
Analysis of Minimax Error Rate for Crowdsourcing and Its Application to Worker Clustering Model

Hideaki Imamura^{1,2} Issei Sato^{1,2} Masashi Sugiyama^{2,1}

Abstract

While crowdsourcing has become an important means to label data, there is great interest in estimating the ground truth from unreliable labels produced by crowdworkers. The Dawid and Skene (DS) model is one of the most well-known models in the study of crowdsourcing. Despite its practical popularity, theoretical error analysis for the DS model has been conducted only under restrictive assumptions on class priors, confusion matrices, or the number of labels each worker provides. In this paper, we derive a minimax error rate under more practical setting for a broader class of crowdsourcing models including the DS model as a special case. We further propose the worker clustering model, which is more practical than the DS model under real crowdsourcing settings. The wide applicability of our theoretical analysis allows us to immediately investigate the behavior of this proposed model, which can not be analyzed by existing studies. Experimental results showed that there is a strong similarity between the lower bound of the minimax error rate derived by our theoretical analysis and the empirical error of the estimated value.

1. Introduction

Crowdsourcing has become an essential tool for large-scale data collection in machine learning. While crowdsourcing provides a less expensive means of labeling data, the data annotated by a crowd can be of low quality because crowd workers are often non-experts and are sometimes even adversarial. Many crowdsourcing services try to

The code for the experiment can be used in <https://github.com/HideakiImamura/MinimaxErrorRate>.

¹The University of Tokyo, Tokyo, Japan ²RIKEN, Tokyo, Japan. Correspondence to: Hideaki Imamura <imamura@ms.k.u-tokyo.ac.jp>.

Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, PMLR 80, 2018. Copyright 2018 by the author(s).

solve this problem by providing redundancy for labeling, i.e., by collecting multiple labels from different workers for each task (Ipeirotis et al., 2010; Welinder et al., 2010; Snow et al., 2008). This strategy raises the question: how can the ground truths be estimated from noisy and redundant labels? Estimating the ground truth is difficult in the *laissez-faire crowdsourcing setting* in which a large number of workers in the world can freely label as many tasks as they want. In this setting, a small number of workers annotate data a large number of times, while most workers perform annotation only a few times. That is, the number of tasks to be labeled per worker typically follows Zipf’s law. In this paper, we focus on this setting because it is more realistic than the *non-laissez-faire crowdsourcing setting*, in which the number of tasks to be labeled per worker is almost constant across all workers and which is assumed in the experiments of many previous works (Welinder et al., 2010; Snow et al., 2008).

In the context of estimating the ground truth of each task from noisy and redundant labels, Dawid & Skene (1979) conducted pioneering research. It is assumed in the Dawid and Skene (DS) model that each worker has his or her own confusion matrix. The ground truth of each task and the confusion matrix of each worker are jointly estimated by the Expectation-Maximization (EM) algorithm (Dempster et al., 1977).

Although the DS model has had empirical success (Welinder et al., 2010; Snow et al., 2008), there are two major problems. First, there are few theoretical error analyses on the performance conducted for the DS model, and the existing theoretical analyses are only valid under strong assumptions. For example, Gao et al. (2016) assumed that the class prior is uniform and Zhang et al. (2014) assumed that the entries of the confusion matrices are strictly positive. Second, the experimental validity of the existing methods is confirmed only when the number of tasks to be labeled per worker is almost constant on both synthetic and real-world data (Welinder et al., 2010; Snow et al., 2008).

To alleviate the first problem, we provide a novel theoretical error analysis under milder assumptions based on Fano’s method, which is a useful minimax lower-bounding technique (Yu, 1997). Our theoretical analysis is applicable

to all models that use the ground truth of each task and the confusion matrix of each worker, including the DS model and its extensions, thanks to the fact that our assumptions are much weaker than those required in previous work.

To alleviate the second problem, we extend the DS model to be able to handle worker clusters. Intuitively, by clustering workers, even when the number of tasks to be labeled per worker is small, the number of tasks to be labeled per *worker cluster* can be increased and thus estimation can be stabilized. Note that our widely applicable theoretical analysis explained above allows us to investigate the behavior of the proposed clustering model, while the existing theoretical analysis methods cannot be used due to their restrictive assumptions. We experimentally show the usefulness of our worker clustering (WC) model for reducing the influence of variations in the number of tasks to be labeled per worker. We also numerically confirm the validity of our theoretical error analysis.

2. Related Work

A large number of studies on the quality assurance of data collected by crowdsourcing have been conducted (Karger et al., 2011; Bachrach et al., 2012; Zhou et al., 2012; Chen et al., 2013; Karger et al., 2013; Parisi et al., 2014; Venanzi et al., 2014; Karger et al., 2014; Tian & Zhu, 2015). One of the most practical and pioneering studies in this field is the Dawid and Skene (DS) model (Dawid & Skene, 1979). The DS model is based on an estimation paradigm that uses the ground truths of all tasks and the confusion matrices of all workers; their inference algorithm is based on the EM algorithm (Dempster et al., 1977). Application research using their method has been actively carried out (Hui & Walter, 1980; Smyth et al., 1995; Albert & Dodd, 2004). As many experiments on synthetic and real-world data have demonstrated, the DS model is practical for estimating the ground truth of each task from noisy and redundant labels.

There is also many studies that uses the EM algorithm (Whitehill et al., 2009; Welinder et al., 2010; Raykar et al., 2010; Welinder & Perona, 2010; Liu et al., 2012; Liu & Wang, 2012; Zhang et al., 2014). Liu & Wang (2012) assume priors over the class prior of the ground truths of tasks and the confusion matrices of workers. Zhang et al. (2014) devised an effective way to initialize the EM algorithm. Specifically, the initial values of the confusion matrices of all workers were estimated using the method of moments. Karger et al. (2011) proposed an iterative algorithm for binary labeling problems which is not based on the DS model, and gave a theoretical analysis under strong assumptions.

