Active Learning from Peers

Keerthiram Murugesan **Jaime Carbonell**
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
{kmuruges,jgc}@cs.cmu.edu

NIPS-2017

Introduction

■ Motivation

This paper addresses the challenge of learning from peers in an online multitask. Instead of always requesting a label from a human oracle, the proposed method first determines if the learner for each task can acquire that label with sufficient confidence from its peers.

1. Receive an example $x^{(t)}$ for the task k
2. If the task k is not confident in the prediction for this example, ask the *peers* or *related tasks* whether they can give a confident label to this example.
3. If the *peers* are not confident enough, ask the oracle for the true label $y^{(t)}$.

Problem Setup

■ Symbolic meaning

$\{(x_k^{(i)}, y_k^{(i)})\}_{i=1}^{N_k}$ K tasks, where the k-th task is associated with N_k training examples

$\{w_k^*\}_{k\in[K]}$ The parameter of the model for k-th task

$$\ell_{kk}^{(t)*} = (1 - y^{(t)}\langle x^{(t)}, w_k^*\rangle)_+$$
$$\ell_{km}^{(t)*} = (1 - y^{(t)}\langle x^{(t)}, w_m^*\rangle)_+$$
hinge losses on example $((x^{(t)}, k), y^{(t)})$

$h_k(x^{(t)}) = \langle w_k^{(t-1)}, x^{(t)}\rangle$ k-th task model`s output for the received $x^{(t)}$

$\hat{p}_{kk} = \langle w_k^{(t-1)}, x^{(t)}\rangle$ and $\hat{p}_{km} = \langle w_m^{(t-1)}, x^{(t)}\rangle$

Method of Learning from Peers

■ Current task model`s confidence

$|h_k(x^{(t)})|$ measure the confidence of the k-*th* task learner on this example

a Bernoulli random variable $P^{(t)}$ for the event $|h_k(x^{(t)})| \leq b_1$ with probability $\frac{b_1}{b_1+|h_k(x^{(t)})|}$

■ Peers` task model`s confidence

a Bernoulli random variable $Q^{(t)}$ for the event $|h_m(x^{(t)})| \leq b_2$ with probability

$$\frac{b_2}{b_2+\sum_{m\in[K],m\neq k} \tau_{km}^{(t-1)} |h_m(x^{(t)})|}$$

τ_{km} relationship between tasks using the cross-task error l_{km}

$$\tau_{km}^{(t)} = \frac{\tau_{km}^{(t-1)} e^{-\frac{Z^{(t)}}{\lambda}\ell_{km}^{(t)}}}{\sum_{\substack{m'\in[K]\\ m'\neq k}} \tau_{km'}^{(t-1)} e^{-\frac{Z^{(t)}}{\lambda}\ell_{km'}^{(t)}}} \quad m\in[K], m\neq k$$

Adaptive Smoothed Online Multi-Task Learning NIPS-2016

Active Learning from Peers

Conventional Uncertainty

$$h_k(x^{(t)}) = \langle w_k^{(t-1)}, x^{(t)} \rangle$$

Current task model`s confidence

Peers` task model`s confidence

$$\tilde{p}^{(t)} = \sum_{m \neq k, m \in [K]} \tau_{km}^{(t-1)} \hat{p}_{km}^{(t)} \text{ and } \tilde{y}^{(t)}$$

Update the model parameter $\boldsymbol{w}$

(when made a mistake)

Update the task relationship $\boldsymbol{\tau}$

Algorithm 1: Active Learning from Peers

Input : $b_1 > 0$, $b_2 > 0$ s.t., $b_2 \geq b_1$, $\lambda > 0$, Number of rounds T

Initialize $w_m^{(0)} = \mathbf{0}\ \forall m \in [K]$, $\boldsymbol{\tau}^{(0)}$.

