

Temporal-aware Language Representation Learning From Crowdsourced Labels

Yang Hao Xiao Zhai Wenbiao Ding Zitao Liu*

TAL Education Group, Beijing, China

{haoyang2, zhaixiao, dingwenbiao, liuzitao}@tal.com

Abstract

Learning effective language representations from crowdsourced labels is crucial for many real-world machine learning tasks. A challenging aspect of this problem is that the quality of crowdsourced labels suffer high intra- and inter-observer variability. Since the high-capacity deep neural networks can easily memorize all disagreements among crowdsourced labels, directly applying existing supervised language representation learning algorithms may yield suboptimal solutions. In this paper, we propose *TACMA*, a temporal-aware language representation learning heuristic for crowdsourced labels with multiple annotators. The proposed approach (1) explicitly models the intra-observer variability with attention mechanism; (2) computes and aggregates per-sample confidence scores from multiple workers to address the inter-observer disagreements. The proposed heuristic is extremely easy to implement in around 5 lines of code. The proposed heuristic is evaluated on four synthetic and four real-world data sets. The results show that our approach outperforms a wide range of state-of-the-art baselines in terms of prediction accuracy and AUC. To encourage the reproducible results, we make our code publicly available at <https://github.com/CrowdsourcingMining/TACMA>.

1 Introduction

Crowdsourcing offers the ability to utilize the power of human computation to generate data annotations that are needed to train various AI systems. For many practical supervised learning applications, it may be infeasible (or very expensive) to obtain objective and reliable labels due to many reasons such as varying skill-levels and biases of crowdsourced workers. Instead, to improve the quality of labels, we can collect subjective

and inconsistent labels from multiple heterogeneous crowdsourced workers. In practice, there is a substantial amount of disagreement between the crowdsourced workers (Nie et al., 2020), i.e., inter-observer variability or even between a worker and the same worker looking at the same example some time later (Guan et al., 2018), i.e., intra-observer variability. Hence, it is of great practical interest to address supervised learning problems in this scenario.

Meanwhile, with the recent advances of deep neural networks (DNNs), supervised representation learning (SRL) has led to rapid improvements in the ability of learning intrinsic nonlinear embeddings using DNNs that preserves the distance between similar examples close and dissimilar examples far on the embedding space. In spite of the significant progress for SRL applications such as face recognition (Schroff et al., 2015), image retrieval (Xia et al., 2014), directly applying existing deep language representation learning approaches on crowdsourced labels may yield poor generalization performance (Han et al., 2018). Because of the high capacity, DNNs could entirely memorize the inconsistency within crowdsourced labels sooner or later during the modeling training process. Besides, this phenomenon does not change with the choice of training optimizations or network architectures (Han et al., 2018).

A large spectrum of approaches have been successfully developed in either estimating true labels from crowdsourced labels, a.k.a., truth inference or label aggregation (Dawid and Skene, 1979; Whitehill et al., 2009), learning via adversarial data generation (Wang et al., 2020a), or learning language representations discriminatively from large-scale consistent labeled data with complicated neural architectures (Rodrigues and Pereira, 2018). However, learning effective neural embeddings directly from crowdsourced labels of real-world data poses

*Corresponding author: Zitao Liu.

numerous challenges. First, crowdsourced workers conduct labeling tasks sequentially, i.e., they label samples one after another. Such sequential labeling behavior is a process of learning, and the expertise of the workers is not stable but gradually changing even without feedback (Elliott and Riach, 1965). According to Miller’s Law (Miller, 1956), humans retain what they just learned in their short-term working memory with a limited span of 7 ± 2 . Temporal factors such as fatigue (Zhang et al., 2018) and intrinsic motivation (Kaufmann et al., 2011) implicitly influence the crowdsourcing quality, which are different from existing well-studied factors, such as the quality of crowdsourced workers, the difficulty of data samples, the price of annotation tasks, etc. In the following, such unconscious temporal behaviors are referred to as “*temporal labeling effects*”. How to model such sample-level temporal information for each individual worker undoubtedly poses a hard modeling problem. Second, a large number of real-world crowdsourced data sets have a substantial amount of disagreement among labels and a relatively small sample size. The majority of existing SRL approaches are discriminatively trained on large-scale consistent labeled data to learn their complicated neural architectures, which may easily overfit the inconsistent crowdsourced data.

In this paper we study and develop solutions that are applicable and can learn effective neural language representations from crowdsourced labels in an end-to-end manner. Our work focuses on the refinements of a popular deep language representation learning paradigm: the deep metric learning (DML) (Koch et al., 2015; Xu et al., 2019; Wang et al., 2020b). We aim to develop an algorithm to automatically learn a nonlinear language representation of the crowdsourced data from multiple workers using DNNs.

