

Learning with Biased Complementary Labels

Xiyu Yu¹, Tongliang Liu¹, Mingming Gong^{2,3}, and Dacheng Tao¹

¹ UBTECH Sydney AI Centre, SIT, FEIT, The University of Sydney
{xiyu0300@uni., tongliang.liu@, dacheng.tao@}sydney.edu.au

² Department of Philosophy, Carnegie Mellon University

³ Department of Biomedical Informatics, University of Pittsburgh
mig73@pitt.edu

Abstract. In this paper, we study the classification problem in which we have access to easily obtainable surrogate for true labels, namely complementary labels, which specify classes that observations do **not** belong to. Let Y and \bar{Y} be the true and complementary labels, respectively. We first model the annotation of complementary labels via transition probabilities $P(\bar{Y} = i|Y = j), i \neq j \in \{1, \dots, c\}$, where c is the number of classes. Previous methods implicitly assume that $P(\bar{Y} = i|Y = j), \forall i \neq j$, are identical, which is not true in practice because humans are biased toward their own experience. For example, as shown in Figure 1, if an annotator is more familiar with monkeys than prairie dogs when providing complementary labels for meerkats, she is more likely to employ “monkey” as a complementary label. We therefore reason that the transition probabilities will be different. In this paper, we propose a framework that contributes three main innovations to learning with **biased** complementary labels: (1) It estimates transition probabilities with no bias. (2) It provides a general method to modify traditional loss functions and extends standard deep neural network classifiers to learn with biased complementary labels. (3) It theoretically ensures that the classifier learned with complementary labels converges to the optimal one learned with true labels. Comprehensive experiments on several benchmark datasets validate the superiority of our method to current state-of-the-art methods.

Keywords: Multi-class classification, biased complementary labels, transition matrix, modified loss function

1 Introduction

Large-scale training datasets translate supervised learning from theories and algorithms to practice, especially in deep supervised learning. One major assumption that guarantees this successful translation is that data are accurately labeled. However, collecting true labels for large-scale datasets is often expensive, time-consuming, and sometimes impossible. For this reason, some weak but cheap supervision information has been exploited to boost learning performance. Such supervision includes side information [33], privileged information [29], and weakly supervised information [15] based on semi-supervised data [37,



Fig. 1. A comparison between true labels (top) and complementary labels (bottom).

9, 6], positive and unlabeled data [23], or noisy labeled data [19, 30, 10, 11, 4, 8]. In this paper, we study another weak supervision: the complementary label which specifies a class that an object does **not** belong to. Complementary labels are sometimes easily obtainable, especially when the class set is relatively large. Given an observation in multi-class classification, identifying a class label that is incorrect for the observation is often much easier than identifying the true label.

Complementary labels carry useful information and are widely used in our daily lives: for example, to identify a language we do not know, we may say “not English”; to categorize a new movie without any fighting, we may say “not action”; and to recognize an image of a previous American president, we may say “not Trump”. Ishida et al. [13] then proposed learning from examples with only complementary labels by assuming that a complementary label is uniformly selected from the $c - 1$ classes other than the true label class ($c > 2$). Specifically, they designed an unbiased estimator such that learning with complementary labels was asymptotically consistent with learning with true labels.

Sometimes, annotators provide complementary labels based on both the content of observations and their own experience, leading to the biases in complementary labels. Thus, complementary labels are mostly non-uniformly selected from the remaining $c - 1$ classes, some of which even have no chance of being selected for certain cases. Regarding the bias governed by the observation content, let us take labeling digits 0-9 as an example. Since digit 1 is much more dissimilar to digit 3 than digit 8, the complementary labels of “3” are more likely to be assigned with “1” rather than “8”. Regarding the bias governed by annotators’ experience, taking our example above, we can see that if one is more familiar with monkeys than other animals, she may be more likely to use “monkey” as a complementary label.



Motivated by the cause of biases, we here model the biased procedure of annotating complementary labels via probabilities $P(\bar{Y} = i | Y = j), i \neq j \in \{1, \dots, c\}$. Note that the assumption that a complementary label is uniformly selected from the remaining $c - 1$ classes implies $P(\bar{Y} = i | Y = j) = 1/(c - 1), i \neq j \in \{1, \dots, c\}$. However, in real applications, the probabilities should not be

$1/(c-1)$ and can differ vastly. How to estimate the probabilities is a key problem for learning with complementary labels.

We therefore address the problem of learning with biased complementary labels. For effective learning, we propose to estimate the probabilities $P(\bar{Y} = i|Y = j), i \neq j \in \{1, \dots, c\}$ without biases. Specifically, we prove that given a clear observation \mathbf{x}_j for the j -th class, i.e., the observation satisfying $P(Y = j|\mathbf{x}_j) = 1$, which can be easily identified by the annotator, it holds that $P(\bar{Y} = i|Y = j) = P(\bar{Y} = i|\mathbf{x}_j), i \in \{1, \dots, j-1, j, \dots, c\}$. This implies that probabilities $P(\bar{Y} = i|Y = j), i \neq j \in \{1, \dots, c\}$ can be estimated without biases by learning $P(\bar{Y} = i|\mathbf{x}_j)$ from the examples with complementary labels. To obtain these clear observations, we assume that a small set of easily distinguishable instances (e.g., 10 instances per class) is usually not expensive to obtain.

Given the probabilities $P(\bar{Y}|Y)$, we modify traditional loss functions proposed for learning with true labels so that the modifications can be employed to efficiently learn with biased complementary labels. We also prove that by exploiting examples with complementary labels, the learned classifier converges to the optimal one learned with true labels with a guaranteed rate. Moreover, we also empirically show that the convergence of our method benefits more from the biased setting than from the uniform assumption, meaning that we can use a small training sample to achieve a high performance.

Comprehensive experiments are conducted on benchmark datasets including UCI, MNIST, CIFAR, and Tiny ImageNet, which verifies that our method significantly outperforms the state-of-the-art methods with accuracy gains of over 10%. We also compare the performance of classifiers learned with complementary labels to those learned with true labels. The results show that our method almost attains the performance of learning with true labels in some situations.