for $t = 1 \ldots T$ **do**

Receive $(x^{(t)}, k)$

Compute $\hat{p}_{kk}^{(t)} = \langle x^{(t)}, w_k^{(t-1)} \rangle$

Predict $\hat{y}^{(t)} = sign(\hat{p}_{kk}^{(t)})$

Draw a Bernoulli random variable $P^{(t)}$ with probability $\frac{b_1}{b_1 + |\hat{p}_{kk}^{(t)}|}$

if $P^{(t)} = 1$ **then**

Compute $\hat{p}_{km}^{(t)} = \langle x^{(t)}, w_m^{(t-1)} \rangle\ \forall m \neq k, m \in [K]$

Compute $\tilde{p}^{(t)} = \sum_{m \neq k, m \in [K]} \tau_{km}^{(t-1)} \hat{p}_{km}^{(t)}$ and $\tilde{y}^{(t)} = sign(\tilde{p}^{(t)})$

Draw a Bernoulli random variable $Q^{(t)}$ with probability $\frac{b_2}{b_2 + |\tilde{p}^{(t)}|}$

end

Set $Z^{(t)} = P^{(t)} Q^{(t)}$ & $\tilde{Z}^{(t)} = P^{(t)}(1 - Q^{(t)})$

Query true label $y^{(t)}$ if $Z^{(t)} = 1$ and set $M^{(t)} = 1$ if $\hat{y}^{(t)} \neq y^{(t)}$

Update $w_k^{(t)} = w_k^{(t-1)} + (M^{(t)} Z^{(t)} y^{(t)} + \tilde{Z}^{(t)} \tilde{y}^{(t)}) x^{(t)}$

Update τ:

$$\tau_{km}^{(t)} = \frac{\tau_{km}^{(t-1)} e^{-\frac{Z^{(t)}}{\lambda} \ell_{km}^{(t)}}}{\sum_{\substack{m' \in [K] \\ m' \neq k}} \tau_{km'}^{(t-1)} e^{-\frac{Z^{(t)}}{\lambda} \ell_{km'}^{(t)}}} \quad m \in [K], m \neq k \tag{1}$$

end

Theoretical Analysis

- **a theoretical mistake bound**

$$\mathbb{E}\Big[\sum_{t\in[T]} M^{(t)}\Big] \leq \frac{b_2}{\gamma}\Big[\frac{(2b_1+X^2)^2}{8b_1\gamma}(\|w_k^*\|^2 + \max_{m\in[K],m\neq k}\|w_m^*\|^2) + (1+\frac{X^2}{2b_1})\Big(\tilde{L}_{kk} + \max_{m\in[K],m\neq k}\tilde{L}_{km}\Big)\Big]$$

- **the expected number of label requests**

$$\sum_t \frac{b_1}{b_1+|h_k(x^{(t)})|}\frac{b_2}{b_2+\max_{\substack{m\in[K]\\ m\neq k}}|h_m(x^{(t)})|}$$

Experiment

■ Dataset

Landmine Detection data

Spam Detection data

Sentiment Analysis data

■ Method baseline

Random

Independent: chooses the examples via selective sampling independently for each task

■ Method of this paper

PEERsum

PEERone

$$\tilde{p}^{(t)} = \sum_{m \neq k, m \in [K]} \tau_{km}^{(t-1)} \hat{p}_{km}^{(t)} \text{ and } \tilde{y}^{(t)} = sign(\tilde{p}^{(t)})$$