Briefly, the DML is a classical and widely used approach for language representation learning that preserves the distance between similar examples close and dissimilar examples far on the embedding space. The majority of existing DML techniques restricted to just noise-free labels appropriately. However, learning effective representation from highly inconsistent crowdsourced data sets from multiple workers gives rise to numerous important questions: (1) since in practice, annotation performance is affected and varied over time (Boksem et al., 2005; Zhang et al., 2018), *how do we capture*

such temporal labeling effects in the DML learning framework? (2) while in some cases the problem may be alleviated by pre-processing methods, such as filtering (Li et al., 2016), label correction (Li et al., 2019a), truth inference (Dawid and Skene, 1979; Raykar et al., 2010), etc., the number of remained instances is often significantly reduced or such pre-processing errors for many problems will be propagated to the downstream representation learning tasks. *How to capture the label uncertainties from multiple workers and at the same time prevent the overfitting problem in an end-to-end framework?*

In this work we address the above issues by presenting a temporal-aware language representation learning heuristic for crowdsourced labels with multiple annotators (TACMA), that

- utilizes the attention mechanism to capture the temporal influence among sequential labeling tasks according to each worker’s short-term working memory.
- estimates and aggregates the annotation confidence from disagreements among multiple workers for each sample.
- supports language representation learning with DML into an end-to-end fashion, and is extremely easy to implement based on existing DML framework with crowdsourced labels i.e., RLL (Xu et al., 2019), in around 5 line of codes.

2 Related Work

2.1 Truth Inference in Crowdsourcing

A large body of research has focused on inferring true labels from crowdsourced labels from multiple workers (Dawid and Skene, 1979; Whitehill et al., 2009; Li et al., 2019c; Rodrigues and Pereira, 2018). The majority of truth inference approaches are inspired by the classic Expectation-Maximization learning paradigm that iterates between estimating the expertise of annotators given true labels inferred and inferring true labels given the expertise of annotators (Dawid and Skene, 1979; Whitehill et al., 2009; Zhang et al., 2014; Li et al., 2019c). Some improvements include modeling the difficulty of items and the expertise of annotators jointly (Whitehill et al., 2009), applying spectral methods to initialize worker confusion matrix (Zhang et al., 2014), and modeling correlations of workers (Li et al., 2019c), etc.

In spite of the successful applications of the truth inference techniques, the majority of aforementioned approaches do not consider the temporal effects of labeling tasks of each individual worker and they cannot seamlessly integrate into deep SRL frameworks.

2.2 Learning from Noisy Labels

Learning with noisy labels has been an important research topic since the beginning of machine learning (Frénay and Verleysen, 2013) and a large spectrum of models have been developed and successfully applied in improving the model prediction performance in noisy settings from different perspectives such as effective label cleaning (Lee et al., 2018), robust model architectures (Vahdat, 2017) and loss functions (Ghosh et al., 2017), sample re-weighting (Ren et al., 2018), and carefully designed training procedures (Zhong et al., 2019).

However, in this work, different from above approaches of robust learning from noisy labels that assume certain percentage of labels are corrupted, our scenario focuses on noisy labels obtained from multiple annotators where the disagreement (corruption) proportion might be surprisingly high and sometimes even 100%, i.e., no completely agreement on every single sample from all crowd workers.

2.3 Deep Metric Learning

DML approaches automatically learn nonlinear metric spaces (Schroff et al., 2015). Many approaches have achieved promising results in many tasks such as face recognition (Schroff et al., 2015), person re-identification (Yi et al., 2014), and collaborative filtering (Hsieh et al., 2017) etc. Recently a body of works have attempted to learn effective embeddings from crowdsourced labels by using DML approaches (Xu et al., 2019; Wang et al., 2020b). For example, Xu et al. estimated crowdsourced label confidence and adjust the DML loss function accordingly (Xu et al., 2019). An exhaustive review of previous work is beyond the scope of this paper. We refer to the survey of (Schroff et al., 2015) on works of DML. Although DML approaches are able to learn effective representations, they heavily rely on comparisons within pairs or triplets, which is very sensitive to ambiguous examples and may be easily misled by inconsistent crowdsourced labels.

Please note that models from the above three categories are complementary and they can be

combined. For example, learning representation from crowdsourced labels can be conducted in two stages where the truth inference algorithms in Section 2.1 is applied to get estimated labels and then the standard DML approaches in Section 2.3 are used to output the learned embeddings. Details are discussed in Section 4.

3 The Proposed Approach

3.1 Notations

Without loss of generality, we consider crowdsourcing scenarios that each data sample is annotated by multiple workers. Following the crowdsourcing practice and to avoid the order effect (Hogarth and Einhorn, 1992) and cheating, each worker will annotate the same set of samples but with shuffled orders. Let α^j be the sample order index set for the j^{th} worker and α_i^j be the index of i^{th} sample for worker j . Let $\mathbf{x}_{\alpha_i^j}$ and $y_{\alpha_i^j}$ be the feature vectors and the worker’s assigned label for sample α_i^j . Let $\mathcal{F}(\cdot)$ represent the learned language representation. Let $(\cdot)^+$ and $(\cdot)^-$ be the indicators of positive and negative examples.