The idea of clustering worker was already proposed by

Venanzi et al. (2015) and Moreno et al. (2015). The problem setting of Venanzi et al. (2015) supposed that the inputs of the model are not labels but real-valued vectors, which is different from our problem setting. Moreno et al. (2015) solved the same problem as ours, but assumed a complex generation process for confusion matrices and labels given by workers. To give the worker clustering structure, their model becomes complex, and difficult to analyze theoretically.

As a number of algorithms have been proposed, statistical understanding of crowdsourcing has been actively researched, such as (Ghosh et al., 2011; Dalvi et al., 2013; Karger et al., 2014; Zhang et al., 2014; Gao et al., 2016; Bonald & Combes, 2017). These studies are divided into two types. One is a theoretical analysis of the estimation paradigm that uses the ground truths and confusion matrices (Zhang et al., 2014; Gao et al., 2016). The other is a theoretical analysis of their own models and algorithms (Ghosh et al., 2011; Dalvi et al., 2013; Karger et al., 2014; Bonald & Combes, 2017). Our theoretical analysis is of the former category.

Zhang et al. (2014) initialized the confusion matrices of all workers using the method of moments and showed convergence of the EM algorithm initialized with their method. However, they assumed that the minimum value of the entries of the confusion matrices is greater than a positive constant. Gao et al. (2016) derived the minimax optimal convergence rate for the DS model, but they assumed that the class prior is uniform.

Bonald & Combes (2017) gave a probabilistic concentration inequality for the estimation error between the ground truth and estimated truth without using the confusion matrices of workers. While their result cannot be calculated empirically, our theoretical result can be calculated using the estimated confusion matrices and the class prior. We describe the usefulness of calculating the bound in Section 6.3 Because of our weaker assumptions, our theoretical analysis is applicable to all methods that use the estimation paradigm based on the ground truths of all tasks and confusion matrices of all workers.

3. Dawid and Skene Model

In this section, we formulate the problem of estimating the ground truths of tasks based on noisy and redundant labels provided by workers and their confusion matrices. The model formulated in this section is the basis of all models to be theoretically analyzed in the next section.

Suppose we have n tasks labeled by m workers on K possible labels. Let $X_{i,j}$ be the label of the i -th task given by the j -th worker. Denote by $X_{i,j} = k$ that the j -th worker labels $k \in \{1, 2, \dots, K\}$ to the i -th task. Denote by $X_{i,j} = 0$

Table 1. Relationship between existing work and this work. DS stands for Dawid and Skene, while WC stands for worker clustering.

	MODEL	THEORETICAL ANALYSIS	LAISSEZ-FAIRE SETTING	CLASS PRIOR
DAWID & SKENE (1979)	DS	-	-	NOT UNIFORM
KARGER ET AL. (2011)	OTHER	✓	-	NO CLASS PRIOR
LIU & WANG (2012)	DS	-	-	NOT UNIFORM
GAO ET AL. (2016)	DS	✓	-	UNIFORM
THIS WORK	DS & WC	✓	✓	NOT UNIFORM

that the j -th worker does not label the i -th task. We use $G_i \in \{1, 2, \dots, K\}$ to denote the ground truth of the i -th task. Here, $\mathbb{S}^{K \times K}$ is the set of right stochastic matrices, that is,

$$\forall P \in \mathbb{S}^{K \times K}, \forall k \in \{1, 2, \dots, K\}, \sum_{k'=1}^K P_{k,k'} = 1,$$

$$\forall P \in \mathbb{S}^{K \times K}, \forall k \in \{1, 2, \dots, K\}, \forall k' \in \{1, 2, \dots, K\}, P_{k,k'} \geq 0.$$

The ability of the j -th worker is measured by a confusion matrix $\pi^j \in \mathbb{S}^{K \times K}$ with its (k, k') -element $\pi_{k,k'}^j$ being the probability that the j -th worker labels k' when the true label is k . For simplicity, we define

$$X = \{X_{i,j}\}_{i=1, j=1}^{n,m}, G = \{G_i\}_{i=1}^n, \pi = \{\pi^j\}_{j=1}^m.$$

We suppose that when $G_i = k$ is the ground truth of the i -th task, the label given by the j -th worker is sampled from a multinomial distribution parametrized by the k -th row of the confusion matrix of the j -th worker, namely, π_k^j . We assume that the ground truth of each task, G_i , is sampled from a multinomial distribution parametrized by $\rho = (\rho_1, \rho_2, \dots, \rho_K)$, where $\rho_k \geq 0$ for any $k \in \{1, \dots, K\}$ and $\sum_{k=1}^K \rho_k = 1$. We call ρ the class prior of all tasks. We have observed variables X , latent variables G , and parameters $\{\pi, \rho\}$. The graphical model is plotted in Figure 1-(a).

Then, the joint distribution of this model is expressed as follows.

$$p(X, G | \rho, \pi) = \left[\prod_{i=1}^n \prod_{j=1}^m p(X_{i,j} | G_i, \pi^j) \right] \left[\prod_{i=1}^n p(G_i | \rho) \right],$$

where each probability distribution is given as follows.

$$p(X_{i,j} | G_i, \pi^j) = \prod_{k'=1}^K (\pi_{G_i, k'}^j)^{\delta(X_{i,j}=k')},$$

$$p(G_i | \rho) = \prod_{k=1}^K \rho_k^{\delta(G_i=k)}.$$

The aim is to predict the ground truths of all tasks, G , from observed labels X . Dawid & Skene (1979) used the EM algorithm to infer the ground truths. However, when we conduct theoretical analysis in the next section, we take an approach that does not depend on a specific inference algorithm.

4. Minimax Error Analysis

We give a lower bound on the minimax error for a class of models that use the ground truths of tasks and confusion matrices of workers. Let us define some concepts and formulas for analysis.