2 Related Work

Learning with complementary labels. To the best of our knowledge, Ishida et al. [13] is the first to study learning with complementary labels. They assumed that the transition probabilities are identical and then proposed modifying traditional one-versus-all (OVA) and pairwise-comparison (PC) losses for learning with complementary labels. The main differences between our method and [13] are: (1) Our work is motivated by the fact that annotating complementary labels are often affected by human biases. Thus, we study a different setting in which transition probabilities are different. (2) In [13], modifying OVA and PC losses is naturally suitable for the uniform setting and provides an unbiased estimator for the expected risk of classification with true labels. In this paper, our method can be generalized to many losses such as cross-entropy loss and directly provides an unbiased estimator for the risk minimizer. Due to these differences, [13] often achieves promising performance in the uniform setting while our method achieves good performance in both the uniform and non-uniform setting.

Learning with noisy labels. In the setting of label noise, transition probabilities are introduced to statistically model the generation of noisy labels. In classifica-

tion and transfer learning, methods [21, 18, 32, 35] employ transition probabilities to modify loss functions such that they can be robust to noisy labels. Similar strategies to modify deep neural networks by adding a transition layer have been proposed in [26, 22]. However, this is the first time that this idea is applied to the new problem of learning with biased complementary labels. Different from label noise, here, all diagonal entries of the transition matrix are zeros and the transition matrix sometimes may be not required to be invertible in empirical.

3 Problem Setup

In multi-class classification, let $\mathcal{X} \in \mathbb{R}^d$ be the feature space and $\mathcal{Y} = [c]$ be the label space, where d is the feature space dimension; $[c] = \{1, \dots, c\}$; and $c > 2$ is the number of classes. We assume that variables (X, Y, \bar{Y}) are defined on the space $\mathcal{X} \times \mathcal{Y} \times \mathcal{Y}$ with a joint probability measure $P(X, Y, \bar{Y})$ ($P_{XY\bar{Y}}$ for short).

In practice, true labels are sometimes expensive but complementary labels are cheap. This work thus studies the setting in which we have a large set of training examples with biased complementary labels and a very small set of correctly labeled examples. The latter is only used for estimating transition probabilities. Our aim is to learn the optimal classifier with respect to the examples with true labels by exploiting the examples with complementary labels.

For each example $(\mathbf{x}, y) \in \mathcal{X} \times \mathcal{Y}$, a complementary label \bar{y} is selected from the complement set $\mathcal{Y} \setminus \{y\}$. We assign a probability for each $\bar{y} \in \mathcal{Y} \setminus \{y\}$ to indicate how likely it can be selected, i.e., $P(\bar{Y} = \bar{y} | X = \mathbf{x}, Y = y)$. In this paper, we assume that \bar{Y} is independent of feature X conditioned on true label Y , i.e., $P(\bar{Y} = \bar{y} | X = \mathbf{x}, Y = y) = P(\bar{Y} = \bar{y} | Y = y)$. This assumption considers the bias which depends only on the classes, e.g., if the annotator is not familiar with the features in a specific class, she is likely to assign complementary labels that she is more familiar with. We summarize all the probabilities into a transition matrix $\mathbf{Q} \in \mathbb{R}^{c \times c}$, where $Q_{ij} = P(\bar{Y} = j | Y = i)$ and $Q_{ii} = 0, \forall i, j \in [c]$. Here, Q_{ij} denotes the entry value in the i -th row and j -th column of \mathbf{Q} . Note that transition matrix is also widely exploited in Markov chains [7] and has many applications in machine learning, such as learning with label noise [21, 26, 22].

If complementary labels are uniformly selected from the complement set, then $\forall i, j \in [c]$ and $i \neq j$, $Q_{ij} = \frac{1}{c-1}$. Previous work [13] has proven that the optimal classifier can be found under the uniform assumption. Sometimes, this is not true in practice due to human biases. Therefore, we focus on situations in which $Q_{ij}, \forall i \neq j$, are different. We mainly study the following problems: how to modify loss functions such that the classifier learned with these biased complementary labels can converge to the optimal one learned with true labels; the speed of the convergence; and how to estimate transition probabilities.

4 Methodology

In this section, we study how to learn with biased complementary labels. We first review how to learn optimal classifiers from examples with true labels. Then,

we modify loss functions for complementary labels and propose a deep learning based model accordingly. Lastly, we theoretically prove that the classifier learned by our method is consistent with the optimal classifier learned with true labels.

4.1 Learning with True Labels

The aim of multi-class classification is to learn a classifier $f(\mathbf{x})$ that predicts a label y for a given observation \mathbf{x} . Typically, the classifier is of the following form:

$$f(X) = \arg \max_{i \in [c]} g_i(X), \quad (1)$$

where $\mathbf{g} : \mathcal{X} \rightarrow \mathbb{R}^c$ and $g_i(X)$ is the estimate of $P(Y = i|X)$.

Various loss functions $\ell(f(X), Y)$ have been proposed to measure the risk of predicting $f(X)$ for Y [1]. Formally, the expected risk is defined as.

$$R(f) = \mathbb{E}_{(X,Y) \sim P_{XY}} [\ell(f(X), Y)]. \quad (2)$$

The optimal classifier is the one that minimizes the expected risk; that is,

$$f^* = \arg \min_{f \in \mathcal{F}} R(f), \quad (3)$$

where \mathcal{F} is the space of f .

However, the distribution P_{XY} is usually unknown. We then approximate $R(f)$ by using its empirical counterpart: $R_n(f) = \frac{1}{n} \sum_{i=1}^n \ell(f(\mathbf{x}_i), y_i)$, where $\{(\mathbf{x}_i, y_i)\}_{1 \leq i \leq n}$ are i.i.d. examples drawn according to P_{XY} .

Similarly, the optimal classifier is approximated by $f_n = \arg \min_{f \in \mathcal{F}} R_n(f)$.

4.2 Learning with Complementary Labels

True labels, especially for large-scale datasets, are often laborious and expensive to obtain. We thus study an easily obtainable surrogate; that is, complementary labels. However, if we still use traditional loss functions ℓ when learning with these complementary labels, similar to Eq.(1), we can only learn a mapping $\mathbf{q} : \mathcal{X} \rightarrow \mathbf{R}^c$ that tries to predict conditional probabilities $P(\bar{Y}|X)$ and the corresponding classifier that predicts a \bar{y} for a given observation \mathbf{x} .

Therefore, we need to modify these loss functions such that the classifier learned with biased complementary labels can converge to the optimal one learned with true labels. Specifically, let $\bar{\ell}$ be the modified loss function. Then, the expected and empirical risks with respect to complementary labels are defined as $\bar{R}(f) = \mathbb{E}_{(X,\bar{Y}) \sim P_{X\bar{Y}}} [\bar{\ell}(f(X), \bar{Y})]$ and $\bar{R}_n(f) = \frac{1}{n} \sum_{i=1}^n \bar{\ell}(f(\mathbf{x}_i), \bar{y}_i)$, respectively. Here, $\{(\mathbf{x}_i, \bar{y}_i)\}_{1 \leq i \leq n}$ are examples with complementary labels.