3.2 Temporal-Aware Memory Confidence

According to Miller’s Law (Miller, 1956), humans can only hold a very limited number of objects in their short-term working memories. When workers conduct labeling tasks, they tend to make relative comparisons in their memory spans and the annotation quality of one sample is largely influenced by its preceding samples. Therefore, in this work, we focus on studying and modeling the effects of unconscious human behaviors during the labeling process that may implicitly influence the overall crowdsourcing quality. We design an approach to explicitly capture such unconscious temporal human behaviors, i.e., temporal labeling effects. We aim to ensure that *the newly annotated samples should obtain the consistent label with similar samples that have already been annotated recently*. Here we first define the short-term labeling memory as follows:

Definition 1. (SHORT-TERM LABELING MEMORY) *A short-term labeling memory of i^{th} sample, i.e., indexed as α_i^j , is composed of a sequence of the current item and k most recent historical items that have been labeled by worker j , i.e.,*

$$\mathbf{M}_i^j = \{ \langle \mathbf{x}_{\alpha_i^j}, y_{\alpha_i^j} \rangle, \langle \mathbf{x}_{\alpha_{i-1}^j}, y_{\alpha_{i-1}^j} \rangle, \dots, \langle \mathbf{x}_{\alpha_{i-k}^j}, y_{\alpha_{i-k}^j} \rangle \}.$$

When the new labeling task arrives, i.e., the i^{th} sample, we compute a weight for every element in worker j 's short-term labeling memory \mathbf{M}_i^j as the dot product of their learned language representations. This weight might be viewed as an attention over the short-term labeling memory per sample per worker.

To form a proper probability distribution over the elements in \mathbf{M}_i^j , we normalize the weights using the softmax function. This way we model probability $s_{\alpha_{i-l}^j}$ that represents the similarity between the i^{th} sample and the sample appears at position l in \mathbf{M}_i^j . In a functional form this is:

$$s_{\alpha_{i-l}^j} \propto \exp\left(\mathcal{F}(\mathbf{x}_{\alpha_{i-l}^j}), \mathcal{F}(\mathbf{x}_{\alpha_i^j})\right), \quad l = 0, \dots, k$$

Then we define a memory confidence score, i.e., c_i^j , to represent the probability that how likely the sample i is positive ($y_{\alpha_i^j} = 1$) solely considering similar samples in the short-term labeling memory. The memory confidence score of c_i^j is computed as follows:

$$c_i^j = \Pr(y_{\alpha_i^j} = 1) \propto \sum_{l=0}^k \mathbb{1}[y_{\alpha_{i-l}^j} = 1] s_{\alpha_{i-l}^j}$$

Please note that our attention based temporal-aware memory confidence scores are not limited to binary crowdsourcing tasks and it can be easily extended to multi-class tasks.

3.3 Multi-Worker Confidence Aggregation

For each sample i , after collecting the memory confidence scores from all workers, we conduct the mean pooling as our aggregation operation, and the final aggregated multi-worker confidence is computed as follows: $c_i = \text{MeanPooling}(c_i^1, c_i^2, \dots, c_i^m)$, where m is the number of workers.

3.4 Representation Learning Framework

We use DML as our representation learning framework. Specifically, following the suggestion of (Xu et al., 2019), instead of using pair and triplet comparisons, we use group, a.k.a., n -tuple, as our comparison unit. A group is made up of two positive and n negative examples. Similar to (Xu et al., 2019), we choose to learn our model parameters by maximizing the conditional likelihood of retrieving

the positive example \mathbf{x}_j^+ given the positive example \mathbf{x}_i^+ from a given group.

Importantly, we do not assume that we know the ground truth label of items in the training set and the validation set. During the training stage of the representation learning framework, after obtaining the aggregated multi-worker confidence c_i of an item with methods introduced in Section 3.3, its label is estimated by $\arg \max c_i$.

Given a collection of groups, we optimize the DML model parameters by maximizing the sum of log conditional likelihood of finding a positive example \mathbf{x}_j^+ given the paired positive example \mathbf{x}_i^+ within every group \mathbf{g} , which will push items of the same class close and items of different classes far in the embedding space. Furthermore, we incorporate the aggregated temporal-aware multi-worker confidence scores from Section 3.3 into the loss function to capture the inconsistency of crowdsourced labels. The loss function is defined as $\mathcal{L}(\Omega) = -\sum \log p(\mathbf{x}_j^+ | \mathbf{x}_i^+)$,

$$p(\mathbf{x}_j^+ | \mathbf{x}_i^+) = \frac{\exp(\eta \cdot c_j \cdot r_{ij})}{\sum_{\mathbf{x}_* \in \mathbf{g}, \mathbf{x}_* \neq \mathbf{x}_i^+} \exp(\eta \cdot c_* \cdot r_{i*})}$$

where Ω is the parameter set of the DNN. r_{i*} represents the cosine similarity score between the representations of \mathbf{x}_i^+ and \mathbf{x}_* in the embedding space. η is a smoothing hyper parameter in the softmax function, which is set empirically on a held-out data set in our experiment. Since $\mathcal{L}(\Omega)$ is differentiable with respect to Ω , we use gradient based optimization approach to train the DNN.