We focus on the model described in Section 3. Note that the model may optionally contain additional structures, such as a prior for the class prior (Liu & Wang, 2012). Moreover, we can use any inference algorithm to estimate the ground truths and class prior of tasks and the confusion matrices of workers. Let \hat{G} be estimated truths of all tasks; let \hat{G}_i be an estimate of the ground truth of the i -th task. The set of $\{1, \dots, K\}$ is denoted by $[K]$. A loss is measured by the error rate given by

$$\mathcal{L}(\hat{G}, G) = \frac{1}{n} \sum_{i=1}^n \delta(\hat{G}_i \neq G_i),$$

where $\delta(\cdot)$ is the indicator function. Let \mathbb{P} be the joint probability distribution of the data $\{X_{i,j}\}$ given π, ρ and G . Let \mathbb{E} be the associated expectation operator. Denote by $H(\rho)$ the entropy of the ground truth of each task with respect to the class prior:

$$H(\rho) = - \sum_{k=1}^K \rho_k \log \rho_k.$$

Let $\text{KL}(\pi_{g_*}^j || \pi_{g'_*}^j)$ be the Kullback-Leibler divergence from one row of the confusion matrix to the other:

$$\text{KL}(\pi_{g_*}^j || \pi_{g'_*}^j) = \sum_{k=1}^K \pi_{g_*}^j \log \frac{\pi_{g_*}^j}{\pi_{g'_*}^j}.$$

We assume the following.

Assumption 1. Given ρ and π , the labels of tasks given by workers are in accordance with the models that use the ground truths of tasks and confusion matrices of workers such as the DS model.

We bound the minimax error rate as follows.

Theorem 1. Under Assumption 1, the minimax error rate is lower bounded as follows.

$$\inf_{\hat{G}} \sup_{G \in [K]^n} \mathbb{E}[\mathcal{L}(\hat{G}, G)] \geq \frac{1}{n \log K} \left(R(\rho, \pi) - \frac{\log 2}{n} \right),$$

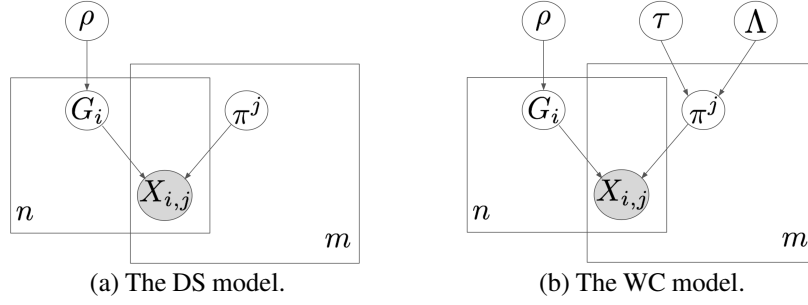


Figure 1. Graphical models.

where

$$R(\rho, \pi) = H(\rho) - \sum_{j=1}^m \sum_{g=1}^K \sum_{g'=1}^K \rho_g \rho_{g'} \text{KL}(\pi_{g^*}^j \| \pi_{g'^*}^j).$$

The details of the proof are given in Appendix A. The proof of this theorem is based on Fano's method by Yu (1997), which is a well-known minimax lower bounding technique.

Assumption 1 is weaker than that of many previous works such as Gao et al. (2016) and Zhang et al. (2014). As we mentioned in the introduction, their approaches make more restrictive assumptions to conduct theoretical analysis. For example, Gao et al. (2016) assumed that class prior ρ is uniform, that is, $\rho_k = \frac{1}{K}$ for any $k \in \{1, \dots, K\}$. Zhang et al. (2014) assumed that all entries $\pi_{g,k}^j$ of the confusion matrix is strictly positive. In contrast, thanks to the weak assumptions, our lower bound is applicable to the DS mode and other models that use the ground truths of tasks and confusion matrices of workers such as Zhang et al. (2014) and Liu & Wang (2012).

Our lower bound can be used to measure the performance of each model. Specifically, the performance of each model can be measured by the value of $R(\rho, \pi)$, which is the main part of the lower bound of the minimax error rate. We conducted numerical experiments to measure the performance of each model, the results of which are given in the experiment section.

In this paper, we provided the lower bound of the minimax error rate but did not give an upper bound. The theoretical analysis of the previous work such as Gao et al. (2016) easily derived an upper bound under mild assumptions. However, the derivation depends not only on the model, but also on an inference algorithm, because $\inf_{\hat{G}} \sup_{G \in [K]^n} \mathbb{E}[\mathcal{L}(\hat{G}, G)]$ includes the infimum over estimate \hat{G} . Therefore, it is not appropriate to use such an upper bound when analyzing the behavior of the model itself.

5. Worker Clustering Model

We propose a model that is more practical than the DS model in the *laissez-faire crowdsourcing setting*, where

workers can label as many tasks as they want. The main idea behind our model is simple: divide all workers into several disjoint clusters. Intuitively, by clustering workers, even if the number of labels provided by each worker is small, the number of labels per *worker cluster* can be increased to stabilize inference.

The definition of the labels given by workers X , the ground truths of all tasks G , the class prior ρ , and the confusion matrices of all workers π are the same as that of the DS model. The proposed model can actually be regarded as an extension of the DS model and the HybridConfusion model (Liu & Wang, 2012). The biggest difference from the previous models is that the proposed model limits the possible values of confusion matrices π to a maximum of L values $\Lambda = \{\Lambda_1, \Lambda_2, \dots, \Lambda_L\}$, where $L \leq m$. This means that we are clustering m workers into a maximum of L groups. We also suppose that the confusion matrix π^j is determined by a multinomial distribution parametrized by $\tau = (\tau_1, \tau_2, \dots, \tau_L)$, that is, π^j is equal to Λ_l with probability τ_l . This means that the j -th worker belongs to the l -th cluster with probability τ_l . We have observed variables X , latent variables $\{G, \pi\}$ and parameters $\{\Lambda, \rho, \tau\}$. The graphical model is shown in Figure 1-(b).