$b_1 = 1$ b_2 is tuned from 20 different values

Experiment

Average rate of mistakes

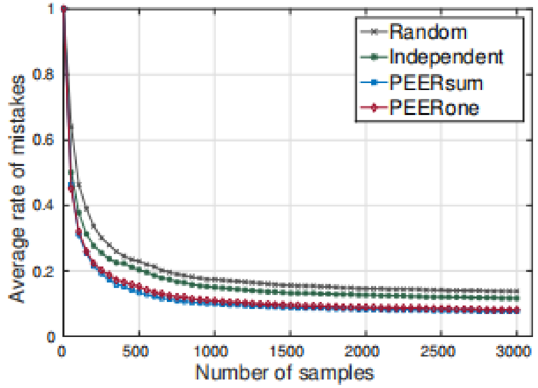

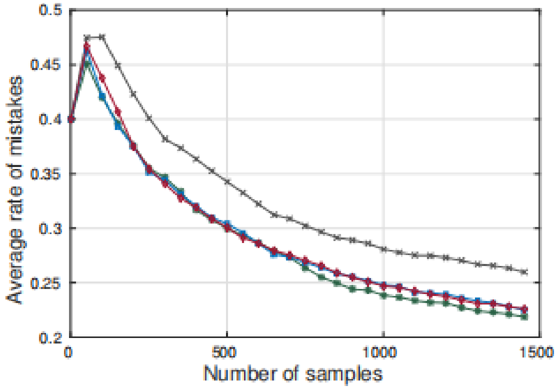

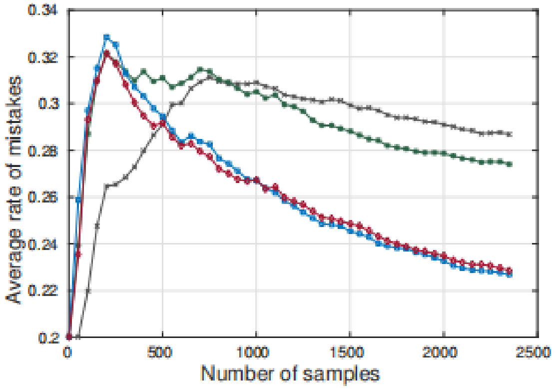

Average rate of label and peer requests

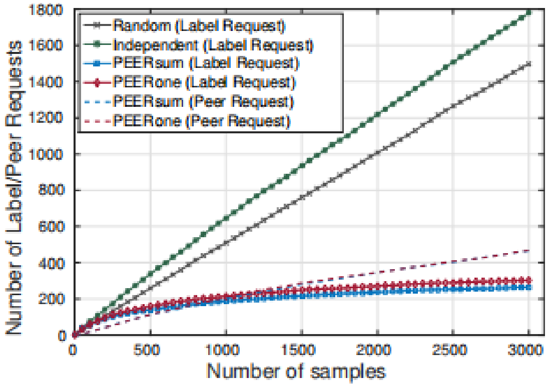

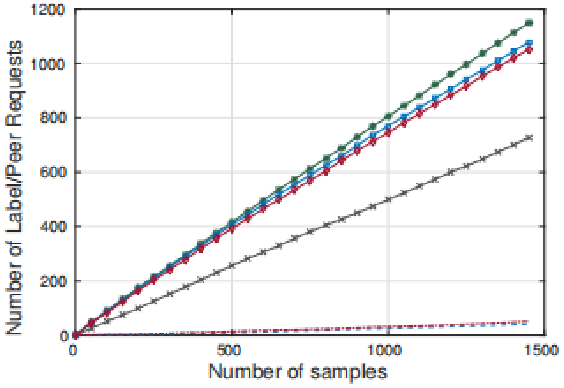

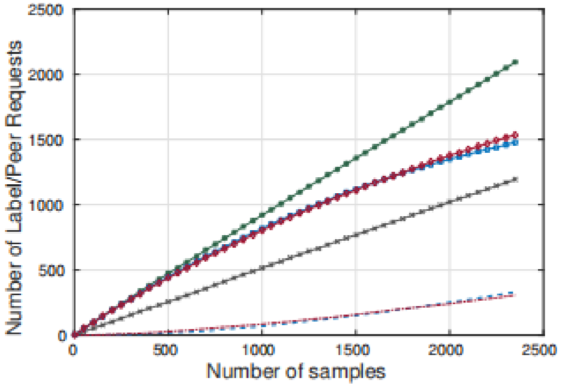

Average rate of (training) mistakes

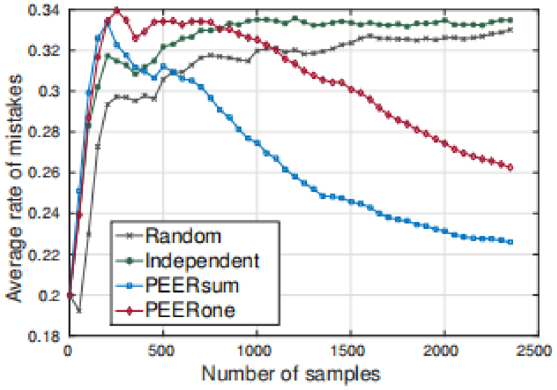

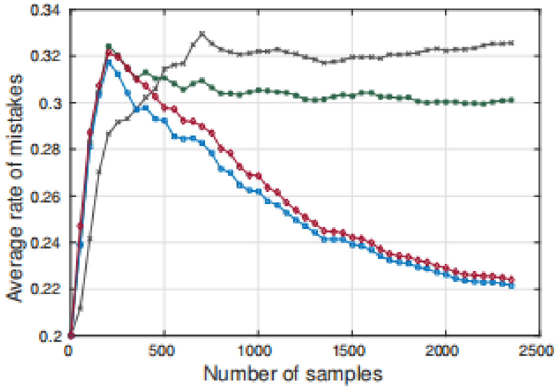

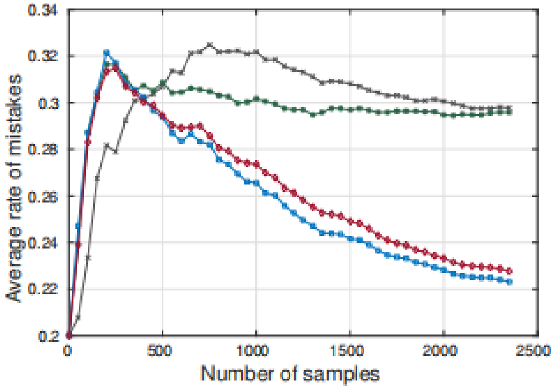

Experiment

Table 1: Average test accuracy on three datasets: means and standard errors over 10 random shuffles.

Models	Landmine Detection			Spam Detection			Sentiment Analysis		
	ACC	*#Queries*	*Time (s)*	*ACC*	*#Queries*	*Time (s)*	*ACC*	*#Queries*	*Time (s)*
Random	0.8905 (0.007)	1519.4 (31.9)	0.38	0.8117 (0.021)	753.4 (29.1)	8	0.7443 (0.028)	1221.8 (22.78)	35.6
Independent	0.9040 (0.016)	1802.8 (35.5)	0.29	0.8309 (0.022)	1186.6 (18.3)	7.9	0.7522 (0.015)	2137.6 (19.1)	35.6
PEERsum	**0.9403** (0.001)	265.6 (18.7)	0.38	**0.8497** (0.007)	1108.8 (32.1)	8	**0.8141** (0.001)	1494.4 (68.59)	36
PEERone	0.9377 (0.003)	303 (17)	1.01	0.8344 (0.018)	1084.2 (24.2)	8.3	0.8120 (0.01)	1554.6 (92.2)	36.3