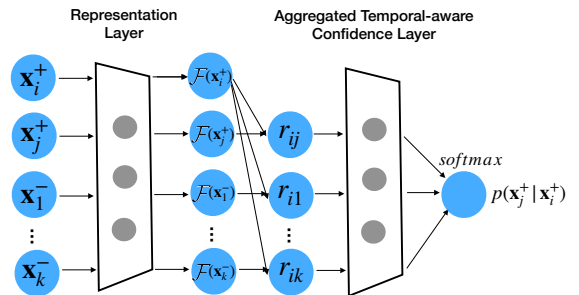


Figure 1: The model structure. Groups made up of two positive and n negative examples are fed into the neural network to obtain their language representations. The cosine similarity scores, i.e., r_{i*} , are calculated between the representations of \mathbf{x}_i^+ and \mathbf{x}_* in the embedding space. Finally, the goal of training is to maximize the conditional likelihood $p(\mathbf{x}_j^+ | \mathbf{x}_i^+)$, which incorporates temporal-aware memory confidence scores c_i^j .

4 Experiments

Experiments are conducted on both real-world and synthetic data sets. The internal cross validation approach is used to select hyper parameters when optimizing models’ predictive performances. Means as well as standard deviations of both accuracy and AUC scores are reported, to comprehensively evaluate the performance of our proposed method, i.e. *TACMA*.

4.1 Real-World Data Sets

Experiments are first conducted on 4 real-world data sets and the corresponding descriptive statistics can be found in Table 1.

- *Emotion*: A vocal emotional speech data set with binary labels indicating whether the voice fragment is exciting or not.
- *Concluding*: A linguistic data set where each item is labeled on whether it is a conclusion of a lesson.
- *Commending*: A linguistic data set of ASR transcripts from real-world classroom recordings. Each item is labeled on whether it’s a commending instruction from the instructors.
- *Question*: A vocal speech data set where each item is labeled on whether it is an interrogative sentence.

Acoustic features of the *Emotion* data set are extracted using OpenSmile¹ with the computational paralinguistic challenge’s (COMPARE-2013) feature set (Schuller et al., 2013). Sentence embedding features are extracted with a Chinese RoBERTa pretrained model². Again we emphasize that the ground truth labels of items in the training and validation set are not observed. In order to evaluate the performance of each model objectively, the labels of items in test sets are labeled by experts and they have reached an agreement on the labels of items.

Inter-observer variability of each data set is measured with Fleiss-kappa score (Fleiss, 1971). Intra-observer variability, i.e., the level of consistency of an annotator when labeling items from the same class, is hard to directly measure without ground truth labels. We will explore the effect of intra-observer variability using temporal-aware memory confidence in Section 4.8.

¹<https://www.audeering.com/opensmile/>

²<https://github.com/ymcui/Chinese-BERT-wwm>

4.2 Synthetic Data Sets

In real-world scenarios, annotators are not guaranteed to be serious about their annotating work, and one may assign random labels in order to get paid quickly. Methods designed for crowdsourcing scenarios should be able to get rid of the influence of these noisy annotations. Hence we build synthetic data sets to evaluate the robustness to irresponsible annotators of each method. Starting from the original *Question* data set, we gradually add 2, 4, 6 and 8 simulated irresponsible annotators. They make random judgments regardless of the features of items. Hence in the worst case, 8 out of 13 workers are making random judgments, resulting in an extreme low kappa of 0.02. Experiments conducted on these synthetic data sets are helpful to examine the robustness of methods.

4.3 Baselines

We carefully selected several groups of baselines as follows:

Group 1: Truth Inference. A wide range of label aggregation methods are chosen as our baselines. Some widely-used methods according to the survey (Zheng et al., 2017) are included, i.e., *EM* (Dawid and Skene, 1979), *Spectral-EM* (Zhang et al., 2014), *GLAD* (Whitehill et al., 2009), *IBCC* (Kim and Ghahramani, 2012), *VI-BP* (Qiang et al., 2012), *VI-MF* (Qiang et al., 2012), *KOS* (Karger et al., 2011), *ZenCrowd* (Demartini et al., 2012), *LFC* (Raykar et al., 2010), *PM* (Li et al., 2014), and the implementation of these algorithms can mostly be found in the website³. Meanwhile some more recent works are also included: *EBCC* (Li et al., 2019c), *BWA* (Li et al., 2019b).

Group 2: Representation Learning. Our proposed method is compared with representation learning methods via deep metric learning, including Triplet with semi-hard example mining (Schroff et al., 2015), i.e., *Triple*, and Triplet networks with Center Loss (He et al., 2018), i.e., *Center*. Recent works of learning effective embeddings from crowdsourced labels using DML approaches are also important baselines: *RLL-MLE* (Xu et al., 2019), *RLL-Bayesian* (Xu et al., 2019), *RECLE* (Wang et al., 2020b).

