

Weakly Supervised Learning Meets Ride-Sharing User Experience Enhancement

Lan-Zhe Guo¹ Feng Kuang² Zhang-Xun Liu² Yu-Feng Li¹ Nan Ma² Xiao-Hu Qie^{2*}

¹National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China

²Didi Chuxing, Beijing, China

guolz@lamda.nju.edu.cn, {kuangfeng, liuzhangxun}@didichuxing.com,

liyf@lamda.nju.edu.cn, {mandyma.edu, tiger.qie}@didiglobal.com

Abstract

Weakly supervised learning aims at coping with scarce labeled data. Previous weakly supervised studies typically assume that there is only one kind of weak supervision in data. In many applications, however, raw data usually contains more than one kind of weak supervision at the same time. For example, in user experience enhancement from Didi, one of the largest online ride-sharing platforms, the ride comment data contains severe label noise (due to the subjective factors of passengers) and severe label distribution bias (due to the sampling bias). We call such a problem as ‘*compound weakly supervised learning*’. In this paper, we propose the CWSL method to address this problem based on Didi ride-sharing comment data. Specifically, an instance reweighting strategy is employed to cope with severe label noise in comment data, where the weights for harmful noisy instances are small. Robust criteria like AUC rather than accuracy and the validation performance are optimized for the correction of biased data label. Alternating optimization and stochastic gradient methods accelerate the optimization on large-scale data. Experiments on Didi ride-sharing comment data clearly validate the effectiveness. We hope this work may shed some light on applying weakly supervised learning to complex real situations.

Introduction

Conventional supervised learning has achieved great success in various applications. Most successful techniques, such as deep learning, require ground-truth labels to be given for a big training data set. In many real-world applications, however, it can be difficult to attain strong supervision due to the fact that the hand-labeled data sets are time-consuming and expensive to collect. Thus, it is desirable for machine learning techniques to be able to work well with weakly supervised data (Zhou 2017).

Compared to the data in traditional supervised learning, weakly supervised data does not have a large amount of precise label information. Typically, three types of weakly supervised data exist: incomplete supervised data (Oliver

et al. 2018), i.e., only a small subset of training data is given with labels whereas the other data remains unlabeled; inaccurate supervised data (Frénay and Verleysen 2014), i.e., the given labels have not always been ground-truth and inexact supervised data (Carbonneau et al. 2018), i.e., only coarse-grained labels are given. Weakly supervised learning (WSL) has attracted considerable attention, which has consequently resulted in a large number of WSL methods, e.g., (Liu and Ziebart 2014; Oliver et al. 2018; Hendrycks et al. 2018).

In previous studies on weakly supervised learning, a basic assumption is that there is only one kind of weak supervision in training data. However, in many real-world applications such as ride-sharing user experience enhancement, such an assumption is difficult to hold. In Didi ride-sharing comment data, it suffers from severe label noises, i.e., collected labels from comment questions in ride-sharing user experience are low-quality since they can easily be influenced by subjective factors (e.g., emotions) and many misoperations of the passengers. Meanwhile, there is also severe label distribution bias, i.e., label distributions between training and testing data are different, because the negative comments in ride-sharing user experience are rare and they are dynamically changing with time.

It is evident that single type weakly supervised learning could not tackle the problem concerned in this paper. For example, label noise learning methods ignore the label distribution bias which may seriously hurt the performance in practice; while label distribution bias correction methods assume that the ground-truth label is accessible to each instance, which is not the case in our situations. Note that the data scenario studied in this paper is quite different from the previous weakly labeled studies. We call such a problem as ‘*compound weakly supervised learning*’.

In this paper, a novel method named CWSL is proposed and verified on Didi ride-sharing comment data, where the goal is to pre-detect negative comments and help improve ride-sharing user experience. It is crucial to decrease the influence of noisy instances in the training stage at the same time correct the label distribution bias to make the model work well in practice. In our method, an instance reweight strategy is employed to cope with severe labeling noise by

*This work was supported by the National Natural Science Foundation of China (61772262) and Didi research fund. Yu-Feng Li is the corresponding author of this work.
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

assigning a small weight to noisy instances. To make the learned model performs well on the test distribution, the instance reweight process is under the guidance of validation performance which is evaluated by a robust AUC criteria. The idea is formulated as a novel bi-level optimization. An alternating optimization and stochastic gradient method is adopted to make the method adapt to large-scale data. Experiments on Didi ride-sharing comment data clearly validate the effectiveness.

Related Work

The problem focused in this paper, i.e. compound weakly supervised learning, more specifically, is the intersection of label noise learning and label distribution bias problem. Both can be regarded as weakly supervised frameworks where label information conveyed by training instances is incorrect or biased with true distribution.

The performance of machine learning models has been shown to degrade noticeably in the presence of label noise (Fréney and Verleysen 2014). Considerable efforts have been made to correct or learn with noisy labels (Natarajan et al. 2013; Biggio, Nelson, and Laskov 2011; Fréney and Verleysen 2014; Li, Guo, and Zhou 2019). (Mnih and Hinton 2012) allow for label noise robustness by modifying the model’s architecture. (Northcutt, Wu, and Chuang 2017) propose to reweight the noisy instances according to the predicted probability. (Patrini et al. 2017) make use of the forward loss correction mechanism, and propose an estimate of the label corruption estimation matrix. There are also methods that correct labels with small clean validation data. For example, (Veit et al. 2017) use validation data to train a label cleaning network by estimating the residuals between the noisy and clean labels. (Hendrycks et al. 2018) use validation data to estimate a label correction matrix. However, these methods are not sufficient to conquer the studied problem well, because they ignore the label distribution bias that may cause performance degradation seriously in practice.

One kind of strategy to correct label distribution bias is to develop bias robust models. (Liu and Ziebart 2014) develops a framework that can learn a robust bias-aware probabilistic classifier using a minimax estimation formulation. (Chen et al. 2016) propose a label distribution bias robust regression method by minimizing conditional Kullback-Leibler divergence and considering the worst-case performance. Another kind of strategy attempts to correct label distribution bias by reweighting training instances. (Huang et al. 2007) proposes to reweight instances by matching distributions between training and testing sets in feature space. (Zhang et al. 2013) proposes to estimate the weights by reweighting data to reproduce the covariate distribution on the test distribution. There are also methods aim to correct label bias via domain adaptation methods (Azizzadenesheli et al. 2019). However, these methods assume that the ground-truth label is accessible to each instance, which is not the case in our situations and can influence the model performance severely.

