

Cleaning Crowdsourced Labels Using Oracles For Statistical Classification

by

Mohamad Dolatshah

B.Sc., Iran University of Science and Technology, 2016

Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of
Master of Science

in the
School of Computing Science
Faculty of Applied Sciences

**© Mohamad Dolatshah 2018
SIMON FRASER UNIVERSITY
Summer 2018**

Copyright in this work rests with the author. Please ensure that any reproduction or re-use
is done in accordance with the relevant national copyright legislation.

Approval

Name: Mohamad Dolatshah

Degree: Master of Science (Computing Science)

Title: Cleaning Crowdsourced Labels Using Oracles
For Statistical Classification

Examining Committee:

Chair: Oliver Schulte
Professor

Jiannan Wang
Senior Supervisor
Assistant Professor

Jian Pei
Supervisor
Professor

Parmit Chilana
Examiner
Assistant Professor

Date Defended: August 21, 2018

Abstract

Nowadays, crowdsourcing is being widely used to collect training data for solving classification problems. However, crowdsourced labels are often noisy, and there is a performance gap between classification with noisy labels and classification with ground-truth labels. In this thesis, we consider how to apply oracle-based label cleaning to reduce the gap. We propose TARS, a label-cleaning advisor that can provide two pieces of valuable advice for data scientists when they need to train or test a model using noisy labels. Firstly, in the model testing stage, given a test dataset with noisy labels, and a classification model, TARS can use the test data to estimate how well the model will perform w.r.t. ground-truth labels. Secondly, in the model training stage, given a training dataset with noisy labels, and a classification algorithm, TARS can determine which label should be sent to an oracle to clean such that the model can be improved the most. For the first advice, we propose an effective estimation technique, and study how to compute confidence intervals to bound its estimation error. For the second advice, we propose a novel cleaning strategy along with two optimization techniques, and illustrate that it is superior to the existing cleaning strategies. We evaluate TARS on both simulated and real-world datasets. The results show that (1) TARS can use noisy test data to accurately estimate a model's true performance for various evaluation metrics; and (2) TARS can improve the model accuracy by a larger margin than the existing cleaning strategies, for the same cleaning budget.

Keywords: data cleaning; classification; label noise

Dedication

To my family.

Acknowledgements

I would like to express my deepest gratitude to my senior supervisor, Dr. Jiannan Wang. It is unimaginable how somebody can be so passionate about research and science and always ready to answer questions and constantly engaged with solving problems and attacking challenges, no matter day or night. There was absolutely no time during my studies that he was not interested in continuing research and there was no challenge for him to be unwilling to solve. I owe him my highest gratefulness for supporting me during this time and helping me to do research no matter what problem happened during the process. I cannot thank him enough for being such a great person he is.

I would also like to thank Dr. Jian Pei for continuous help during the process of writing the thesis. In spite of being very busy, he was always available for helping and suggesting. Without his extremely helpful insights this work would not have been possible.

I am also grateful to my lab mates for their kind support. Especially, Mathew Teoh, who had contributed to very important parts of the work. I would not have been able to complete this work without his contribution. It was a wonderful experience to be able to work with such an enthusiastic and hard working group of young people, hungry for doing something wonderful every day.

Last but not least, my warmest and most sincere gratitude goes to my family, especially my mother who is the reason for whatever I have today and gave me unforgettable and unconditional love despite all the hardships which is the most valuable thing in my life. They emotionally supported me and comforted me from the other side of the world. There was nothing other than the love of my family, encouraging me to carry on.

Table of Contents

Approval	ii
Abstract	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vi
List of Tables	viii
List of Figures	ix
1 Introduction	1
2 Related Work	6
3 Problem Definition	8
3.1 Background	8
3.2 Advice 1: How Good is a Model?	9
3.3 Advice 2: Which Label Should be Cleaned?	11
4 Preliminary: Label Consolidation	12
4.1 The Need For Label Consolidation	12
4.2 Computing Consolidated Labels	12
4.3 Quantify Consolidated Label’s Uncertainty	13
5 Model Evaluation	15
5.1 Estimating Model’s True Performance	15
5.2 Bounding Estimation Error	17
5.2.1 Analytical Confidence Interval	17
5.2.2 Empirical Confidence Interval	19
6 Cleaning Strategy	22

6.1	Limitations of Existing Strategies	22
6.2	Main Idea: Expected Model Improvement	24
6.3	Further Optimization	26
7	Experiments	29
7.1	Experimental Settings	29
7.2	Evaluation of Advice 1	31
7.2.1	Sensitivity Analysis	31
7.2.2	Performance on Real-world Datasets	32
7.3	Evaluation of Advice 2	32
7.3.1	Optimization Techniques	32
7.3.2	Efficiency Study	33
7.3.3	Cleaning Strategies	34
8	Conclusion	37
9	Future Work	38
	Bibliography	39
	Appendix A Multi Class Generalization	44
A.0.1	Computing Consolidated Labels	44
A.0.2	Estimating Model's True Performance	45
A.0.3	Expected Model Improvement	45
A.0.4	Empirical Experiment	46

List of Tables

Table 7.1	Dataset statistics (SS: synthetic data with simulated noisy labels; RS: real-world data with simulated noisy labels; RR: real-world data with real-world crowdsourced labels)	29
-----------	---	----

List of Figures

Figure 1.1	TARS can provide two pieces of valuable advice when data scientists need to train or test a model using noisy labels	2
Figure 1.2	Datasets with ground-truth labels, crowdsourced labels, and a model’s predicted labels.	3
Figure 3.1	An illustration of (a) model’s confusion matrix and (b) four evaluation metrics derived from the confusion matrix.	10
Figure 5.1	The relationships between sample size (N), label noise (β), and estimation error (half the width of the 95% confidence interval), for model accuracy $\theta = 0.8$. (For simplicity, we do not show the estimation error for $\beta \in (0.5, 1]$ because based on Equation 5.10, it will be the same as $1 - \beta$). . .	19
Figure 5.2	An illustration of using bootstrap to compute empirical confidence intervals for complex estimators.	21
Figure 6.1	Illustrating the limitation of uncertain sampling.	23
Figure 6.2	Illustrating the limitation of sorting by noise rate.	23
Figure 7.1	Comparison of the estimation error between TARS and DirtyEval by varying sample size (Gaussian).	35
Figure 7.2	Comparison of the estimation error between TARS and DirtyEval by varying noise rates (Gaussian).	35
Figure 7.3	Comparison of the estimation error between TARS and DirtyEval by varying the number of votes per instance (Gaussian).	35
Figure 7.4	Comparison of the estimation error between TARS and DirtyEval by varying the percentage of good workers (Gaussian).	35
Figure 7.5	Comparison of the estimation error of TARS and DirtyEval on real-world datasets.	35
Figure 7.6	Evaluating the effectiveness of the proposed optimization techniques of the TARS’s cleaning strategy with different noise rates (Gaussian).	36
Figure 7.7	Evaluating the impact of the pruning technique on the efficiency and effectiveness of TARS.	36

Figure 7.8	Comparing the effectiveness of different cleaning strategies by varying the number of cleaned instances.	36
Figure A.1	Comparison of cleaning strategies on the multi-class dataset	46

Chapter 1

Introduction

Classification is a fundamental problem in machine learning and statistics. It has achieved great success in many real-world applications, such as entity resolution, disease prediction, and fraud detection. The goal of classification is to train a *classifier* (a.k.a. *model*) from a collection of $\langle \text{instance}, \text{label} \rangle$ pairs such that the classifier can predict the labels of unseen instances. The labeled pairs are split into two parts: training data and test data, where the training data is for model training and the test data is for model evaluation.

Despite the great success that classification has achieved, it suffers from the high cost of data labeling. Crowdsourcing is a promising way to overcome this limitation. Crowdsourcing platforms such as Amazon Mechanical Turk have a large pool of crowd workers available. These workers can be used to label data at low cost and fast speed. But at the same time, crowd workers are not accurate. They often provide *noisy labels* with certain probabilities to be wrong. Although this issue can be mitigated by assigning an instance to multiple workers and then infer the instance’s true label using a truth-inference algorithm like Majority Vote, the state-of-the-art truth-inference algorithms are still far from perfect [73].

Having noisy labels in training data will negatively affect the performance of a classification algorithm because the algorithm tries to predict noisy labels rather than ground-truth labels. In the Machine Learning community, there has already been some work that studies how to clean noisy labels to solve this problem [8]. However, the noisy labels are cleaned by heuristic algorithms, which have no guarantee on cleaning accuracy and may even mess up a lot of correct labels [21].

Unlike the existing work, our work draws some inspiration from the recent progress in the data cleaning community [3, 20, 38], and focuses on a *different data-cleaning scenario*. We consider that there exists an oracle who can be queried to clean noisy labels. Each query is to ask the oracle to verify whether a training example, $\langle \text{instance}, \text{label} \rangle$, is correctly labeled or not. If not, replace its label with the true label. This scenario more often than not holds in reality. Imagine a data scientist needs to train a good classifier on noisy data. In this situation, she can ask internal experts from her company to serve as oracles to clean the noisy labels.

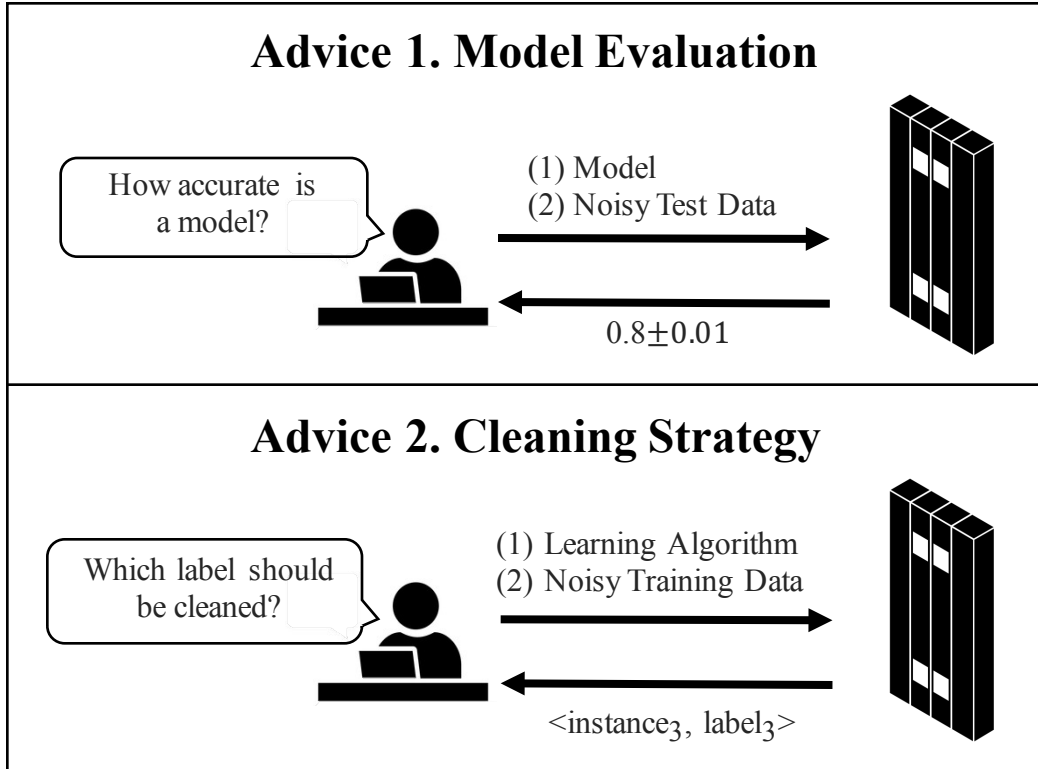


Figure 1.1: TARS can provide two pieces of valuable advice when data scientists need to train or test a model using noisy labels

It is worth noting that the goal of this thesis is *not* to develop yet another classification algorithm for noisy labels. Instead, we aim to develop a *label cleaning advisor*, named as TARS¹, that can advise a data scientist on how to make the best use of oracle-based cleaning for classification. As shown in Figure 1.1, suppose a data scientist has already trained a model on a noisy dataset. The first piece of advice she may ask for is how accurate the model is. If the current model is already good enough, there is no need to further spend effort on label cleaning. However, if she finds that the current model does not meet her need, the next question she may ask is which label should be cleaned such that the model can be improved the most. TARS can provide the two pieces of advice:

- *Advice 1. Model Evaluation.* Given a model, and a test dataset labeled by the crowd, TARS can tell a data scientist how well the model will perform w.r.t. *ground-truth labels* (rather than w.r.t. noisy labels). One may argue that since the test dataset is relatively small, it may be affordable to clean the entire test dataset. This is true if there is only a single test dataset. In reality, however, a data scientist often needs to construct multiple test datasets to avoid overfitting. For example, suppose a 10-fold cross validation is applied for this purpose. The dataset will be divided into ten

¹TARS is named after an intelligent robot in *Interstellar* who can provide insightful advice for human beings.

(a) Ground-truth Labels		(b) Crowdsourced (noisy) labels		(c) Model's Prediction	
Instance	True Label	Instance	Noisy Label	Instance	Predicted Label
x_1	+1	x_1	$(w_1, +1)$	x_1	+1
x_2	+1	x_2	$(w_1, +1)$	x_2	+1
x_3	-1	x_3	$(w_1, +1)$	x_3	-1
x_4	-1	x_4	$(w_1, -1)$	x_4	-1
x_5	-1	x_5	$(w_2, -1)$	x_5	+1

Figure 1.2: Datasets with ground-truth labels, crowdsourced labels, and a model's predicted labels.

subsets, and each of the ten subsets will be treated as a test dataset. While each subset could be small, the union of the ten subsets is large.

- *Advice 2. Cleaning Strategy.* Given a learning algorithm, and a training dataset labeled by the crowd, TARS can tell a data scientist which label in the training dataset should be cleaned such that the new model, re-trained on the cleaned training dataset using the learning algorithm, has the best performance. The size of a training dataset is typically in the range from a few hundred to tens of thousands [1,2]. It is expensive to ask oracles to clean them all. With the help of Advice 2, data scientists are able to apply oracles to clean them *progressively*. For example, suppose there is a small cleaning budget (e.g., 50), but a training dataset is much bigger (e.g., 1000). In this situation, a data scientist can ask for Advice 2 iteratively for 50 times and identify which label should be cleaned at each iteration.