The joint distribution of this model is expressed as follows.

$$p(X, G, \pi | \rho, \tau, \Lambda) = \left[\prod_{i=1}^n \prod_{j=1}^m p(X_{i,j} | G_i, \pi^j) \right] \left[\prod_{i=1}^n p(G_i | \rho) \right] \left[\prod_{j=1}^m p(\pi^j | \tau, \Lambda) \right],$$

where each probability distribution is given as follows.

$$p(X_{i,j} | G_i, \pi^j) = \prod_{k'=1}^K (\pi_{G_i, k'}^j)^{\delta(X_{i,j}=k')},$$

$$p(G_i | \rho) = \prod_{k=1}^K \rho_k^{\delta(G_i=k)},$$

$$p(\pi^j | \tau, \Lambda) = \prod_{l=1}^L \tau_l^{\delta(\pi^j=\Lambda_l)}.$$

Note that our theoretical analysis given in Section 4 is also applicable to our proposed WC model. To apply the lower

bound to our model, we use Λ instead of the confusion matrices as follows. We assume that Λ , class prior ρ , and a group ℓ_j which each worker belongs to are known. Under this assumption, $\hat{\phi}_{j,l} = 1$ when $l = \ell_j$ and otherwise $\hat{\phi}_{j,l} = 0$. Our lower bound of the minimax error rate is then established with Λ_{ℓ_j} instead of a confusion matrix π^j .

To estimate latent variables and optimize parameters, we adopt the strategy of *empirical variational inference* (Robbins, 1956). The details of the derivations are in Appendix B in the supplementary material.

The behavior of the WC model when L gradually increases is as follows. Since L is the maximum number of worker clusters, even if $L = m$, not all m workers belong to different clusters. Some clusters may include many workers, but others may not include any workers. The actual number of clusters is adaptively determined for each target dataset. Therefore, it is not good to increase the value of L , but there should be an appropriate value of L . The method of determining the value of L using our theoretical analysis given in Section 4 is described in Section 6.3.

6. Experiments

We empirically analyzed the proposed method on synthetic and real-world data. First, we analyzed that the WC model proposed in Section 5 performs better than the existing method in the *laissez-faire crowdsourcing setting* where the number of labels per worker is in accordance with Zipf’s law. Second, we investigated that there is a strong similarity between the lower bound of the minimax error rate derived by our theoretical analysis in Section 4 and the empirical error of the estimated value.

We compared our proposed method, the WC model, with the classical Majority Voting (MV) scheme and the DS model. The experiments were conducted on four synthetic and three real-world datasets.

6.1. Synthetic Data

For synthetic data, we generated $m = 100$ workers and $n = 1000$ binary annotation tasks. The ground truth of each task was sampled from the binomial distribution with the class prior $\rho = (0.3, 0.7)$. For each worker, a confusion matrix was generated as follows. Workers were divided into two types: honest and adversarial. Each row of each confusion matrix was sampled from the Dirichlet distribution. The concentration parameter of the Dirichlet distribution for the k -th row was

$$\alpha = (\underbrace{10, \dots, 10}_{k-1}, 100, \underbrace{10, \dots, 10}_{K-k})$$

for honest workers and

$$\alpha = (\underbrace{10, \dots, 10}_{k-1}, 1, \underbrace{10, \dots, 10}_{K-k})$$

for adversarial workers. To simulate the labeling process of a crowd, we consider two settings: The first setting is the “Constant Labeling Number (CLN)” setting, in which we determine the number N of tasks that all workers commonly annotate, and then each worker is given N tasks randomly and label those tasks in accordance with his or her own confusion matrix. The second setting is the “Zipf Labeling Number (ZLN)” setting, in which, for any $j \in \{1, \dots, m\}$, we determine the number N_j of tasks that the j -th worker annotates in accordance with Zipf’s law, and then the j -th worker is randomly given N_j tasks and labels those tasks according to his or her own confusion matrix. The CLN setting was assumed in the experiments of many existing studies (e.g. Welinder et al., 2010; Snow et al., 2008). The ZLN setting is suitable for the *laissez-faire crowdsourcing setting*.

The input of each model was a set of triplets (task, worker, and label) generated by this procedure. The outputs of the WC and DS models were estimated values of the ground truths and class prior of all tasks and the confusion matrices of all workers. In order to measure how robust each method is to the proportion of adversaries, we showed the change in accuracy as a function of the proportion of adversaries (see Figure 2). We changed the proportion of adversaries from 0 to 50%.

The accuracy of MV monotonically decreases as the proportion of adversaries increases in all cases. In the CLN setting, when the number of times each worker labels is sufficiently large ($N = 100$), the accuracies of DS and WC remain high even if the proportion of adversaries is high (nearly 0.5) (see Figure 2-(c)). In contrast, when the number of times each worker labels is small ($N = 2, 10$), the accuracy of DS is low (almost the chance level), but that of WC remains high (see Figure 2-(a), (b)). Intuitively, the reason for this phenomenon is that WC effectively increased the number of labeled tasks per cluster by clustering workers and thus the estimation became more reliable. In particular, this phenomenon is noticeable when the number of tasks to be labeled per worker follows Zipf’s law, i.e., in the ZLN setting (see Figure 2-(d)). Even if the average or the maximum value of the number of tasks to be labeled per worker is large, when most workers only label a small number, the accuracy of DS is not high. When the number of times most workers label is small, the proposed method WC is more effective than MV and DS.