Group 3: Learning from Noisy Data. Group 3 contains methods of learning with noisy labels: *LC* (Arazo et al., 2019) use a two-component beta mixture model to perform unsupervised noise mod-

³https://zhydhkcws.github.io/crowd_truth_inference/index.html

Table 1: Data sets statistics. Data sets statistics. It should be noted that the class ratio of each training set is estimated by majority voting since the ground truth labels are not observed. The labels of items in each test set are annotated by experts and they have reached an agreement on the label of each item.

Data Sets	Emotion	Commending	Question	Concluding	Syn-2	Syn-4	Syn-6	Syn-8
# of annotators	5	7	5	5	7	9	11	13
# of train samples	3067	1200	3140	1208	3140	3140	3140	3140
# of validation samples	766	299	785	302	785	785	785	785
# of test samples	800	1300	2000	648	2000	2000	2000	2000
kappa	0.84	0.69	0.82	0.37	0.35	0.2	0.12	0.08
train class ratio (majority voting)	0.42	0.50	0.63	0.42	0.63	0.63	0.63	0.63

eling, and *DivideMix* (Li et al., 2019a) leverages semi-supervised learning techniques. *CrowdLayer* (Rodrigues and Pereira, 2018) is an end-to-end approach learning a DNN from noisy labels with a crowd layer.

Group 4: Combining Group 1 with Groups 2 & 3. Some methods of Group 2 & 3, i.e. *Triple*, *Center*, *LC*, *DivideMix*, are not specifically designed for crowdsourcing scenarios. Although majority-voting labels are served as a default choice, these models should be trained with labels inferred by methods of Group 1 as stronger baselines, since methods of Group 1 are likely to provide more accurate inferred labels than majority voting. These methods are therefore trained with labels inferred by *EBCC*, which achieves the best performances of Group 1 in all data sets.

4.4 Setup and Implementation Details

Experimental codes are implemented in Tensorflow 1.8 available at <https://github.com/CrowdsourcingMining/TACMA>.

Experiments are conducted on a server with a GTX 1080 Ti GPU. We set the tuple size n to 5 for all the experiments, as suggested in (Xu et al., 2019). The representation learning network has a simple structure, i.e., 2 fully-connected layers with a drop-out rate of 0.2, a learning rate of $1e-3$, and hyper-parameters including sizes of each layer and scale of ℓ_2 regularization searched via grid searching with cross validation. The network weights are initialized with a normal distribution initializer and updated with Adadelta optimizer (Zeiler, 2012). For all the representation learning methods, the downstream classifier is set to be a logistic regression classifier with ℓ_2 penalty containing the only hyper-parameter C as penalty strength ranging from $1e-2$ to $1e4$.

4.5 Performance Comparison

We compare performance of *TACMA* with existing methods on 4 real-world data sets and the results are summarized in Table 2. *TACMA* outperforms all the 4 groups of baselines, and here are some observations:

- The advantage of *TACMA* over truth inference methods gets bigger on the *Concluding* data set than other data sets. The *Concluding* data set has a low kappa score of 0.37, indicating that there are more disagreements among workers, which makes it hard to inference correct labels regardless of items’ features. By contrast, *TACMA* makes full use of representations of items to gain more information resulting in the best performance.
- Although labels inferred by *EBCC* boost the performances of representation learning models, e.g., *Triple+EBCC*, they still perform inferior to *TACMA*, a possible explanation is that these two-stage methods give equal weight to each item and ignores temporal labeling effects. *TACMA* is able to discover potential conflicts in the short-term working memory, by applying the attention mechanism and gives low weights to the conflicting judgments.
- *TACMA* shares the same representation network structures with other methods of representation learning with crowdsourced labels i.e., *RLL-MLE*, *RLL-bayesian* and *RECLE*. The learned representations are compared in Figure 2 by feeding the raw features into representation network and performing dimension reduction into 2-dimensional space with t-SNE method (Van and Hinton, 2008). In the raw feature space, items of different classes are interleaved with each other. By contrast, learned representations of *TACMA* are more

Table 2: Prediction accuracy and AUC scores on 4 real-world data sets. The experiments are repeated 5 times and the means and standard deviations are reported.