There are some straightforward solutions to the above two problems. In the following, we will use the simple example in Figure 1.2 to illustrate the limitations of these solutions, and then demonstrate the contributions made in this thesis to overcome the limitations.

Let us first consider Advice 1. Figure 1.2(b) shows a test dataset labeled by two crowd workers w_1 and w_2 . For simplicity, we assume that the workers w_1 and w_2 have the same noise rate = 0.2, which means that each of them has a probability of 0.2 to give a different label from the true label. We apply a given model to the test dataset and obtain the predicted label of each instance (see Table (c)). Since ground-truth labels are unknown, the model's true accuracy cannot be directly derived. One naive approach is to treat noisy labels as ground-truth labels and then compute the accuracy based on the noisy labels. However, this approach is biased because it ignores the workers' noise rates. In this example, the accuracies computed based on the noisy labels and the ground-truth labels are $\frac{3}{5}$ and $\frac{4}{5}$, respectively, and the difference is $\frac{1}{5}$.

To overcome the limitation, we present a new estimator that estimates the model's true performance (e.g., accuracy, precision, recall, or F-score) by considering not only noisy labels but also noise rates. We prove that our estimator is unbiased (i.e., in expectation the estimator's estimated value is equal to the true value). We further study how to compute a confidence interval for the esti-

mator in order to bound its estimation error (i.e., bound the difference between the estimated value and the true value). This turns out to be a challenging problem because the estimator’s error comes from two sources: sampling and labeling. We theoretically analyze how each source contributes to the overall estimation error, and show that the contribution of the first (second) source is controlled by test data size N (a noise rate β). Interestingly, they will not be affected by each other: (i) as N increases, the overall estimation error will decrease at a rate of $\mathcal{O}(\frac{1}{\sqrt{N}})$, regardless of what β is, and (ii) as β decreases, the overall estimation error will decrease at a rate of $\mathcal{O}(\frac{1}{|0.5-\beta|})$, regardless of what N is. In other words, the estimation error can be decreased by either increasing test data size or improving worker quality; the above theoretical results show the trade-off of each choice.

Next, let us consider Advice 2. Suppose Figure 1.2(b) represents a training dataset, where w_1 ’s noise rate is 0.4 and w_2 ’s noise rate is 0.01. There are five noisy labels in the dataset, and TARS needs to decide which one should be sent to an oracle to clean. Please note that this problem is different from active learning [57] because active learning assumes that data is unlabeled but here data has been labeled by crowd workers. We need to incorporate label noise into our cleaning strategy otherwise an oracle may clean many instances that have already been correctly labeled. For example, consider the data in Figure 1.2(b), the first four instances have the label noise of 0.4 and the label noise of the last instance is only 0.01. A good cleaning strategy should try to avoid sending the last instance to an oracle because it has a probability of 0.99 to be correct. (In addition to active learning, there are some other cleaning strategies proposed in the literature. Please refer to Section 6.1 for a more detailed discussion.)

To this end, we propose a new cleaning strategy, called expected model improvement (EMI). EMI estimates the expected model improvement of cleaning each noisy label and then selects the noisy label with the largest estimated value to clean. We illustrate the limitations of the existing cleaning strategies and explain why EMI can overcome the limitations. While the idea of EMI sounds promising in theory, we need to address some practical issues. The first issue is which data should be used to estimate the expected model improvement. If we choose the data improperly, EMI may end up training a model that performs well on the training data but not on the test data. The second issue is how to break the tie when two noisy labels have the same expected model improvement. We propose optimization techniques to address the issues and demonstrate their effectiveness experimentally. In the end, we analyze the time complexity of our approach, and study how to improve its efficiency in practice.

Figure 1.2 only shows a simplified version of our problem. In the thesis, we study a more general version of the problem, where an instance can be labeled by multiple workers, a confusion matrix is used to model worker quality, and various metrics such as accuracy, precision, recall, and F-score can be chosen for model evaluation. In summary, our thesis makes the following contributions:

- To the best of our knowledge, we are the first to study how to use an oracle to clean crowdsourced labels for classification. We identify two challenging problems (model evaluation and cleaning strategy) and present the formal problem definitions.

- We propose an estimator that can estimate a model’s true accuracy based on noisy data. We prove that the estimator is unbiased and we compute a confidence interval to bound its estimation error. We also discuss how to extend our solution to other evaluation metrics such as precision, recall, F-score.
- We develop a new cleaning strategy, called EMI, that can effectively decide which label should be cleaned. We explain why EMI is superior to the existing cleaning strategies to solve our problem. We further improve the effectiveness of EMI by developing two optimization techniques, and discuss how to make it run efficiently in practice.
- We evaluate TARS on both simulated and real-world datasets. The results show that (1) TARS can use noisy labels to accurately estimate how well a model will perform w.r.t. ground-truth labels; and (2) TARS can improve the model accuracy by a larger margin than the existing cleaning strategies, for the same cleaning budget.

The remainder of this thesis is organized as follows. Chapter 2 discusses about existing methods in the literature. Chapter 3 formally defines the model evaluation and cleaning strategy problems. Since each instance may be labeled by multiple workers, we introduce how label consolidation works in Chapter 4. After that, we discuss how to solve the model evaluation problem in Chapter 5 and the cleaning strategy problem in Chapter 6. Experimental results are presented in Chapter 7, followed by conclusion and future work (Chapter 8).

Chapter 2

Related Work

Crowdsourcing. There are three research topics in crowdsourcing related to our work: task assignment [5, 10, 19, 28, 44, 51, 74], truth inference [18, 27, 35, 53, 73], and active learning (from crowds) [29, 42, 47, 70]. Task assignment studies how to determine which task should be assigned to an incoming crowd worker. A number of task assignment algorithms have been proposed [5, 44, 74], but their objectives are not to maximize the performance of a classification model. Truth inference studies the problem of inferring the ground-truth label of each instance based on (inconsistent) labels from different workers. There are many interesting ideas proposed to solve this problem [27, 35, 53]. For traditional truth-inference techniques where no oracle-based cleaning is involved [73], this is orthogonal to our problem because we can first use them to get a noisy label along with noise rate for each instance, and then apply TARS to train or test a model using the noisy labels. The oracle-based cleaning was adopted by a recent truth-inference study [35]. However, as shown in the experiment, TARS outperformed this approach because it aims to maximize label quality rather than model quality. Active learning (from crowds) determines which unlabeled instance should be sent to a crowd worker to label. Since a crowd worker may make mistakes, there are some studies on the trade-off between asking another crowd worker to relabel an instance or label a new instance [42, 58]. In our problem, we consider that all the instances have been labeled by the crowd and an oracle can be used to clean the instances.

Machine-based Data Cleaning. There is an extensive study on the methods of data cleaning without human intervention [9, 30, 32, 56]. These methods can be categorized into three types: threshold based, classification filtering and model based. In threshold based methods, outliers can be identified and removed when the measurement for anomaly exceeds a threshold [22, 23, 24, 25, 60]. Predictions of a classifier is used in the classification filtering method to identify mislabeled records in order to remove them from data [26, 37, 61]. And in model based algorithms, noisy instances are detected by their impact on learning models that are sensitive to label noise [13, 15, 31, 45, 49, 55, 68]. These strategies are usually based on heuristics and also suffer from a chicken-or-egg dilemma, where good models are needed for filtering noise and learning with noisy data results in poor classifiers.

Human-guided Data Cleaning. Algorithmic data cleaning approaches have been improving in quality, but still far from perfect [11]. In view of the challenge, human-guided data cleaning has

recently attracted a lot of attention [3, 12, 14, 20, 38, 50, 64, 65, 67, 69, 71]. The existing studies can be broadly divided into two categories. One category is to leverage humans (either crowd workers or experts) to solve a particular data-cleaning problem, such as entity resolution [14, 20, 64, 67], missing value imputation [50], and data repairing [12, 69]. The other category is to clean data for a particular data analysis task, such as building a machine-learning model [38] and answering SQL queries [3, 65]. Our work belongs to the second category. In particular, we are focused on cleaning crowdsourced labels for statistical classification, which is a problem that has not been explored before.

Learning with Noisy Labels. There is a large body of work in the Machine Learning community on learning with noisy labels (see [21] for a survey). Some existing approaches aim to develop a robust algorithm to tolerate label noise [43, 59]. Liu et al. [43] use importance reweighting to ensure that any surrogate loss function can be used for classification, proving that the label noise involved in training will not affect the search for an optimal classifier. Sukhbaatar et al. [59] train neural networks on images with noisy labels directly, but add an extra layer designed to model the label noise. There are also some works [8, 62] that seek to leverage data cleaning for model training. Brodley and Friedl [8] use ensembles of classifiers to identify mislabeled instances and then remove them from the training data. Veit et al. [62] use a small sample of clean data to reduce the severity of label noise on the (much larger) noisy dataset. In contrast to these works, this thesis is focused on a different data-cleaning scenario, i.e., oracle-based label cleaning.

Chapter 3

Problem Definition

We first provide some background knowledge in Section 3.1, and then formally define our problems in Sections 3.2 and 3.3.

3.1 Background

Classification With ground-truth labels. For ease of presentation, we will use binary classification to illustrate our approach, but the proposed approach can be easily extended to multiclass classification (see [17] for more detail about the extension). Let \mathbb{G} be the joint distribution on $(x, y) \in X \times Y$, where x represents an input instance (typically a vector) and $y \in \{-1, +1\}$ represents the true label. Denote a sample drawn i.i.d. from \mathbb{G} as $S = \{(x_i, y_i)\}_{i=1}^N$. A classification algorithm aims to train a *predictive model* (i.e., a *classifier*) f based on S and then make predictions over the unseen instances in \mathbb{G} , where f can be thought of as a decision function that takes an instance x as input and outputs a real value $t = f(x)$. The instance will be classified as +1 if $t > 0$, and -1 otherwise.

Crowdsourced Data. Obtaining ground truth labels can be expensive, but labeling instances with imprecise crowd workers can be cheap. Let $\mathcal{W} = \{w_j\}_{j=1}^K$ denote a set of workers. We assume that each instance x_i is labeled by a subset of k_i workers; let $\mathcal{L}^i = \{(w_j, l_{i,j}) \mid w_j \text{ labels } x_i\}$ denote the corresponding labels, where $l_{i,j} \in \{-1, +1\}$ represents the label given by worker w_j . Let $\mathcal{C} = \{(x_i, \mathcal{L}^i)\}_{i=1}^N$, which we call the *crowdsourced data*.

Worker Model. Since crowd workers are not perfect, existing crowdsourcing work typically uses *confusion matrix* to model worker quality. The confusion matrix of each worker w_j is a 2×2 matrix,

$$q^{(j)} = \begin{bmatrix} q_{-1,-1}^{(j)} & q_{-1,+1}^{(j)} \\ q_{+1,-1}^{(j)} & q_{+1,+1}^{(j)} \end{bmatrix},$$

where each row represents a true label, each column represents a worker's provided label, and $q_{y,l}^{(j)}$ ($y \in \{\pm 1\}, l \in \{\pm 1\}$) means that given an instance with true label y , worker w_j provides label l with probability of $q_{y,l}^{(j)}$. For example, the following illustrates the confusion matrices of a perfect worker w_1 and an imperfect worker w_2 (with the noise rates of 0.2 given $y = -1$ and 0.3 given $y = +1$):

$$q^{(1)} = \begin{bmatrix} 1, & 0 \\ 0, & 1 \end{bmatrix}, \quad q^{(2)} = \begin{bmatrix} 0.8, & 0.2 \\ 0.3, & 0.7 \end{bmatrix}.$$

Prior Probability. Each cell in a confusion matrix is a conditional probability, $q_{y,l}^{(j)} = P(L = l \mid Y = y)$. To compute the joint probability distribution $P(L, Y)$, we also need to know a *prior probability* (or simply called the prior), denoted by $P(Y)$. Intuitively, the prior is the probability of an instance having a label of +1 or -1.

Computing Confusion Matrices and Prior. An important problem is how to compute workers' confusion matrices $q^{(j)}$ (for all $j \in [1, K]$) and the prior $P(Y)$ in practice. This problem has been extensively studied in the crowdsourcing literature [41]. One simple idea is to manually label a small sample of instances upfront, and then mix these instances with other unlabeled instances and ask workers to label them all. Since workers do not know which instances have been pre-labeled, these pre-labeled instances can be used to compute workers' confusion matrices as well as the prior.

Another common idea is to leverage label redundancy. To improve quality, each instance is often labeled by multiple workers. If a worker often provides inconsistent labels with the majority of other workers, then the worker is very likely to be a low-quality worker. Based on this idea, existing work treats confusion matrices and the prior as unknown parameters and adopts an EM algorithm [16] to iteratively estimate their values.

If a worker only labels a few instances, it is hard to estimate its confusion matrix accurately. To avoid this to happen, we can batch multiple instances into a task. In this way, each worker will complete at least a single batch of instances. In fact, this batching idea has been widely adopted by existing crowdsourcing studies to save cost (e.g., [46, 64]). For example, consider the *Restaurant* dataset collected by [64] (see the description of the dataset in Section 7.1). Since the batch size was set to 18, all the workers in this dataset labeled at least 18 instances.

There are certainly many other approaches to improve the computation of workers' confusion matrices and the prior. In this thesis, we treat this as an orthogonal problem, and assume that they are given as input of our problems.

Classification With Noisy Labels. Let \mathcal{A} denote a classification algorithm that learns a classifier f on crowdsourced data \mathcal{C} . This problem has been well studied in the Machine Learning community [21, 52]. Their basic ideas are either to develop a robust classification algorithm to tolerate label noise or to adopt automatic cleaning algorithms to filter/correct noisy labels. In this thesis, we treat \mathcal{A} as a black box, which takes crowdsourced data \mathcal{C} as input and outputs a classifier.

