NeuCrowd: Neural Sampling Network for Representation Learning with Crowdsourced Labels

Yang Hao
TAL Education Group
Beijing, China 100080

Wenbiao Ding
TAL Education Group
Beijing, China 100080

Zitao Liu[*]
TAL Education Group
Beijing, China 100080

Abstract

Representation learning approaches require a massive amount of discriminative training data, which is unavailable in many scenarios, such as healthcare, small city, education, etc. In practice, people refer to crowdsourcing to get annotated labels. However, due to issues like data privacy, budget limitation, shortage of domain-specific annotators, the number of crowdsourced labels are still very limited. Moreover, because of annotators' diverse expertises, crowdsourced labels are often inconsistent. Thus, directly applying existing representation learning algorithms may easily get the overfitting problem and yield suboptimal solutions. In this paper, we propose *NeuCrowd*, a unified framework for representation learning from crowdsourced labels. The proposed framework (1) creates a sufficient number of high-quality n-tuplet training samples by utilizing safety-aware sampling and robust anchor generation; and (2) automatically learns a neural sampling network that adaptively learns to select effective samples for representation learning network. The proposed framework is evaluated on both synthetic and 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[1].

1 INTRODUCTION

Representation learning, especially deep learning, has been proven to be crucial in many different domains such as information retrieval (Grbovic & Cheng, 2018), recommender systems (Xue, Dai, Zhang, Huang, & Chen, 2017), compute vision (Duan, Chen, Lu, & Zhou, 2019; Sohn, 2016), etc. Such approaches are usually discriminatively trained on massive labeled data, which are mostly generated from explicit or implicit online user engagement, like ratings, comments, clicks, hides, etc (Bengio, Courville, & Vincent, 2013).

However, in many real-world scenarios such as healthcare, small city, education, finance, etc., labeled data sets are typically insufficient or unavailable. To alleviate this problem, human efforts can be involved to acquire labeled data manually and crowdsourcing provides a flexible solution (Whitehill, Wu, Bergsma, Movellan, & Ruvolo, 2009; Raykar et al., 2010; Rodrigues, Pereira, & Ribeiro, 2014; Soto & Hirschberg, 2017). Theoretically, we could annotated data sets as large as we want via crowdsourcing platforms such as Amazon Mechanical Turk [2], CrowdTruth[3], etc. Unfortunately, the number of crowdsourced labels are still very limited due to a variety of reasons as follows:

- **data privacy**: data sets in many offline scenarios are usually difficult to collect due to privacy concerns. For example, in medical diagnostic imaging, patients' privacy data sets are prohibited to be opened to the public by applicable laws.

- **specialist shortage**: the crowdsourced tasks may require strong domain specialties. For instance, in educational data mining, student assessments require pedagogical specialties from annotators, in which doesn't scale by nature.

---

[*]The corresponding author.

[1]To encourage the reproducible results, we make our code public on a github repository, i.e., `https://github.com/crowd-data-mining/NeuCrowd`

[2]https://www.mturk.com/
[3]http://crowdtruth.org/

- **high cost**: the labeling tasks may require excessive budgets or tedious and time-consuming efforts. For example, it may take a crowd worker less than 1 second to annotate an image while the worker has to watch a 60-min classroom recording before determining the class quality, i.e., whether the class is good or bad.

- **label inconsistency**: due to different levels of expertises, one object may get distinct labels from multiple annotators, which makes crowdsourced labels very inconsistent or noisy.

Recent years have witnessed great efforts on learning with small labeled data (Fei-Fei, Fergus, & Perona, 2006; Ravi & Larochelle, 2016; Vinyals, Blundell, Lillicrap, Wierstra, et al., 2016). Also inferring true labels from inconsistent crowdsourced labels has been studied for decades (Whitehill et al., 2009; Raykar et al., 2010; Rodrigues et al., 2014). However, research on representation learning with small and inconsistent crowdsourced labels is rather limited. Therefore, the objective of this work is to study and develop approaches that can be used for learning representation from crowdsourced labels. More specifically, we target on answering two questions: (1) Since the annotated samples in healthcare, education and many other domains are usually in an incredibly smaller order of magnitude (a few hundreds or less), compared to web-scale data sets, how do we take advantage of deep representation learning under the limited sample setting? and (2) Due to the fact that crowdsourced labels may be highly inconsistent, how do we handle such uncertainty and make the learning procedure more efficient?

In this work we address above issues by presenting a unified framework *NeuCrowd* that is applicable to learn effective representations from very limited crowdsourced data. Specifically, we propose a scheme of generating hundreds of thousands of safety-aware and robust training instances from a limited amount of inconsistent crowdsourced labeled data.

Our data augmentation approach generalizes the deep triplet embedding learning in computer vision into crowdsourcing settings with multiple negative examples, a.k.a., *n*-tuplet, where each *n*-tuplet is consist of an anchor, a positive example and *n-2* negative examples (Sohn, 2016; Xu et al., 2019). Furthermore, in order to expedite the learning process and improve the quality of the learned representations, we specifically design a neural sampling network to adaptively select "hard" *n*-tuplet training samples. Different from most existing hard example mining heuristics (Shrivastava, Gupta, & Girshick, 2016), our framework is able to train both the representa-tion learning network and the sampling network simultaneously. Hence, the sampling network is able to dynamically exploit relations among *n*-tuplet samples without any hard-coding heuristic.

