模式分析与机器智能
工业和信息化部重点实验室
MIIT Key Laboratory of
Pattern Analysis & Machine Intelligence

ParNeC 模式识别与神经计算研究组
PAttern Recognition and NEural Computing

An Unbiased Risk Estimator for Learning with Augmented Classes

Yu-Jie Zhang, Peng Zhao, Lanjihong Ma, Zhi-Hua Zhou
National Key Laboratory for Novel Software Technology,
Nanjing University, Nanjing 210023, China
{zhangyj, zhaop, maljh, zhouzh}@lamda.nju.edu.cn

NeurIPS 2020

Introduction

- Open-set problem

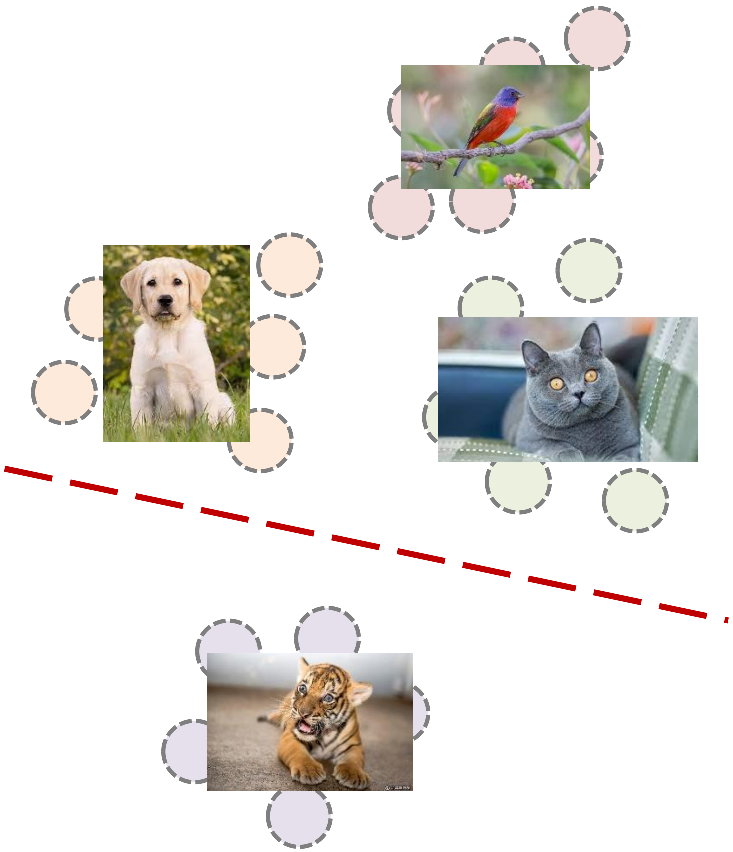

- **The real world is 'open'.**
- **Open-set issue**: models always misclassifies unseen class into one of seen class, which makes its predictions unreliable.
- **Goal**: making the learning system robust to identify unseen classes in the non-stationary environments.

Learning with Augmented Classes

- LAC problem: augmented classes unobserved in training data might emerge in testing