3.2 Advice 1: How Good is a Model?

The first piece of advice from TARS is focused on the *model testing* stage. It considers the situation when a user has already trained a model and wants to evaluate the model's performance using crowdsourced data.

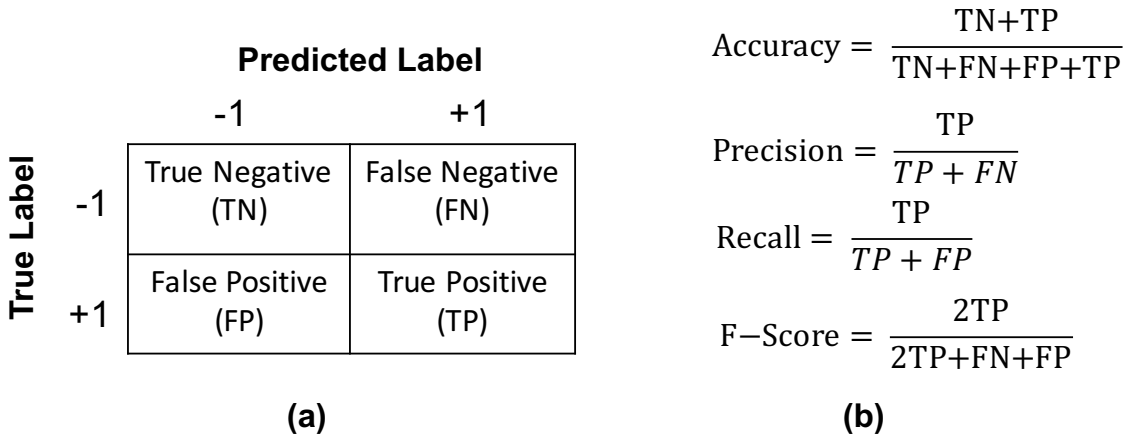


Figure 3.1: An illustration of (a) model’s confusion matrix and (b) four evaluation metrics derived from the confusion matrix.

Evaluation Metrics. To evaluate a model’s performance, people often compute a confusion matrix for the model and then derive different types of evaluation metrics from the matrix. Figure 3.1(a) illustrates a model’s confusion matrix. We can see that it is similar to a worker’s confusion matrix, where each row also represents a true label, but the difference is that each column represents a model’s predicted label rather than a worker’s provided label. The matrix has four cells: *True Positive (TP)*, *False Positive (FP)*, *True Negative (TN)*, and *False Negative (FN)*. TARS aims to estimate the value of each cell. In this way, any evaluation metric computed based on these cells can be derived in a straightforward manner. Figure 3.1(b) shows the definitions of four representative evaluation metrics.

Let eval denote a user-specified evaluation metric. The model’s *true performance* is denoted by $\text{eval}(\mathbb{G}, f)$. For example, suppose eval is accuracy, then $\text{eval}(\mathbb{G}, f)$ computes the f ’s accuracy on \mathbb{G} . Since we do not have access to \mathbb{G} but only crowdsourced data \mathcal{C} , TARS aims to use \mathcal{C} to estimate $\text{eval}(\mathbb{G}, f)$. Let $\widetilde{\text{eval}}(\mathcal{C}, f)$ denote an estimation of $\text{eval}(\mathbb{G}, f)$. We say $\widetilde{\text{eval}}(\mathcal{C}, f)$ to be unbiased if the expected value of the estimation is equal to the true value, i.e., $\mathbb{E}[\widetilde{\text{eval}}(\mathcal{C}, f)] = \text{eval}(\mathbb{G}, f)$.

TARS computes a *confidence interval* to bound the estimation error of $\widetilde{\text{eval}}(\mathcal{C}, f)$. Suppose the estimated value is $\widetilde{\text{eval}}(\mathcal{C}, f) = 0.8$. Given a confidence level (e.g., 95%), a confidence interval (e.g., 0.8 ± 0.01) indicates that the difference between the estimated value and the true value is within ± 0.01 with 95% probability. The wider the confidence interval, the larger the estimation error.

Problem 1 (Model Evaluation). *Given crowdsourced data \mathcal{C} , a model f , and an evaluation metric eval , TARS aims to determine (1) an unbiased estimator, $\widetilde{\text{eval}}(\mathcal{C}, f)$, of the model’s true performance, and (2) a confidence interval, $[\widetilde{\text{eval}}(\mathcal{C}, f) - \epsilon_1, \widetilde{\text{eval}}(\mathcal{C}, f) + \epsilon_2]$, at a given confidence level.*

3.3 Advice 2: Which Label Should be Cleaned?

The second piece of advice that TARS can provide is focused on the *model training* stage. It considers the situation when a user trains a model using noisy data, but the model is not good enough. The user wants to know which instance-label pair should be sent to an *oracle* to clean such that the model's true performance can be improved the most.

Oracle Labeller. Cleaning, or ground truth labeling, can be thought of as querying a perfect worker w_Ω called an *oracle*. It follows that the noise rates for w_Ω are $P(L = -1 | Y = +1) = P(L = +1 | Y = -1) = 0$. We assume that queries are expensive, and thus calls to the oracle are constrained by a budget.

Cleaning Data. Suppose we query an oracle to clean an instance x_i and obtain ground truth label y_i . One could update crowdsourced data \mathcal{C} with this new knowledge by replacing \mathcal{L}^i with $\{(w_\Omega, y_i)\}$. We say that we are *cleaning* instance x_i , because we are substituting a set of imprecise labels with the ground truth label provided by the oracle.

Cleaning Strategy. Recall that a classification algorithm \mathcal{A} takes crowdsourced data \mathcal{C} as input and outputs a model f . By cleaning the labels in \mathcal{C} , we hope \mathcal{A} can produce a better model. It is expensive to query an oracle to get ground truth labels, thus the goal is to strategically choose x_i to clean. The best choice of x_i would result in a dataset which leads to a model more accurate than cleaning any other label. Based on the notation above, we have a formal definition of the problem below:

Problem 2 (Cleaning Strategy). *Given crowdsourced data \mathcal{C} , a classification algorithm \mathcal{A} , and an evaluation metric eval , let f_i be the model resulting from cleaning instance x_i , then training on \mathcal{C} with \mathcal{A} . TARS aims to determine which instance should be cleaned such that the model's true performance can be improved the most:*

$$i^* = \operatorname{argmax}_{i=1,\dots,N} \text{eval}(\mathbb{G}, f_i).$$

Chapter 4

Preliminary: Label Consolidation

Given crowdsourced data, if there are instances in the data having multiple worker labels, TARS will first consolidate these labels into a single label. In other words, after this process, each instance in the crowdsourced data will have a single consolidated label along with a confusion matrix that quantifies the uncertainty of the consolidated label. In this section, we will start by showing why there is a need for this process, and then present how to get the consolidated label as well as the consolidated confusion matrix.

4.1 The Need For Label Consolidation

Both Problems 1 and 2 require estimating the performance of some model which is defined based on the ground-truth labels in G . Rather than ground-truth labels, we only have worker labels, which are noisy. The trick to bridging the gap between worker labels and ground-truth labels is noticing that there are some relationships between them, which are captured by the workers' confusion matrices.

Therefore, we can use worker labels along with workers' confusion matrices to infer the most likely true label for each instance, and then figure out how to quantify the uncertainty of each inferred label. Specifically, given a crowdsourced dataset $\mathcal{C} = \{(x_i, \mathcal{L}^i)\}_{i=1}^N$, we aim to get a new dataset, denoted by $\mathcal{D} = \{(x_i, y'_i, r^{(i)})\}_{i=1}^N$, where y'_i and $r^{(i)}$ represent the inferred label and the uncertainty of the inferred label mentioned above. We will present how to compute y'_i and $r^{(i)}$ in Sections 4.2 and 4.3, respectively.

4.2 Computing Consolidated Labels

Given an instance x_i with worker labels \mathcal{L}^i , the basic idea of getting the x_i 's most likely label is to compare the values of two conditional probabilities: $P(Y_i = +1 \mid \mathcal{L}^i)$ and $P(Y_i = -1 \mid \mathcal{L}^i)$, the probability that instance x_i has a true label of +1 (resp. -1), conditioned on the labels that the workers provide. If the former (latter) is larger, it means that the instance's label is more likely to be +1 (resp. -1).

We use the work of Dawid and Skene [16] to compute the conditional probabilities. Below is an explanation of their approach, adapted to our notation. Assuming that workers provide labels independently of one another, we have:

$$\begin{aligned}
P(Y_i = +1 | \mathcal{L}^i) &= \frac{P(\mathcal{L}^i | Y_i = +1)P(Y_i = +1)}{P(\mathcal{L}^i)} \\
&\propto P(\mathcal{L}^i | Y_i = +1)P(Y_i = +1) \\
&= P(Y_i = +1) \prod_{l_{i,j} \in \mathcal{L}^i} P(L_j = l_{i,j} | Y_i = +1) \\
&= P(Y = +1) \prod_{l_{i,j} \in \mathcal{L}^i} q_{+1, l_{i,j}}^{(j)} \tag{4.1}
\end{aligned}$$

Similarly, we can compute $P(Y_i = -1 | \mathcal{L}^i)$. By comparing their values, we obtain the consolidated label y'_i (if there is a tie, we break the tie randomly).

$$y'_i = \operatorname{argmax}_{\bar{y} \in \{-1, 1\}} P(Y = \bar{y}) \prod_{l_{i,j} \in \mathcal{L}^i} q_{\bar{y}, l_{i,j}}^{(j)} \tag{4.2}$$

Thus, the chance of instance x_i having a ground truth label of, say, +1 is influenced by two main factors. If the instances for which $Y = +1$ are extremely common (i.e. $P(Y = +1)$ is very close to 1), this increases our belief that $Y_i = +1$. Likewise, if the labels that workers provide are likely to happen provided that $Y = +1$ were true (i.e., $q_{+1, l}^k$ is very close to +1), then this also increases our belief that $Y_i = +1$.

4.3 Quantify Consolidated Label's Uncertainty

Suppose an instance x_i is labeled by a group of k_i workers, denoted by \mathcal{W}_i . Let y'_i denote the consolidated label inferred from worker labels using the above method. Like the definition of a worker's confusion matrix, we define the *consolidated confusion matrix* associated with a group of workers as a 2×2 matrix:

$$r^{(\mathcal{W}_i)} = \begin{bmatrix} r_{-1, -1}^{(\mathcal{W}_i)} & r_{-1, +1}^{(\mathcal{W}_i)} \\ r_{+1, -1}^{(\mathcal{W}_i)} & r_{+1, +1}^{(\mathcal{W}_i)} \end{bmatrix},$$

where each row represents a true label, each column represents a consolidated label, and $r_{y, y'}^{(\mathcal{W}_i)}$ ($y \in \{\pm 1\}, y' \in \{\pm 1\}$) means that given an instance with true label y , the consolidated label is y' with probability of $r_{y, y'}^{(\mathcal{W}_i)}$. In other words, each cell in the matrix is a conditional probability $r_{y, y'}^{(\mathcal{W}_i)} = P(Y'_i = y' | Y = y)$ ($y', y \in \{\pm 1\}$). If the context is clear, we abbreviate $r^{(\mathcal{W}_i)}$ as $r^{(i)}$ or r .

We use the Law of Total Probability to compute the conditional probability.

$$P(Y'_i | Y) = \sum_n P(Y'_i | \bar{\mathcal{L}}_n^i, Y) P(\bar{\mathcal{L}}_n^i | Y), \tag{4.3}$$

where $\{\bar{\mathcal{L}}_n^i : n = 1, 2, \dots, 2^{k_i}\}$ represents all combinations of the labels that the group of k_i workers from \mathcal{W}_i can provide. For example, suppose there are two workers, w_1, w_2 . Then, there will be four combinations of worker labels: $\bar{\mathcal{L}}_1^i = \{(w_1, -1), (w_2, -1)\}$, $\bar{\mathcal{L}}_2^i = \{(w_1, -1), (w_2, +1)\}$, $\bar{\mathcal{L}}_3^i = \{(w_1, +1), (w_2, -1)\}$, and $\bar{\mathcal{L}}_4^i = \{(w_1, +1), (w_2, +1)\}$, where, e.g., $\bar{\mathcal{L}}_1^i$ means that w_1 provides -1 and w_2 provides -1.

The equation depends on two forms of conditional probabilities: $P(Y' | \bar{\mathcal{L}}_n^i, Y)$ and $P(\bar{\mathcal{L}}_n^i | Y)$. For the latter, we have already discussed how to compute it in Equation A.1.

$$P(\bar{\mathcal{L}}_n^i | Y) = \prod_{(w_j, l) \in \bar{\mathcal{L}}_n^i} P(\bar{\mathcal{L}}_n^i = (w_j, l) | Y) = \prod_{(w_j, l) \in \bar{\mathcal{L}}_n^i} q_{y, l}^{(j)} \quad (4.4)$$

For the former, since y'_i depends only on $\bar{\mathcal{L}}_n^i$ (i.e., the workers and the labels that they provide), we have that:

$$P(Y' = y' | \bar{\mathcal{L}}_n^i, Y) = \begin{cases} 1 & \text{if } y' = \bar{y} \\ 0 & \text{otherwise} \end{cases} \quad (4.5)$$

where \bar{y} represents the consolidated label inferred from $\bar{\mathcal{L}}_n^i$.

Chapter 5

Model Evaluation

Now we have a noisy dataset $\mathcal{D} = \{(x_i, y'_i, r^{(i)})\}_{i=1}^N$, where each instance is associated with a single noisy label y'_i along with a noise rate $r^{(i)}$. There are a number of challenging problems that need to be addressed: (1) how to develop a unified estimation framework that works for different evaluation metrics (Section 5.1); (2) how to bound the difference between the estimated value and the true value under the new framework (Section 5.2); (3) how to give a quantitative analysis on how each factor (e.g., sample vs. population, noisy labels vs. ground-truth labels) contributes to the bound (Section 5.2.1). In this section, we propose novel solutions to these non-trivial problems.