Overall this paper makes the following contributions:

- We propose a safety-aware and robust data augmentation technique that considers the inconsistence and uncertainty between annotators and creates a sufficient number of robust *n*-tuplet training samples.

- We design a sampling network to automatically and adaptively select optimized (a.k.a., hard) *n*-tuplet samples for the representation learning framework. The sampling network doesn't rely on any pre-fixed heuristic and both the embedding network and the sampling network are optimized simultaneously.

- We conduct a detailed and comprehensive experimental comparison of the proposed framework on multiple data sets from different domains. To encourage the reproducible results, we make our code and data publicly available on a github repository.

2 RELATED WORK

In this section, we briefly review related works. Our work lies at the intersection of limited labeled data learning, crowdsourced labels learning, and metric learning with hard example mining.

2.1 Learning with Limited Data

Both few/zero-short learning and semi/weakly supervised learning approaches have been developed to enable learning with limited labeled data in different ways. Few/zero-shot learning, which is motivated by the fact that humans can learn new concepts with very little supervision, aims to learn new concepts from very small amount of labeled examples (Fei-Fei et al., 2006; Snell, Swersky, & Zemel, 2017; F. Sung et al., 2018). While semi/weakly supervised learning makes use of the large amount of unlabeled data to learn better predictors (Takamatsu, Sato, & Nakagawa, 2012; A. J. Ratner, De Sa, Wu, Selsam, & Ré, 2016; A. Ratner et al., 2017).

Although few-shot learning methods yield promising results on unseen categories, they demands large data sets from other categories. This may be infeasible in many real-world domains other than computer vision. In spite of the successful applications of the semi-supervised or weekly supervised approaches, they may not work when the total available data is limited.

2.2 Learning with Crowdsourced Labels

Truth inference is well studied in crowdsourcing research (Whitehill et al., 2009; Raykar et al., 2010; Rodrigues et al., 2014), which aims at directly inferring the ground truth from workers' annotations. Whitehill et al. propose a probabilistic framework that iteratively adjust the inferred ground truth estimates based on the the performance of the annotators (Whitehill et al., 2009). Raykar et al. proposed an EM algorithm to jointly learn the levels of annotators and the regression models (Raykar et al., 2010). Rodrigues et al. generalize Gaussian process classification to consider multiple annotators with diverse expertise (Rodrigues et al., 2014).

The majority of aforementioned algorithms have been designed to address the label inconsistency problem and they cannot work as expected when labels are limited. While in this work, we aim to develop algorithms which can jointly solve the representation learning challenges from limited and inconsistent labels.

2.3 Metric Learning with Hard Example Mining

Deep metric learning approaches construct pairs (Koch, Zemel, & Salakhutdinov, 2015; Sohn, 2016) or triplets (Schroff, Kalenichenko, & Philbin, 2015; He, Zhou, Zhou, Bai, & Bai, 2018) with different objective functions. Consequently, various hard example mining techniques are developed to select "hard" training samples to expedite the optimization convergence (K.-K. Sung, 1996). Many approaches along this direction have achieved promising results in many tasks such as object detection (Shrivastava et al., 2016), face recognition (Sohn, 2016; Schroff et al., 2015), etc.

Although deep metric learning with hard example mining is able to learn effective representations, they require a large amount of data. Moreover, deep metric learning approaches heavily rely on the comparisons within pairs or triplets, which are very sensitive to outliers or ambiguous examples and may be easily misled by inconsistent crowdsourced labels.

3 METHODOLOGY

3.1 Notation

Following conventions, we use bold upper case for collections or sets, bold lower case letters for vectors and calligraphic typeface for functions. We use $(\cdot)^+$ and $(\cdot)^-$ to indicate positive and negative examples. More specifically, let $\mathbf{D}$ be the original crowdsourced data set, i.e., $\mathbf{D} = \{e_i\}_{i=1}^N = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$, where e_i is the ith crowdsourced example and $\mathbf{x}_i$ denotes the raw features of e_i

and $\mathbf{y}_i \in \mathbb{R}^{d \times 1}$ denotes the crowdsourced labels of e_i. N is the size of $\mathbf{D}$ and d is the number of crowd workers. Let $y_{ij} \in \{0, 1\}$ be the binary label from the jth worker for e_i, i.e., $\mathbf{y}_i = [y_{i1}, \cdots, y_{id}]$.

3.2 N-tuplet

Similar to (Sohn, 2016; Xu et al., 2019), we define n-tuplet as follows:

Definition 1 *(n-TUPLET) An* n-*tuplet T is an n-element collection that consists of an anchor, a positive example (to the anchor) and n-2 negative examples, i.e.,*

$$T \stackrel{\text{def}}{=} (e^*, e^+, e_1^-, \cdots, e_{n-2}^-)$$

where e^* is an anchor and e^+ is a positive example to e^* and $\{e_i^-\}_{i=1}^{n-2}$ are negative examples.

The n-tuplet is a generalization of triplet (when $n = 3$) where it contains more than one negative examples. In terms of model learning, different from triplets that only support learning from negative examples one at a time, the n-tuplets try to maximize the distances between positive examples and all the other $n - 2$ negative examples all at once (Sohn, 2016).