	Commending		Emotion		Question		Concluding	
	ACC	AUC	ACC	AUC	ACC	AUC	ACC	AUC
EM	0.794±0.019	0.871±0.008	0.883±0.012	0.967±0.005	0.877±0.010	0.941±0.005	0.681±0.004	0.720±0.015
Spectral-EM	0.794±0.017	0.870±0.007	0.886±0.010	0.964±0.003	0.876±0.009	0.941±0.004	0.681±0.004	0.720±0.013
GLAD	0.794±0.017	0.870±0.007	0.886±0.010	0.964±0.003	0.878±0.009	0.942±0.004	0.689±0.009	0.742±0.014
IBCC	0.794±0.017	0.870±0.007	0.889±0.004	0.964±0.004	0.876±0.009	0.941±0.004	0.681±0.004	0.720±0.013
VI-BP	0.794±0.017	0.870±0.007	0.892±0.012	0.968±0.005	0.877±0.008	0.941±0.004	0.681±0.004	0.720±0.013
VI-MF	0.799±0.013	0.874±0.003	0.786±0.000	0.898±0.000	0.876±0.009	0.941±0.004	0.685±0.003	0.725±0.009
KOS	0.799±0.013	0.874±0.003	0.786±0.000	0.898±0.000	0.878±0.009	0.942±0.004	0.694±0.005	0.747±0.010
ZenCrowd	0.794±0.019	0.871±0.008	0.895±0.011	0.971±0.004	0.877±0.010	0.941±0.005	0.689±0.010	0.742±0.016
LFC	0.794±0.019	0.871±0.008	0.883±0.009	0.967±0.004	0.877±0.010	0.941±0.005	0.681±0.004	0.720±0.015
PM	0.799±0.014	0.867±0.008	0.887±0.010	0.966±0.003	0.874±0.010	0.940±0.004	0.677±0.009	0.730±0.013
EBCC	0.812±0.006	0.874±0.003	0.895±0.012	0.970±0.005	0.878±0.007	0.941±0.008	0.694±0.003	0.748±0.006
BWA	0.794±0.020	0.867±0.008	0.888±0.005	0.965±0.004	0.875±0.009	0.939±0.004	0.689±0.007	0.741±0.013
Triple	0.793±0.012	0.871±0.006	0.804±0.005	0.876±0.002	0.888±0.002	0.941±0.001	0.725±0.014	0.821±0.008
Center	0.806±0.002	0.859±0.001	0.701±0.007	0.780±0.007	0.840±0.006	0.905±0.008	0.705±0.015	0.797±0.007
RLL-MLE	0.800±0.008	0.866±0.001	0.854±0.016	0.961±0.008	0.853±0.013	0.919±0.006	0.735±0.004	0.828±0.009
RLL-Bayesian	0.816±0.000	0.861±0.001	0.877±0.006	0.954±0.004	0.877±0.004	0.932±0.003	0.725±0.001	0.839±0.001
RECLE	0.812±0.002	0.858±0.000	0.746±0.001	0.836±0.001	0.880±0.024	0.934±0.012	0.729±0.003	0.838±0.005
LC	0.560±0.085	0.700±0.028	0.611±0.046	0.715±0.007	0.715±0.018	0.720±0.007	0.701±0.018	0.790±0.011
DivideMix	0.515±0.016	0.733±0.014	0.535±0.000	0.730±0.000	0.734±0.009	0.720±0.014	0.654±0.025	0.710±0.007
CrowdLayer	0.802±0.008	0.878±0.007	0.757±0.008	0.798±0.009	0.852±0.003	0.920±0.002	0.676±0.014	0.722±0.011
LC+EBCC	0.581±0.070	0.687±0.029	0.825±0.024	0.845±0.019	0.758±0.010	0.830±0.012	0.705±0.018	0.784±0.004
DivideMix+EBCC	0.515±0.018	0.730±0.014	0.726±0.039	0.832±0.028	0.760±0.012	0.833±0.014	0.659±0.006	0.720±0.005
Triple+EBCC	0.814±0.004	0.872±0.000	0.893±0.003	0.968±0.004	0.890±0.001	0.938±0.003	0.737±0.003	0.825±0.007
Center+EBCC	0.814±0.004	0.866±0.003	0.826±0.016	0.884±0.018	0.844±0.005	0.909±0.005	0.742±0.006	0.848±0.003
TACMA	0.831±0.002	0.882±0.004	0.904±0.002	0.973±0.001	0.899±0.005	0.945±0.003	0.765±0.006	0.855±0.010

separated than the other methods, reducing the difficulty of downstream classification tasks.

4.6 Robustness to Irresponsible Workers

We select some representatives from Groups 1-4 and draw the curves of accuracy on synthetic data sets containing different number of irresponsible workers in Figure 3. We can find that:

- Truth inference methods such as *EBCC* stay stable facing different numbers of irresponsible workers. On the other hand, the accuracy of other methods decreases when increasing the number of irresponsible workers. This result may be explained by the fact that for methods including *RLL-Bayesian*, *Triple*, learning effective representations of items heavily relies on correct labels, and hence becomes harder as the labels become more noisy.
- *TACMA* maintains the highest accuracy of all the methods. Unlike the two-stage method i.e., *Triplet + EBCC*, which gives equal weight to each item and ignores *temporal labeling effects*, *TACMA* is able to discover potential conflicts in the short-term working memory using the attention mechanism, and give low training weights to the conflicting judgments.

4.7 Effect of Working Memory Sizes

We set the working memory size ranging from 3 to 11 to find the optimized length and at the same time explore its influence on performance, shown in Figure 4. The accuracy of our proposed method goes up at the beginning with the increasing working memory size, and the standard deviations gradually become smaller at the same time. It is reasonable because potential inconsistent judgments among similar items cannot be found without observing enough historical annotations. As the working memory size continues extending, the accuracy scores become relatively stable, indicating that there is sufficient evidence to estimate the time-aware confidence of the current annotation.