Our method is also related to other ‘mixed’ cases under weakly supervised learning frameworks such as *multi-instance multi-label learning* (Zhou et al. 2012), *semi-supervised label noise learning* (Zhang et al. 2019), *semi-*

Notation	Meaning
$\mathbf{x} \in \mathbb{R}^d$	Feature vector of instance
$y \in \{0, 1\}$	Label of instance
n	Number of training instances
m	Number of validation instances, $m \ll n$
$\{\mathbf{x}_i^t, y_i^t\}_{i=1}^n$	Training instances
$\{\mathbf{x}_i^v, y_i^v\}_{i=1}^m$	Validation instances
θ	Model parameters
$\hat{y} = f(\mathbf{x}, \theta)$	Prediction of model θ on instance \mathbf{x}
$\mathbf{w} \in \mathcal{W}$	Weights for training instances
w_i	Weight for the i -th training instance
$\ell(\cdot, \cdot)$	Training loss function
$\mathcal{L}(\cdot, \cdot)$	Validation loss function
λ_θ	Step size for θ
$\lambda_{\mathbf{w}}$	Step size for \mathbf{w}

Table 1: Summary of important notations used in this paper.

supervised multi-label learning (Wei et al. 2018), *semi-supervised weak label learning* (Dong, Li, and Zhou 2018) and *semi-supervised partial label learning* (Wang, Li, and Zhou 2019). It is worth noting that our paper is different from these works as we focus on the mix of severe label noise and biased label distribution.

The Proposed Method

Problem Setting and Notations

Let $(\mathbf{x}, y) \in \mathcal{X} \times \mathcal{Y}$ be the feature-label pair where $\mathbf{x} \in \mathbb{R}^d$ is the feature vector including driver features, passenger features, order specific features and comment question, $y \in \{0, 1\}^n$ is the label where 1 indicates positive comment and 0 indicates negative comment. $\{(\mathbf{x}_i^t, y_i^t), 1 \leq i \leq n\}$ is the training set and $\{(\mathbf{x}_i^v, y_i^v), 1 \leq i \leq m\}$ is a small clean unbiased validation set verified by domain experts. In reality, $m \ll n$. Hereafter, we will use superscript t to denote the training set and superscript v to denote validation set.

Let $\hat{y} = f(\theta, \mathbf{x})$ be the prediction of model θ on instance \mathbf{x} , $\ell(\cdot, \cdot)$ be a loss function to be minimized during the training procedure. Following the classical ERM (Empirical Risk Minimization) framework, we aim to minimize the expected loss for the training set: $\frac{1}{n} \sum_{i=1}^n \ell(f(\theta, \mathbf{x}_i^t), y_i^t) = \frac{1}{n} \sum_{i=1}^n \ell_i(\theta)$ where each training instance is weighted equally, and $\ell_i(\theta)$ stands for the loss function associating with instances (\mathbf{x}_i^t, y_i^t) . The important notations and meanings used in this paper are summarized in Table 1.

Instance Reweight Mechanism

The difficulty of CWSL is to learn with severe noisy labels and at the same time making the learned model works well on real distribution. One intuitive solution to handle noisy instances is to assign a small weight to noisy instances to decrease their influence (Northcutt, Wu, and Chuang 2017), i.e., a weighted loss is minimized during the training procedure:

$$\theta^* = \arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n w_i \ell_i(\theta) \quad (1)$$

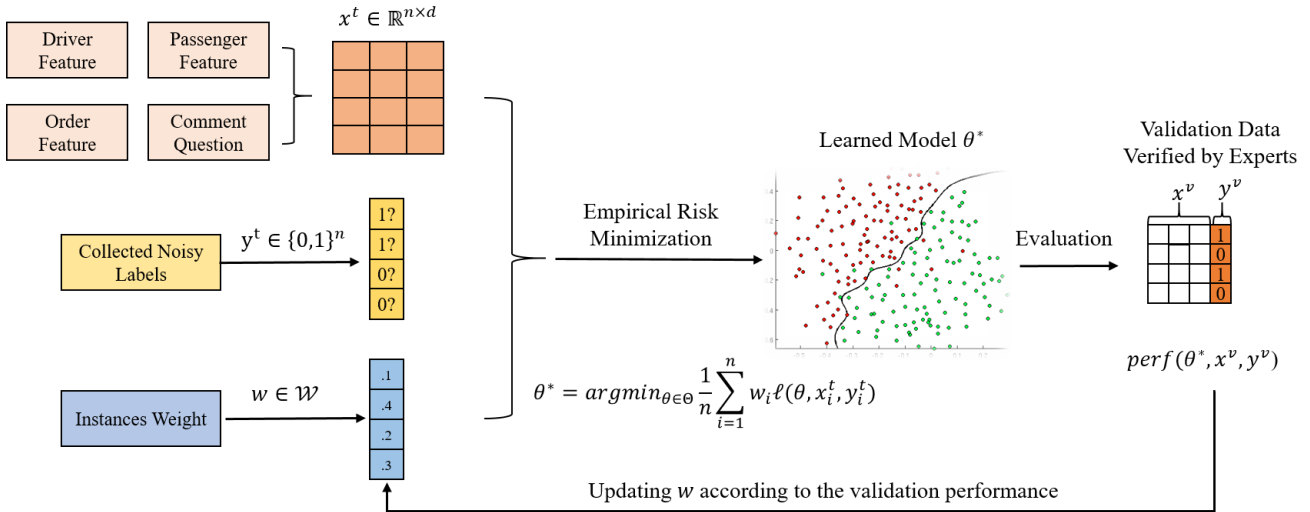


Figure 1: Illustration of the proposed algorithm.

where w_i is the weight for each instance, Θ is the space of model parameter, θ^* is the optimal model learned by the ERM procedure.