5.1 Estimating Model's True Performance

We first present our unified estimation framework. Recall that a model's performance (e.g., accuracy, F-score) is determined by its confusion matrix: $\begin{bmatrix} \text{TN} & \text{FN} \\ \text{FP} & \text{TP} \end{bmatrix}$. To estimate a model's performance, the key is to figure out how to estimate the values of the four cells in the confusion matrix. We will use TP as an example to illustrate this estimation process.

Overview. The estimation framework consists of two steps. The first step is to write TP in a form of the sum of *loss*, and the second step is to use an existing approach [48] to estimate the loss.

Step 1. A loss function, denoted by $\text{Loss}(t, y)$, measures the difference between a model's prediction t and a true label y . When ground-truth labels are accessible, we can represent TP as follows:

$$\text{TP} = \sum_{i=1}^N \text{Loss}(t_i, y_i), \quad (5.1)$$

where $\text{Loss}(t_i, y_i) = 1$ if t_i and y_i are both positive; 0, otherwise.

Step 2. In reality, however, we do not have access to ground-truth labels but only noisy labels. Thus, we need to use a noisy label y' along with its noise rate r to estimate the true loss $\text{Loss}(t, y)$. Natarajan et al. [48] proposed an unbiased estimator:

$$\widetilde{\text{Loss}}(t, y') = \frac{(1 - r_{-y', y'}) \cdot \text{Loss}(t, y') - r_{y', -y'} \cdot \text{Loss}(t, -y')}{1 - r_{+1, -1} - r_{-1, +1}}. \quad (5.2)$$

Please note that this estimator works for any bounded loss function. It is defined based on a noisy label y' , thus does not require knowing the true label y .

Unbiased Estimator of TP. By plugging $\widetilde{\text{Loss}}(t, y')$ into Equation 5.1, we obtain an estimator of TP:

$$\widetilde{\text{TP}} = \sum_{i=1}^N \widetilde{\text{Loss}}(t_i, y'_i). \quad (5.3)$$

Since $\widetilde{\text{Loss}}(t, y')$ is unbiased, due to the linearity of expectation, we can easily prove that $\widetilde{\text{TP}}$ is unbiased, i.e., $\mathbb{E}[\widetilde{\text{TP}}] = \text{TP}$.

Unbiased Estimator of TN, FN, FP. Using a similar approach, we can get an unbiased estimator for TN, FN, and FP. The only difference from TP is that they need to choose a different loss function. For example, suppose we want to estimate TN. Then, the loss function w.r.t. TN should be defined as $\text{Loss}(t_i, y_i) = 1$ if t_i and y_i are both negative; 0, otherwise.

Estimating Accuracy, Precision, Recall, F-Score. Now we have known how to estimate each value in the model's confusion matrix. These estimated values can be composed to get the model's performance w.r.t. each evaluation metric that is defined in Figure 3.1. For example, by plugging the estimated values of TP and TN into the accuracy's definition, we can get the estimated value of accuracy: $\frac{\widetilde{\text{TP}} + \widetilde{\text{TN}}}{N}$. Similarly, we can get the estimated values for precision: $\frac{\widetilde{\text{TP}}}{\widetilde{\text{TP}} + \widetilde{\text{FP}}}$, recall:

$$\frac{\widetilde{\text{TP}}}{\widetilde{\text{TP}} + \widetilde{\text{FN}}}, \text{ and F-score: } \frac{2\widetilde{\text{TP}}}{2\widetilde{\text{TP}} + \widetilde{\text{FP}} + \widetilde{\text{FN}}}.$$

For accuracy, we can prove that the estimation is unbiased, i.e.,

$$\mathbb{E}\left[\frac{\widetilde{\text{TP}} + \widetilde{\text{TN}}}{N}\right] = \frac{\mathbb{E}[\widetilde{\text{TP}}] + \mathbb{E}[\widetilde{\text{TN}}]}{N} = \frac{\text{TP} + \text{TN}}{N}.$$

However, for the other three evaluation metrics, their estimators are not unbiased. This is because that when both numerator X and denominator Y are random variables, we do *not* have $\mathbb{E}\left[\frac{X}{Y}\right] = \frac{\mathbb{E}[X]}{\mathbb{E}[Y]}$. This type of estimator is often called conditionally unbiased given Y . Despite that they are not unbiased in theory, we validate their effectiveness in the experiments, and find that they perform very well on both synthetic and real-world datasets.

5.2 Bounding Estimation Error

In this section, we study how to compute confidence intervals for these estimators. This problem is challenging because there are two sources of error involved and a confidence interval has to take both of them into consideration.

Sample vs. Population. The first source of error comes from sampling. Since the entire population is not accessible, our estimator can only look at a sample of data and use it to estimate how well a model will perform over the entire population.

Noisy Labels vs. ground-truth labels. The other source of error comes from noisy labels. Since ground-truth labels are not accessible, our estimator can only look at noisy labels and use them to estimate how well a model will perform w.r.t. ground-truth labels.

To address this challenge, we develop an analytical confidence interval based on the central limit theorem (CLT). From the analytical confidence interval, we can easily see how each source of error contributes to the overall estimation error (i.e., half the width of the confidence interval). However, the analytical confidence interval only works for accuracy. For the other evaluation metrics, we show how to compute their empirical confidence intervals using bootstrapping and conduct experiments to explore the impact of the two sources of error on the overall estimation error in various situations.

5.2.1 Analytical Confidence Interval

We first introduce some background knowledge about CLT, then present an analytical confidence interval for accuracy, and finally dive into the confidence interval to gain more insights.

Central Limit Theorem (CLT). Consider a population with mean μ and variance σ^2 . Given a random sample of size N from the population, $\{X_1, X_2, \dots, X_N\}$, CLT states that the sample mean $\tilde{\mu} = \frac{1}{N} \sum_{i=1}^N X_i$ follows a normal distribution with mean μ and variance $\frac{\sigma^2}{N}$. Note that CLT does not require that the original population has to follow a normal distribution.

Suppose that we treat the sample mean as an estimator of the population mean. Based on CLT, the confidence interval for the estimator is

$$\mu \pm \lambda \sqrt{\frac{\sigma^2}{N}}, \quad (5.4)$$

where λ is a parameter determined by a confidence level (e.g., $\lambda = 1.96$ for 95% confidence interval, $\lambda = 2.58$ for 99% confidence interval). Since the population mean μ and variance σ^2 are unknown, they can be replaced by the estimated mean $\tilde{\mu}$ and variance $\tilde{\sigma}^2$ based on a sample.

Analytical Confidence Interval. The estimator of accuracy can be represented as follows:

$$\text{accuracy} \approx \frac{\widetilde{\text{TP}} + \widetilde{\text{TN}}}{N} = \frac{1}{N} \sum_{i=1}^N \text{Loss}_{0/1}(t_i, y'_i),$$

where $\widetilde{\text{Loss}}_{0/1}(t_i, y'_i)$ is an unbiased estimator of $\text{Loss}_{0/1}(t_i, y_i)$, and $\text{Loss}_{0/1}(t_i, y_i) = 1$ if t_i and y_i have the same sign (i.e., either both positive or both negative); $\text{Loss}_{0/1}(t_i, y_i) = 0$, otherwise. We can see that the estimator is in the form of mean. Let $X_i = \widetilde{\text{Loss}}_{0/1}(t_i, y'_i)$ for each $i \in [1, N]$. The confidence interval for the estimator can be directly derived from Equation 5.5.

$$\mathbb{E}[X] \pm \lambda \sqrt{\frac{\text{var}(X)}{N}} \quad (5.5)$$

In-Depth Analysis. We now provide an in-depth analysis of the confidence interval. As mentioned in the beginning of this section, there are two sources of error. Our analysis aims to answer two questions: (1) *how does sample size affect the confidence interval?* (2) *how does label noise affect the confidence interval?*

For simplicity, we assume that each instance has the same label noise of $r_{-1,+1} = r_{+1,-1} = \beta$. Based on Equation 5.2, we find that $\widetilde{\text{Loss}}_{0/1}(t, y')$ can only take two possible values:

$$\widetilde{\text{Loss}}_{0/1}(t, y') = \begin{cases} \frac{1-\beta}{1-2\beta} & \text{if } t \text{ and } y' \text{ have the same sign} \\ \frac{-\beta}{1-2\beta} & \text{otherwise} \end{cases} \quad (5.6)$$

Suppose that $\widetilde{\text{Loss}}_{0/1}(t, y')$ has a probability of p being $a = \frac{1-\beta}{1-2\beta}$ and $1-p$ being $b = \frac{-\beta}{1-2\beta}$. It is easy to see that $a + b = 1$. The expected value of $\widetilde{\text{Loss}}_{0/1}(t, y')$ is:

$$\mathbb{E}[X] = pa + (1-p)b \quad (5.7)$$

The variance of $\widetilde{\text{Loss}}_{0/1}(t, y')$ is:

$$\begin{aligned} \text{Var}[X] &= \mathbb{E}[X^2] - \mathbb{E}[X]^2 \\ &= pa^2 + (1-p)b^2 - \mathbb{E}[X]^2 \\ &= pa^2 + (1-p)b^2 - (pa + (1-p)b - \mathbb{E}[X]) - \mathbb{E}[X]^2 \\ &= -ab + \mathbb{E}[X] - \mathbb{E}[X]^2 \end{aligned} \quad (5.8)$$

Let θ denote a model's true accuracy. Since $\frac{1}{N} \sum \widetilde{\text{Loss}}_{0/1}(t, y')$ is an unbiased estimator of the true accuracy, then we have $\mathbb{E}[X] = \theta$.

$$\text{Var}[X] = -ab + \theta - \theta^2 \quad (5.9)$$

By plugging Equation 5.9 into Equation 5.5, we obtain a closed form confidence interval of our estimator:

$$\theta \pm \lambda \sqrt{\frac{-ab + \theta - \theta^2}{N}} \quad (5.10)$$

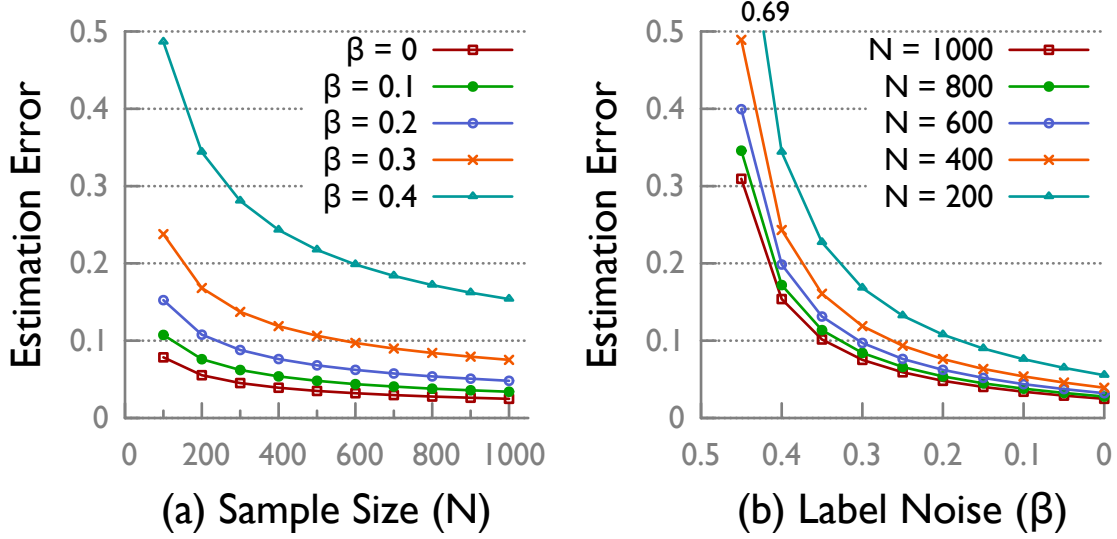


Figure 5.1: The relationships between sample size (N), label noise (β), and estimation error (half the width of the 95% confidence interval), for model accuracy $\theta = 0.8$. (For simplicity, we do not show the estimation error for $\beta \in (0.5, 1]$ because based on Equation 5.10, it will be the same as $1 - \beta$).

where $ab = \frac{-\beta(1-\beta)}{(1-2\beta)^2}$, β is label noise, θ is a model's true accuracy, N is sample size, and λ is constant determined by a confidence level.

From this equation, we can analyze how each source of error contributes to the overall estimation error.

Insight 1. The first source of error (sample vs. population) is controlled by sample size N . It only affects the denominator of the confidence interval. Figure 5.1(a) shows the relationship between sample size and estimation error, for different label noise β . We can see that as sample size N increases, regardless of what β is, estimation error will decrease at a rate of $\mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$. For example, when sample size is increased from $N = 100$ to 1000, estimation error will decrease by about $\mathcal{O}\left(\sqrt{\frac{1000}{100}}\right) = 3$ times.

Insight 2. The second source of error (noisy label vs. true label) is controlled by label noise β . It only affects the numerator of the confidence interval. Figure 5.1(b) demonstrates the relationship between label noise and estimation error, for different sample size N . We can see that as noise decreases, regardless of what N is, estimation error will decrease at a rate of $\mathcal{O}\left(\frac{1}{|0.5-\beta|}\right)$. For example, when β is decreased from $\beta = 0.4$ to $\beta = 0.1$, estimation error will decrease by about $\mathcal{O}\left(\frac{|0.5-0.1|}{|0.5-0.4|}\right) = 4$ times.

5.2.2 Empirical Confidence Interval

We first introduce some background knowledge about bootstrapping, and then present our approach to compute empirical confidence intervals for other estimators than accuracy.