Denote f^* and \bar{f}_n as the optimal solution obtained by minimizing $\bar{R}(f)$ and $\bar{R}_n(f)$, respectively. They are $f^* = \arg \min_{f \in \mathcal{F}} \bar{R}(f)$ and $\bar{f}_n = \arg \min_{f \in \mathcal{F}} \bar{R}_n(f)$.

We hope that the modified loss function $\bar{\ell}$ can ensure that $\bar{f}_n \xrightarrow{n} f^*$, which implies that by learning with complementary labels, the classifier we obtain can also approach to the optimal one defined in (3).




Fig. 2. An overview of our method. We modify the deep neural network by adding a layer that multiplies the output of the softmax function by \mathbf{Q}^\top .

Recall that in transition matrix \mathbf{Q} , $Q_{ij} = P(\bar{Y} = j|Y = i)$ and $Q_{ii} = P(\bar{Y} = i|Y = i) = 0, \forall i \in [c]$. We observe that $P(Y|X)$ can be transferred to $P(\bar{Y}|X)$ by using the transition matrix \mathbf{Q} ; that is, $\forall j \in [c]$,

$$\begin{aligned} P(\bar{Y} = j|X) &= \sum_{i \neq j} P(\bar{Y} = j, Y = i|X) \\ &= \sum_{i \neq j} P(\bar{Y} = j|Y = i, X)P(Y = i|X) \\ &= \sum_{i \neq j} P(\bar{Y} = j|Y = i)P(Y = i|X). \end{aligned} \quad (4)$$

Intuitively, if $q_i(X)$ tries to predict the probability $P(\bar{Y} = i|X)$, $\forall i \in [c]$, then $\mathbf{Q}^{-\top} \mathbf{q}$ can predict the probability $P(Y|X)$. To enable end-to-end learning rather than transferring after training, we let

$$\mathbf{q}(X) = \mathbf{Q}^\top \mathbf{g}(X), \quad (5)$$

where $\mathbf{g}(X)$ is now an intermediate output, and $f(X) = \arg \max_{i \in [c]} g_i(X)$.

Then, the modified loss function $\bar{\ell}$ is

$$\bar{\ell}(f(X), \bar{Y}) = \ell(\mathbf{q}(X), \bar{Y}). \quad (6)$$

In this way, if we can learn an optimal \mathbf{q}^* such that $q_i^*(X) = P(\bar{Y} = i|X), \forall i \in [c]$, meanwhile, we can also find the optimal \mathbf{g}^* and the classifier f^* .

This loss modification method can be easily applied to deep learning. As shown in Figure 2, we achieve this simply by adding a linear layer to the deep neural network. This layer outputs $\mathbf{q}(X)$ by multiplying the output of the softmax function (i.e., $\mathbf{g}(X)$) by the transposed transition matrix \mathbf{Q}^\top . With sufficient training examples with complementary labels, this deep neural network often simultaneously learns good classifiers for both (X, \bar{Y}) and (X, Y) .

Note that, in our modification, the forward process does not need to compute $\mathbf{Q}^{-\top}$. Even though the subsequent analysis for identification requires the transition matrix to be invertible, sometimes, we may have no such requirement in practice. We also show an example in the Supplementary Material that even with singular transition matrices, high classification performance can also be achieved if no column of \mathbf{Q} is all-zero.

5 Identification of the Optimal Classifier

In this section, we aim to prove that the proposed loss modification method ensures the identifiability of the optimal classifier under a reasonable assumption:

Assumption 1 *By minimizing the expected risk $R(f)$, the optimal mapping \mathbf{g}^* satisfies $g_i^*(X) = P(Y = i|X), \forall i \in [c]$.*

Based on Assumption 1, we can prove that $\bar{f}^* = f^*$ by the following theorem:

Theorem 1 *Suppose that \mathbf{Q} is invertible and Assumption 1 is satisfied, then the minimizer \bar{f}^* of $\bar{R}(f)$ is also the minimizer f^* of $R(f)$; that is, $\bar{f}^* = f^*$.*

Please find the detailed proof in the Supplementary Material. Given sufficient training data with complementary labels, \bar{f}_n can converge to \bar{f}^* , which can be proved in the next section. According to Theorem 1, this also implies that \bar{f}_n also converges to the optimal classifier f^* .

Examples of Loss Functions. The proof of Theorem 1 relies on Assumption 1. However, for many loss functions, Assumption 1 can be provably satisfied. Here, we take the cross-entropy loss as an example to demonstrate this fact. The cross-entropy loss is widely used in deep supervised learning and is defined as

$$\ell(f(X), Y) = - \sum_{i=1}^c 1(Y = i) \log(g_i(X)), \quad (7)$$

where $1(\cdot)$ is an indicator function; that is, if the input statement is true, it outputs 1; otherwise, 0. For the cross-entropy loss, we have the following lemma:

Lemma 1 *Suppose ℓ is the cross-entropy loss and $\mathbf{g}(X) \in \Delta^{c-1}$, where Δ^{c-1} refers to a standard simplex in \mathbb{R}^c ; that is, $\forall \mathbf{x} \in \Delta^{c-1}, x_i \geq 0, \forall i \in [c]$ and $\sum_{i=1}^c x_i = 1$. By minimizing the expected risk $R(f)$, we have $g_i^*(X) = P(Y = i|X), \forall i \in [c]$.*

Please see the detailed proof in the Supplementary Material. In fact, losses such as square-error loss $\ell(f(X), Y) = \sum_{j=1}^c (1(Y = j) - g_j(X))^2$, also satisfy Assumption 1. The readers can prove it themselves using similar strategy. Combined with Theorem 1, we can see, by applying the proposed method to loss functions such as cross-entropy loss, we can prove that the optimal classifier f^* can be found even when learning with biased complementary labels.

6 Convergence Analysis

In this section, we show an upper bound for the estimation error of our method. This upper bound illustrates a convergence rate for the classifier learned with complementary labels to the optimal one learned with true labels. Moreover, with the derived bound, we can clearly see that the estimation error could further benefit from the setting of biased complementary labels under mild conditions.