However, w_i is unknown upon beginning. Previous label noise learning studies typically calculate the instances weights based on the training loss for the reason that when the noisy instances are fewer than clean instances, the clean instance is more likely to have a smaller training loss (Frénay and Verleysen 2014). In our method, to make the learned model performs well on test distribution, a small clean unbiased validation data is introduced. It is usually acceptable to collect small validation data with small expense in real applications. For example, in Didi, an annotation team can collect accurate labels for a small dataset via telephone interview.

We argue that the optimal weight is that the model trained with this weight should maximize the validation performance since the validation performance acts as a good approximation of the generalization performance. Specifically, once the model trained with \mathbf{w} owns a good validation performance, it can possibly perform well on the general data distribution. Formally, we consider the following optimization,

$$\mathbf{w}^* = \arg \max_{\mathbf{w}} \text{perf}(\theta^*, \mathbf{x}^v, \mathbf{y}^v) \quad (2)$$

where $\text{perf}(\theta^*, \mathbf{x}^v, \mathbf{y}^v)$ is the task specific performance measure (e.g., AUC, Accuracy, F1-Score) of model θ^* on validation set $(\mathbf{x}^v, \mathbf{y}^v)$.

Therefore, by combining the above considerations, the objective of our weight learning algorithm is formulated as:

$$\begin{aligned} \max_{\mathbf{w} \in \mathcal{W}} \quad & \text{perf}(\theta^*, \mathbf{x}^v, \mathbf{y}^v) \\ \text{s.t.} \quad & \theta^* = \arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n w_i \ell_i(\theta) \end{aligned} \quad (3)$$

where the weight \mathbf{w} is from a convex set \mathcal{W} that is typically set as $\mathcal{W} = \{\mathbf{w} | 0 \leq \mathbf{w} \leq 1\}$. Figure 1 illustrates the framework of our method.

The proposed formulation Eq.(3) is a bi-level optimization problem where one optimization problem is nested within another problem (Bard 2013). The lower-level optimization is to find a weighted empirical risk minimizer model given the training set whereas the upper-level optimization is to maximize the validation performance given the learned model. The effectiveness of bi-level optimization (Bard 2013) has been recently demonstrated in the machine learning community. For example, (Franceschi et al. 2017; 2018) propose hyper-parameter optimization methods based on bi-level formulation. (Zhang, Zhu, and Wright 2018) introduce validation data and debug the training set labels according to validation performance. (Ren et al. 2018) propose to reweight examples for robust deep learning based on bi-level optimization meta-learning.

Robust AUC Criteria

To conquer the bias label distribution, where the label distributions in the training data and testing data are different since negative comments in ride-sharing user experience are much fewer than that of positive comments while they are dynamically changing with time in test data. This means directly training a binary classifier in terms of accuracy will predict almost all incoming instances to be positive ones, which obviously does not make sense. So we adopt AUC as the performance measure as AUC is not sensitive to the class ratio (Liu, Wu, and Zhou 2008), which has been shown effective for many applications, such as website ad click prediction (McMahan et al. 2013) where only a very small fraction of web history contains ads clicked by visitors.

Specifically, let \mathbf{x}^+ , \mathbf{x}^- be positive and negative instances generated from distribution \mathcal{P}^+ and \mathcal{P}^- . The AUC is defined as:

$$\text{AUC} = \mathbb{E}_{\substack{\mathbf{x}^+ \sim \mathcal{P}^+ \\ \mathbf{x}^- \sim \mathcal{P}^-}} [\mathbb{I}\{f(\theta, \mathbf{x}^+) - f(\theta, \mathbf{x}^-) > 0\}] \quad (4)$$

where $\mathbb{I}(\cdot)$ is the indicator function which returns 1 if the argument is true and 0 otherwise. This expectation is the

probability that a positive instance is ranked higher than a negative instance.

Owing to its non-convexity and discontinuous, direct optimization of AUC often leads to a NP-hard problem (Gao et al. 2013). To make a compromise for avoiding computational difficulties, we propose to optimize a convex surrogate loss function (Gultekin et al. 2018).

Replacing $\mathbb{I}\{f(\theta, \mathbf{x}^+) - f(\theta, \mathbf{x}^-) > 0\}$ in Eq.(4) with the surrogate loss $\phi(f(\theta, \mathbf{x}^+) - f(\theta, \mathbf{x}^-))$, we now aim to minimize:

$$\mathbb{E}_{\substack{\mathbf{x}^+ \sim \mathcal{P}^+ \\ \mathbf{x}^- \sim \mathcal{P}^-}} [\phi(f(\theta, \mathbf{x}^+) - f(\theta, \mathbf{x}^-))] \quad (5)$$

In this paper, we adopt the pairwise squared loss $\phi(t) = (1-t)^2$ as the surrogate loss function, for the reason that it is convex differentiable and we can easily prove its consistency with the original AUC according to Theorem 2 in (Gao and Zhou 2015).

Denote the positive and negative instances in validation set \mathbf{x}^v as $\mathcal{S}^+ = \{\mathbf{x}_1^{v+}, \dots, \mathbf{x}_{m_+}^{v+}\}$ and $\mathcal{S}^- = \{\mathbf{x}_1^{v-}, \dots, \mathbf{x}_{m_-}^{v-}\}$ where m_+ and m_- are number of positive and negative instances. For a model θ , the surrogate validation loss can be written as:

$$\begin{aligned} \mathcal{L}(\theta) &= \frac{1}{m^+m^-} \sum_{i=1}^{m^+} \sum_{j=1}^{m^-} [(1 - (f(\theta, \mathbf{x}_i^{v+}) - f(\theta, \mathbf{x}_j^{v-})))^2] \\ &= \frac{1}{m^+m^-} \sum_{i=1}^{m^+} \sum_{j=1}^{m^-} [(1 - (\theta^\top \mathbf{x}_i^{v+} - \theta^\top \mathbf{x}_j^{v-}))^2] \\ &= 1 - 2\theta^\top \left[\frac{1}{m^+m^-} \sum_{i=1}^{m^+} \sum_{j=1}^{m^-} (\mathbf{x}_i^{v+} - \mathbf{x}_j^{v-}) \right] \\ &\quad + \theta^\top \left[\frac{1}{m^+m^-} \sum_{i=1}^{m^+} \sum_{j=1}^{m^-} (\mathbf{x}_i^{v+} - \mathbf{x}_j^{v-})(\mathbf{x}_i^{v+} - \mathbf{x}_j^{v-})^\top \right] \theta \\ &= 1 - 2\theta^\top \mu_m + \theta^\top \Sigma_m \theta \end{aligned} \quad (6)$$