As we discussed, the limited number of labeled data in certain domains may easily lead to the overfitting problems for many supervised representation learning approaches. Fortunately, this issue can be largely alleviated by exploring the trick of n-tuplets. Theoretically, we could create a size of $\mathcal{O}(P^2 Q^{n-2})$ n-tuplets where P and Q are the number of positive and negative examples. By sampling and reassembling from original data set $\mathbf{D}$, we are able to significantly augment the training data size compared to the one that directly training models from individual examples, i.e., $\mathcal{O}(P + Q)$. Therefore, in this work, we develop our representation learning framework that is optimized on batches of n-tuplets instead of individual examples.

3.3 The NeuCrowd Model

Although the n-tuplet based paradigm sheds light on the problem of learning from small data, building end-to-end representation learning solutions from crowdsourced labels is still challenging and gives rise to the following important questions:

Q1 : How do we effectively construct n-tuplets from highly inconsistent crowdsourced data?

Q2 : How do we improve the efficiency of the embedding model training when using a set of n-tuplets?

In this work, we address *Q1* by proposing (1) a safety-aware sampling strategy to "clean up" the *n*-tuplet construction space by dynamically identifying inconsistent crowdsourced examples along the embedding model learning (See Section 3.3.1); and (2) a robust anchor generation method to artificially create anchors that reduce ambiguity and chances of outliers within *n*-tuplets (See Section 3.3.2). We answer *Q2*, we develop an deep embedding network that is able to learn from *n*-tuplets (See Section 3.3.3) and explicitly design a sampling network, which is able to adaptively select the "hardest" *n*-tuplets and automatically co-learn its parameters with the representation learning network without any heuristic (See Section 3.3.4). The iterative joint learning paradigm is described in Section 3.4. The entire *NeuCrowd* framework is illustrated in Figure 1.
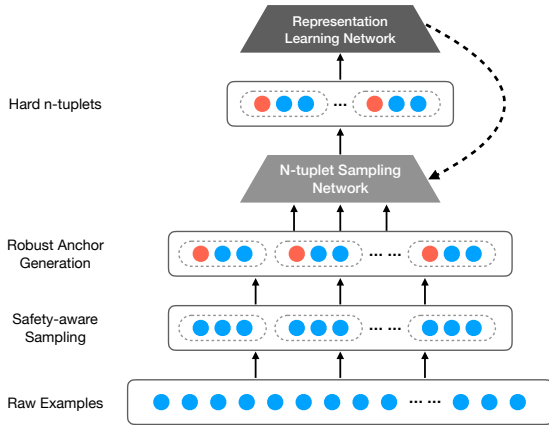

Figure 1: The overview of the *NeuCrowd* framework. The dash line indicates that the representation learning network passes its training loss to the *n*-tuplet sampling network. The blue dots represent original crowdsourced examples and red dots represent the robust anchors.

3.3.1 Safety-aware Sampling

When obtaining reliable labels from crowds, an example is usually annotated by multiple workers (Raykar et al., 2010; Yan, Rosales, Fung, Subramanian, & Dy, 2014; Yan et al., 2010). Consequentially, we may get different votes for the same example. For example, assuming that e_i^+ and e_j^+ are two positive examples whose corresponding 5-person crowdsourced binary labels are (1, 1, 1, 1, 1) and (1, 1, 1, 0, 0), our assurance of the labels of e_i^+ and e_j^+ should be different. Here we refer to label assurance as the measurement of the degree of disagreement of annotated labels within the same example, which is defined as follows:

Definition 2 (LABEL ASSURANCE) *Given a crowdsourced example e_i, its label assurance, i.e., $\mathcal{A}(e_i)$ is*

computed as follows:

$$\mathcal{A}(e_i) = \left| 1 - \frac{2}{d} \|\mathbf{y_i}\|_1 \right|$$

where $|\cdot|$ denotes the absolute value and $\|\cdot\|_1$ represents the vector ℓ_1 norm.

The label assurance measures the disagreement degree among workers and reaches to its minimum value[4] of 0 when a tie or a draw happens and goes to its maximum value of 1 when all workers' labels are consistent.

Since the representation learning network optimizes its parameters purely from *n*-tuplets and it tries to push the $n - 2$ negative examples all at once within each *n*-tuplet, incorporating unsure labels will easily confuse the objective function and lead to inferior representations. Therefore, it is necessary to exclude those ambiguous examples when constructing the training set. Due to the fact that such ambiguous instances may make up 50% of all labels, simply discarding all the ambiguous example directly doesn't work (Zhong, Tang, & Zhou, 2015; Takeoka, Dong, & Oyamada, 2020).

Therefore, in this work, we propose a safety-aware sampling technique to dynamically sample safe examples in the embedding space when constructing the *n*-tuplets. Here at each training iteration *t*, we define the safe example as follows:

Definition 3 (SAFE EXAMPLE) *Let $\mathbf{N}_i(t)^+$ and $\mathbf{N}_i(t)^-$ denote the index sets of k nearest neighbors for e_i at iteration t in the embedding space. Without loss of generality, assume e_i is positive, an example e_i is safe if and only if*

$$\sum_{p \in \mathbf{N}_i(t)^+} \mathcal{A}(e_p) > \sum_{q \in \mathbf{N}_i(t)^-} \mathcal{A}(e_q) + \delta$$

where $k = \sqrt{b}$ and b is the batch size and δ is the safety margin. $\mathbf{N}_i(t)^+$ and $\mathbf{N}_i(t)^-$ are computed by using the learned embedding at iteration t with the ℓ_2 distance function.