Since $\bar{f}^* = f^*$, we have $|\bar{f}_n - f^*| = |\bar{f}_n - \bar{f}^*|$. We will upper bound the error $|\bar{f}_n - \bar{f}^*|$ via upper bounding $\bar{R}(\bar{f}_n) - \bar{R}(\bar{f}^*)$; that is, when $\bar{R}(\bar{f}_n) - \bar{R}(\bar{f}^*) \rightarrow 0$, $|\bar{f}_n - \bar{f}^*| \rightarrow 0$. Specifically, it has been proven that

$$\begin{aligned} \bar{R}(\bar{f}_n) - \bar{R}(\bar{f}^*) &= \bar{R}(\bar{f}_n) - \bar{R}_n(\bar{f}_n) + \bar{R}_n(\bar{f}_n) - \bar{R}_n(\bar{f}^*) + \bar{R}_n(\bar{f}^*) - \bar{R}(\bar{f}^*) \\ &\leq \bar{R}(\bar{f}_n) - \bar{R}_n(\bar{f}_n) + \bar{R}_n(\bar{f}^*) - \bar{R}(\bar{f}^*) \\ &\leq 2 \sup_{f \in \mathcal{F}} |\bar{R}(f) - \bar{R}_n(f)|, \end{aligned} \quad (8)$$

where the first inequality holds because $\bar{R}_n(\bar{f}_n) - \bar{R}_n(\bar{f}^*) \leq 0$ and the error in the last line is called the generalization error.

Let $(X_1, \bar{Y}_1), \dots, (X_n, \bar{Y}_n)$ be independent variables. By employing the concentration inequality [3], the generalization error can be upper bounded by using the method of Rademacher complexity [2].

Theorem 2 ([2]) *Let the loss function be upper bounded by M . Then, for any $\delta > 0$, with the probability $1 - \delta$, we have*

$$\sup_{f \in \mathcal{F}} |\bar{R}(f) - \bar{R}_n(f)| \leq 2\mathfrak{R}_n(\bar{\ell} \circ \mathcal{F}) + M \sqrt{\frac{\log 1/\delta}{2n}}, \quad (9)$$

where $\mathfrak{R}_n(\bar{\ell} \circ \mathcal{F}) = \mathbb{E} [\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \sigma_i \bar{\ell}(f(X_i), \bar{Y}_i)]$ is the Rademacher complexity; $\{\sigma_1, \dots, \sigma_n\}$ are Rademacher variables uniformly distributed from $\{-1, 1\}$.

Before upper bounding $\mathfrak{R}_n(\bar{\ell} \circ \mathcal{F})$, we need to discuss the specific form of the employed loss function $\bar{\ell}$. By exploiting the well-defined binary loss functions, one-versus-all and pairwise-comparison loss functions [36] have been proposed for multi-class learning. In this section, we discuss the modified loss function $\bar{\ell}$ defined by Eqs. (6) and (7), which can be rewritten as,

$$\begin{aligned} \bar{\ell}(f(X), \bar{Y}) &= - \sum_{i=1}^c 1(\bar{Y} = i) \log((\mathbf{Q}^\top \mathbf{g})_i(X)) \\ &= - \sum_{i=1}^c 1(\bar{Y} = i) \log \left(\frac{\sum_{j=1}^c Q_{ji} \exp(h_j(X))}{\sum_{k=1}^c \exp(h_k(X))} \right), \end{aligned} \quad (10)$$

where $(\mathbf{Q}^\top \mathbf{g})_i$ denotes the i -th entry of $\mathbf{Q}^\top \mathbf{g}$; $\mathbf{h} : \mathcal{X} \rightarrow \mathbb{R}^c$, $h_i(X) \in \mathcal{H}$, $\forall i \in [c]$; and $g_i(X) = \frac{\exp(h_i(X))}{\sum_{k=1}^c \exp(h_k(X))}$.

Usually, the convergence rates of generalization bounds of multi-class learning are at most $O(c^2/\sqrt{n})$ with respect to c and n [13, 20]. To reduce the dependence on c of our derived convergence rate, we rewrite $\bar{R}(f)$ as follows:

$$\begin{aligned} \bar{R}(f) &= \int_X \sum_{i=1}^c P(\bar{Y} = i) P(X|\bar{Y} = i) \bar{\ell}(f(X), \bar{Y} = i) dX \\ &= \sum_{i=1}^c P(\bar{Y} = i) \int_X P(X|\bar{Y} = i) \bar{\ell}(f(X), \bar{Y} = i) dX \\ &= \sum_{i=1}^c \bar{\pi}_i \bar{R}_i(f), \end{aligned} \quad (11)$$

where $\bar{R}_i(f) = \mathbb{E}_{X \sim P(X|\bar{Y}=i)} \bar{\ell}(f(X), \bar{Y} = i)$ and $\bar{\pi}_i = P(\bar{Y} = i)$.

Similar to Theorem 2, we have the following theorem.

Theorem 3 *Suppose $\bar{\pi}_i = P(\bar{Y} = i)$ is given. Let the loss function be upper bounded by M . Then, for any $\delta > 0$, with the probability $1 - c\delta$, we have*

$$\begin{aligned} \bar{R}(\bar{f}_n) - \bar{R}(\bar{f}^*) &\leq 2 \sup_{f \in \mathcal{F}} |\bar{R}(f) - \bar{R}_n(f)| \\ &\leq 2 \sum_{i=1}^c \bar{\pi}_i \sup_{f \in \mathcal{F}} |\bar{R}_i(f) - \bar{R}_{i,n_i}(f)| \\ &\leq 2 \sum_{i=1}^c \bar{\pi}_i \left(2\mathfrak{R}_{n_i}(\bar{\ell} \circ \mathcal{F}) + M \sqrt{\frac{\log 1/\delta}{2n_i}} \right) \\ &= \sum_{i=1}^c \left(4\bar{\pi}_i \mathfrak{R}_{n_i}(\bar{\ell} \circ \mathcal{F}) + 2\bar{\pi}_i M \sqrt{\frac{\log 1/\delta}{2n_i}} \right), \end{aligned} \quad (12)$$

where $\mathfrak{R}_{n_i}(\bar{\ell} \circ \mathcal{F}) = \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n_i} \sum_{j=1}^{n_i} \sigma_j \bar{\ell}(f(X_j), \bar{Y}_j = i) \right]$ and $\bar{R}_{i,n_i}(f)$ is the empirical counterpart of $\bar{R}_i(f)$, and $n_i, i \in [c]$, represents the numbers of X whose complementary labels are $\bar{Y} = i$.