4.8 Relations between Temporal-aware Memory Confidence and Worker’s Expertise

In this part we further explore the relations between worker’s expertise and temporal-aware memory confidence. To evaluate a worker’s expertise, a Logistic Regression classifier is trained with labels annotated by this same person, and the accuracy on the corresponding test set is recorded. On the other hand, the temporal-aware confidence of all




Figure 2: Visualization of learned representations on the test set of Question data. The raw features of items are fed into the representation network to obtain the semantic representations, and dimension reduction using t-SNE method is performed for visualization.




Figure 3: Accuracy curves on synthetic data sets containing different number of irresponsible annotators who make random decisions.




Figure 4: The effect of different working memory sizes on prediction accuracy on real-world data sets.

the judgments made by this worker is averaged.

We perform standardization on both accuracy scores and the averaged temporal-aware confidence scores within the corresponding data set, and put the standardized values of all the 62 workers from 4 real-world data sets and 4 synthetic data sets together in Figure 5, to reveal the universal relation between temporal-aware confidence and the worker’s expertise. We can find a wide range of intra-observer variability among different workers, estimated by their temporal-aware confidence scores. A strong positive correlation is found between averaged confidence and prediction accuracy (*pearson r* = 0.844). Specifically, synthetic irresponsible annotators, colored in blue, are automatically clustered in the lower left corner, indicating that the poor performances of the classifiers trained with their labels derive from huge inner inconsistencies in their judgments.




Figure 5: The relations between standardized temporal-aware memory confidence and standardized prediction accuracy of annotators in both real and synthetic data sets. Most of the irresponsible annotators appear in the lower left corner, indicating that there are internal conflicts in their judgments (low confidence), and therefore LR models trained with these labels perform worse than average.

5 Conclusion

We presented TACMA, an end-to-end framework for language representation learning from crowd-sourced labels. Comparing with traditional SRL approaches, the advantages of our framework are: (1) it is able to consider temporal labeling effects within sequences of sample-level labeling tasks for each worker; (2) it automatically computes and aggregates sample-level confidence scores from multiple workers which makes the training process more effective. Experimental results on both synthetic and real-world data sets demonstrates that our approach outperforms other state-of-the-art baselines in terms of accuracy and AUC scores.

References

- Eric Arazo, Diego Ortego, Paul Albert, Noel O’Connor, and Kevin Mcguinness. 2019. Unsupervised label noise modeling and loss correction. In *ICML*, pages 312–321.
- Maarten AS Boksem, Theo F Meijman, and Monique M Lorist. 2005. Effects of mental fatigue on attention: an erp study. *Cognitive brain research*, 25(1):107–116.
- Alexander Philip Dawid and Allan M Skene. 1979. Maximum likelihood estimation of observer error-rates using the em algorithm. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 28(1):20–28.
- Gianluca Demartini, Djellel Difallah, and Philippe Cudre-Mauroux. 2012. *Zencrowd: Leveraging probabilistic reasoning and crowdsourcing techniques for large-scale entity linking*. pages 469–478.
- DN Elliott and WD Riach. 1965. Effect of repeated practice with and without feedback upon discrimination performance. *The Journal of the Acoustical Society of America*, 37(6):1194–1194.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5):378.
- Benoît Fréney and Michel Verleysen. 2013. Classification in the presence of label noise: a survey. *IEEE transactions on neural networks and learning systems*, 25(5):845–869.
- Aritra Ghosh, Himanshu Kumar, and PS Sastry. 2017. Robust loss functions under label noise for deep neural networks. In *AAAI*, volume 31.
- Melody Guan, Varun Gulshan, Andrew Dai, and Geoffrey Hinton. 2018. Who said what: Modeling individual labelers improves classification. In *AAAI*, volume 32.
- Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, pages 8527–8537.
- Xinwei He, Yang Zhou, Zhichao Zhou, Song Bai, and Xiang Bai. 2018. Triplet-center loss for multi-view 3d object retrieval. In *CVPR*, pages 1945–1954.
- Robin M Hogarth and Hillel J Einhorn. 1992. Order effects in belief updating: The belief-adjustment model. *Cognitive psychology*, 24(1):1–55.
- Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. 2017. *Collaborative metric learning*. WWW ’17, page 193–201, Republic and Canton of Geneva, CHE. WWW Steering Committee.
- David R. Karger, Sewoong Oh, and Devavrat Shah. 2011. Iterative learning for reliable crowdsourcing systems. In *NIPS*, pages 1953–1961.
- Nicolas Kaufmann, Thimo Schulze, and Daniel Veit. 2011. More than fun and money. worker motivation in crowdsourcing-a study on mechanical turk. In *Amcis*, volume 11, pages 1–11. Detroit, Michigan, USA.
- Hyun-Chul Kim and Zoubin Ghahramani. 2012. Bayesian classifier combination. In *Artificial Intelligence and Statistics*, pages 619–627.
- Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. 2015. Siamese neural networks for one-shot image recognition. In *ICML Deep Learning Workshop*, volume 2. Lille.
- Kuang-Huei Lee, Xiaodong He, Lei Zhang, and Linjun Yang. 2018. Cleannet: Transfer learning for scalable image classifier training with label noise. In *CVPR*, pages 5447–5456.
- Chaoqun Li, Victor S Sheng, Liangxiao Jiang, and Hongwei Li. 2016. Noise filtering to improve data and model quality for crowdsourcing. *Knowledge-Based Systems*, 107:96–103.
- Junnan Li, Richard Socher, and Steven CH Hoi. 2019a. Dividemix: Learning with noisy labels as semi-supervised learning. In *ICLR*.
- Qi Li, Yaliang Li, Jing Gao, Bo Zhao, Wei Fan, and Jiawei Han. 2014. Resolving conflicts in heterogeneous data by truth discovery and source reliability estimation. In *SIGMOD*, pages 1187–1198.
- Yuan Li, Benjamin IP Rubinstein, and Trevor Cohn. 2019b. Truth inference at scale: A bayesian model for adjudicating highly redundant crowd annotations. In *WWW*, pages 1028–1038.
- Yuan Li, Benjamin Rubinstein, and Trevor Cohn. 2019c. Exploiting worker correlation for label aggregation in crowdsourcing. In *ICML*, pages 3886–3895.