Inspired by safe borderline SMOTE (Han, Wang, & Mao, 2005), we consider a positive (negative) example is safe if the sum of assurance of its *k* nearest positive (negative) neighbors are larger than the one from its negative (positive) neighbors. Please note that the neighbors are defined in the embedding space at each iteration and will dynamically change along the model training. When constructing the *n*-tuplets, we conduct the safety-aware

---

[4]The minimum value goes to $1/d$ when d is odd.

sampling by giving safe examples higher probabilities compared to unsafe examples.

3.3.2 Robust Anchor Generation

Anchors play one of the most important roles in the *n*-tuplets based model learning. When learning from *n*-tuplets, both the positive example and multiple negative examples are compared with the corresponding anchor in each *n*-tuplet. The result of model learning highly relies on the quality of those anchors and any ambiguous anchor will mislead the optimization to a suboptimal solution. Therefore, to reduce such inferior effect, we develop a robust anchor generation approach that artificially creates a batch-level "gold standard" anchor by summarizing the all the anchors' information by their label assurance scores, i.e.,

Definition 4 *(ROBUST ANCHOR)*

Let $c_1, \cdots, c_m$ be the example indices of selected anchors and m be the total number of n-tuplets we generated within each training batch. The batch-level robust anchor is computed by:

$$e_r^* = \sum_{j=1}^{m} \mathcal{A}(e_{c_j}) e_{c_j}^*$$

After creating the robust anchor e_r^*, we replace all the original anchors from *n*-tuplets within this batch with e_r^*. The robust anchor e_r^* is more closer to the center of cluster formed by the high consistent crowdsourced examples and the influence of ambiguous anchors is significantly reduced. It is worth noting that the robust anchor's calculation is computationally friendly and easy to implement by adding a robust anchor generation layer in the representation learning network, depicted in Figure 1.

3.3.3 Representation Learning Network

Inspired by the discriminative training approaches widely used in information retrieval (Huang et al., 2013; Palangi et al., 2016) and natural language processing (Dos Santos & Gatti, 2014), we present a supervised training approach to learn the representation network by maximizing the conditional likelihood of retrieving positive example e^+ given our robust anchor e_r^* from the corresponding *n*-tuplet with the rest $n - 2$ negative examples. Similar to (Xu et al., 2019), we design a weight-sharing deep neural networks (DNN) for each example within the *n*-tuplets.

Formally, give a embedding network $\Omega(\Theta)$, let $\mathcal{F}_\Theta(e_i)$ be the learned representation of example e_i, we compute

the posterior probability of e^+ given e_r^* through a soft-max function, i.e.,

$$p\big(e^+|e_r^*\big) = \frac{\exp\Big(\eta \cdot \mathcal{A}(e^+) \cdot \mathcal{R}(\mathcal{F}_\Theta(e^+), \mathcal{F}_\Theta(e_r^*))\Big)}{\sum\limits_{e_j \in \mathcal{T}} \exp\Big(\eta \cdot \mathcal{A}(e_j) \cdot \mathcal{R}(\mathcal{F}_\Theta(e_j), \mathcal{F}_\Theta(e_r^*))\Big)}$$

where $\mathcal{R}(\cdot, \cdot)$ is the similarity function, i.e., $\mathcal{R}(\mathcal{F}_\Theta(e_1), \mathcal{F}_\Theta(e_2)) \overset{\text{def}}{=} C - \|\mathcal{F}_\Theta(e_1) - \mathcal{F}_\Theta(e_2)\|_{\ell_2}$. C and η are smoothing hyper-parameters in the softmax function, which are set empirically on a held-out data set in our experiments.

To maximize the posterior, we would like to maximize the relevance between two positive embeddings $\mathcal{F}_\Theta(e_r^*)$ and $\mathcal{F}_\Theta(e^+)$, in the meanwhile, minimize the relevance between the robust anchor embedding $\mathcal{F}_\Theta(e_r^*)$ and all the other negative embeddings, i.e., $\{\mathcal{F}_\Theta(e_j^-)\}_{j=1}^{n-2}$. As distance is proportional to the inverse of relevance, similar data examples are pull closer while dissimilar examples are pushed away in the embedding space.

Hence, given a collection of *n*-tuplets, we optimize parameters of the embedding network $\Omega(\Theta)$ by maximizing the sum of log conditional likelihood of finding a positive example e^+ given the robust anchor e_r^* from *n*-tuplet $\mathcal{T}$, i.e., $\mathcal{L}(\Omega(\Theta)) = -\sum \log p\big(e^+|e_r^*\big)$. Since $\mathcal{L}(\Omega(\Theta))$ is differentiable with respect to $\Omega(\Theta)$, we use gradient based optimization approach to train the embedding network.

3.3.4 *N*-tuplet Sampling Network

By constructing the training *n*-tuplets with safety-aware sampling and robust anchors, we are able to get quadratic or cubic training sample size compared to the original data set. On the one hand, we provide the embedding network sufficient training data and avoid the overfitting problem, but on the other, the training process may become extremely long and may not guarantee an optimal performance. Therefore, we explicitly design a sampling network to adaptively select effective training examples from a massive number of *n*-tuplets.