6.2. Real-World Data

In the real-world data experiments, we compared crowdsourcing algorithms on three datasets: two binary tasks and one multi-class task. The two binary tasks were labeling bird species (Welinder et al., 2010) (the Bird dataset) and recognizing textual entailment (Snow et al., 2008) (the

Table 2. Statistics of real-world datasets.

DATASET	#CLASSES	#TASKS	#WORKERS	#LABELS
BIRD	2	108	39	4212
RTE	2	800	164	8000
DOG	4	807	52	7354

Table 3. Error rate (%) of estimating the ground truths of tasks on real-world datasets. MV, DS, and WC stand for majority voting, Dawid and Skene, and worker clustering, respectively.

MODEL \ DATASET	RTE	BIRD	DOG
MV	8.12	24.07	18.46
DS	7.62	15.74	16.60
WC ($L = 1$)	11.37	19.44	18.71
WC ($L = 2$)	7.00	17.59	18.59
WC ($L = 10$)	6.75	11.11	16.98
WC ($\min R(\rho, \pi)$)	6.75	11.11	16.36

RTE dataset). The multi-class task was labeling the breed of dogs (Zhou et al., 2012) (the Dog dataset). The statistics for the datasets were summarized in Table 2.

We calculated the error rate of the WC model by changing the value of parameter L . In Table 3, we reported the values for $L = 1, 2, 10$ and the value of the error rate when $R(\rho, \pi)$ attains its minimum. Table 3 also shows the error rate calculated by MV and DS as baselines. It can be seen that the error rate decreases as the value of L increases in the WC model. For the RTE dataset, the WC model is superior to both MV and DS when $L \geq 2$. For the Bird dataset, the WC model outperforms MV for all L and DS when $L \geq 10$. For the Dog dataset, the WC model does not outperform DS even when $L = 10$; as the value of L is increased, the WC model outperforms DS when $R(\rho, \pi)$ is minimum.

As expected, the performance of the WC model depends heavily on the choice of L . However, the experimental results show that larger L always yields better performance. We will investigate this issue more systematically in Section 6.3.

6.3. Similarity between Lower Bound of Minimax Error Rate and Empirical Error

We experimentally investigated the similarity between the lower bound of the minimax error rate and the empirical error of the ground truths of tasks estimated by the WC model.

It is important to note that the lower bound of the minimax error rate derived by our theoretical analysis can be approximately calculated from the class prior and confusion matrices estimated by each method, but the empirical error of the ground truths of tasks estimated by each method can not be calculated without the ground truths of those tasks. In practice, we want to adopt a model that makes the empirical error as small as possible, but its calculation is impos-

sible because of the necessity of the ground truths of tasks. If there is some similarity between the empirical error and lower bound of the minimax error rate, we can estimate the behavior of each model by the theoretical lower bound even in a more realistic situation in which we do not know the ground truths of tasks.

Here we experimentally investigate the behavior of $R(\rho, \pi)$ and $\mathcal{L}(\hat{G}, G)$ against the change in the value of the maximum number of worker clusters, L for the WC model. $R(\rho, \pi)$ is the main part of the lower bound of the minimax error rate derived by our theoretical analysis introduced in Section 4, while $\mathcal{L}(\hat{G}, G)$ is the empirical error of the estimated truths of tasks by the WC model.

The experiment was conducted on three synthetic datasets and three real-world datasets. For each dataset, we plotted in Figure 3 the change in $R(\rho, \pi)$ and $\mathcal{L}(\hat{G}, G)$ as a function of the parameter L . The value of L was changed from 1 to m . In Figure 3, the results of $\mathcal{L}(\hat{G}, G)$ and $R(\rho, \pi)$ are shown in the left and right columns, respectively.

For all cases, we see that increasing the value of L tends to decrease the values of $R(\rho, \pi)$ and $\mathcal{L}(\hat{G}, G)$ similarly. This means that the behavior of $\mathcal{L}(\hat{G}, G)$ can be well predicted by $R(\rho, \pi)$. Looking at the change in $\mathcal{L}(\hat{G}, G)$, to achieve better performance, we can see that it is sufficient to set the value of L to be large to some extent. This behavior can be well captured by looking at the change in $R(\rho, \pi)$ which can be computed in practice. Thus, our theoretical lower bound may be useful to investigate the behavior of different models.

7. Conclusions

In this paper, we gave a novel theoretical error analysis by using Fano’s method for any models based on the ground truths of tasks and the confusion matrices of workers. We have considered a realistic *laissez-faire crowdsourcing setting* and proposed clustering workers to reduce the effect of the bias in the number of tasks to be labeled per worker. Our theoretical analysis is applicable to all models that use the ground truths of tasks and the confusion matrices of workers, including Dawid & Skene (1979) and its variations, thanks to the weak assumptions required in our analysis. Through experiments on synthetic and real-world data, we have found that there is a strong similarity between the lower bound of the minimax error rate derived by our theoretical analysis and the empirical error of the estimated value. Thus, our theoretical lower bound, which can be approximately calculated from an established class prior and estimated confusion matrices, can be useful to investigate the behavior of different models in practice.

Acknowledgement IS was supported by JST CREST Grant Number JPMJCR17A1, Japan. and MS was supported by KAKENHI 17H00757.

A. Proof of Theorem 1

Fano's method (Yu, 1997) uses Fano's inequality (Fano, 1949) to give a lower bound for the minimax error rate. In the method, the uniformity of a class prior is assumed. However, it is not realistic for crowdsourcing. Therefore, we do not use Fano's method directly but improve the method to be suitable for crowdsourcing. We use the following inequality between the error probability and a conditional entropy.

Lemma 2 (Fano's inequality Fano (1949)).
For any Markov chain $V \rightarrow X \rightarrow \hat{V}$, we have

$$h(P(\hat{V} \neq V)) + P(\hat{V} \neq V)(\log(|\mathcal{V}|) - 1) \geq H(V|\hat{V}),$$

where $h(p) = -p \log p - (1-p) \log(1-p)$, \mathcal{V} is the set of possible value of V , and $H(V|\hat{V})$ is the entropy of V conditioned on \hat{V} .