- George A Miller. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2):81.
- Yixin Nie, Xiang Zhou, and Mohit Bansal. 2020. What can we learn from collective human opinions on natural language inference data? *EMNLP*.
- Liu Qiang, Peng Jian, and Alexander Ihler. 2012. Variational inference for crowdsourcing. *NIPS*, 25.
- Raykar, Vikas, C., Shipeng, Yu, Zhao, Linda, H., Valadez, and Gerardo. 2010. Learning from crowds. *JMLR*.
- Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. 2018. Learning to reweight examples for robust deep learning. *arXiv preprint arXiv:1803.09050*.
- Filipe Rodrigues and Francisco Pereira. 2018. Deep learning from crowds. In *AAAI*, volume 32.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823.
- Björn Schuller, Stefan Steidl, Anton Batliner, Alessandro Vinciarelli, Klaus Scherer, Fabien Ringeval, Mohamed Chetouani, Felix Weninger, Florian Eyben, Erik Marchi, et al. 2013. The interspeech 2013 computational paralinguistics challenge: Social signals, conflict, emotion, autism. In *Proceedings INTERSPEECH 2013, 14th Annual Conference of the International Speech Communication Association, Lyon, France*.
- Arash Vahdat. 2017. Toward robustness against label noise in training deep discriminative neural networks. In *NIPS*, pages 5596–5605.
- L Van, der Maaten and Ge Hinton. 2008. (visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research*, 9(2):2579–2605.
- Wentao Wang, Tyler Derr, Yao Ma, Suhang Wang, Hui Liu, Zitao Liu, and Jiliang Tang. 2020a. Learning from incomplete labeled data via adversarial data generation. In *2020 IEEE International Conference on Data Mining (ICDM)*, pages 1316–1321. IEEE.
- Wentao Wang, Guowei Xu, Wenbiao Ding, Yan Huang, Guoliang Li, Jiliang Tang, and Zitao Liu. 2020b. Representation learning from limited educational data with crowdsourced labels. *TKDE*.
- Jacob Whitehill, Ting-fan Wu, Jacob Bergsma, Javier R Movellan, and Paul L Ruvolo. 2009. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. In *NIPS*, pages 2035–2043.
- Rongkai Xia, Yan Pan, Hanjiang Lai, Cong Liu, and Shuicheng Yan. 2014. Supervised hashing for image retrieval via image representation learning. In *AAAI*.
- Guowei Xu, Wenbiao Ding, Jiliang Tang, Songfan Yang, Gale Yan Huang, and Zitao Liu. 2019. Learning effective embeddings from crowdsourced labels: An educational case study. In *ICDE*, pages 1922–1927. IEEE.
- Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. 2014. Deep metric learning for person re-identification. In *ICPR*, pages 34–39. IEEE.
- Matthew D Zeiler. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Ying Zhang, Xianghua Ding, and Ning Gu. 2018. Understanding fatigue and its impact in crowdsourcing. In *CSCWD*, pages 57–62. IEEE.
- Yuchen Zhang, Xi Chen, Dengyong Zhou, and Michael I Jordan. 2014. Spectral methods meet em: A provably optimal algorithm for crowdsourcing. In *NIPS*, pages 1260–1268.
- Yudian Zheng, Guoliang Li, Yuanbing Li, Caihua Shan, and Reynold Cheng. 2017. Truth inference in crowdsourcing: Is the problem solved? *PVLDB*, 10(5):541–552.
- Yaoyao Zhong, Weihong Deng, Mei Wang, Jiani Hu, Jianteng Peng, Xunqiang Tao, and Yaohai Huang. 2019. Unequal-training for deep face recognition with long-tailed noisy data. In *CVPR*, pages 7812–7821.