Similar to Section 3.3.3, we design our sampling network by *n* parameter-sharing DNNs with multi-layer fully-connected projections. Every example in the *n*-tuplet is passed to its corresponding DNN to get its non-linear embedding. These *n* embeddings are concatenated and used as features to optimize the "hardness" score prediction. In general, the "hardness" scores can be calculated by any real-valued function. Here, we choose to use the training loss from representation learning network as our

surrogates of "hardness" scores. We use the square loss as the objective loss function of sampling network.

3.4 Joint Learning Paradigm

We jointly optimize the representation learning network and the n-tuplet sampler network in an iterative manner, which is described as the following repeated steps:

Step 1.: The sampling network scores every element in the n-tuplet set, i.e., $\mathbf{T}$, which is constructed by using safety-aware sampling and robust anchor generation. Those n-tuplets with higher "hardness" scores are selected into set $\mathbf{H}$

Step 2.: The representation network uses $\mathbf{H}$ for its parameter optimization. It passes the training loss L to the n-tuple sampling network through forward-propagation.

Step 3.: The sampling network fine-tunes its parameters by utilizing the training loss L from the representation learning network.

4 EXPERIMENTS

In this section, we will (1) demonstrate the effectiveness of the proposed framework by comprehensively comparing with representative state-of-the-art baselines; and (2) conduct component analysis to examine the effect of key components of the proposed framework. In the following, the proposed method is referred to as *NeuCrowd*.

Experiments are conducted on both synthetic and real-world crowdsourced data sets from different domains. We would also like to note that the hyper parameters used in our methods are selected (in all experiments) by the internal cross validation approach while optimizing models' predictive performances. We report accuracy and AUC scores to comprehensively evaluate the performance of our proposed method. To encourage the reproducible results, we make our code public on a github repository, i.e., `https://github.com/crowd-data-mining/NeuCrowd`.

4.1 Data

Synthetic Data To get a good understanding of our approach, we first test it on two synthetic data sets. Here we use the same simulation approaches as Guyon et al. used in the NIPS 2003 variable selection task (Guyon, Gunn, Ben-Hur, & Dror, 2005) for generating synthetic samples. Briefly, we create 4 clusters of points normally distributed (std=1) about vertices of a multi-dimensional hypercube with sides of length 2 and assigns 2 clusters

to each class (positive or negative) and sample features from these Gaussian distributions. Detailed can be found in *scikit-learn.org* docs[5]. To generate crowdsourced labels, we consider two settings of different number of crowd workers, i.e., 7 workers (*Syn1*) and 11 workers (*Syn2*). Here, crowdsourced labels are simulated by assigning a mislabeling probability for each worker. The mislabeling probability is obtained from a truncated normally distribution with upper bound of 0.5, lower bound of 0.01, mean of 0.1 and std of 0.1.

Pre-K Children Speech Data We test our framework on a data set of speech contest of pre-K children, i.e., *Pre-K*. The contest examines the ability of addressing speech in front of public audience. Each example is a 1-min video with binary labels indicating speech fluency. We extract both the linguistic features and acoustic features from the videos[6].

Hotel Review Data We use a hotel comment data, i.e., *hotel*, as a benchmark data set. The data is collected from a third party rating website. Each comment has a binary label indicating its positive or negative rating sentiment. The goal is to learn the language embedding to distinguish positive and negative comments.

Oral Language Skill Data We also test our *NeuCrowd* on a student oral language skill data set, i.e., *Oral*, which is obtained from a third party online training platform. Each example is a 2-min video clip of students practicing their oral language skills. The predictive task is to determine whether the students' spoken language proficiency is fluent enough or not. Similar to (Xu et al., 2019), a wide range of linguistic features from the raw texts after having automatic speech recognition on the clips.

4.1.1 Data Statistics

We summarize the crowdsourcing settings and data statistics in Table 1. In Table 1, *kappa* represents the value of Fleiss' kappa measurement (Fleiss, 1971), which is a statistical measure for assessing the reliability of agreement between a fixed number of raters when assigning categorical ratings to a number of items or classifying items. *class ratio* denotes the class label ratio that is computed by # of positive samples divided by # of total samples.

---

[5] `https://scikit-learn.org/stable/modules/generated/sklearn.datasets.make_classification.html`

[6] Acoustic features are extracted by using OpenSmile, i.e., `https://www.audeering.com/opensmile/`

Table 1: Data sets statistics.

	# of annotators	# of features	train size	validation size	test size	kappa	class ratio
Syn1	7	1200	800	200	500	0.52	0.50
Syn2	11	1200	800	200	500	0.44	0.50
Pre-K	11	1632	940	247	300	0.60	0.65
Hotel	7	300	447	112	140	0.80	0.43
Oral	11	50	908	384	200	0.41	0.70

4.2 Baselines

We carefully choose three groups of state-of-the-art as our baselines to comprehensively assess the effectiveness of the proposed framework.