Due to the fact that $\bar{\ell}$ is actually defined with respect to \mathbf{h} rather than f , we would like to bound the error by the Rademacher complexity of \mathcal{H} . We observe that the relationship between $\mathfrak{R}_{n_i}(\bar{\ell} \circ \mathcal{F})$ and $\mathfrak{R}_{n_i}(\mathcal{H})$ is:

Lemma 2 *Let $\bar{\ell}(f(X), \bar{Y} = i) = -\log \left(\frac{\sum_{k=1}^c Q_{ki} \exp(h_k(X))}{\sum_{k=1}^c \exp(h_k(X))} \right)$ and suppose that $h_i(X) \in \mathcal{H}, \forall i \in [c]$, we have $\mathfrak{R}_{n_i}(\bar{\ell} \circ \mathcal{F}) \leq c\mathfrak{R}_{n_i}(\mathcal{H})$.*

The detailed proof can be found in the Supplementary Material. Combine Theorem 3 and Lemma 2, we have the final result:

Corollary 1 *Suppose $\bar{\pi}_i = P(\bar{Y} = i)$ is given. Let the loss function be upper bounded by M . Then, for any $\delta > 0$, with the probability $1 - c\delta$, we have*

$$\bar{R}(\bar{f}_n) - \bar{R}(\bar{f}^*) \leq \sum_{i=1}^c \left(4c\bar{\pi}_i \mathfrak{R}_{n_i}(\mathcal{H}) + 2\bar{\pi}_i M \sqrt{\frac{\log 1/\delta}{2n_i}} \right). \quad (13)$$

In current state-of-the-art methods [13], the convergence rate of $\mathfrak{R}_n(\bar{\ell} \circ \mathcal{F})$ is of order $O(c^2/\sqrt{n})$ with respect to c and n while our derived bound $\sum_{i=1}^c 4c\bar{\pi}_i \mathfrak{R}_{n_i}(\mathcal{H})$ is of order $\max_{i \in [c]} O(c/\sqrt{n_i})$. Since our error bound depends on n_i , the bound would be loose if n_i (or $\bar{\pi}_i$) is small. However, if $\bar{\pi}_i$ is balanced and n_i is about n/c , our convergence rate is of order $O(c\sqrt{c}/\sqrt{n})$, which is smaller than the error bounds provided by previous methods if c is very large.

Remark. Theorem 3 and Corollary 1 aim to provide the proof of uniform convergence for general losses and show how the convergence rate can benefit

from the biased setting under mild conditions. Thus, assuming the loss is upper-bounded is reasonable for many loss functions such as the square-error loss. If the readers would like to derive specific error bound for the cross-entropy loss, strategies in [31] can be employed. If we assume that the transition matrix \mathbf{Q} is invertible, we can derive similar results as those in Lemma 1-3 [31] for the modified loss function, which can be finally deployed to derive generalization error bound similar to Corollary 1.

7 Estimating \mathbf{Q}

In the aforementioned method, transition matrix \mathbf{Q} is assumed to be known, which is not true. Here, we thus provide an efficient method to estimate \mathbf{Q} .

When learning with complementary labels, we completely lose the information of true labels. Without any auxiliary information, it is impossible to estimate the transition matrix which is associated with the class priors of true labels. On the other hand, although it is costly to annotate a very large-scale dataset, a small set of easily distinguishable observations are assumed to be available in practice. This assumption is also widely used in estimating transition probabilities in label noise problem [28] and class priors in semi-supervised learning [34]. Therefore, in order to estimate \mathbf{Q} , we manually assign true labels to 5 or 10 observations in each class. Since these selected observations are often easy to classify, we further assume that they satisfy the anchor set condition [18]:

Assumption 2 (Anchor Set Condition) *For each class y , there exists an anchor set $\mathcal{S}_{\mathbf{x}|y} \subset \mathcal{X}$ such that $P(Y = y|X = \mathbf{x}) = 1$ and $P(Y = y'|X = \mathbf{x}) = 0$, $\forall y' \in \mathcal{Y} \setminus \{y\}$, $\mathbf{x} \in \mathcal{S}_{\mathbf{x}|y}$.*

Here, $\mathcal{S}_{\mathbf{x}|y}$ is a subset of features in class y . Given several observations in $\mathcal{S}_{\mathbf{x}|y}$, $y \in [c]$, we are ready to estimate the transition matrix \mathbf{Q} . According to Eq. (4),

$$P(\bar{Y} = \bar{y}|X) = \sum_{y' \neq \bar{y}} P(\bar{Y} = \bar{y}|Y = y')P(Y = y'|X). \quad (14)$$

Suppose $\mathbf{x} \in \mathcal{S}_{\mathbf{x}|y}$, then $P(Y = y|X = \mathbf{x}) = 1$ and $P(Y = y'|X = \mathbf{x}) = 0$, $\forall y' \in \mathcal{Y} \setminus \{y\}$. We have

$$P(\bar{Y} = \bar{y}|X = \mathbf{x}) = P(\bar{Y} = \bar{y}|Y = y). \quad (15)$$

That is, the probabilities in \mathbf{Q} can be obtained via $P(\bar{Y}|X)$ given the observations in the anchor set of each class. Thus, we need only to estimate this conditional probability, which has been proven to be achievable in Lemma 1. In this paper, with the training sample $\{(x_i, \bar{y}_i)\}_{1 \leq i \leq n}$, we estimate $P(\bar{Y}|X)$ by training a deep neural network with the softmax function and cross-entropy loss. After obtaining these conditional probabilities, each probability $P(\bar{Y} = \bar{y}|Y = y)$ in the transition matrix can be estimated by averaging the conditional probabilities $P(\bar{Y} = \bar{y}|X = \mathbf{x})$ on the anchor data \mathbf{x} in class y .

8 Experiments

We evaluate our algorithm on several benchmark datasets including the UCI datasets, USPS, MNIST [16], CIFAR10, CIFAR100 [14], and Tiny ImageNet⁴. All our experiments are trained on neural networks. For USPS and UCI datasets, we employ a one-hidden-layer neural network ($d-3-c$) [13]. For MNIST, LeNet-5 [17] is deployed, and ResNet [12] is exploited for the other datasets. All models are implemented in PyTorch⁵.