where

$$\mu_m = \frac{1}{m^+m^-} \sum_{i=1}^{m^+} \sum_{j=1}^{m^-} (\mathbf{x}_i^{v+} - \mathbf{x}_j^{v-}) \quad (7)$$

$$\Sigma_m = \frac{1}{m^+m^-} \sum_{i=1}^{m^+} \sum_{j=1}^{m^-} (\mathbf{x}_i^{v+} - \mathbf{x}_j^{v-})(\mathbf{x}_i^{v+} - \mathbf{x}_j^{v-})^\top \quad (8)$$

Gradient Descent Updating

By combining the forms Eq.(3) and Eq.(6), the final objective of the optimization is formulated as:

$$\begin{aligned} \min_{\theta \in \Omega} \quad & \mathcal{L}(\theta^*) \\ \text{s.t.} \quad & \theta^* = \arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n w_i \ell_i(\theta) \end{aligned} \quad (9)$$

where $\mathcal{L}(\theta^*) = -2\theta^{*\top} \mu_m + \theta^{*\top} \Sigma_m \theta^*$ and the negative log-likelihood is adopted as the loss function in our task: $\ell_i(\theta) = -y_i^t \theta^\top \mathbf{x}_i^t + \log(1 + e^{\theta^\top \mathbf{x}_i^t})$.

Notice that the lower-level problem can be replaced equivalently with its Karush-Kuhn-Tucker (KKT) condition (Boyd and Vandenberghe 2004):

$$g(\mathbf{w}, \theta) \equiv \frac{1}{n} \sum_{i=1}^n w_i \frac{\partial \ell_i(\theta)}{\partial \theta} = 0 \quad (10)$$

According to the implicit function theorem (Nakasho, Futa, and Shidama 2017), we have the following Jacobian matrix:

$$J = \frac{\partial \theta}{\partial \mathbf{w}} = - \begin{bmatrix} \frac{\partial g_1}{\partial \theta_1} & \dots & \frac{\partial g_1}{\partial \theta_d} \\ \vdots & & \vdots \\ \frac{\partial g_d}{\partial \theta_1} & \dots & \frac{\partial g_d}{\partial \theta_d} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial g_1}{\partial w_1} & \dots & \frac{\partial g_1}{\partial w_n} \\ \vdots & & \vdots \\ \frac{\partial g_d}{\partial w_1} & \dots & \frac{\partial g_d}{\partial w_n} \end{bmatrix} \quad (11)$$

Referring to Eq.(10), we can compute $\frac{\partial g}{\partial \theta}$ and $\frac{\partial g}{\partial \mathbf{w}}$ as:

$$\frac{\partial g}{\partial \theta} = \mathbf{w}^\top \left(\frac{\partial^2 \ell(\theta)}{\partial \theta \partial \theta^\top} \right), \quad \frac{\partial g}{\partial \mathbf{w}} = \frac{\partial \ell(\theta)}{\partial \theta} \quad (12)$$

The gradient and Hessian of the training loss are posed as:

$$\frac{\partial \ell(\theta)}{\partial \theta} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^t (-y_i^t + p(\mathbf{x}_i^t; \theta)) \quad (13)$$

$$\frac{\partial^2 \ell(\theta)}{\partial \theta \partial \theta^\top} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^t \mathbf{x}_i^{t\top} p(\mathbf{x}_i^t; \theta)(1 - p(\mathbf{x}_i^t; \theta)) \quad (14)$$

where $p(\mathbf{x}_i^t; \theta) = \Pr\{\hat{y} = 1 | \mathbf{x}_i^t, \theta\}$.

The matrix J tells us how the model parameters θ varies with respect to an infinitesimal change to \mathbf{w} . Now, we can apply the chain rule to get the gradient of the whole optimization problem w.r.t. \mathbf{w} :

$$\frac{\partial \mathcal{L}(\theta^*)}{\partial \mathbf{w}} = \frac{\partial \mathcal{L}(\theta^*)}{\partial \theta} \frac{\partial \theta}{\partial \mathbf{w}} \quad (15)$$

Finally, we can solve the bi-level optimization problem with gradient-based methods and the iteration rule is:

$$\mathbf{w}_t = \mathbf{w}_{t-1} - \lambda_{\mathbf{w}} \left. \frac{\partial \mathcal{L}(\theta^*)}{\partial \mathbf{w}} \right|_{\mathbf{w}=\mathbf{w}_{t-1}} \quad (16)$$

where $\lambda_{\mathbf{w}}$ is the step size for \mathbf{w} .

Notice that to solve the gradient w.r.t. \mathbf{w} , we need to compute the inverse of the Hessian matrix and in practice, the matrix inverses is not pleasing. Here, we propose an alternative method. We compute $-(\frac{\partial g}{\partial \theta})^{-1} \frac{\partial g}{\partial \mathbf{w}}$ as it is the solution to $(\frac{\partial g}{\partial \theta}) \mathbf{x} = -\frac{\partial g}{\partial \mathbf{w}}$. The linear system can be solved conveniently and only requires matrix-vector products, that is we do not have to materialize the Hessian.