Group 1: True Label Inference from Crowdsourcing

The first group contains methods inferring true labels from crowdsourced labels. They are listed as follows:

- Logistic regression with every pair (instance, label) provided by each crowd worker as a separate example. Note that this amounts to using a soft probabilistic estimate of the actual ground truth to learn the classifier, i.e., *SoftProb* (Raykar et al., 2010).

- Logistic regression with EM labels, i.e., *EM* (Dempster, Laird, & Rubin, 1977). The labels are treated as hidden variables and inferred by expectation-maximization. The EM algorithm iteratively estimates worker's accuracy and exploits the estimated accuracy to compute the aggregated result.

- Logistic regression with GLAD labels, i.e., *GLAD*. GLAD infers the true labels by jointly inferring the true label, worker's expertise and the difficulty of each data instance (Whitehill et al., 2009).

Group 2: Representation Learning with Limited Labels

The second group includes representation learning methods designed for limited labels. They are listed as follows:

- Contrastive Loss, i.e., *Contrastive* (Koch et al., 2015) We train a Siamese network that learns an embedding with pairs of examples to minimize distance between intra-class instances.

- Triplet networks, i.e., *Triple* (Schroff et al., 2015). We train a triplet network that takes an anchor, a positive (of same class as an anchor) and negative (of different class than an anchor) examples. The

objective is to learn embeddings such that the anchor is closer to the positive example than it is to the negative example by some margin value.

- Triplet networks with semi-hard example mining, i.e., *TripleSemi* (Schroff et al., 2015). The triplet with smallest distance between anchor and negative example in the embedding space is chosen, from triplets where the anchor is farther from the negative example than the positive one.

- Triplet networks with lifted structured loss, i.e., *Lifted* (Oh Song, Xiang, Jegelka, & Savarese, 2016). Lifted structured loss is based on all the pairwise edges among positive and negative pairs of samples in the training set, which fully explores the relations of instances.

- Triplet networks with center Loss, i.e., *Center* (He et al., 2018). Distance between each instance and the center (not weighted by vote confidence) is learned for each category, with the goal to minimize intra-class variations and maximize inter-class distances at the same time.

Group 3: Two-stage Models by Combining Group 1 and Group 2

The third group are methods combining baselines from the first (i.e., inferring true labels) and second groups (i.e., learning embedding with limited labels). They solve the problems of the limited and inconsistent labels in two stages. Due to the page limit, we only combine the best approaches from first and second groups.

4.3 Implementation Details

Experimental codes are implemented in tensorflow 1.8 https://www.tensorflow.org/ and executed on a server with Ubuntu 14.04 LTS and a GTX 1080 Ti GPU. We set n to 5 for all the following experiments. We use a weight-sharing deep neural networks with 2 fully-connected layers as the representation learning network and the sampling network. We set dropout rate to 0.2. We initialize the network weights with normal distribution initializer. We use Adadelta as our optimizer (Zeiler,

Table 2: Prediction results on both synthetic and real-world data sets. Acc is short for accuracy. "-" represents the algorithm never converges.

Method	Group	Syn1		Syn2		Pre-K		Hotel		Oral	
		Acc	AUC	Acc	AUC	Acc	AUC	Acc	AUC	Acc	AUC
SoftProb	group 1	0.666	0.720	0.664	0.734	-	-	0.850	0.921	0.825	0.940
EM	group 1	0.652	0.717	0.666	0.732	0.830	**0.923**	0.857	0.918	0.815	0.949
GLAD	group 1	0.614	0.664	0.636	0.701	0.830	**0.923**	0.850	0.911	0.805	0.948
Contrastive	group 2	0.580	0.590	0.600	0.579	0.82	0.873	0.850	0.908	0.855	0.905
Triple	group 2	0.654	0.695	0.644	0.683	0.850	0.883	0.864	0.916	0.795	0.927
TripleSemi	group 2	0.638	0.678	0.626	0.641	0.650	0.628	0.771	0.795	0.84	0.920
Centered	group 2	0.594	0.612	0.594	0.605	0.760	0.801	0.843	0.912	0.850	**0.957**
Lifted	group 2	0.588	0.596	0.598	0.614	0.740	0.784	0.836	0.887	0.860	0.940
Contrastive+EM	group 3	0.596	0.605	0.612	0.622	0.800	0.855	0.850	0.906	**0.865**	0.950
Triple+EM	group 3	0.650	0.661	0.672	**0.772**	0.860	0.918	0.864	0.918	**0.865**	0.954
TripleSemi+EM	group 3	0.616	0.653	0.640	0.685	0.745	0.725	0.750	0.792	0.850	0.908
Centered+EM	group 3	0.594	0.600	0.592	0.620	0.790	0.794	0.850	0.918	**0.865**	0.948
Lifted+EM	group 3	0.588	0.600	0.600	0.616	0.720	0.740	0.857	0.911	0.860	0.944
NeuCrowd	our	**0.678**	**0.729**	**0.688**	0.751	**0.870**	0.898	**0.871**	**0.928**	**0.865**	0.927

2012). The learning rate for both embedding network and sampling network is set to 1e-3. Sizes of each layer and scale of ℓ_2 regularization are hyper-parameters that are set by grid searching with cross validation. Downstream logistic regression classifier is trained with inverse of ℓ_2 regularization strength C as the only hyper-parameter ranging from 1e-4 to 1e4.