UCI and USPS. We first evaluate our method on USPS and six UCI datasets: WAVEFORM1, WAVEFORM2, SATIMAGE, PENDIGITS, DRIVE, and LETTER, downloaded from the UCI machine learning repository. We apply the same strategies of annotating complementary labels, standardization, validation, and optimization with those in [13]. The learning rate is chosen from $\{10^{-5}, \dots, 10^{-1}\}$, weight decay from $\{10^{-7}, 10^{-4}, 10^{-1}\}$, batch size 100.

For fair comparison in these experiments, we assume the transition probabilities are identical and known as prior. Thus, no examples with true labels are required here. All results are shown in Table 1. Our loss modification (LM) method is compared to a partial label (PL) method [5], a multi-label (ML) method [24], and “PC/S” (the pairwise-comparison formulation with sigmoid loss), which achieved the best performance in [12]. We can see, “PC/S” achieves very good performances. The relatively higher performance of our method may be due to that our method provides an unbiased estimator for risk minimizer.

MNIST. MNIST is a handwritten digit dataset including 60,000 training images and 10,000 test images from 10 classes. To evaluate the effectiveness of our method, we consider the following three settings: (1) for each image in class y , the complementary label is uniformly selected from $\mathcal{Y} \setminus \{y\}$ (“**uniform**”); (2) the complementary label is non-uniformly selected, but each label in $\mathcal{Y} \setminus \{y\}$ has non-zero probability to be selected (“**without0**”); (3) the complementary label is non-uniformly selected from a small subset of $\mathcal{Y} \setminus \{y\}$ (“**with0**”).

To generate complementary labels, we first give the probability of each complementary label to be selected. In the “uniform” setting, $P(\bar{Y} = j|Y = i) = \frac{1}{9}, \forall i \neq j$. In the “without0” setting, for each class y , we first randomly split $\mathcal{Y} \setminus \{y\}$ to three subsets, each containing three elements. Then, for each complementary label in these three subsets, the probabilities are set to $\frac{0.6}{3}, \frac{0.3}{3}$, and $\frac{0.1}{3}$, respectively. In the “with0” setting, for each class y , we first randomly selected three labels in $\mathcal{Y} \setminus \{y\}$, and then randomly assign them with three probabilities whose summation is 1. After \mathbf{Q} is given, we assign complementary label to each image based on these probabilities. Finally, we randomly set aside 10% of training data as validation set.

In all experiments, the learning rate is fixed to $1e-4$; batch size 128; weight decay $1e-4$; maximum iterations 60,000; and stochastic gradient descend (SGD) with momentum $\gamma = 0.9$ [27] is applied to optimize deep models. Note that, as shown in [13] and previous experiments, [13] and our method have surpassed

⁴ The dataset is available at <http://cs231n.stanford.edu/tiny-imagenet-200.zip>

⁵ <http://pytorch.org>

Table 1. Classification accuracy on USPS and UCI datasets: the means and standard deviations of classification accuracy over 20 trials in percentages are reported. “#train” is the number of training and validation examples in each class. “#test” is the number of test examples in each class.

| Dataset | c | d | #train | #test | PC/S | PL | ML | LM (ours) |
|-----------|---------|-----|--------|-------|-------------------|-------------------|------------------|-------------------|
| WAVEFORM1 | 1 ~ 3 | 21 | 1226 | 398 | 85.8 (0.5) | 85.7 (0.9) | 79.3 (4.8) | 85.1 (0.6) |
| WAVEFORM2 | 1 ~ 3 | 40 | 1227 | 408 | 84.7 (1.3) | 84.6 (0.8) | 74.9 (5.2) | 85.5 (1.1) |
| SATIMAGE | 1 ~ 7 | 36 | 415 | 211 | 68.7 (5.4) | 60.7 (3.7) | 33.6 (6.2) | 69.3 (3.6) |
| PENDIGITS | 1 ~ 5 | 16 | 719 | 336 | 87.0 (2.9) | 76.2 (3.3) | 44.7 (9.6) | 92.7 (3.7) |
| | 6 ~ 10 | | 719 | 335 | 78.4 (4.6) | 71.1 (3.3) | 38.4 (9.6) | 85.8 (1.3) |
| | even # | | 719 | 336 | 90.8 (2.4) | 76.8 (1.6) | 43.8 (5.1) | 90.0 (1.0) |
| | odd # | | 719 | 335 | 76.0 (5.4) | 67.4 (2.6) | 40.2 (8.0) | 86.5 (0.5) |
| | 1 ~ 10 | | 719 | 335 | 38.0 (4.3) | 33.2 (3.8) | 16.1 (4.6) | 62.8 (5.6) |
| DRIVE | 1 ~ 5 | 48 | 3955 | 1326 | 89.1 (4.0) | 77.7 (1.5) | 31.1 (3.5) | 93.3 (4.6) |
| | 6 ~ 10 | | 3923 | 1313 | 88.8 (1.8) | 78.5 (2.6) | 30.4 (7.2) | 92.8 (0.9) |
| | even # | | 3925 | 1283 | 81.8 (3.4) | 63.9 (1.8) | 29.7 (6.3) | 84.3 (0.7) |
| | odd # | | 3939 | 1278 | 85.4 (4.2) | 74.9 (3.2) | 27.6 (5.8) | 85.9 (2.1) |
| | 1 ~ 10 | | 3925 | 1269 | 40.8 (4.3) | 32.0 (4.1) | 12.7 (3.1) | 75.1 (3.2) |
| LETTER | 1 ~ 5 | 16 | 565 | 171 | 79.7 (5.4) | 75.1 (4.4) | 28.3 (10.4) | 84.3 (1.5) |
| | 6 ~ 10 | | 550 | 178 | 76.2 (6.2) | 66.8 (2.5) | 34.0 (6.9) | 84.4 (1.0) |
| | 11 ~ 15 | | 556 | 177 | 78.3 (4.1) | 67.4 (3.4) | 28.6 (5.0) | 88.3 (1.9) |
| | 16 ~ 20 | | 550 | 184 | 77.2 (3.2) | 68.4 (2.1) | 32.7 (6.4) | 85.2 (0.7) |
| | 21 ~ 25 | | 585 | 167 | 80.4 (4.2) | 75.1 (1.9) | 32.0 (5.7) | 82.5 (1.0) |
| | 1 ~ 25 | | 550 | 167 | 5.1 (2.1) | 5.0 (1.0) | 5.2 (1.1) | 7.0 (3.6) |
| USPS | 1 ~ 5 | 256 | 652 | 166 | 79.1 (3.1) | 70.3 (3.2) | 44.4 (8.9) | 86.4 (4.5) |
| | 6 ~ 10 | | 542 | 147 | 69.5 (6.5) | 66.1 (2.4) | 37.3 (8.8) | 88.1 (2.7) |
| | even # | | 556 | 147 | 67.4 (5.4) | 66.2 (2.3) | 35.7 (6.6) | 79.5 (5.4) |
| | odd # | | 542 | 147 | 77.5 (4.5) | 69.3 (3.1) | 36.6 (7.5) | 86.3 (3.1) |
| | 1 ~ 10 | | 542 | 127 | 30.7 (4.4) | 26.0 (3.5) | 13.3 (5.4) | 37.2 (5.4) |