Bootstrapping. Consider a population P and a random sample of the population S . Given an estimator est , suppose the estimator can use the sample to estimate a *complex* parameter of the population. Once the estimate $\text{est}(S)$ is derived, we want to compute its confidence interval. If we could create multiple samples from the population, S_1, S_2, \dots, S_n , then we would get a distribution of the estimates, $\text{est}(S_1), \text{est}(S_2), \dots, \text{est}(S_n)$. Based on the distribution, we can find an interval that covers 95% of the estimates and use it as a 95% confidence interval of the estimation. In practice, however, the population is not available, thus it is impossible to draw another sample from P . The basic idea of bootstrapping is to simulate this procedure by constructing a number of resamples of S . Specifically, to get a resample, it samples the data from S with replacement. Let $S^{(1)}, S^{(2)}, \dots, S^{(n)}$ denote n resamples. Then, it can get a distribution of the estimates based on resamples, $\text{est}(S^{(1)}), \text{est}(S^{(2)}), \dots, \text{est}(S^{(n)})$, from which we can compute a 95% confidence interval.

Empirical Confidence Interval. Now we present how to use the bootstrap to compute empirical confidence intervals for our estimators. Consider a crowdsourced dataset D . We can think of D as a random sample of a population. While it is feasible to get multiple crowdsourced datasets from the population, the monetary cost and the time for doing so can be quite high. For example, suppose each instance needs 0.5 dollar and 5 seconds to label on average. Getting one crowdsourced dataset of size $|D| = 1000$ will cost us \$500 dollars and 1.4 hours. Repeat this for 1000 times will cost us as high as 0.5 million dollars and 58 days.

We use bootstrapping to avoid the need to repeatedly draw samples from the population. Figure 5.2 illustrates the procedure. Given D and an estimator $\widetilde{\text{eval}}(D, f)$, the first step is to construct n resamples of D , denoted by $D^{(1)}, D^{(2)}, \dots, D^{(n)}$. Then, we apply the estimator to each resample to get n estimates,

$$\widetilde{\text{eval}}(D^{(1)}, f), \quad \widetilde{\text{eval}}(D^{(2)}, f), \quad \dots, \quad \widetilde{\text{eval}}(D^{(n)}, f).$$

Given a confidence level (e.g., 95%), let $\widetilde{\text{eval}}_{2.5\%}$ and $\widetilde{\text{eval}}_{97.5\%}$ denote the 2.5th and 97.5th percentile of the distribution, respectively. Then, the 95% confidence interval is denoted by $[\widetilde{\text{eval}}_{2.5\%}, \widetilde{\text{eval}}_{97.5\%}]$.

This approach works for all the estimators developed in Section 5.1 including precision, recall, F-score. Like accuracy, the estimation error of other estimators also come from two sources. We empirically study its relationship with sample size and label noise for these estimators in the experiments.




Figure 5.2: An illustration of using bootstrap to compute empirical confidence intervals for complex estimators.

Chapter 6

Cleaning Strategy

Now we present how TARS provides the second piece of advice: which label should be cleaned? Please note that, unlike the previous section, here we turn our focus to the *model training* stage.

Recall that we are given $S = \{(x_i, y_i)\}_{i=1}^N$ drawn from \mathbb{G} , and aim to make predictions on the dataset $S_{\text{test}} = \{(x_j, y_j)\}_{j=1}^M$, also drawn from \mathbb{G} . We could train a model f on S , predict labels on instances x_j from S_{test} , and compare the predicted labels with each y_j . Ideally, the predicted labels should be “fairly close” to the actual labels S_{test} . One reason why this procedure works is because both S and S_{test} are drawn from \mathbb{G} ; in other words, the data we train on is an acceptable representation of the data we are expected to make predictions on.

Labels in noisy dataset $\mathcal{D} = \{(x_i, y'_i, r^{(i)})\}_{i=1}^N$ are not guaranteed to be correct, so \mathcal{D} might not adequately represent the data we ultimately have to make predictions on. This means if we train directly on the pairs $\{(x_i, y'_i)\}_{i=1}^N$ and are asked to predict the labels from a testing set S_{test} , there could be unacceptably many errors. Therefore, we study the cleaning strategy problem, aiming to choose the “best” instance to clean that would bring the greatest benefit to the resulting model.

In the following, we first explain why the existing cleaning strategies do not work in Section 6.1, and then present the main idea of our cleaning strategy in Section 6.2. We find that the naive implementation of this idea did not work very well in the experiments. We discuss the issues and propose effective solutions in Section 6.3.

6.1 Limitations of Existing Strategies

Below are three classes of existing cleaning strategies, and explanations as to why they are not the perfect solution to our problem.

Active Learning. Typical active learning settings start with a small pool of (cleanly) labeled data, and a large pool of unlabeled data. Similar to our problem setting, an oracle is available to obtain ground truth labels, and the goal is to choose the most informative instances to label under a budget [57]. Active learning literature refers to approaches to this problem as query strategies. There are many query strategies proposed in the literature, such as uncertain sampling [40], expected error reduction [54]. However, since active learning involves labeling *unlabeled* rather than *noisy* data,

Consider an example in Figure 6.2. The above cleaning strategy will select x_i for cleaning because it has the greatest noise. As shown in Figure 6.2(b), even if x_1 's label was flipped, the model would still keep unchanged. In comparison, x_j 's label has a slightly smaller noise rate, but if its label was flipped, the model's decision boundary would move from the left side of x_j to its right side, leading to a big change of the model. Therefore, x_j should have a higher priority to be selected, for this example.

ActiveClean. As we can see, a good cleaning strategy should be aware of both noise rate as well as model changes. A recent paper takes these two factors into consideration and proposes a new cleaning strategy, called ActiveClean [38]. ActiveClean predicts the true label of each instance, and then estimates how cleaning each instance would change the model based on predicated ground-truth labels, and finally selects the instance that would lead to the biggest change. However, ActiveClean has two limitations to solve our problem. First, it does not leverage the given noise rates to predict the true label of each instance. Second, it uses stochastic gradient descent to update the model after cleaning each batch of instances, thus the model's performance may not be very stable for our oracle-based cleaning scenario, where only a small number of instances (e.g., hundreds of instances) can be cleaned.

6.2 Main Idea: Expected Model Improvement

To inform our decision of which instance to clean, we wish to understand how cleaning an instance could improve the current model. Let f be the model trained without cleaning any data. Let f_i be the model trained after cleaning instance x_i . There are two possible cases about f_i .

Case 1: If the label is *not* flipped (i.e., $y_i = y'_i$), the model will stay the same, i.e., $f_i = f$.

Case 2: If the label is flipped (i.e., $y_i = -y'_i$), the model's performance will change by $\text{eval}(\mathbb{G}, f_i) - \text{eval}(\mathbb{G}, f)$.

Please note that cleaning an instance is not always guaranteed to improve the model. For example, in Case 2, if $\text{eval}(\mathbb{G}, f_i) < \text{eval}(\mathbb{G}, f)$, the model's performance will get worse. Since we do not know whether the label would be flipped or not until cleaning the instance, the model's true improvement cannot be obtained. Nevertheless, it could be possible to compute the model's expected improvement based on the probabilities that each case may happen.

Let $P(\text{case}_1)$ and $P(\text{case}_2)$ denote the probabilities that Case 1 and Case 2 happen, respectively. Then, the *expected model improvement* (EMI) is defined as:

$$\begin{aligned} \text{EMI}(i) &= P(\text{case}_1) \cdot 0 + P(\text{case}_2) \cdot (\text{eval}(\mathbb{G}, f_i) - \text{eval}(\mathbb{G}, f)) \\ &= P(\text{case}_2) \cdot (\text{eval}(\mathbb{G}, f_i) - \text{eval}(\mathbb{G}, f)) \end{aligned} \quad (6.1)$$

Our cleaning strategy computes $\text{EMI}(i)$ for each instance x_i , and then selects the instance x_{i^*} with the largest value and sends it to an oracle to clean.

$$i^* = \underset{i}{\operatorname{argmax}} \text{EMI}(i) \quad (6.2)$$

Next, we discuss how to compute $\text{EMI}(i)$, which has three parts:

Computing $P(\text{case}_2)$. $P(\text{case}_2)$ represents the probability that the noisy label is flipped after cleaning. Note that we have already known that the noisy label is y'_i . Thus, $P(\text{case}_2)$ represents the conditional probability of the true label being $-y'_i$ given the noisy label y'_i , i.e., $P(\text{case}_2) = P(Y_i = -y'_i | Y'_i = y'_i)$, where $y'_i \in \{-1, +1\}$ is constant. Based on the Bayes' rule, we can easily derive that

$$\begin{aligned} P(\text{case}_2) &= \frac{P(Y'_i = y'_i | Y_i = -y'_i) \cdot P(Y_i = -y'_i)}{P(Y'_i = y'_i)} \\ &\propto P(Y'_i = y'_i | Y_i = -y'_i) \cdot P(Y_i = -y'_i) \end{aligned} \quad (6.3)$$

$P(Y'_i = y'_i | Y_i = -y'_i)$ is equal to the noise rate $r_{-y'_i, y'_i}^{(i)}$ and $P(Y_i = -y'_i)$ is equal to the prior $P(Y = -y'_i)$. Therefore, we obtain

$$P(\text{case}_2) \propto r_{-y'_i, y'_i}^{(i)} \cdot P(Y = -y'_i) \quad (6.4)$$

Similarly, we can obtain

$$P(\text{case}_1) \propto r_{y'_i, y'_i}^{(i)} \cdot P(Y = y'_i) \quad (6.5)$$

Since $P(\text{case}_1) + P(\text{case}_2) = 1$, we have

$$P(\text{case}_2) = \frac{r_{-y'_i, y'_i}^{(i)} \cdot P(Y = -y'_i)}{r_{-y'_i, y'_i}^{(i)} \cdot P(Y = -y'_i) + r_{y'_i, y'_i}^{(i)} \cdot P(Y = y'_i)} \quad (6.6)$$

Computing $\text{eval}(\mathbb{G}, f)$. Since \mathbb{G} is not available, we cannot compute $\text{eval}(\mathbb{G}, f)$ directly. Fortunately, in Section 5.1, we have discussed a way to estimate it based on noisy data \mathcal{D} , which is accessible. Thus, we can use $\widetilde{\text{eval}}(\mathcal{D}, f)$ to approximate $\text{eval}(\mathbb{G}, f)$.

Computing $\text{eval}(\mathbb{G}, f_i)$. If we knew f_i , $\text{eval}(\mathbb{G}, f_i)$ could be estimated similarly as above. Recall that f_i denotes the resulting model from training on D after cleaning instance x_i . If x_i 's label is not flipped, we do not need to consider this case because the model stays the same as f ; if x_i 's label is flipped, we can retrain a model f_i on the new dataset \mathcal{D}_i , where \mathcal{D}_i represents the dataset resulting from flipping the label of instance x_i of \mathcal{D} .

Remarks. EMI incorporates both noise rates and model changes, thus overcomes the limitations of the existing cleaning strategies. For example, consider x_i in Figure 6.1. Since it has a small noise rate, leading to a small value of $P(\text{case}_2)$, EMI tends to not select x_i . Consider x_i in Fig-

ure 6.2. Since the model would not change after flipping the label of x_i , leading to a zero value of $\text{eval}(\mathbb{G}, f_i) - \text{eval}(\mathbb{G}, f)$, EMI will not select x_i .

Time Complexity. Let $N = |\mathcal{D}|$ denote the size of training data. EMI needs to retrain N models for each iteration. Let $\text{cost}_{\text{train}}$ denote the average cost of training a model on \mathcal{D} . The cost of each iteration is $\mathcal{O}(N \cdot \text{cost}_{\text{train}})$. Suppose there are t iterations. Then, the total cost of our approach will be $\mathcal{O}(t \cdot N \cdot \text{cost}_{\text{train}})$. Based on our observation of real-world datasets [1,2], the training data size N is typically in the range from a few hundred to tens of thousands. Without improving the efficiency, it is expensive to apply our approach to a training dataset with thousands of instances.

In practice, there are two simple but effective ideas that can be used to improve the efficiency.

(1) *Pruning.* Rather than retrain N models for every iteration, we can prune $p\%$ of the iterations, where $p\%$ is a parameter to balance the trade-off between efficiency and effectiveness. For example, 50% means that we prune 50% of the iterations and retrain models for every other iteration. This pruning idea works well in practice because if $\text{EMI}(i)$ is high in one iteration, since the noise rate keeps unchanged, its value tends to be high in later iterations. As a result, we can use the value of $\text{EMI}(i)$ from a previous iteration to approximate its value in a later iteration. For the above example, if $p\% = 90\%$, then the cost will be reduced by $10\times$.

(2) *Parallel Computing.* At each iteration, we can retrain N models in parallel. Recall that these models are trained on $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_N$, where \mathcal{D}_i only has one instance different from \mathcal{D} for all $i \in [1, N]$. We keep one copy of training data \mathcal{D} in memory. When training a model on \mathcal{D}_i , we read \mathcal{D} and flip the label of the i -th instance on-the-fly. In this way, the training data \mathcal{D} is read-only, and there is no coordination when training multiple models in parallel. For the above example, suppose CPU has 16 cores. The total cost will be reduced by $16\times$. In addition to parallel computing, we can leverage distributed computing frameworks (e.g., Spark) or GPU computing to further reduce the running time. In this thesis, we treat this as an orthogonal problem and defer additional exploration to future work.

6.3 Further Optimization

We find that the naive implementation of EMI did not perform very well in the experiments. The general reason might be the usage of simple models with simple decision boundary. Because we are cleaning only one instance in each iteration, this instance does not have a big impact on the model and the EMI for a set of instances might be very similar. Another issue might be related to train/test split and the fact that choosing the same data for training and testing might cause overfitting in machine learning models. We discuss the reasons that cause the problem and propose effective techniques to optimize EMI.

Splitting the noisy data The first reason is related to which noisy dataset should be used to estimate $\text{eval}(\mathbb{G}, f_i)$ and $\text{eval}(\mathbb{G}, f)$. One natural idea is to use \mathcal{D}_i for $\text{eval}(\mathbb{G}, f_i)$ and \mathcal{D} for $\text{eval}(\mathbb{G}, f)$

because f_i is trained on \mathcal{D}_i and f is trained on \mathcal{D} :

$$\text{eval}(\mathbb{G}, f_i) \approx \widetilde{\text{eval}}(\mathcal{D}_i, f_i), \quad \text{eval}(\mathbb{G}, f) \approx \widetilde{\text{eval}}(\mathcal{D}, f).$$

However, there are two issues about this idea.