4.4 Performance Comparison

From these results, we make the following observations to compare performance of existing methods and *Neu-Crowd*:

- Methods in **group 3**, which combine EM with methods in **group 2** to solve the problem of limited data and crowdsourced label inconsistencies at the same time, outperform baseline methods in **group 2** on most of the data sets, which suggests that it's necessary to get rid of noise when using crowdsourced labels. In fact, simply training embeddings with majority-vote labels results in inferiority compared with classic methods.

- To get a robust anchor in each batch of training, our proposed framework makes full use of label assurance as weights, rather than calculates the arithmetic average as center. These assurance-aware anchors are closer to instances with high confidence instead of ones with low confidence in random positions, helping determine which groups are hard due

to inefficient embedding training instead of mislabeling. As a result, *NeuCrowd* performs better than other center based loss structures i.e. *Center*.

- Another component to guarantee the sampler network to select safe hard examples is the safety-aware sampling in *NeuCrowd*. Since a mislabeled instance would participate in many n-tuplets throughout the training process, we exclude them dynamically during the overall learning process. This process takes advantage of the learned high-quality non-linear representations from the embedding network. Therefor our framework performs better than networks trained with fixed EM estimated labels in **group 3**.

4.5 Component Analysis

We systematically examine the effect of key components by constructing following model variants:

- NeuCrowd-SA :it eliminates the contribution of safety-aware sampling.

- NeuCrowd-SA-RA: it eliminates the contributions of both safety-aware sampling and robust anchors.

- NeuCrowd-SA-RA-SN: it eliminates the contributions of safety-aware sampling, robust anchors and n-tuple sampling network and only the n-tuple based representation learning model remains, which

Table 3: Key component prediction analysis of *NeuCrowd* on both synthetic and real-world data sets.

Method	Syn1		Syn2		Pre-K		Hotel		Oral	
	Acc	AUC	Acc	AUC	Acc	AUC	Acc	AUC	Acc	AUC
NeuCrowd-SA-RA-SN	0.65	0.712	0.674	0.662	0.806	0.87	0.77	0.838	0.837	0.908
NeuCrowd-SA-RA	0.665	0.713	0.664	0.72	0.806	0.882	0.850	0.921	0.81	0.904
NeuCrowd-SA	0.67	0.714	0.68	0.723	0.822	0.869	0.857	0.823	0.884	**0.948**
NeuCrowd	**0.678**	**0.729**	**0.688**	**0.751**	**0.870**	**0.898**	**0.871**	**0.928**	**0.865**	0.927

is equivalent to the *RLL* framework proposed by Xu et al. (Xu et al., 2019).

The prediction performance with different components on both synthetic and real-world data sets is shown in Table 3. As we can see, the results of our experiments show that our *NeuCrowd* model outperforms all other variants in terms of prediction errors on all data sets. As we keep incorporating the proposed components, the performance boosts on every data set. This suggests that prediction performance degrades when ignoring any type of key components. Thus, it is important to incorporate them together when build the end-to-end solutions of representation learning from crowdsourced labels.

Here we also demonstrate the changes of training loss with different components in the proposed framework in Figure 2. Since we observe similar results on other data sets and the space limit, we only show that on *Hotel* data set to save space. First, when comparing *NeuCrowd-SA-RA-SN* to other methods with sampling network, we can see that the optimization converges faster and achieves a much lower training loss. We believe this is because the sampling network is able to appropriately select the hard examples for modeling training. Second, when comparing *NewCrowd* and *NeuCrowd-SA-RA* and *NeuCrowd-SA*, even through they almost have the same converge speed, with safety-aware sampling and robust anchors, *NewCrowd* is able to get the lowest training loss with less than 50 epochs.

5 CONCLUSION

In this paper, we presented a representation learning framework for learning embeddings from limited crowdsourced labels. Comparing with traditional discriminative representation learning approaches, the advantages of our framework are: (1) it is able to learn effective embeddings from very limited data; (2) it automatically selects effective *n*-tuplet training samples which makes the training process more effective. Experimental results on both synthetic and real-world data sets demonstrated that our approach outperforms other state-of-the-art baselines
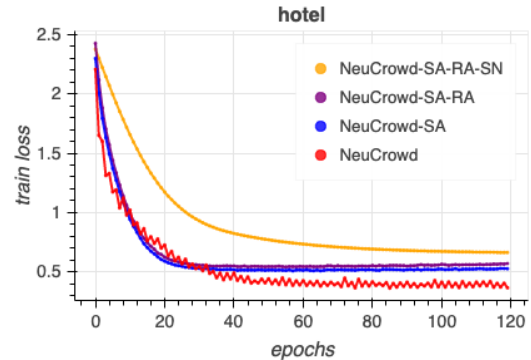

Figure 2: The training loss of NeuCrowd with different component on *Hotel* data set.

in terms of accuracy and AUC score. In the future, we plan to study how to model the quality of crowd workers and incorporate such information into the representation learning framework.

References

Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *TPAMI*, *35*(8), 1798–1828.

Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, *39*(1), 1–22.

Dos Santos, C., & Gatti, M. (2014). Deep convolutional neural networks for sentiment analysis of short texts. In *COLING* (pp. 69–78).

Duan, Y., Chen, L., Lu, J., & Zhou, J. (2019). Deep embedding learning with discriminative sampling policy. In *CVPR* (pp. 4964–4973).