Table 2. Classification accuracy on MNIST: the means and standard deviations of classification accuracy over five trials in percentages are reported. “TL” denotes the result of learning with true labels. “LM/T” and “LM/E” refer to our method with the true \mathbf{Q} and the estimated one, respectively.

| Method | Uniform | Without0 | With0 |
|--------|--------------|--------------|--------------|
| TL | 99.12 | 99.12 | 99.12 |
| PC/S | 86.59 ± 3.99 | 76.03 ± 3.34 | 29.12 ± 1.94 |
| LM/T | 97.18 ± 0.45 | 97.65 ± 0.15 | 98.63 ± 0.05 |
| LM/E | 96.33 ± 0.31 | 97.04 ± 0.31 | 98.61 ± 0.05 |

baseline methods such as PL and ML. In the following experiments, we will not again make comparisons with these baselines.

The results are shown in Table 2. The means and standard deviations of classification accuracy over five trials are reported. Note that the digit data features are not too entangled, making it easier to learn a good classifier. However, we can still see the differences in the performance caused by the change of settings for annotating complementary labels. According to the results shown in Table 2, “PC/S” [13] works relatively well under the uniform assumption but the accuracy deteriorates in other settings. Our method performs well in all settings. It can also be seen that due to the accurate estimates of these probabilities, “LM/E” with the estimated transition matrix \mathbf{Q} is competitive with “LM/T” which exploits the true one.

Table 3. Classification accuracy on CIFAR10: the means and standard deviations of classification accuracy over five trials in percentages are reported. “TL” denotes the result of learning with true labels. “LM/T” and “LM/E” refer to our method with the true \mathbf{Q} and the estimated one, respectively.

| Method | Uniform | Without0 | With0 |
|--------|------------------|------------------|------------------|
| TL | 90.78 | 90.78 | 90.78 |
| PC/S | 41.19 ± 0.04 | 42.97 ± 3.00 | 18.12 ± 1.45 |
| LM/T | 73.38 ± 1.06 | 78.80 ± 0.45 | 85.32 ± 1.11 |
| LM/E | 42.96 ± 0.76 | 70.56 ± 0.34 | 84.60 ± 0.14 |

CIFAR10. We evaluate our method on the CIFAR10 dataset under the aforementioned three settings. CIFAR10 has totally 10 classes of tiny images, which includes 50,000 training images and 10,000 test images. We leave out 10% of the training data as validation set. In these experiments, ResNet-18 [12] is deployed. We start with an initial learning rate 0.01 and divide it by 10 after 40 and 80 epochs. The weight decay is set to $5e - 4$, and other settings are the same as those for MNIST. Early stopping is applied to avoid overfitting.

We apply the same process as MNIST to generate complementary labels. The results in Table 3 verify the effectiveness of our method. “PC/S” achieves promising performance when complementary labels are uniformly selected, and our method outperforms “PC/S” in other settings. In the “uniform” setting, $P(\bar{Y}|X)$ is not well estimated. As a result, the transition matrix is also poorly estimated. “LM/E” thus performs relatively badly.

The results of our method under the “uniform” and “without0” settings (shown in Table 3) are usually worse than that of “with0”. For a certain amount of training images, the empirical results show that in the “uniform” and “without0” setting, the proposed method converges at a slower rate than in the “with0” setting. This phenomenon may be caused by the fact that the uncertainty involved with the transition procedure in the “with0” setting is less than that in “uniform” and “without0” settings, making it easier to learn in the former setting. This phenomenon also indicates that, for images in each class, annotators need not to assign all possible complementary labels, but can provide the labels following the criteria, i.e., each label in the label space should be assigned as complementary label for images in at least one class. In this way, we can reduce the number of training examples to achieve high performance.

CIFAR100. CIFAR100 also presents a collection of tiny images including 50,000 training images and 10,000 test images. But CIFAR100 has totally 100 classes, each with only 500 training images. Due to the label space being very large and the number of training data being limited, in both “uniform” and “without0” settings, few training data are assigned as j for images in each class i , $\forall i \neq j$. Both the proposed method and “PC/S” cannot converge. Here, we only conduct the experiments under the “with0” setting. To generate complementary labels, for each class y , we randomly selected 5 labels from $\mathcal{Y} \setminus \{y\}$, and assign them with non-zero probabilities. Others have no chance to be selected.

Table 4. Classification accuracy on CIFAR100 and Tiny ImageNet under the setting “with0”: the means and standard deviations of classification accuracy over five trials in percentages are reported. “TL” denotes the result of learning with true labels. “LM/T” and “LM/E” refer to our method with the true \mathbf{Q} and the estimated one, respectively.

| Method | CIFAR100 | Tiny ImageNet |
|--------|------------------|------------------|
| TL | 69.55 | 63.26 |
| PC/S | 8.95 ± 1.47 | N/A |
| LM/T | 62.84 ± 0.30 | 52.71 ± 0.71 |
| LM/E | 60.27 ± 0.28 | 49.70 ± 0.78 |

In these experiments, ResNet-34 is deployed. Other experimental settings are the same with those in CIFAR10. Results are shown in the second column of Table 4. “PC/S” can hardly obtain a good classifier, but our method achieves high accuracies that are comparable to learning with true labels.

Tiny ImageNet. Tiny ImageNet represents 200 classes with 500 images in each class from ImageNet dataset [25]. Images are cropped to 64×64 . Detailed information is lost during the down-sampling process, making it more difficult to learn. ResNet-18 for ImageNet [12] is deployed. Instead of using the original first convolutional layer with a 7×7 kernel and the subsequent max pooling layer, we replace them with a convolutional layer with a 3×3 kernel, stride=1, and no padding. The initial learning rate is 0.1, divided by 10 after 20,000 and 40,000 iterations. The batch size is 256 and weight decay is $5e - 4$. Other settings are the same as CIFAR100. The experimental results are shown in the third column of Table 4. We also only test our method under the setting “with0”. “PC/S” cannot converge here, but our method still achieves promising performance.