First, as shown in Equation A.6, the goal is to estimate the difference between $\text{eval}(\mathbb{G}, f_i)$ and $\text{eval}(\mathbb{G}, f)$ as more accurate as possible. Let X and Y denote the estimators of $\text{eval}(\mathbb{G}, f_i)$ and $\text{eval}(\mathbb{G}, f)$, respectively. That is, we aim to minimize $\text{var}(X - Y) = \text{var}(X) + \text{var}(Y) - \text{cov}(X, Y)$. In order to minimize $\text{var}(X - Y)$, we need to increase $\text{cov}(X, Y)$ as more as possible, i.e., making X and Y as more correlated as possible. If X and Y are estimated based on the same noisy data, it will make them much more correlated than be estimated on two different ones. Another issue is about overfitting. If we train a model on a dataset and then use the same dataset to evaluate it, the model may suffer from overfitting. In other words, the model may perform well on the current dataset, but not learn to generalize to unseen data.

To address these issues, we split \mathcal{D} into a training dataset $\mathcal{D}_{\text{train}}$ and a validation dataset \mathcal{D}_{vdn} . Only the instances in $\mathcal{D}_{\text{train}}$ can be cleaned and be used to train a model; the instances in \mathcal{D}_{vdn} cannot be cleaned or train a model, and their job is to estimate $\text{eval}(\mathbb{G}, f_i)$ and $\text{eval}(\mathbb{G}, f)$:

$$\text{eval}(\mathbb{G}, f_i) \approx \widetilde{\text{eval}}(\mathcal{D}_{\text{vdn}}, f_i), \quad \text{eval}(\mathbb{G}, f) \approx \widetilde{\text{eval}}(\mathcal{D}_{\text{vdn}}, f).$$

It is worth noting that this idea has been widely adopted in machine learning, where a validation set is often used for hyperparameter tuning and has shown to be very effective to avoid overfitting.

Weighing with Model Uncertainty Even after splitting noisy data into validation and training sets, another challenge that remains is that the values for $\text{EMI}(i)$ for different instances i can be very similar. To see why this is the case, consider a simple situation where $\text{eval}(\mathbb{G}, f)$ measures the percent of instances classified correctly by f , and the data is labeled by a single worker. Then, the $r^{(i)}$ constant across all instances i , as is $P(\text{case}_2)$.

This means if we have two instances i_1 and i_2 , for which $\text{eval}(\mathbb{G}, f_{i_1}) - \text{eval}(\mathbb{G}, f)$ and $\text{eval}(\mathbb{G}, f_{i_2}) - \text{eval}(\mathbb{G}, f)$ are very similar, the resulting values $\text{EMI}(i_1)$ and $\text{EMI}(i_2)$ will be very similar. In the worst case, $\text{EMI}(i_1) = \text{EMI}(i_2)$, which makes it impossible to distinguish which instance would be a better candidate for cleaning.

To address this issue, we combine EMI with the uncertainty of model f . More specifically, let $u(x_i) = 1 - P(f(x_i) | x_i)$ measure the uncertainty of f 's prediction on x_i . For example, we can interpret small $P(f(x_i) | x_i)$ as a “less confident” prediction, which corresponds to large $u(x_i)$.

Intuitively, if we have two instances whose EMI values are similar, we would like to defer the decision of “which is better” (i.e. which instance is better to clean) to the model’s uncertainty. In this case, the instance for which f more uncertain (i.e. larger $u(x_i)$) should be considered a better candidate for cleaning.

At first glance, it might be tempting to compare instances using $u(x_i) \cdot \text{EMI}(i)$. Suppose x_{i^*} is the best instance to clean. Since $\text{EMI}(i^*)$ is computed using estimators, it's possible that $\text{EMI}(i^*) < 0$. Furthermore, if $u(x_{i^*})$ is large, then the product $u(x_{i^*}) \cdot \text{EMI}(i^*)$ could be very negative, which would rank x_{i^*} below other instances.

Before multiplying by $u(x_i)$, we need to transform $\text{EMI}(i)$ into a positive value, using some function $\sigma: \mathbb{R} \rightarrow \mathbb{R}^+$. In order to preserve the relative ordering of $\text{EMI}(i)$, σ needs to be monotonically increasing. A convenient choice is the sigmoid function $\sigma(t) = \frac{1}{1+e^{-t}}$. To weigh EMI with model uncertainty, we compute:

$$\text{MU}(i) = u(x_i) \cdot \sigma(\text{EMI}(i)), \quad (6.7)$$

Thus, with this optimization, we choose instance i^* to clean by computing:

$$i^* = \underset{i}{\operatorname{argmax}} \text{MU}(i) \quad (6.8)$$

Chapter 7

Experiments

We conduct extensive experiments to evaluate the effectiveness of TARS on synthetic and real-world datasets with simulated noisy labels and crowdsourced noisy labels. The experiments aim to answer five questions. (1) Can we accurately estimate a model’s true performance from noisy labels? (2) How does the estimation error change by varying different parameters (e.g., sample size, noise rates)? (3) Are the proposed optimization techniques effective for EMI? (4) How does the optimized EMI perform compared to the existing cleaning strategies?

7.1 Experimental Settings

Datasets. We used a synthetic dataset and four real-world datasets to evaluate our method. (1) Gaussian contains instances randomly drawn from two different 2D Gaussian distributions with the parameters of $(x_1, y_1) = (0, 0)$ and $\sigma_{x_1} = \sigma_{y_1} = 1$, and $(x_2, y_2) = (1, 0)$ and $\sigma_{x_2} = \sigma_{y_2} = 1$. Gaussian (big) is generated using the same process but the size is ten times larger. (2) Heart, Diabetes, and Cancer are three real-world datasets downloaded from the UCI Machine Learning Repository¹. They are widely used to evaluate classification algorithms in the Machine Learning community. The Heart dataset contains 13 attributes such as age, sex, cholesterol level and etc. while the target value represents if the patient has any sort of heart disease. The study resulted in

Table 7.1: Dataset statistics (SS: synthetic data with simulated noisy labels; RS: real-world data with simulated noisy labels; RR: real-world data with real-world crowdsourced labels)

Dataset	Size	Positive%	#Dimension	Type
Gaussian	1000	50%	2	SS
Gaussian (big)	10000	50%	2	SS
Heart	270	44%	13	RS
Cancer	699	31%	9	RS
Diabetes	768	65%	20	RS
Restaurant	2004	5%	8	RR

¹<http://archive.ics.uci.edu/ml/index.php>

the Cancer dataset was conducted on attributes of breast cancer cells to identify its class as benign or malignant. Diabetes dataset was acquired from diabetes patient records using an electronic device recording with fixed timestamps and paper records capturing logical time slots. (3) Restaurant is a real-world dataset widely used to evaluate entity resolution [64]. Table 7.1 illustrates the detailed statistical information of the six data sets.

Noisy Labels. The noisy labels for the Gaussian, Heart, German, and Cancer datasets were randomly generated, controlled by two parameters, $r_{-1,+1}$ and $r_{+1,-1}$. For example, given $r_{-1,+1} = 0.2$ and $r_{+1,-1} = 0.1$, to generate the noisy labels for a dataset, we do the following for each instance. If the instance’s true label is -1, then it will be flipped with a probability of 0.2; if its true label is +1, then it will be flipped with a probability of 0.1. When an instance needs to be labeled by multiple workers, we will apply this process to generate multiple noisy labels for the instance and get its consolidated label using the method described in Chapter 4.

The noisy labels for the Restaurant dataset were collected from the real crowd in Amazon Mechanical Turk (AMT). Each instance was labeled by a single worker, and the entire dataset was labeled by 24 different workers in total. The noise rates of the workers were in the ranges of $r_{-1,+1} \in [0, 0.4]$ and $r_{+1,-1} \in [0, 0.2]$.

Cleaning Strategies. We compared TARS with six existing cleaning strategies.

- Random cleans the instance randomly selected from the uncleaned instances.
- SortNoise cleans the instance whose noisy label is most likely to be wrong, without considering the impact on the current model.
- ActiveClean cleans the instance which, if the instance was cleaned, would impart the greatest change to the current model.
- Uncertainty Sampling (USample) [39] cleans the instance that the current model is most uncertain about, without considering the noise rate.
- Expected Error Reduction (ExpectError) [54] cleans the instance such that the current model’s error can be reduced the most, where the error is computed on noisy labels rather than estimated w.r.t. ground-truth labels. This is similar (in spirit) to our strategy.
- Hung [35] is a state-of-the-art truth-inference approach. Unlike other approaches, it allows an oracle to validate crowd answers. However, its strategy is to clean the instance in order to maximize label quality rather than model quality.

All the code was written in Python 2.7. We trained logistic regression models on all datasets using scikit-learn². Each dataset was randomly divided into a training set and a test set with the ratio of 2 to 1. The experiments were run on a Windows machine with an Intel Core 8 i7-6700 3.40GHz processor, 16GB of RAM, and 1TB HDD.

²<http://scikit-learn.org/>

7.2 Evaluation of Advice 1

In this section, we first conduct sensitivity analysis on the Advice 1 provided by TARS in order to gain a deep understanding of its performance, and then examine its performance on real-world datasets.

7.2.1 Sensitivity Analysis

We evaluate the estimation error of TARS on the Gaussian dataset by varying the sample size, the noise rate, the number of votes, and the percentage of good workers which are widely used in crowdsourcing literature [4, 6, 7, 34, 36, 63, 66, 72] for evaluating effect of different attributes of noisy data on learning. For performance metrics we use accuracy, precision, recall, and F-score which are also very common in capturing performance of applied machine learning models as a way of quantifying inter-rater reliability. The machine learning community is vastly using these metrics especially in information retrieval researches mainly because they do not need counting of negative cases [33]. When varying one parameter, we set the other parameters with their default values. By default, the sample size is 1000, the noise rate is 0.2, the number of votes is 1, and the percentage of good workers is 0%. We define the estimation error of TARS as half the width of the 95% confidence interval of its estimated value. We compared TARS with DirtyEval, which is a naive estimator presented in the Introduction section. It simply treats the noisy labels as the ground-truth labels without considering noise rates, thus leading to biased estimated results.

Sample Size. Figure 7.1 compares the estimation error of TARS and DirtyEval by varying the dataset size from 100 to 1000, w.r.t. different evaluation metrics. We can see that as the sample size was increased by 10 times, the estimation error of TARS was reduced by about 3 times, but DirtyEval did not change so much. The reason is that TARS takes into account noise rates, which converts label error (i.e., bias) into sampling error (i.e., variance). Increasing the sample size can reduce the sampling error at the rate of $\mathcal{O}(\frac{1}{\sqrt{N}})$. In comparison, DirtyEval has a much larger label error but with a smaller sampling error. While the sampling error can still be decreased, since the overall estimation error was dominated by the label error, the improvement was marginal.

Noise Rate. Figure 7.2 compares the estimation error of TARS and DirtyEval by varying the noise rate from 0.45 to 0.1, w.r.t. different evaluation metrics. We can see that TARS follows a similar trend for all the figures: as the decrease of the noise rate, the estimation error will first decrease dramatically, quickly reaching an estimation error of less than 0.1 at the noise rate of about 0.35. After that, the decreasing speed tends to get slower. Our theoretical analysis in Section 5.2.1 has shown that for accuracy, the estimation error decreases at the rate of $\mathcal{O}(\frac{1}{|0.5-\beta|})$. This experiment validated that it holds for the other evaluation metrics empirically.

Number of Votes. We investigate how adding the number of votes will affect the estimation error. We varied the number of votes from 1 to 15. Figure 7.3 reports the comparison results between TARS and DirtyEval w.r.t. different evaluation metrics. We can see that increasing the number of votes reduced the estimation error exponentially. This is because that as the increase of the

number of votes, the noise rate of a consolidated label will decrease exponentially. We also see that TARS outperformed DirtyEval for not only a single-vote situation but also multiple-vote situations. Eventually, their estimation error will both converge to zero.

Percentage of Good Workers. In a real-world crowdsourcing setting, there could be a mix of high-quality and low-quality workers. We designed an experiment aiming to examine whether TARS can still achieve good performance in this situation. We consider that a “good” worker has the noise rate of 0.9, and a “bad” worker has the noise rate of 0.4. We varied the percentage of the instances labeled by good workers. Figure 7.4 reports the result. We can see that TARS achieved much smaller estimation error than DirtyEval for all situations. For example, when 50% of the instances were labeled by good workers, TARS got an estimation error of about 0.1 (w.r.t. accuracy and F-score), while the estimation error of DirtyEval was more than $2\times$ larger.

7.2.2 Performance on Real-world Datasets

We evaluate TARS on the four real-world datasets. We used the noise rate of 0.2 for Heart, German, and Cancer, and used the real-world crowd workers to label Restaurant. We aim to answer two questions in this experiment: (i) can TARS get an accurate estimate of the model’s true performance on real-world datasets? and (ii) can TARS bound the difference between the estimated value and the true value?

We used TARS and DirtyEval to estimate the model’s true accuracy on the Heart, German, and Cancer datasets. We estimated F-score for the Restaurant dataset because the class labels of the dataset are highly imbalanced. We computed the true accuracy (f-score) using the ground-truth labels. Note that since the entire clean population \mathbb{G} is not available, we can only compute the true accuracy (F-score) based on the clean sample S (i.e., the labeled datasets). Figure 7.5 shows the result. We can see that on the Diabetes and Cancer datasets, TARS returned almost the same accuracy as the true accuracy. On the Heart dataset, although the TARS’s performance was not as good as Diabetes’s or Cancer’s, it was still better than DirtyEval. On the Restaurant dataset, TARS returned an estimated value of F-score with the error more than $2\times$ smaller than DirtyEval. More importantly, on all the datasets, TARS was able to bound the estimation error. Being able to do so is essential for enabling reliable decisions in the real world.