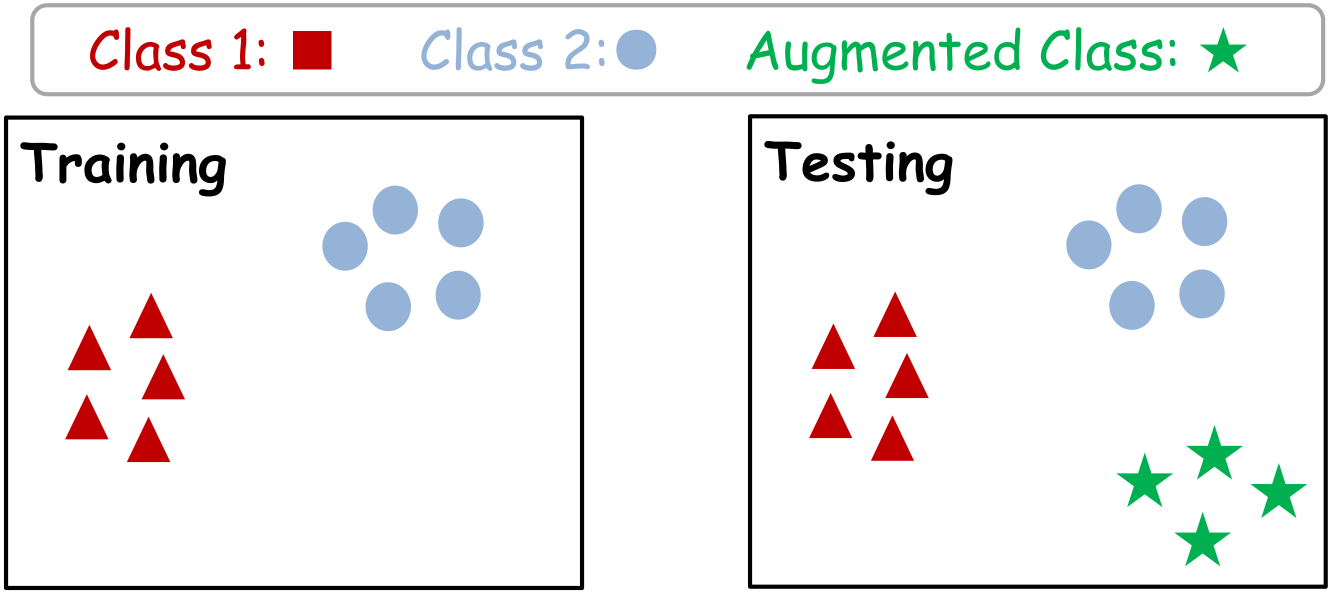

Previous Attempts

- LAC problem: augmented classes unobserved in training data might emerge in testing

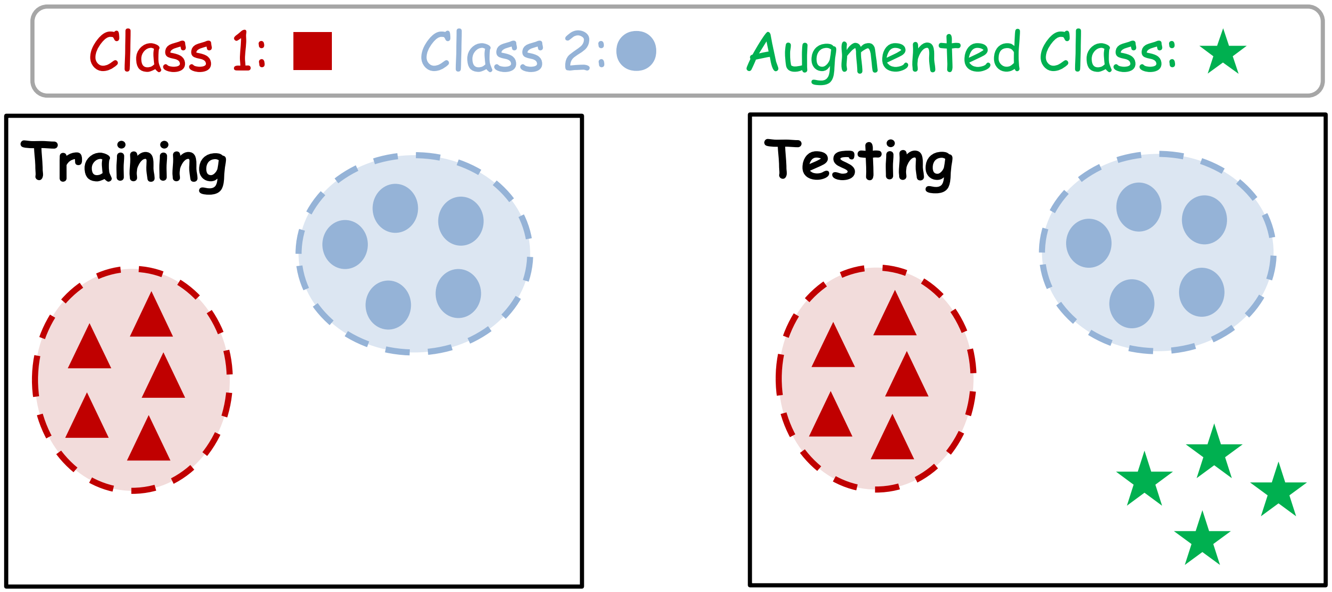

- **Potential Limitation:**
  - Existing methods hardly explore the generalization ability of the model.