9 Conclusion

We address the problem of learning with biased complementary labels. Specifically, we consider the setting that the transition probabilities $P(\bar{Y} = j|Y = i), \forall i \neq j$ vary and most of them are zeros. We devise an effective method to estimate the transition matrix given a small amount of data in the anchor set. Based on the transition matrix, we proposed to modify traditional loss functions such that learning with complementary labels can theoretically converge to the optimal classifier learned from examples with true labels. Comprehensive experiments on a wide range of datasets verify that the proposed method is superior to the current state-of-the-art methods.

Acknowledgement. This work was supported by Australian Research Council Projects FL-170100117, DP-180103424, and LP-150100671. This work was partially supported by SAP SE and research grant from Pfizer titled “Developing Statistical Method to Jointly Model Genotype and High Dimensional Imaging Endophenotype”. We are also grateful for the computational resources provided by Pittsburgh Super Computing grant number TG-ASC170024.

References

1. Bartlett, P.L., Jordan, M.I., McAuliffe, J.D.: Convexity, classification, and risk bounds. *Journal of the American Statistical Association* **101**(473), 138–156 (2006)
2. Bartlett, P.L., Mendelson, S.: Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research* **3**(Nov), 463–482 (2002)
3. Boucheron, S., Lugosi, G., Massart, P.: Concentration inequalities: A nonasymptotic theory of independence. Oxford university press (2013)
4. Cheng, J., Liu, T., Ramamohanarao, K., Tao, D.: Learning with bounded instance- and label-dependent label noise. arXiv preprint arXiv:1709.03768 (2017)
5. Cour, T., Sapp, B., Taskar, B.: Learning from partial labels. *Journal of Machine Learning Research* **12**(May), 1501–1536 (2011)
6. Ehsan Abbasnejad, M., Dick, A., van den Hengel, A.: Infinite variational autoencoder for semi-supervised learning. In: CVPR (July 2017)
7. Gagniuc, P.A.: Markov Chains: From Theory to Implementation and Experimentation. John Wiley & Sons (2017)
8. Gong, C., Zhang, H., Yang, J., Tao, D.: Learning with inadequate and incorrect supervision. In: ICDM. pp. 889–894. IEEE (2017)
9. Haeusser, P., Mordvintsev, A., Cremers, D.: Learning by association: A versatile semi-supervised training method for neural networks. In: CVPR (2017)
10. Han, B., Tsang, I.W., Chen, L., Celina, P.Y., Fung, S.F.: Progressive stochastic learning for noisy labels. *IEEE Transactions on Neural Networks and Learning Systems* (99), 1–13 (2018)
11. Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I., Sugiyama, M.: Co-teaching: Robust training deep neural networks with extremely noisy labels. arXiv preprint arXiv:1804.06872 (2018)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
13. Ishida, T., Niu, G., Sugiyama, M.: Learning from complementary labels. In: NIPS (2017)
14. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images (2009)
15. Law, M.T., Yu, Y., Urtasun, R., Zemel, R.S., Xing, E.P.: Efficient multiple instance metric learning using weakly supervised data. In: CVPR (July 2017)
16. LeCun, Y., Corinna, C., Christopher, B.J.: The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>
17. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
18. Liu, T., Tao, D.: Classification with noisy labels by importance reweighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(3), 447–461 (March 2016)
19. Misra, I., Lawrence Zitnick, C., Mitchell, M., Girshick, R.: Seeing through the Human reporting bias: Visual classifiers from noisy Human-centric labels. In: CVPR. pp. 2930–2939 (2016)
20. Mohri, M., Rostamizadeh, A., Talwalkar, A.: Foundations of machine learning. MIT press (2012)
21. Natarajan, N., Dhillon, I.S., Ravikumar, P.K., Tewari, A.: Learning with noisy labels. In: NIPS. pp. 1196–1204 (2013)

22. Patrini, G., Rozza, A., Menon, A., Nock, R., Qu, L.: Making neural networks robust to label noise: A loss correction approach. In: CVPR (2017)
23. du Plessis, M.C., Niu, G., Sugiyama, M.: Analysis of learning from positive and unlabeled data. In: NIPS. pp. 703–711 (2014)
24. Read, J., Pfahringer, B., Holmes, G., Frank, E.: Classifier chains for multi-label classification. *Machine learning* **85**(3), 333–359 (2011)
25. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* **115**(3), 211–252 (2015). <https://doi.org/10.1007/s11263-015-0816-y>
26. Sukhbaatar, S., Bruna, J., Paluri, M., Bourdev, L., Fergus, R.: Training convolutional networks with noisy labels. arXiv preprint arXiv:1406.2080 (2014)
27. Sutskever, I., Martens, J., Dahl, G., Hinton, G.: On the importance of initialization and momentum in deep learning. In: ICML. pp. 1139–1147 (2013)
28. Vahdat, A.: Toward robustness against label noise in training deep discriminative neural networks. In: NIPS. pp. 5596–5605 (2017)
29. Vapnik, V., Vashist, A.: A new learning paradigm: Learning using privileged information. *Neural networks* **22**(5), 544–557 (2009)
30. Veit, A., Alldrin, N., Chechik, G., Krasin, I., Gupta, A., Belongie, S.: Learning from noisy large-scale datasets with minimal supervision. In: CVPR (July 2017)
31. Wan, L., Zeiler, M., Zhang, S., Le Cun, Y., Fergus, R.: Regularization of neural networks using dropconnect. In: ICML. pp. 1058–1066 (2013)
32. Wang, R., Liu, T., Tao, D.: Multiclass learning with partially corrupted labels. *IEEE Transactions on Neural Networks and Learning Systems* **29**(6), 2568–2580 (2018)
33. Xing, E.P., Jordan, M.I., Russell, S.J., Ng, A.Y.: Distance metric learning with application to clustering with side-information. In: NIPS. pp. 521–528 (2003)
34. Yu, X., Liu, T., Gong, M., Batmanghelich, K., Tao, D.: An efficient and provable approach for mixture proportion estimation using linear independence assumption. In: CVPR. pp. 4480–4489 (2018)
35. Yu, X., Liu, T., Gong, M., Zhang, K., Tao, D.: Transfer learning with label noise. arXiv preprint arXiv:1707.09724 (2017)
36. Zhang, T.: Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research* **5**(Oct), 1225–1251 (2004)
37. Zhu, X.: Semi-supervised learning literature survey (2005)