7.3 Evaluation of Advice 2

In this section, we first examine how the proposed optimization techniques can improve the effectiveness of the naive implementation of EMI, and then compared TARS (i.e., EMI with all the optimization techniques) with the state-of-the-art cleaning strategies.

7.3.1 Optimization Techniques

In Section 6.3, we identified the possible issues when applying EMI in practice, and proposed two optimization techniques, denoted by `op1` and `op2`, where `op1` represents the optimization of split-

ting the noisy data and `op1` represents the optimization of weighing with model uncertainty. Figure 7.6 compares the three variants of EMI on the Gaussian dataset in various settings of noise rates.

In Figure 7.6(a), we set the noise rates to $r_{-1,+1} = 0.4$ and $r_{+1,-1} = 0.2$. We can see that the noisy labels had a significant negative impact on the model. After cleaning 100 instances (10% of the data), EMI +`op1`+`op2` improved the model’s accuracy from 0.77 to 0.94. However, the other two cleaning strategies did not help to improve the model so much.

In Figure 7.6(b), we set the noise rates to $r_{-1,+1} = 0.2$ and $r_{+1,-1} = 0.2$. In this setting, we can see that the model’s accuracy was (almost) not affected by the noisy labels. Therefore, all three cleaning strategies started with a very accurate model. After sending some instances to an oracle to clean, EMI and EMI +`op1` sometimes led to a much worse model, but EMI +`op1`+`op2` avoided this kind of situation to happen.

In Figure 7.6(c) and (d), we evaluated the optimization techniques on a mix of noise rates, where the former has 50% of the instances with (0.4, 0) and 50% with (0.2, 0.2); the latter has 90% with (0.4, 0) and 10% with (0.2, 0.2). We can see that in both settings, EMI +`op1`+`op2` outperformed the other two variants, further validating the effectiveness of the proposed optimization techniques.

Since EMI +`op1`+`op2` achieved the best performance, we will only use it to compare with the existing cleaning strategies.

7.3.2 Efficiency Study

We proposed two techniques, parallel computing and pruning, in Section 6.2 to optimize the efficiency of our cleaning strategy. The result for the use of parallel computing is obvious, which will improve the efficiency by c times without hurting the effectiveness, where c is the number of CPU cores. So, we only show the results for the pruning technique. We used the Gaussian (big) dataset and set the noise rates to (0.4, 0). Let $\text{TARS}(p\%)$ denote our cleaning strategy with the pruning percentage of $p\%$. For example, $\text{TARS}(0\%)$ means that no pruning is adopted; $\text{TARS}(50\%)$ means that 50% of the iterations are pruned; $\text{TARS}(100\%)$ means that all the iterations (except for the first one) are pruned.

We varied the data size from 2000 to 10,000, and calculated the average time that each approach spent over 1000 iterations. Figure 7.7(a) shows the result. We can see that the pruning technique can significantly improve the efficiency. For example, without pruning, $\text{TARS}(0\%)$ spent about 13 seconds in deciding which instance to clean; with 80% pruning, $\text{TARS}(0\%)$ reduced the average time to less than 4 seconds. Next, we evaluate the impact of the pruning technique on the effectiveness of TARS. Figure 7.7(b) shows the result. We can see that using pruning only had a slightly negative impact on the accuracy. This is because that when data is large, cleaning a few instances will not change the current model a lot, thus there is no need to retrain all the models for every iteration.

7.3.3 Cleaning Strategies

We compare the effectiveness of different cleaning strategies. We constructed a mix of noise rates: $(0.4, 0.1)$ and $(0.3, 0.4)$, and applied the first noise rate to two thirds of the instances, and applied the second one to the remaining instances. We varied the number of cleaned instances, and compared the model's accuracy w.r.t. different cleaning strategies on the Gaussian, Heart, German, and Cancer datasets.

Figure 7.8 shows the results. We have two interesting observations from the results. First, TARS outperformed Random, SortNoise, USample, ExpectError, Hung on all four datasets. The reason for this result is that TARS considers both label noise and model changes in its cleaning strategy while Random does not consider either of them, SortNoise and Hung do not consider model changes, and USample and ExpectError do not consider label noise. Second, we find that the performance of ActiveClean sometimes was not very stable (see Figure 7.8(b) and Figure 7.8(c)). This is because that ActiveClean is focused on a different cleaning scenario, where there is a large dirty dataset (both features and labels are dirty), and an oracle cleans the data in batches. It is common that each batch contains at least 100 instances, but in this experiment, the entire cleaning budget was only 100.




Figure 7.1: Comparison of the estimation error between TARS and DirtyEval by varying sample size (Gaussian).




Figure 7.2: Comparison of the estimation error between TARS and DirtyEval by varying noise rates (Gaussian).




Figure 7.3: Comparison of the estimation error between TARS and DirtyEval by varying the number of votes per instance (Gaussian).




Figure 7.4: Comparison of the estimation error between TARS and DirtyEval by varying the percentage of good workers (Gaussian).




Figure 7.5: Comparison of the estimation error of TARS and DirtyEval on real-world datasets.




Figure 7.6: Evaluating the effectiveness of the proposed optimization techniques of the TARS’s cleaning strategy with different noise rates (Gaussian).



Figure 7.7: Evaluating the impact of the pruning technique on the efficiency and effectiveness of TARS.




Figure 7.8: Comparing the effectiveness of different cleaning strategies by varying the number of cleaned instances.

Chapter 8

Conclusion

In this thesis, we have studied the problem of cleaning crowdsourced labels using oracles for statistical classification. We developed TARS, a label-cleaning advisor that can provide data scientists with two pieces of advice when they need to train or test a model using noisy labels. We formally defined the corresponding problems: model evaluation and cleaning strategy. For the first problem, we described effective techniques to estimate the model’s true performance as well as bound the estimation error, for different evaluation metrics (accuracy, precision, recall, F-score). For the second problem, we devised a new cleaning strategy, called EMI, to overcome the limitations of the existing cleaning strategies. We developed two techniques to further optimize its effectiveness and proposed two ideas to improve its efficiency. The experimental results show that (1) TARS can accurately estimate the model’s true performance, with the estimation error up to $3\times$ smaller than DirtyEval; (2) TARS can improve the model accuracy by a larger margin than the state-of-the-art cleaning strategies, for the same cleaning budget.

Chapter 9

Future Work

We assumed that we already have the workers' confusion matrices/the prior as input and we considered this as an orthogonal problem. A future question might consider the effect of using many approaches to solve this problem, such as EM algorithm, on the performance of the estimator and query strategy.

Another improvement that can be considered is performance evaluation of different types of machine learning models including noise tolerant and robust to noise models using bigger crowd-sourced and real datasets. More comprehensive experiments are also needed in order to compare with existing methods including methods that identify mislabeled data instances and evaluate the model accuracy over the remaining data and methods that adopt automatic cleaning algorithms to filter/correct noisy labels.

In our study a uniform probability distribution of making label noises by a worker is assumed, which may not be true in practice because some instances difficult to label obviously may lead to a higher probability of returning noisy labels by a worker. Solving this problem along with introducing unbiased estimator for f-measure metric is on the agenda for further improvement of the work.

Bibliography

- [1] Crowdfunder: Data For Everyone. <https://www.figure-eight.com/data-for-everyone/>. Accessed: 2018-05-29.
- [2] UCI Machine Learning Repository: Classification Datasets. <https://archive.ics.uci.edu/ml/datasets.html?format=&task=cla>. Accessed: 2018-05-29.
- [3] Moria Bergman, Tova Milo, Slava Novgorodov, and Wang Chiew Tan. Query-oriented data cleaning with oracles. In *ACM SIGMOD*, pages 1199–1214, 2015.
- [4] Avrim Blum and Shuchi Chawla. Learning from labeled and unlabeled data using graph min-cuts. 2001.
- [5] Rubi Boim, Ohad Greenspan, Tova Milo, Slava Novgorodov, Neoklis Polyzotis, and Wang Chiew Tan. Asking the right questions in crowd data sourcing. In *IEEE (ICDE)*.
- [6] Ioannis Boutsis and Vana Kalogeraki. On task assignment for real-time reliable crowdsourcing.
- [7] Jonathan Bragg, Daniel S Weld, et al. Crowdsourcing multi-label classification for taxonomy creation. In *First AAAI conference on human computation and crowdsourcing*, 2013.
- [8] Carla E. Brodley and Mark A. Friedl. Identifying mislabeled training data. *J. Artif. Intell. Res.*, 11:131–167, 1999.
- [9] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3):15, 2009.
- [10] Peng Cheng, Xiang Lian, Lei Chen, Jinsong Han, and Jizhong Zhao. Task assignment on multi-skill oriented spatial crowdsourcing. *IEEE Trans. Knowl. Data Eng.*, 28(8):2201–2215, 2016.
- [11] Xu Chu, Ihab F. Ilyas, Sanjay Krishnan, and Jiannan Wang. Data cleaning: Overview and emerging challenges. In *SIGMOD*, pages 2201–2206, 2016.
- [12] Xu Chu, John Morcos, Ihab F. Ilyas, Mourad Ouzzani, Paolo Papotti, Nan Tang, and Yin Ye. KATARA: A data cleaning system powered by knowledge bases and crowdsourcing. In *ACM SIGMOD*, pages 1247–1261, 2015.
- [13] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.

- [14] Sanjib Das, Paul Suganthan G. C., AnHai Doan, Jeffrey F. Naughton, Ganesh Krishnan, Rohit Deep, Esteban Arcaute, Vijay Raghavendra, and Youngchoon Park. Falcon: Scaling up hands-off crowdsourced entity matching to build cloud services. In *ACM SIGMOD*, pages 1431–1446, 2017.
- [15] Belur V Dasarathy. Nearest neighbor ($\{NN\}$) norms: $\{NN\}$ pattern classification techniques. 1991.
- [16] Alexander Philip Dawid and Allan M Skene. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Applied statistics*, pages 20–28, 1979.
- [17] Mohamad Dolatshah, Mathew Teoh, Jiannan Wang, and Jan Pei. Cleaning crowdsourced labels using oracles for statistical classification. <http://www.cs.sfu.ca/~jnwang/papers/TR-TAR.pdf>.
- [18] Xin Luna Dong and Felix Naumann. Data fusion - resolving data conflicts for integration. *PVLDB*, 2(2):1654–1655, 2009.
- [19] Ju Fan, Guoliang Li, Beng Chin Ooi, Kian-Lee Tan, and Jianhua Feng. iCrowd: An adaptive crowdsourcing framework. In *ACM SIGMOD*, pages 1015–1030, 2015.
- [20] Donatella Firmani, Barna Saha, and Divesh Srivastava. Online entity resolution using an oracle. *PVLDB*, 9(5):384–395, 2016.
- [21] Benoît Frénay and Michel Verleysen. Classification in the presence of label noise: A survey. *IEEE Trans. Neural Netw. Learning Syst.*, 25(5):845–869, 2014.
- [22] Dragan Gamberger and Nada Lavrač. Noise detection and elimination applied to noise handling in a krk chess endgame. In *International Conference on Inductive Logic Programming*, pages 72–88. Springer, 1996.
- [23] Dragan Gamberger and Nada Lavrač. Conditions for occam’s razor applicability and noise elimination. In *European Conference on Machine Learning*, pages 108–123. Springer, 1997.
- [24] Dragan Gamberger, Nada Lavrač, and Sašo Džeroski. Noise elimination in inductive concept learning: A case study in medical diagnosis. In *International Workshop on Algorithmic Learning Theory*, pages 199–212. Springer, 1996.
- [25] Dragan Gamberger, Nada Lavrac, and Saso Dzeroski. Noise detection and elimination in data preprocessing: experiments in medical domains. *Applied Artificial Intelligence*, 14(2):205–223, 2000.
- [26] Dragan Gamberger, Nada Lavrac, and Ciril Groselj. Experiments with noise filtering in a medical domain. In *ICML*, pages 143–151, 1999.
- [27] Jing Gao, Qi Li, Bo Zhao, Wei Fan, and Jiawei Han. Truth discovery and crowdsourcing aggregation: A unified perspective. *PVLDB*, 8(12):2048–2049, 2015.
- [28] Jinyang Gao, Xuan Liu, Beng Chin Ooi, Haixun Wang, and Gang Chen. An online cost sensitive decision-making method in crowdsourcing systems. In *ACM SIGMOD*, pages 217–228, 2013.

- [29] Daniel Haas, Jiannan Wang, Eugene Wu, and Michael J. Franklin. Clamshell: Speeding up crowds for low-latency data labeling. *PVLDB*, 9(4):372–383, 2015.
- [30] Douglas M Hawkins. *Identification of outliers*, volume 11. Springer, 1980.
- [31] Tom Heskes. The use of being stubborn and introspective. In *Prerational Intelligence: Adaptive Behavior and Intelligent Systems Without Symbols and Logic, Volume 1, Volume 2 Prerational Intelligence: Interdisciplinary Perspectives on the Behavior of Natural and Artificial Systems, Volume 3*, pages 1184–1200. Springer, 2000.
- [32] Victoria Hodge and Jim Austin. A survey of outlier detection methodologies. *Artificial intelligence review*, 22(2):85–126, 2004.
- [33] George Hripcsak and Adam S Rothschild. Agreement, the f-measure, and reliability in information retrieval. 12:296–8, 01 2005.
- [34] Pei-Yun Hsueh, Prem Melville, and Vikas Sindhwani. Data quality from crowdsourcing: a study of annotation selection criteria. In *Proceedings of the NAACL HLT 2009 workshop on active learning for natural language processing*, pages 27–35. Association for Computational Linguistics, 2009.
- [35] Nguyen Quoc Viet Hung, Duong Chi Thang, Matthias Weidlich, and Karl Aberer. Minimizing efforts in validating crowd answers. In *ACM SIGMOD*, pages 999–1014, 2015.
- [36] Ece Kamar, Severin Hacker, and Eric Horvitz. Combining human and machine intelligence in large-scale crowdsourcing. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 467–474. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [37] Taghi M Khoshgoftaar and Pierre Rebour. Generating multiple noise elimination filters with the ensemble-partitioning filter. In *Information Reuse and Integration, 2004. IRI 2004. Proceedings of the 2004 IEEE International Conference on*, pages 369–375. IEEE, 2004.
- [38] Sanjay Krishnan, Jiannan Wang, Eugene Wu, Michael J. Franklin, and Ken Goldberg. Active-clean: Interactive data cleaning for statistical modeling. *PVLDB*, 9(12):948–959, 2016.
- [39] David D. Lewis and Jason Catlett. Heterogenous uncertainty sampling for supervised learning. In *ICML*, pages 148–156, 1994.
- [40] David D. Lewis and William A. Gale. A sequential algorithm for training text classifiers. In *ACM-SIGIR*, pages 3–12, 1994.
- [41] Guoliang Li, Jiannan Wang, Yudian Zheng, and Michael J. Franklin. Crowdsourced data management: A survey. *IEEE Trans. Knowl. Data Eng.*, 28(9):2296–2319, 2016.
- [42] Christopher H. Lin, Mausam, and Daniel S. Weld. Re-active learning: Active learning with relabeling. In *AAAI*, pages 1845–1852, 2016.
- [43] Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence*, 38(3):447–461, 2016.
- [44] Xuan Liu, Meiyu Lu, Beng Chin Ooi, Yanyan Shen, Sai Wu, and Meihui Zhang. CDAS: A crowdsourcing data analytics system. *PVLDB*, 5(10):1040–1051, 2012.