Exploiting Unlabeled data for LAC

- LAC problem: augmented classes unobserved in training data might emerge in testing

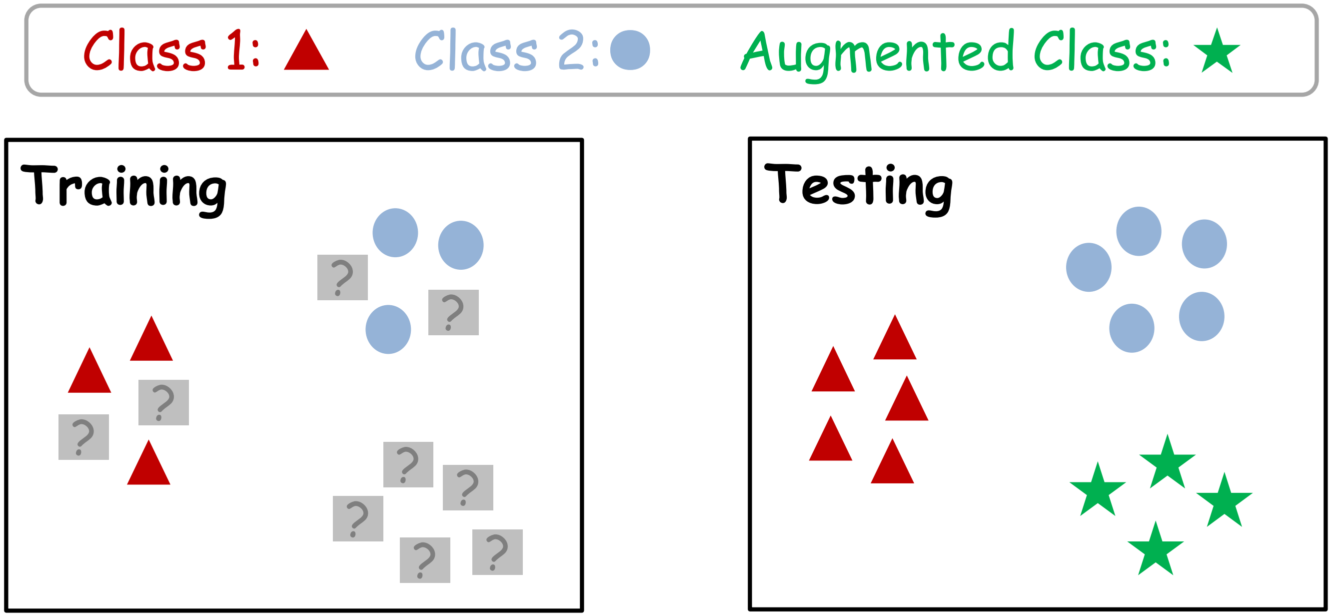

- **Solution in this paper:**
  - Propose an approach with theoretical guarantee by exploiting **unlabeled data**.

Exploiting Unlabeled data for LAC

- Intuition: the distribution of augmented class can be approximated by separating the distribution of known class from that of unlabeled data.

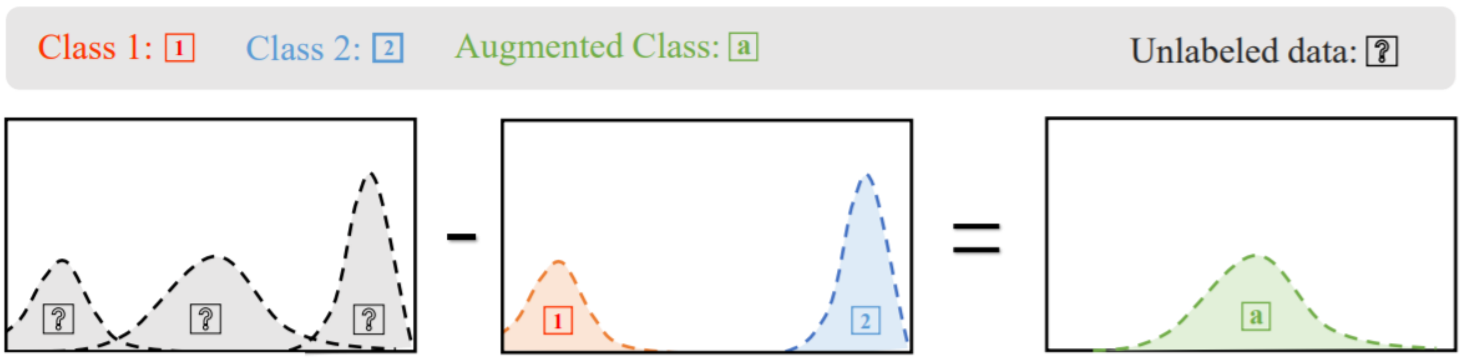

distribution of *unlabeled* data　distribution of *labeled* data　distribution of the class absent in labeled data