The details of the proof are given in Appendix A in supplementary material. Using this inequality, we prove Theorem 1 as follows.

Proof of Theorem 1.

First, using Markov's inequality, we have

$$\begin{aligned} \inf_{\hat{G}} \sup_{G \in [K]^n} \mathbb{E}[\mathcal{L}(\hat{G}, G)] &\geq \inf_{\hat{G}} \sup_{G \in [K]^n} \frac{1}{n} \mathbb{P} \left[\left(\mathcal{L}(\hat{G}, G) \geq \frac{1}{n} \right) \right] \\ &= \inf_{\hat{G}} \left[\frac{1}{n} \sup_{G \in [K]^n} \mathbb{P} \left(\mathcal{L}(\hat{G}, G) \geq \frac{1}{n} \right) \right] \end{aligned}$$

Since $\hat{G} \neq G$ and $\mathcal{L}(\hat{G}, G) = \frac{1}{n} \sum_{i=1}^n \delta(\hat{G}_i \neq G_i) \geq \frac{1}{n}$ are equivalent, we have

$$\begin{aligned} \inf_{\hat{G}} \sup_{G \in [K]^n} \mathbb{E}[\mathcal{L}(\hat{G}, G)] &\geq \inf_{\hat{G}} \left[\frac{1}{n} \sup_{G \in [K]^n} \mathbb{P} \left(\mathcal{L}(\hat{G}, G) \geq \frac{1}{n} \right) \right] \\ &= \inf_{\hat{G}} \left[\frac{1}{n} \sup_{G \in [K]^n} \mathbb{P}(\hat{G} \neq G) \right]. \quad (1) \end{aligned}$$

Second, we evaluate $\mathbb{P}(\hat{G} \neq G)$ by using Lemma (2), i.e.,

$$\begin{aligned} \mathbb{P}(\hat{G} \neq G) &\geq \mathbb{P}(\hat{G} \neq G) \frac{\log(K^n - 1)}{\log K^n} \\ &\geq \frac{\mathbb{P}(\hat{G} \neq G) \log(K^n - 1) + h(P(\hat{G} \neq G)) - \log 2}{\log K^n} \\ &\geq \frac{H(G|\hat{G}) - \log 2}{n \log K}. \quad (2) \end{aligned}$$

Third, we evaluate $H(G|\hat{G})$ by using the relationship between conditional entropy and mutual information, and the

data processing inequality, i.e.,

$$\begin{aligned} H(G|\hat{G}) &= H(G) - I(G; \hat{G}) \\ &\geq H(G) - I(G; X). \end{aligned}$$

Since the elements of $G = (G_1, \dots, G_n)$ are independent and identically distributed, we have

$$H(G|\hat{G}) \geq nH(\rho) - I(G; X) \quad (3)$$

Then the mutual information of G and X is evaluated as follows.

$$\begin{aligned} I(G; X) &= KL(P(G, X) \| P(G)P(X)) \\ &= KL(P(X|G)P(G) \| P(G)P(X)) \\ &= \sum_{G \in [K]^n, X \in [K]^m} P(G)P(X|G) \log \frac{P(G)P(X|G)}{P(G)P(X)} \\ &= \sum_{G \in [K]^n} P(G) KL(P(X|G) \| P(X)) \\ &= \sum_{G \in [K]^n} P(G) KL \left(P(X|G) \left\| \sum_{G' \in [K]^n} P(X|G')P(G') \right. \right) \\ &\leq \sum_{G \in [K]^n} P(G) \sum_{G' \in [K]^n} P(G') KL(P(X|G) \| P(X|G')) \\ &\quad (\text{by convexity of } -\log) \\ &= \sum_{G \in [K]^n} \sum_{G' \in [K]^n} P(G)P(G') \left(\sum_{i=1}^n \sum_{j=1}^m KL(P(X_{i,j}|G) \| P(X_{i,j}|G')) \right) \\ &\quad (\text{by independence of } X) \\ &= \sum_{G \in [K]^n} \sum_{G' \in [K]^n} \sum_{i=1}^n \sum_{j=1}^m P(G)P(G') KL(\pi_{G_i^*}^j \| \pi_{G'_i^*}^j) \\ &= \sum_{i=1}^n \sum_{j=1}^m \sum_{g=1}^K \sum_{g'=1}^K P(G_i = g)P(G'_i = g') KL(\pi_{g^*}^j \| \pi_{g'^*}^j) \\ &= n \sum_{j=1}^m \sum_{g=1}^K \sum_{g'=1}^K \rho_g \rho_{g'} KL(\pi_{g^*}^j \| \pi_{g'^*}^j). \end{aligned}$$

Combining this inequality and Inequality (3), we have

$$\begin{aligned} H(G|\hat{G}) &\geq n \left(H(\rho) - \sum_{j=1}^m \sum_{g=1}^K \sum_{g'=1}^K \rho_g \rho_{g'} KL(\pi_{g^*}^j \| \pi_{g'^*}^j) \right) \\ &= nR(\rho, \pi). \quad (4) \end{aligned}$$

Finally, using Inequality (1), (2), and (4), we get

$$\begin{aligned} \inf_{\hat{G}} \sup_{G \in [K]^n} \mathbb{E}[\mathcal{L}(\hat{G}, G)] &\geq \inf_{\hat{G}} \left[\frac{1}{n} \sup_{G \in [K]^n} \frac{1}{\log K} \left(R(\rho, \pi) - \frac{\log 2}{n} \right) \right] \\ &= \frac{1}{n \log K} \left(R(\rho, \pi) - \frac{\log 2}{n} \right) \end{aligned}$$