Efficient Alternating Optimization

Calculating the optimal \mathbf{w} requires two nested loops of optimization, to further accelerate the optimization for large-scale data, we propose an alternating optimization method by updating \mathbf{w} and θ iteratively. Instead of updating \mathbf{w} based on the optimal θ^* , we update \mathbf{w} after each iteration of the optimization of θ_t , i.e.,

$$\theta_t = \theta_{t-1} - \lambda_{\theta} \mathbf{w}_{t-1} \left. \frac{\partial \ell(\theta)}{\partial \theta} \right|_{\theta=\theta_{t-1}} \quad (17)$$

$$\mathbf{w}_t = \mathbf{w}_{t-1} - \lambda_{\mathbf{w}} \left. \frac{\partial \mathcal{L}(\theta_t)}{\partial \mathbf{w}} \right|_{\mathbf{w}=\mathbf{w}_{t-1}} \quad (18)$$

Then, compared to regular optimization methods, our method only needs $2 \times$ training time. The overall optimization procedure is summarized in Algorithm 1.

Algorithm 1 Alternating Optimization Procedure of CWSL.

Input: Training set $\{\mathbf{x}_i^t, y_i^t\}_{i=1}^n$, validation set $\{\mathbf{x}_i^v, y_i^v\}_{i=1}^m$, current values of weights \mathbf{w}_0 and model parameters θ_0 .

Output: Learned instance weights \mathbf{w} and model parameters θ .

- 1: **for** $t = 1$ to T **do**
 - 2: $\theta_t = \theta_{t-1} - \lambda_\theta \mathbf{w}_{t-1} \left. \frac{\partial \ell(\theta)}{\partial \theta} \right|_{\theta=\theta_{t-1}}$
 - 3: Compute $\frac{\partial g}{\partial \theta}$, $\frac{\partial g}{\partial \mathbf{w}}$ with Eq.(12)
 - 4: Compute J with Eq.(11).
 - 5: Compute $\left. \frac{\partial \mathcal{L}(\theta_t)}{\partial \mathbf{w}} \right|_{\theta=\theta_t}$ with Eq.(15).
 - 6: $\mathbf{w}_t = \mathbf{w}_{t-1} - \lambda_{\mathbf{w}} \left. \frac{\partial \mathcal{L}(\theta_t)}{\partial \mathbf{w}} \right|_{\mathbf{w}=\mathbf{w}_{t-1}}$.
 - 7: Project \mathbf{w}_t to the constrained set \mathcal{W} .
 - 8: **end for**
 - 9: **return** \mathbf{w}_T and θ_T .
-

Experiments

In this section, we conduct experiments to demonstrate the effectiveness of the proposed method. We first conduct experiments on UCI dataset *breast_cancer* with synthetic noisy labels to show the instance weights evolution during the training procedure. Then, we present the evaluation settings on real ride comment user experience data in Didi ride-sharing platform include data set preparation, feature transformation, a brief description of the related comparison method and the detail parameter settings. Finally, extensive results of different methods are presented along with the analysis.

Synthetic Data

To understand how our proposed algorithm contributes to decreasing the influence of noisy labels, we conduct an experiment on UCI dataset *breast_cancer*¹ with synthetic noisy labels to show how the instance weights evolved during the training procedure.

Breast_cancer is a classic binary classification dataset with 569 instances and 32 attributes. We split the dataset into 3 parts: 469 training data, 50 validation data, and 50 test data. For the training data, we randomly select 100 instances and reverse their labels as noisy instances.

For this experiment, we adopt a two-layer neural network as the classification model and binary cross-entropy loss to evaluate the validation performance instead of the AUC we described before, as this is no label distribution bias problem. The training epoch is set to 100 and the initial weights of all instances are set to 0.5.

Figure 2 shows the instance weights evolution during the training procedure. Figure 2(a) shows that our method can separate the noisy and clean instances after only 1 epoch. After 10 epochs, the weight of clean instances has been mainly concentrated in 1.0. After 100 epochs, as expected the algorithm nicely separates the two sets and correctly pushes all 100 noisy instances into zero weight while keeping the clean instances into one weight. Finally, the algorithm can achieve 100% accuracy on the test data. These demonstrate that our proposal is able to reliably decrease the influence of noisy instances by assigning them small weights.

Real Didi Ride-Sharing Comment Data

Didi is one of the largest online ride-sharing platforms that offer peer-to-peer ride-sharing services. Once passengers enter start location and destination, the platform will match a driver nearby to

pick up the passenger. The platform provides a comment question to passengers after each order, e.g., “Was the car smelly?” or “Did your driver detour?”. The goal is to rank these candidate comment questions according to the probability that may receive negative comments because negative comments can help discover the needs of users and the deficiency of services, thus improving service quality and user experience.

The benchmark ride-sharing comment user experience data set was constructed from the real comments in the main city zone of ride-sharing orders within the time period from Mar 1st, 2019 to Apr 1st, 2019. We sample 600,000 instances as the training data and the training data have severe label noise as we mentioned before. To construct validation and test data, we sample 15,000 data from Apr 2st, 2019 according to the real distribution, and verified their labels by domain experts. Class ratio of negative to positive in the real distribution is approximately 1:30. The verified data is equally split into the validation set and test set. Thus, we have prepared the training data set and clean unbiased validation and test data set.

Feature Engineering Each instance in the data set is described with hundreds of raw features including driver features, passenger features, order specific features and comment question. For example, driver and passenger features include age, gender, number of negative reviews for a comment question, etc. Order specific features include cities, price, destination, etc. Different comment question is also described with different numerical features.

The raw feature contains both continuous and categorical features, simply train models on the raw feature will leading to poor performance and manual feature engineering is time-consuming. Inspired by (He et al. 2014), we propose to transform the input features of the classifier in order to improve its performance by boosted decision trees. The boosted tree is a powerful and very convenient way to implement non-linear and tuple transformations and has been widely used in Click-Through Rate (CTR) prediction (Richardson, Dominowska, and Ragno 2007).

Specifically, we treat each individual tree as a categorical feature that takes as value the index of the leaf an instance ends up falling in, then we normalize the transformed feature into zero mean and unit variance. The boosted decision tree we used follows the XGBoost (Chen and Guestrin 2016).