- Class shift condition

Definition 1 (Class Shift Condition). The testing distribution P_{te}, the distribution of known classes P_{kc} and the distribution of augmented classes P_{ac} are under the *class shift condition*, if

$$P_{te} = \theta \cdot P_{\mathsf{kc}} + (1-\theta) \cdot P_{\mathsf{ac}}, \tag{1}$$

where $\theta \in [0,1]$ is a certain mixture proportion.[1]

Relation to PU Learning

- PU learning: learning from **positive** and **unlabeled** examples
  - One-sample assumption: both P and U data is drawn from the identical distribution p(x).
  - **Two-sample assumption**: P data is drawn from the positive marginal density $\underline{p(x|Y=+1)}$ and U data is drawn from p(x).

$$P_X = \pi P_1 + (1-\pi)P_{-1}.$$

$$\pi = p(Y=+1)$$

$$P_{te} = \theta \cdot P_{\mathrm{kc}} + (1-\theta) \cdot P_{\mathrm{ac}};$$

  - Ordinary classification risk (P and N data is both accesible)

$$R(f) := \pi R_1(f) + (1-\pi)\boxed{R_{-1}(f)}.$$

$$R_1(f) = P_1(f(X) \neq 1);$$

treat all unlabeled data as negative,

and then sub the loss of positive data.

  - Everything is 'known'.

EULAC

Definition 1 (Class Shift Condition)**.** The testing distribution P_{te}, the distribution of known classes P_{kc} and the distribution of augmented classes P_{ac} are under the *class shift condition*, if

$$P_{te} = \theta \cdot P_{\mathsf{kc}} + (1-\theta) \cdot P_{\mathsf{ac}}, \tag{1}$$

where $\theta \in [0,1]$ is a certain mixture proportion.[1]

- Class shift condition can re-written as

$$\begin{aligned} p_{XY}^{te}(\mathbf{x}, y) &\overset{(1)}{=} \theta \cdot p_{XY}^{\mathsf{kc}}(\mathbf{x}, y) + (1-\theta) \cdot p_{XY}^{\mathsf{ac}}(\mathbf{x}, y) \\ &= \theta \cdot p_{XY}^{\mathsf{kc}}(\mathbf{x}, y) + \mathbb{1}(y = \mathsf{ac}) \cdot (1-\theta) \cdot p_{X}^{\mathsf{ac}}(\mathbf{x}), \end{aligned}$$

- Then, we have

$$R_\psi(\boldsymbol{f}) = \mathbb{E}_{(x,y)\sim P_{te}}[\psi(\boldsymbol{f}(\boldsymbol{x}), y)]$$

Classification risk over testing distribution

$$\Updownarrow$$

$$R_{LAC} = \underbrace{\theta \mathbb{E}_{(x,y)\sim P_{\mathrm{kc}}}[\psi(\boldsymbol{f}(\boldsymbol{x}), y)]}_{\text{risk on known classes}} + \underbrace{\mathbb{I}(y = ac)(1-\theta)\mathbb{E}_{x\sim P_{\mathrm{ac}}}[\psi(\boldsymbol{f}(\boldsymbol{x}), ac)]}_{\text{risk on augmented classes}}$$

EULAC

Definition 1 (Class Shift Condition)**.** The testing distribution P_{te}, the distribution of known classes P_{kc} and the distribution of augmented classes P_{ac} are under the *class shift condition*, if

$$P_{te} = \theta \cdot P_{\mathsf{kc}} + (1-\theta) \cdot P_{\mathsf{ac}}, \tag{1}$$

where $\theta \in [0,1]$ is a certain mixture proportion.[1]