This completes the proof. \square

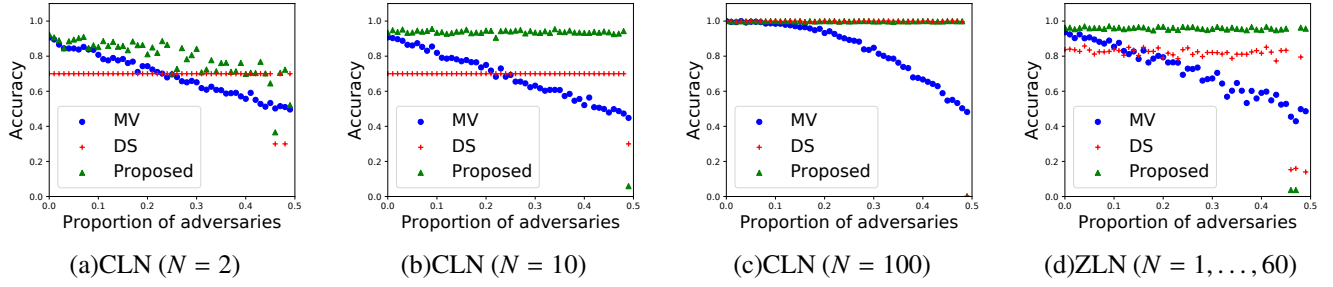


Figure 2. Change of accuracy as a function of the proportion of adversaries. MV, DS, and WC stand for majority voting, Dawid and Skene, and worker clustering, respectively.

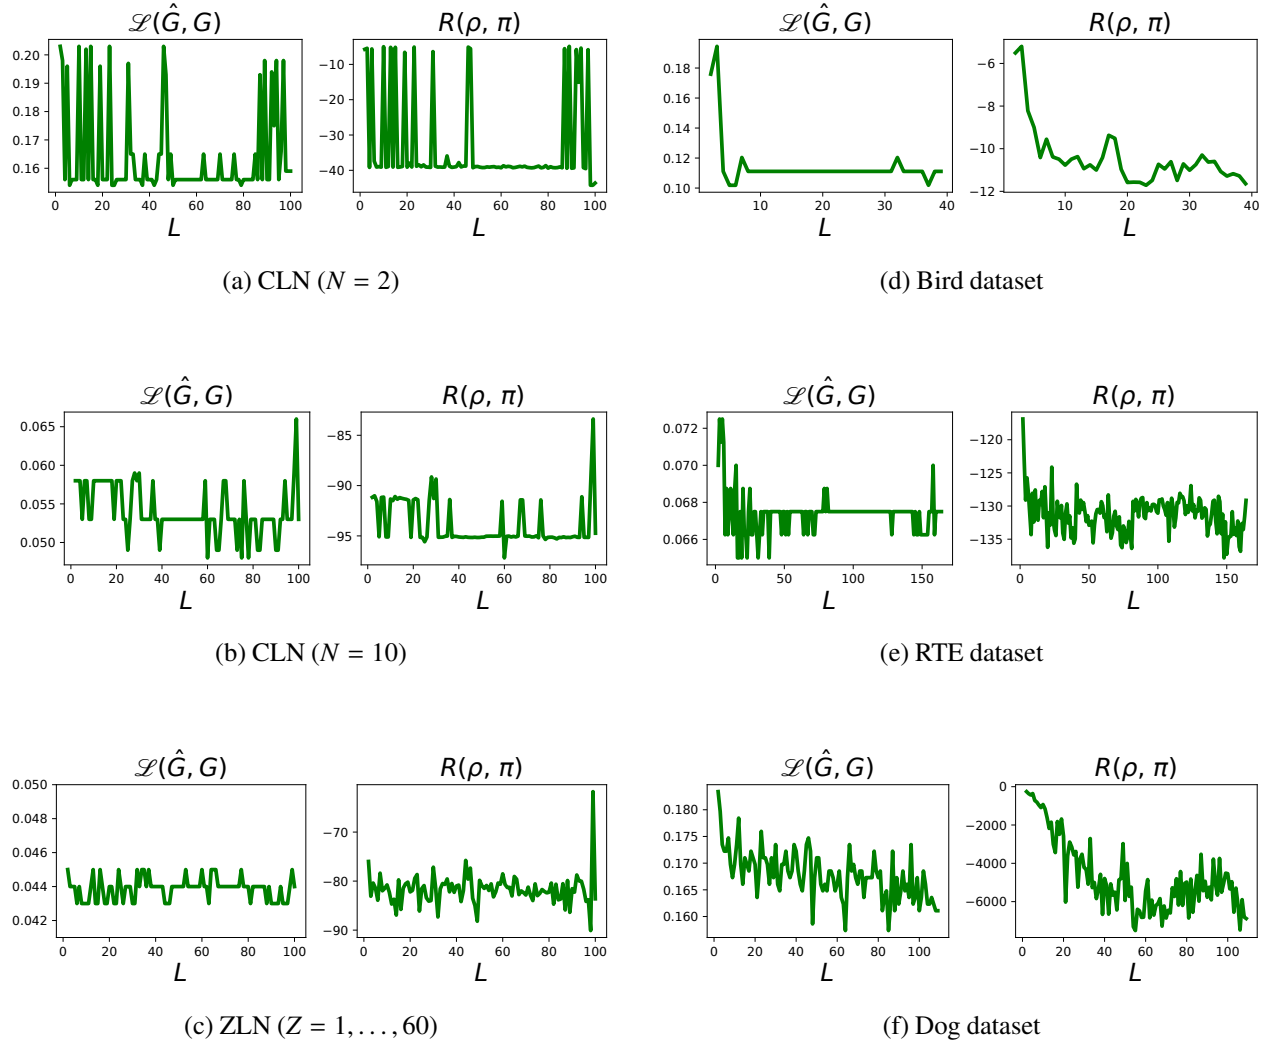


Figure 3. Change in the empirical error $\mathcal{L}(\hat{G}, G)$ and the main part of the lower bound of the minimax error rate $R(\rho, \pi)$ for the maximum number of worker clusters L .