Label Bias Learning Methods To demonstrate the effectiveness of our proposal, we compare the proposed method with three classical classification methods with label distribution bias correction:

- **XGBoost** (Chen and Guestrin 2016): which is a scalable machine learning system for tree boosting (Friedman 2001). The effectiveness of XGBoost has been widely recognized in a number of machine learning and data mining challenges, like Kaggle competitions. For XGBoost, we set the weight of positive class into 4 and other parameters are set as default.
- **LR (Logistic Regression)**: which is a baseline method that commonly used in classification tasks. For the Logistic Regression, we adopt the implementation in Sklearn² and the parameters are set as default.
- **DNN (Deep Neural Network)**: We also design a DNN structure including three hidden layers and a prediction layer, where ReLU (Glorot, Bordes, and Bengio 2011) is adopted as the activation function for each hidden layer. The hyper-parameters are tuned based on the validation set. For instance, the number of units for each hidden layer is set to 64. Dropout rate is set as

¹<https://archive.ics.uci.edu>

²<https://scikit-learn.org>

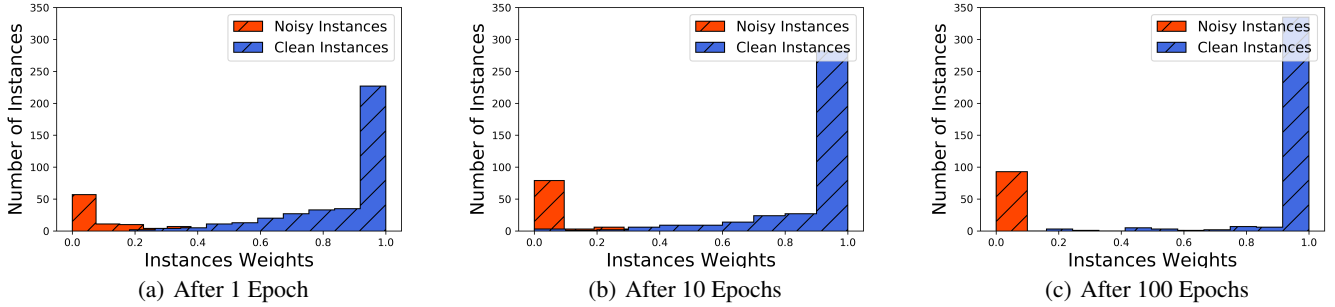


Figure 2: Instances weights distribution after 1 Epoch, 10 Epochs, and 100 Epochs. The instances weights evolution shows that our proposal successfully identifies noisy instances and can reduce their influences by assigning them low weights.

0.5. We adopt the binary cross-entropy loss and SGD algorithm with learning 0.01 to train the DNN model.

For these methods, we correct the predicted probability to deal with the label distribution bias problem. According to **probability calibrating method** (Dal Pozzolo et al. 2015) which is a commonly used method to correct label distribution bias, let β be the sampling ratio of the positive instances in training data, $0 < \beta < 1$, the predicted probability of instance x belong to negative class is p_s , the probability can be corrected with $p = \frac{p_s \beta}{p_s \beta - p_s + 1}$.

Label Noise Learning Methods We also compare our proposal with three state-of-the-art label noise learning methods:

- **Rank Pruning** (Northcutt, Wu, and Chuang 2017): Rank pruning is a novel label noise learning method that can estimate the fraction of mislabeling in both the positive and negative training instances and can be applied to any probabilistic classifier. Rank Pruning achieves perfect noise estimation and equivalent expected risk as learning with correct labels in theoretical and achieves nearly the same performance as learning with correct labels. For rank pruning method, we adopt the XGBoost as the basic classification model.
- **GLC (Gloden Loss Correction)** (Hendrycks et al. 2018): Similar to our proposal, GLC also introduces a small set of trusted data and estimate a noise correction matrix based on the validation data. The GLC method achieves state-of-the-art performance on varies real-world dataset, e.g., CIFAR-10 and IMDB. For GLC, we adopt the DNN as the basic model with structure as we described before.
- **LTR (Learning to Reweight Examples)** (Ren et al. 2018): LTR aims to reweight training instances based on the validation loss. Different from our proposal, LTR simply optimizes the cross-entropy loss on validation data based on a trivial optimization method while our method optimizes the robust AUC. For LTR method, we adopt the same architecture as the DNN method. However, due to the validation data is class imbalance, simply using LTR will obtain a degenerate solution, i.e., predict the labels of all instances as positive. Therefore, we rescale the validation loss based on ratio 1:30 according to the idea of cost-sensitive learning (Zhou and Liu 2006).

For methods that fail to utilize validation data (e.g., XGBoost, LR, DNN, Rank Pruning), we add the validation data into the training set to ensure fairness. All methods are trained on the transformed feature. The model structure we adopted is same with DNN and the hyper-parameters of the proposed method are tuned according to the AUC score on the validation set. The initial weights of

all instances are set to 0.5, and the training epochs are set to 100. The other two hyper-parameters (step sizes λ_θ and λ_w), are fixed to 0.1 and 0.4, respectively.

Table 2: Comparison results of the compared methods. For each method, we report the average AUC score and the running time of 10 runs on the test set.

Methods	AUC	Time Cost (s)
LR	78.64 ± 1.23	47.8
DNN	77.85 ± 1.17	90.4
XGBoost	79.55 ± 1.13	324.3
Rank Pruning	82.22 ± 0.53	601.4
GLC	82.63 ± 0.49	300.5
LTR	85.45 ± 0.51	211.6
Proposal	91.24 ± 0.47	107.1

Comparison Results Table 2 shows the experimental results of our proposal and compared methods. From Table 2, we can see that our proposal significantly outperforms all the compared methods. Compared with XGBoost, we achieve more than 12% performance gain and compared with LTR we also achieve nearly 6% performance gain. These show that the previous works that consider only one weak supervision can not deal with the complex compound situations well. Moreover, training models using validation data only cause poor performance (AUC 60.98 with XGBoost) as the validation data is insufficient to train a good model, so we do not report the result in Table 2. We also try to correct label bias for label noise learning methods but do not achieves performance improvement. One possible reason is that these methods modify the priors of the training set and consequently simply adopt the label bias correction method is not useful. These results demonstrate the effectiveness of our proposal.