Fei-Fei, L., Fergus, R., & Perona, P. (2006). One-shot learning of object categories. *TPAMI*, *28*(4), 594–611.

Fleiss, J. L. (1971). Measuring nominal scale agreement among many raters. *Psychological Bulletin*, *76*(5), 378.

Grbovic, M., & Cheng, H. (2018). Real-time personalization using embeddings for search ranking at airbnb. In *SIGKDD* (pp. 311–320).

Guyon, I., Gunn, S., Ben-Hur, A., & Dror, G. (2005). Result analysis of the nips 2003 feature selection challenge. In *NIPS* (pp. 545–552).

Han, H., Wang, W.-Y., & Mao, B.-H. (2005). Borderline-smote: a new over-sampling method in imbalanced data sets learning. In *International conference on intelligent computing* (pp. 878–887).

He, X., Zhou, Y., Zhou, Z., Bai, S., & Bai, X. (2018). Triplet-center loss for multi-view 3d object retrieval. In *CVPR* (pp. 1945–1954).

Huang, P.-S., He, X., Gao, J., Deng, L., Acero, A., & Heck, L. (2013). Learning deep structured semantic models for web search using clickthrough data. In *CIKM* (pp. 2333–2338).

Koch, G., Zemel, R., & Salakhutdinov, R. (2015). Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop* (Vol. 2).

Oh Song, H., Xiang, Y., Jegelka, S., & Savarese, S. (2016). Deep metric learning via lifted structured feature embedding. In *CVPR* (pp. 4004–4012).

Palangi, H., Deng, L., Shen, Y., Gao, J., He, X., Chen, J., . . . Ward, R. (2016). Deep sentence embedding using long short-term memory networks: Analysis and application to information retrieval. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, *24*(4), 694–707.

Ratner, A., Bach, S. H., Ehrenberg, H., Fries, J., Wu, S., & Ré, C. (2017). Snorkel: Rapid training data creation with weak supervision. *VLDB*, *11*(3), 269–282.

Ratner, A. J., De Sa, C. M., Wu, S., Selsam, D., & Ré, C. (2016). Data programming: Creating large training sets, quickly. In *NIPS* (pp. 3567–3575).

Ravi, S., & Larochelle, H. (2016). Optimization as a model for few-shot learning.

Raykar, V. C., Yu, S., Zhao, L. H., Valadez, G. H., Florin, C., Bogoni, L., & Moy, L. (2010). Learning from crowds. *JMLR*, *11*(Apr), 1297–1322.

Rodrigues, F., Pereira, F., & Ribeiro, B. (2014). Gaussian process classification and active learning with multiple annotators. In *ICML* (pp. 433–441).

Schroff, F., Kalenichenko, D., & Philbin, J. (2015). Facenet: A unified embedding for face recognition and clustering. In *CVPR* (pp. 815–823).

Shrivastava, A., Gupta, A., & Girshick, R. (2016). Training region-based object detectors with online hard example mining. In *CVPR* (pp. 761–769).

Snell, J., Swersky, K., & Zemel, R. (2017). Prototypical networks for few-shot learning. In *NIPS* (pp. 4077–4087).

Sohn, K. (2016). Improved deep metric learning with multi-class n-pair loss objective. In *NIPS* (pp. 1857–1865).

Soto, V., & Hirschberg, J. (2017). Crowdsourcing universal part-of-speech tags for code-switching. *arXiv preprint arXiv:1703.08537*.

Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P. H., & Hospedales, T. M. (2018). Learning to compare: Relation network for few-shot learning. In *CVPR* (pp. 1199–1208).

Sung, K.-K. (1996). Learning and example selection for object and pattern detection.

Takamatsu, S., Sato, I., & Nakagawa, H. (2012). Reducing wrong labels in distant supervision for relation extraction. In *ACL* (pp. 721–729).

Takeoka, K., Dong, Y., & Oyamada, M. (2020). Learning with unsure responses. In *AAAI*.

Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. (2016). Matching networks for one shot learning. In *NIPS* (pp. 3630–3638).

Whitehill, J., Wu, T.-f., Bergsma, J., Movellan, J. R., & Ruvolo, P. L. (2009). Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. In *NIPS* (pp. 2035–2043).

Xu, G., Ding, W., Tang, J., Yang, S., Huang, G. Y., & Liu, Z. (2019). Learning effective embeddings from crowdsourced labels: An educational case study. In *ICDE* (pp. 1922–1927).

Xue, H.-J., Dai, X., Zhang, J., Huang, S., & Chen, J. (2017). Deep matrix factorization models for recommender systems. In *IJCAI* (pp. 3203–3209).

Yan, Y., Rosales, R., Fung, G., Schmidt, M., Hermosillo, G., Bogoni, L., . . . Dy, J. (2010). Modeling annotator expertise: Learning when everybody knows a bit of something. In *AISTATS* (pp. 932–939).

Yan, Y., Rosales, R., Fung, G., Subramanian, R., & Dy, J. (2014). Learning from multiple annotators with varying expertise. *Machine learning*, *95*(3), 291–327.

Zeiler, M. D. (2012). Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.

Zhong, J., Tang, K., & Zhou, Z.-H. (2015). Active learning from crowds with unsure option. In *IJCAI*.