References

- Albert, P. S. and Dodd, L. E. A cautionary note on the robustness of latent class models for estimating diagnostic error without a gold standard. *Biometrics*, 60(2):427–435, 2004.
- Bachrach, Y., Minka, T., Guiver, J., and Graepel, T. How to grade a test without knowing the answers — a Bayesian graphical model for adaptive crowdsourcing and aptitude testing. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, pp. 1183–1190, 2012.
- Bonald, T. and Combes, R. A minimax optimal algorithm for crowdsourcing. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 4355–4363, 2017.
- Chen, X., Lin, Q., and Zhou, D. Optimistic knowledge gradient policy for optimal budget allocation in crowdsourcing. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 2013.
- Dalvi, N., Dasgupta, A., Kumar, R., and Rastogi, V. Aggregating crowdsourced binary ratings. In *Proceedings of the 22nd International Conference on World Wide Web (WWW)*, 2013.
- Dawid, A. and Skene, A. Maximum likelihood estimation of observer error-rates using the em algorithm. *Applied Statistics*, 28(1):20–28, 1979.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.
- Fano, R. M. *Transmission of Information*. M.I.T. Press, 1949.
- Gao, C., Lu, Y., and Zhou, D. Exact exponent in optimal rates for crowdsourcing. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pp. 603–611, 2016.
- Ghosh, A., Kale, S., and McAfee, R. P. Who moderates the moderators?: crowdsourcing abuse detection in user-generated content. In *Proceedings of the 12th ACM conference on Electronic commerce of ACM EC*, pp. 167–176, 2011.
- Hui, S. L. and Walter, S. D. Estimating the error rates of diagnostic tests. *Biometrics*, pp. 167–171, 1980.
- Ipeirotis, P. G., Provost, F., and Wang, J. Quality management on amazon mechanical turk. In *Proceedings of the ACM SIGKDD Workshop on Human Computation*, pp. 64–67, 2010.
- Karger, D., Oh, S., and Shah, D. Iterative learning for reliable crowdsourcing systems. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1953–1961, 2011.
- Karger, D., Oh, S., and Shah, D. Budget-optimal task allocation for reliable crowdsourcing systems. *Operations Research*, 62(1):1–24, 2014.
- Karger, D. R., Oh, S., and Shah, D. Efficient crowdsourcing for multi-class labeling. *ACM SIGMETRICS Performance Evaluation Review*, 41(1):81–92, 2013.
- Liu, C. and Wang, Y. Truelabel + confusions: A spectrum of probabilistic models in analyzing multiple ratings. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, 2012.
- Liu, Q., Peng, J., and Ihler, A. T. Variational inference for crowdsourcing. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 692–700, 2012.
- Moreno, P. G., Teh, Y. W., and Perez-Cruz, F. Bayesian nonparametric crowdsourcing. *Journal of Machine Learning Research*, 16(1):1607–1627, 2015.
- Parisi, F., Strino, F., Nadler, B., and Kluger, Y. Ranking and combining multiple predictors without labeled data. *Proceedings of the National Academy of Sciences*, 111(4):1253–1258, 2014.
- Raykar, V. C., Yu, S., Zhao, L. H., Valadez, G. H., Florin, C., Bogoni, L., and Moy, L. Learning from crowds. *The Journal of Machine Learning Research*, 11:1297–1322, 2010.
- Robbins, H. An empirical bayes approach to statistics. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, 1: Contributions to the Theory of Statistics:157–163, 1956.
- Smyth, P., Fayyad, U., Burl, M., Perona, P., and Baldi, P. Inferring ground truth from subjective labelling of venus images. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1085–1092, 1995.

- Snow, R., O'Connor, B., Jurafsky, D., and Ng, A. Y. Cheap and fast - but is it good? evaluating non-expert annotations for natural language tasks. *In Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2008.
- Tian, T. and Zhu, J. Max-margin majority voting for learning from crowds. *In Advances in Neural Information Processing Systems (NIPS)*, pp. 1612–1620, 2015.
- Venanzi, M., Guiver, J., Kazai, G., Kohli, P., and Shokouhi, M. Community-based bayesian aggregation models for crowdsourcing. *In Proceedings of the 23rd International Conference on World Wide Web*, pp. 155–164, 2014.
- Venanzi, M., Teacy, W. T. L., Rogers, A., and Jennings, N. R. Bayesian modelling of community-based multidimensional trust in participatory sensing under data sparsity. 2015.
- Welinder, P. and Perona, P. Online crowdsourcing: rating annotators and obtaining cost-effective labels. *In IEEE Conference on Computer Vision and Pattern Recognition Workshops (ACVHL)*, 2010.
- Welinder, P., Branson, S., Belongie, S., and Perona, P. The multidimensional wisdom of crowds. *In Advances in Neural Information Processing Systems (NIPS)*, pp. 2424–2432, 2010.
- Whitehill, J., Ruvolo, P., Wu, T., Bergsma, J., and Movellan, J. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. *In Advances in Neural Information Processing Systems (NIPS)*, pp. 2035–2043, 2009.
- Yu, B. Assouad, fano, and le cam. *In Festschrift for Lucien Le Cam*, pp. 423–435, 1997.
- Zhang, Y., Chen, X., Zhou, D., and Jordan, M. Spectral methods meet EM: A provably optimal algorithm for crowdsourcing. *In Advances in Neural Information Processing Systems (NIPS)*, pp. 1260–1268, 2014.
- Zhou, D., Platt, J. C., Basu, S., and Mao, Y. Learning from the wisdom of crowds by minimax entropy. *In Advances in Neural Information Processing Systems (NIPS)*, pp. 2195–2203, 2012.