Validation Set Previously, experiments are conducted based on clean unbiased validation data. It is interesting to study how the performance influenced by the validation data bias. We conduct experiments with the ratio of positive to negative class in the validation set varied from 1:1 to more than 40:1 on three methods based on the validation set, i.e., LTR, GLC, and our proposed method. The AUC on test data is reported in Figure 3. From Figure 3, we can see that, as the change of validation set bias, our proposal is consistently better than the other two well-performed methods. This demonstrates that our proposal is robust to the validation set bias.

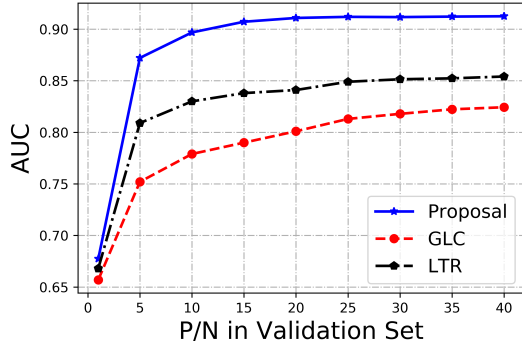


Figure 3: AUC on test data with the varied ratio between positive to negative class in the validation set

Running Time We also report the running time of compared methods in Table 2. The training of our proposal is implemented in Pytorch³ and all the code are run on a single machine with 10cores 2.20GHz Intel Xeon(R) CPU and 32GB main memory.

From the results, we can see that our proposal is the fastest method compared with label noise learning methods whereas Rank Pruning that based on XGBoost is the slowest algorithm and our proposal only cost a half time compared with LTR, though these two algorithms are all based on bi-level optimization. Compared with the three baseline methods, our proposal is also comparable with DNN and faster than XGBoost. These demonstrate that our proposal is efficient to handle large-scale dataset.

Table 3: Performance results of our proposal and compared methods on Precision, Recall and F1-Score metrics.

Methods	Precision	Recall	F1-Score
LR	24.08%	34.69%	26.43%
DNN	23.26%	35.34%	25.71%
XGBoost	9.51%	69.80%	16.73%
Rank Pruning	12.26%	54.73%	18.27%
GLC	24.13%	34.24%	26.79%
LTR	8.63%	87.35%	15.70%
Proposal	16.50%	80.82 %	27.41%

Performance Measures It is meaningful to study whether the proposal is effective in other performance measures. Table 3 presents the experimental comparison on Precision, Recall, and F1-Score metrics, that are commonly used measures to evaluate the ranking quality.

Specifically, denoting all ground-truth positive instances as G and all predicted positive instances as T , then precision and recall are defined as follows:

$$precision = \frac{|G \cap T|}{|T|}, \quad recall = \frac{|T \cap G|}{|G|}$$

and the F1-Score is defined as:

$$F1 = \frac{2 * precision * recall}{precision + recall}$$

³<https://pytorch.org/>

From the results, we can see that LTR achieve the best recall while the lowest precision, DNN, LR, and GLC achieve similar results on these three metrics. Our proposal trades off precision and recall best and achieves the best result on F1-Score. These results reveal that although the proposal is designed to optimize AUC, it has a certain degree of robustness to the change of performance measures.

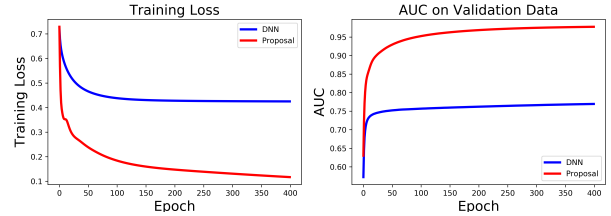


Figure 4: Training Loss v.s. Epoch (left) and Validation AUC v.s. Epoch (right) of our proposal and DNN.

Convergence We also conduct experiments to study the convergence property of our proposed alternating optimization method compared with the normal gradient method. The training loss and validation AUC v.s. training epoch of our proposal and deep neural network are reported in Figure 4. From Figure 4, we can see that the number of epoch required to converge of our proposal is comparable with the naive DNN based on single-level gradient descent method. This demonstrates the efficiency of our proposed optimization method.

Conclusion

In this paper, we study compound weakly supervised learning motivated by real application, i.e., user experience enhancement in Didi, one of the largest online ride-sharing platforms. This is an intersection of two kinds of weak supervision, i.e., severe noisy label and distribution biased label happen simultaneously. This is a new kind of weakly supervised learning problem that to the best of our knowledge, has not been thoroughly studied before. To address this problem, we propose the CWSL method. An instance reweighting strategy is employed to cope with severe label noise. Robust criteria AUC and the validated performance are optimized for the correction of biased data label. Efficient algorithms accelerate the optimization on large-scale data. Experiments on real data sets validate the effectiveness of CWSL in handling with compound weakly supervised data. In the future, we will consider extending this work to dynamic scenarios with an automatic model update.

References

- Azizzadenesheli, K.; Liu, A.; Yang, F.; and Anandkumar, A. 2019. Regularized learning for domain adaptation under label shifts. In *Proceedings of the 7th International Conference on Learning Representations*.
- Bard, J. F. 2013. *Practical Bilevel Optimization: Algorithms and Applications*. Springer Science & Business Media.
- Biggio, B.; Nelson, B.; and Laskov, P. 2011. Support vector machines under adversarial label noise. In *Proceedings of the 3rd Asian Conference on Machine Learning*, 97–112.
- Boyd, S., and Vandenberghe, L. 2004. *Convex optimization*. Cambridge University Press.
- Carbonneau, M.-A.; Cheplygina, V.; Granger, E.; and Gagnon, G. 2018. Multiple instance learning: A survey of problem characteristics and applications. *Pattern Recognition* 329–353.