Experiment

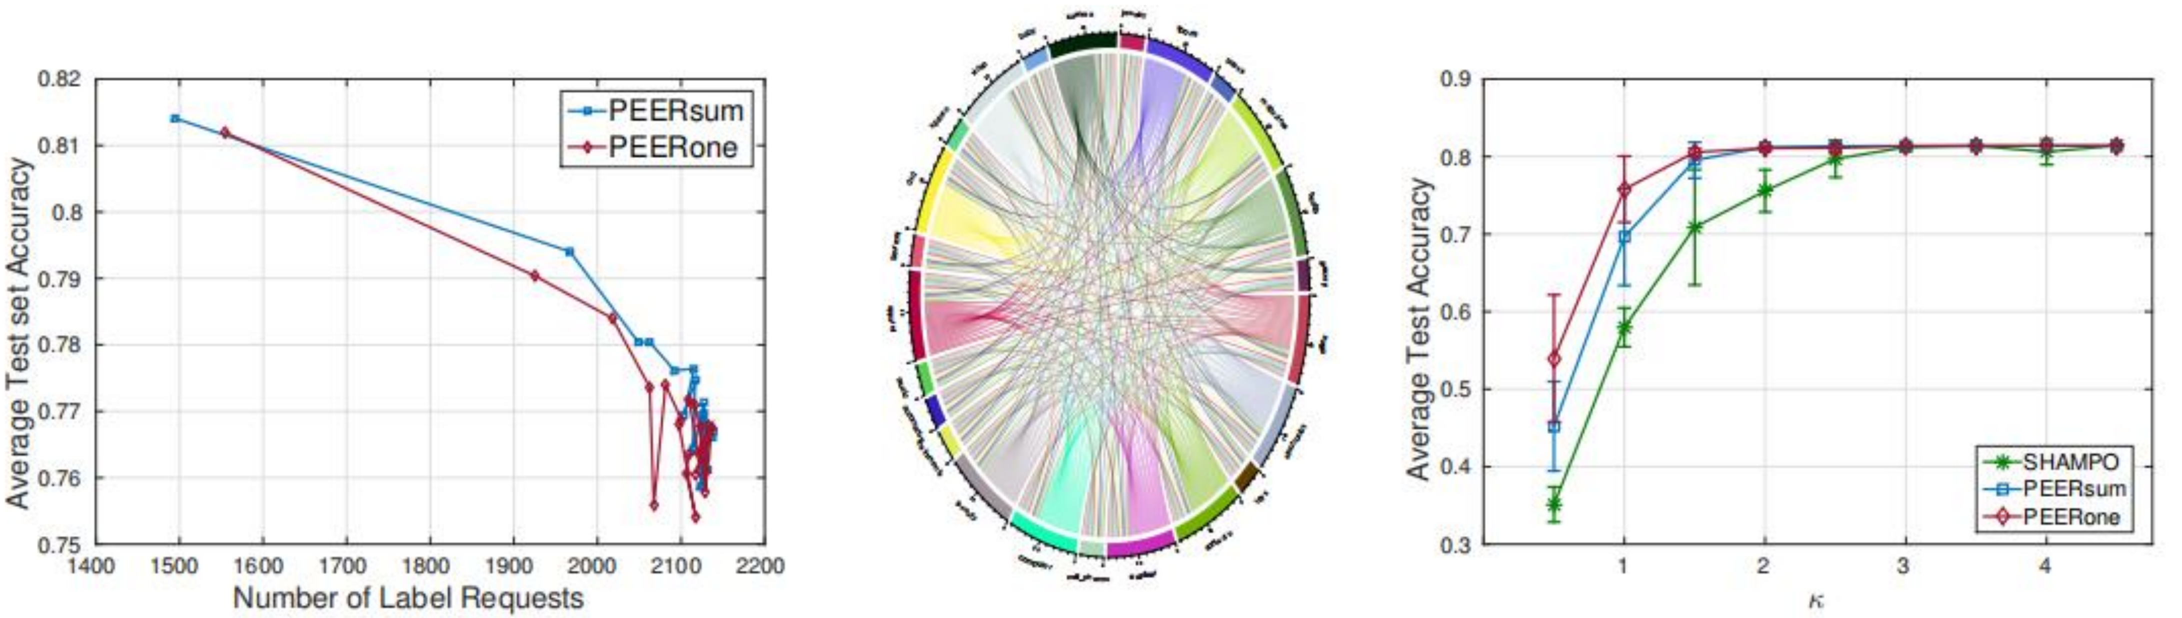

Figure 3: Average test set ACC calculated for different values of b_2 (left). A visualization of the peer query requests among the tasks in *sentiment* learned by PEERone (middle) and comparison of proposed methods against SHAMPO in parallel setting. We report the average test set accuracy (right).

Authors and related papers

Keerthiram Murugesan Jaime Carbonell

Adaptive Smoothed Online Multi-Task Learning	NIPS-2016
Self-Paced Multitask Learning with Shared Knowledge	IJCAI-2017
Multi-Task Multiple Kernel Relationship Learning	SDM-2017