- [45] Andrea Malossini, Enrico Blanzieri, and Raymond T Ng. Detecting potential labeling errors in microarrays by data perturbation. *Bioinformatics*, 22(17):2114–2121, 2006.
- [46] Adam Marcus, Eugene Wu, David R. Karger, Samuel Madden, and Robert C. Miller. Human-powered sorts and joins. *PVLDB*, 5(1):13–24, 2011.
- [47] Barzan Mozafari, Purnamrita Sarkar, Michael J. Franklin, Michael I. Jordan, and Samuel Madden. Scaling up crowd-sourcing to very large datasets: A case for active learning. *PVLDB*, 8(2):125–136, 2014.
- [48] Nagarajan Natarajan, Inderjit S. Dhillon, Pradeep Ravikumar, and Ambuj Tewari. Learning with noisy labels. In *NIPS*, pages 1196–1204, 2013.
- [49] Seishi Okamoto and Nobuhiro Yugami. An average-case analysis of the k-nearest neighbor classifier for noisy domains. In *IJCAI (1)*, pages 238–245, 1997.
- [50] Hyunjung Park and Jennifer Widom. Crowdfill: collecting structured data from the crowd. In *ACM SIGMOD*, pages 577–588, 2014.
- [51] Julien Pilourdault, Sihem Amer-Yahia, Dongwon Lee, and Senjuti Basu Roy. Motivation-aware task assignment in crowdsourcing. In *EDBT*, pages 246–257, 2017.
- [52] Vikas C. Raykar, Shipeng Yu, Linda H. Zhao, Gerardo Hermosillo Valadez, Charles Florin, Luca Bogoni, and Linda Moy. Learning from crowds. *Journal of Machine Learning Research*, 11:1297–1322, 2010.
- [53] Theodoros Rekatsinas, Manas Joglekar, Hector Garcia-Molina, Aditya G. Parameswaran, and Christopher Ré. Slimfast: Guaranteed results for data fusion and source reliability. In *SIGMOD*, pages 1399–1414, 2017.
- [54] Nicholas Roy and Andrew McCallum. Toward optimal active learning through sampling estimation of error reduction. In *ICML*, pages 441–448, 2001.
- [55] José Salvador Sánchez, Filiberto Pla, and Francesc J Ferri. Prototype selection for the nearest neighbour rule through proximity graphs. *Pattern Recognition Letters*, 18(6):507–513, 1997.
- [56] Bernhard Schölkopf, Robert C Williamson, Alex J Smola, John Shawe-Taylor, and John C Platt. Support vector method for novelty detection. In *Advances in neural information processing systems*, pages 582–588, 2000.
- [57] Burr Settles. Active learning literature survey. *University of Wisconsin, Madison*, 52(55-66):11, 2010.
- [58] Victor S. Sheng, Foster J. Provost, and Panagiotis G. Ipeirotis. Get another label? improving data quality and data mining using multiple, noisy labelers. In *ACM SIGKDD*, pages 614–622, 2008.
- [59] Sainbayar Sukhbaatar and Rob Fergus. Learning from noisy labels with deep neural networks. *arXiv preprint arXiv:1406.2080*, 2(3):4, 2014.
- [60] Jiang-wen Sun, Feng-ying Zhao, Chong-jun Wang, and Shi-fu Chen. Identifying and correcting mislabeled training instances. In *Future generation communication and networking (FGCN 2007)*, volume 1, pages 244–250. IEEE, 2007.

- [61] Jaree Thongkam, Guandong Xu, Yanchun Zhang, and Fuchun Huang. Support vector machine for outlier detection in breast cancer survivability prediction. In *Asia-Pacific Web Conference*, pages 99–109. Springer, 2008.
- [62] Andreas Veit, Neil Alldrin, Gal Chechik, Ivan Krasin, Abhinav Gupta, and Serge J. Belongie. Learning from noisy large-scale datasets with minimal supervision. In *IEEE CVPR*, pages 6575–6583, 2017.
- [63] Jeroen Vuurens, Arjen P de Vries, and Carsten Eickhoff. How much spam can you take? an analysis of crowdsourcing results to increase accuracy. In *Proc. ACM SIGIR Workshop on Crowdsourcing for Information Retrieval (CIR’11)*, pages 21–26, 2011.
- [64] Jiannan Wang, Tim Kraska, Michael J. Franklin, and Jianhua Feng. Crowder: Crowdsourcing entity resolution. *PVLDB*, 5(11):1483–1494, 2012.
- [65] Jiannan Wang, Sanjay Krishnan, Michael J. Franklin, Ken Goldberg, Tim Kraska, and Tova Milo. A sample-and-clean framework for fast and accurate query processing on dirty data. In *ACM SIGMOD*, pages 469–480, 2014.
- [66] Lu Wang and Zhi-Hua Zhou. Cost-saving effect of crowdsourcing learning. In *IJCAI*, pages 2111–2117, 2016.
- [67] Steven Euijong Whang, Peter Lofgren, and Hector Garcia-Molina. Question selection for crowd entity resolution. *PVLDB*, 6(6):349–360, 2013.
- [68] D Randall Wilson and Tony R Martinez. Reduction techniques for instance-based learning algorithms. *Machine learning*, 38(3):257–286, 2000.
- [69] Mohamed Yakout, Ahmed K. Elmagarmid, Jennifer Neville, Mourad Ouzzani, and Ihab F. Ilyas. Guided data repair. *PVLDB*, 4(5):279–289, 2011.
- [70] Yan Yan, Rómer Rosales, Glenn Fung, and Jennifer G. Dy. Active learning from crowds. In *ICML*, pages 1161–1168, 2011.
- [71] Chen Jason Zhang, Lei Chen, Yongxin Tong, and Zheng Liu. Cleaning uncertain data with a noisy crowd. In *ICDE*, pages 6–17, 2015.
- [72] Liyue Zhao, Gita Sukthankar, and Rahul Sukthankar. Incremental relabeling for active learning with noisy crowdsourced annotations. In *SocialCom/PASSAT*, pages 728–733, 2011.
- [73] Yudian Zheng, Guoliang Li, Yuanbing Li, Caihua Shan, and Reynold Cheng. Truth inference in crowdsourcing: is the problem solved? *PVLDB*, 10(5):541–552, 2017.
- [74] Yudian Zheng, Jiannan Wang, Guoliang Li, Reynold Cheng, and Jianhua Feng. QASCA: A quality-aware task assignment system for crowdsourcing applications. In *ACM SIGMOD*, pages 1031–1046, 2015.

Appendix A

Multi Class Generalization

In this section, we talk about how we can generalize our approach to the multi-class classification problem.

Classification With ground-truth labels. In order to be able to handle multi-class classification, we have to support all class labels for y . Therefore, our definition would be $y = l_k$ where $l_k \in Y = \{l_1, \dots, l_n\}$ representing the set of all possible ground-truth labels.

Crowdsourced Data. Corresponding to the ground-truth labels we also have to change the definition for noisy labels. Let $\mathcal{L}^i = \{(w_j, l_{i,j}) \mid w_j \text{ labels } x_i\}$ denote the corresponding labels, where $l_{i,j} \in Y$ represents the label given by worker w_j .

Worker Model. The confusion matrix of each worker w_j is a $n \times n$ matrix, where n represents the number of classes. For example, for 3-class classification we have:

$$q^{(j)} = \begin{bmatrix} q_{l_1, l_1}^{(j)}, & q_{l_1, l_2}^{(j)}, & q_{l_1, l_3}^{(j)} \\ q_{l_2, l_1}^{(j)}, & q_{l_2, l_2}^{(j)}, & q_{l_2, l_3}^{(j)} \\ q_{l_3, l_1}^{(j)}, & q_{l_3, l_2}^{(j)}, & q_{l_3, l_3}^{(j)} \end{bmatrix}$$

where each row represents a true label, each column represents a worker's provided label, and $q_{y,l}^{(j)}$ ($y \in Y, l \in Y$) means that given an instance with true label y , worker w_j provides label l with probability of $q_{y,l}^{(j)}$.

A.0.1 Computing Consolidated Labels

Given an instance x_i with worker labels \mathcal{L}^i , the basic idea of getting the x_i 's most likely label is to compare the values of n conditional probabilities: $P(Y_i = l_k \mid \mathcal{L}^i)$, the probability that instance x_i has a true label of l_k , conditioned on the labels that the workers provide. We will choose whatever label in this set that has the biggest conditional probability.

$$\begin{aligned}
P(Y_i = l_k | \mathcal{L}^i) &= \frac{P(\mathcal{L}^i | Y_i = l_k)P(Y_i = l_k)}{P(\mathcal{L}^i)} \\
&\propto P(\mathcal{L}^i | Y_i = l_k)P(Y_i = l_k) \\
&= P(Y_i = l_k) \prod_{l_{i,j} \in \mathcal{L}^i} P(L_j = l_{i,j} | Y_i = l_k) \\
&= P(Y = l_k) \prod_{l_{i,j} \in \mathcal{L}^i} q_{l_k, l_{i,j}}^{(j)} \tag{A.1}
\end{aligned}$$

A.0.2 Estimating Model's True Performance

For estimating model's performance we have to generalize the proposed unbiased estimator to handle multiple classes. The way to do it is to consider the loss function for each of the possible classes:

$$\text{Loss}(t, l_k) = \sum_n P(Y' = l_i | Y = l_k) \cdot \widetilde{\text{Loss}}(t, l_i) \tag{A.2}$$

Therefore we will have n equations and n unknown variables. By solving these n equations we can have the estimated loss function for each class $\widetilde{\text{Loss}}(t, l_i)$.

A.0.3 Expected Model Improvement

As discussed before we have the notion of improvement signal as follow:

$$\text{EMI}(i)_k = P(\text{case}_k) \cdot (\text{eval}(\mathbb{G}, f_i) - \text{eval}(\mathbb{G}, f)) \tag{A.3}$$

Then we can generalize this concept for each of the class labels:

$$\begin{aligned}
P(\text{case}_k) &= \frac{P(Y'_i = y'_i | Y_i = l_k) \cdot P(Y_i = l_k)}{P(Y'_i = y'_i)} \\
&\propto P(Y'_i = y'_i | Y_i = l_k) \cdot P(Y_i = l_k) \tag{A.4}
\end{aligned}$$

$P(Y'_i = y'_i | Y_i = l_k)$ is equal to the noise rate $r_{l_k, y'_i}^{(i)}$ and $P(Y_i = l_k)$ is equal to the prior $P(Y = l_k)$. Therefore, we obtain

$$P(\text{case}_k) \propto r_{l_k, y'_i}^{(i)} \cdot P(Y = l_k) \tag{A.5}$$

We can get the improvement signal for all of the possible classes and then for each record we have the final improvement signal as the maximum improvement signal among class label.

$$\text{EMI}(i) = \max \text{EMI}(i)_k \quad (\text{A.6})$$

A.0.4 Empirical Experiment

We evaluated our method on multi class Thyroid Disease dataset from UCI archive containing 3 classes and 215 instances. For simulating error in real data, noise rate of (0.1,0.1,0.3) corresponding to each class, was applied to the dataset. We varied the number of cleaned instances, and compared the model's accuracy w.r.t. different cleaning strategies on this dataset. The result is shown in Figure A.1.

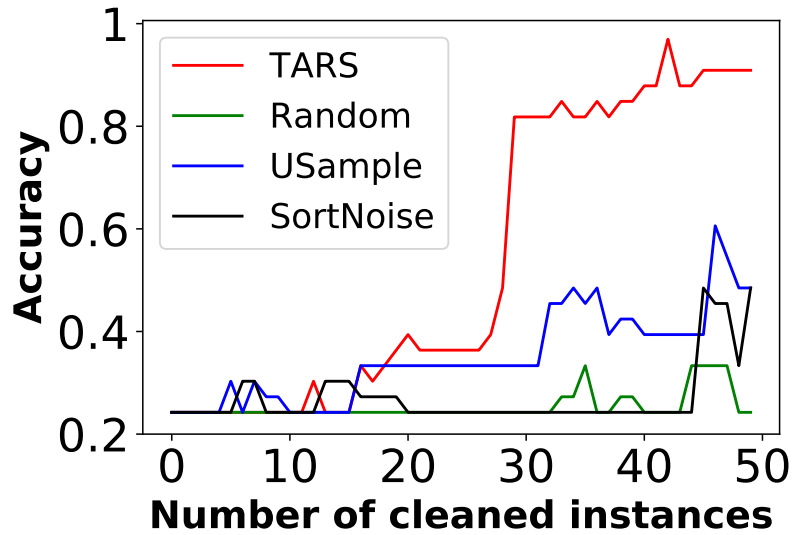


Figure A.1: Comparison of cleaning strategies on the multi-class dataset