- Chen, T., and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.
- Chen, X.; Monfort, M.; Liu, A.; and Ziebart, B. D. 2016. Robust covariate shift regression. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, 1270–1279.
- Dal Pozzolo, A.; Caelen, O.; Johnson, R. A.; and Bontempi, G. 2015. Calibrating probability with undersampling for unbalanced classification. In *IEEE Symposium Series on Computational Intelligence*, 159–166.
- Dong, H.-C.; Li, Y.-F.; and Zhou, Z.-H. 2018. Learning from semi-supervised weak-label data. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
- Franceschi, L.; Donini, M.; Frasconi, P.; and Pontil, M. 2017. Forward and reverse gradient-based hyperparameter optimization. In *Proceedings of the 34th International Conference on Machine Learning*, 1165–1173.
- Franceschi, L.; Frasconi, P.; Salzo, S.; and Pontil, M. 2018. Bilevel programming for hyperparameter optimization and meta-learning. In *Proceedings of the 35th International Conference on Machine Learning*, 1563–1572.
- Frénay, B., and Verleysen, M. 2014. Classification in the presence of label noise: a survey. *IEEE Transactions on Neural Networks and Learning Systems* 845–869.
- Friedman, J. H. 2001. Greedy function approximation: a gradient boosting machine. *Annals of Statistics* 1189–1232.
- Gao, W., and Zhou, Z. 2015. On the consistency of AUC pairwise optimization. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, 939–945.
- Gao, W.; Jin, R.; Zhu, S.; and Zhou, Z.-H. 2013. One-pass auc optimization. In *Proceedings of the 30th International Conference on Machine Learning*, 906–914.
- Glorot, X.; Bordes, A.; and Bengio, Y. 2011. Deep sparse rectifier neural networks. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, 315–323.
- Gultekin, S.; Saha, A.; Ratnaparkhi, A.; and Paisley, J. W. 2018. MBA: mini-batch AUC optimization. *CoRR* abs/1805.11221.
- He, X.; Pan, J.; Jin, O.; Xu, T.; Liu, B.; Xu, T.; Shi, Y.; Atallah, A.; Herbrich, R.; Bowers, S.; et al. 2014. Practical lessons from predicting clicks on ads at facebook. In *Proceedings of the 8th International Workshop on Data Mining for Online Advertising*, 1–9.
- Hendrycks, D.; Mazeika, M.; Wilson, D.; and Gimpel, K. 2018. Using trusted data to train deep networks on labels corrupted by severe noise. In *Advances in Neural Information Processing Systems*, 10456–10465.
- Huang, J.; Gretton, A.; Borgwardt, K.; Schölkopf, B.; and Smola, A. J. 2007. Correcting sample selection bias by unlabeled data. In *Advances in Neural Information Processing Systems*, 601–608.
- Li, Y.-F.; Guo, L.-Z.; and Zhou, Z.-H. 2019. Towards safe weakly supervised learning. *IEEE transactions on Pattern Analysis and Machine Intelligence*.
- Liu, A., and Ziebart, B. 2014. Robust classification under sample selection bias. In *Advances in Neural Information Processing Systems*, 37–45.
- Liu, X.-Y.; Wu, J.; and Zhou, Z.-H. 2008. Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics* 539–550.
- McMahan, H. B.; Holt, G.; Sculley, D.; Young, M.; Ebner, D.; Grady, J.; Nie, L.; Phillips, T.; Davydov, E.; Golovin, D.; et al. 2013. Ad click prediction: a view from the trenches. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1222–1230.
- Mnih, V., and Hinton, G. E. 2012. Learning to label aerial images from noisy data. In *Proceedings of the 29th International Conference on Machine Learning*, 567–574.
- Nakasho, K.; Futa, Y.; and Shidama, Y. 2017. Implicit function theorem. part I. *Formalized Mathematics* 269–281.
- Natarajan, N.; Dhillon, I. S.; Ravikumar, P. K.; and Tewari, A. 2013. Learning with noisy labels. In *Advances in Neural Information Processing Systems*, 1196–1204.
- Northcutt, C. G.; Wu, T.; and Chuang, I. L. 2017. Learning with confident examples: Rank pruning for robust classification with noisy labels. In *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence*.
- Oliver, A.; Odena, A.; Raffel, C. A.; Cubuk, E. D.; and Goodfellow, I. 2018. Realistic evaluation of deep semi-supervised learning algorithms. In *Advances in Neural Information Processing Systems*, 3235–3246.
- Patrini, G.; Rozza, A.; Krishna Menon, A.; Nock, R.; and Qu, L. 2017. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1944–1952.
- Ren, M.; Zeng, W.; Yang, B.; and Urtasun, R. 2018. Learning to reweight examples for robust deep learning. In *Proceedings of the 35th International Conference on Machine Learning*, 4331–4340.
- Richardson, M.; Dominowska, E.; and Ragno, R. 2007. Predicting clicks: estimating the click-through rate for new ads. In *Proceedings of the 16th International Conference on World Wide Web*, 521–530.
- Veit, A.; Alldrin, N.; Chechik, G.; Krasin, I.; Gupta, A.; and Belongie, S. 2017. Learning from noisy large-scale datasets with minimal supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 839–847.
- Wang, Q.; Li, Y.; and Zhou, Z. 2019. Partial label learning with unlabeled data. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 3755–3761.
- Wei, T.; Guo, L.-Z.; Li, Y.-F.; and Gao, W. 2018. Learning safe multi-label prediction for weakly labeled data. *Machine Learning* 107(4):703–725.
- Zhang, K.; Schölkopf, B.; Muandet, K.; and Wang, Z. 2013. Domain adaptation under target and conditional shift. In *Proceedings of the 30th International Conference on Machine Learning*, 819–827.
- Zhang, Z.; Zhao, P.; Jiang, Y.; and Zhou, Z. 2019. Learning from incomplete and inaccurate supervision. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD*, 1017–1025.
- Zhang, X.; Zhu, X.; and Wright, S. J. 2018. Training set debugging using trusted items. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, 4482–4489.
- Zhou, Z.-H., and Liu, X.-Y. 2006. On multi-class cost-sensitive learning. In *Proceedings of the 21st National Conference on Artificial Intelligence*, 567–572.
- Zhou, Z.-H.; Zhang, M.-L.; Huang, S.-J.; and Li, Y.-F. 2012. Multi-instance multi-label learning. *Artificial Intelligence* 2291–2320.
- Zhou, Z.-H. 2017. A brief introduction to weakly supervised learning. *National Science Review* 44–53.