- The unbiased estimator can be written as

$$R_{LAC} = \theta \mathbb{E}_{(x,y)\sim P_{\mathrm{kc}}}[\psi(\boldsymbol{f}(\boldsymbol{x}), y)] + \mathbb{I}(y = ac)\boxed{(1-\theta)\mathbb{E}_{x\sim P_{\mathrm{ac}}}[\psi(\boldsymbol{f}(\boldsymbol{x}), ac)]}$$

$$\Updownarrow$$

$$(1-\theta)\cdot p_X^{\mathsf{ac}}(\mathbf{x}) = p_X^{te}(\mathbf{x}) - \theta \cdot p_X^{\mathsf{kc}}(\mathbf{x}).$$

$$R_{LAC} = \underbrace{\theta \mathbb{E}_{(x,y)\sim P_{\mathrm{kc}}}[\psi(\boldsymbol{f}(\boldsymbol{x}), y)]}_{\text{labeled training data}} + \underbrace{\mathbb{E}_{x\sim p_x^{te}(x)}[\psi(\boldsymbol{f}(\boldsymbol{x}), ac)]}_{\text{unlabeled training data}}$$

Experiments: Comparison Results

- Performance of classifying known classes and identifying augmented classes
  - Comparison on RKHS-based EULAC

Table 1: Macro-F1 scores on benchmark datasets. The best method is emphasized in bold. Besides, • indicates that EULAC is significantly better than others (paired t-tests at 5% significance level).

Dataset	OVR-SVM	W-SVM	OSNN	EVM	LACU-SVM	PAC-iForest	EULAC
usps	75.42 ± 4.87 •	79.77 ± 4.97 •	63.14 ± 8.91 •	61.14 ± 6.27 •	69.20 ± 8.34 •	55.69 ± 13.3 •	**86.52 ± 2.72**
segment	71.78 ± 5.12 •	80.82 ± 9.38 •	85.10 ± 5.98	82.13 ± 5.88 •	40.69 ± 12.5 •	63.64 ± 13.1 •	**86.17 ± 5.80**
satimage	54.67 ± 9.80 •	76.29 ± 13.2 •	62.48 ± 11.2 •	72.10 ± 8.16 •	51.56 ± 17.3 •	60.76 ± 7.79 •	**81.25 ± 6.18**
optdigits	80.11 ± 3.80 •	87.82 ± 4.64 •	86.97 ± 3.79 •	72.00 ± 8.33 •	80.92 ± 3.68 •	71.65 ± 5.46 •	**91.54 ± 2.95**
pendigits	72.78 ± 5.19 •	87.79 ± 3.95	86.69 ± 3.39 •	**89.94 ± 1.30**	70.66 ± 6.18 •	73.21 ± 4.52 •	88.41 ± 4.81
SenseVeh	48.07 ± 3.80 •	45.96 ± 2.32 •	49.91 ± 6.88 •	51.24 ± 3.91 •	51.61 ± 3.31 •	54.12 ± 7.19 •	**77.33 ± 2.17**
landset	60.43 ± 7.65 •	68.91 ± 17.0 •	73.25 ± 9.23 •	76.00 ± 7.79 •	53.59 ± 9.88 •	70.50 ± 7.16 •	**85.70 ± 4.46**
mnist	66.74 ± 2.76 •	75.38 ± 4.62 •	57.75 ± 10.9 •	58.39 ± 5.94 •	63.53 ± 7.58 •	48.31 ± 9.62 •	**80.66 ± 5.38**
shuttle	37.39 ± 14.1 •	58.48 ± 34.5 •	48.21 ± 16.4 •	–	34.18 ± 13.4 •	29.36 ± 8.70 •	**66.49 ± 17.9**
EULAC w/ t/ l	9/ 0/ 0	8/ 1/ 0	8/ 1/ 0	8/ 1/ 0	9/ 0/ 0	9/ 0/ 0	rank first 8/ 9

Experiments: Comparison Results

- Performance of classifying known classes and identifying augmented classes
  - Comparison on deep models

Table 2: AUC for DNN-based EULAC

Methods	mnist	Cifar-10	SVHN
SoftMax	97.8 ± 0.6	67.7 ± 3.8	88.6 ± 1.4
OpenMax	98.1 ± 0.5	69.5 ± 4.4	89.4 ± 1.3
G-OpenMax	98.4 ± 0.5	67.5 ± 4.4	89.6 ± 1.7
OSRCI	$\mathbf{98.8} \pm \mathbf{0.4}$	69.9 ± 3.8	$\mathbf{91.0} \pm \mathbf{1.0}$
EULAC	98.6 ± 0.4	$\mathbf{85.2} \pm \mathbf{2.0}$	$\mathbf{91.2} \pm \mathbf{2.8}$

Study on the Size of Unlabeled Data

- Performance of classifying known classes and identifying augmented classes
  - Influence on the size of unlabeled data

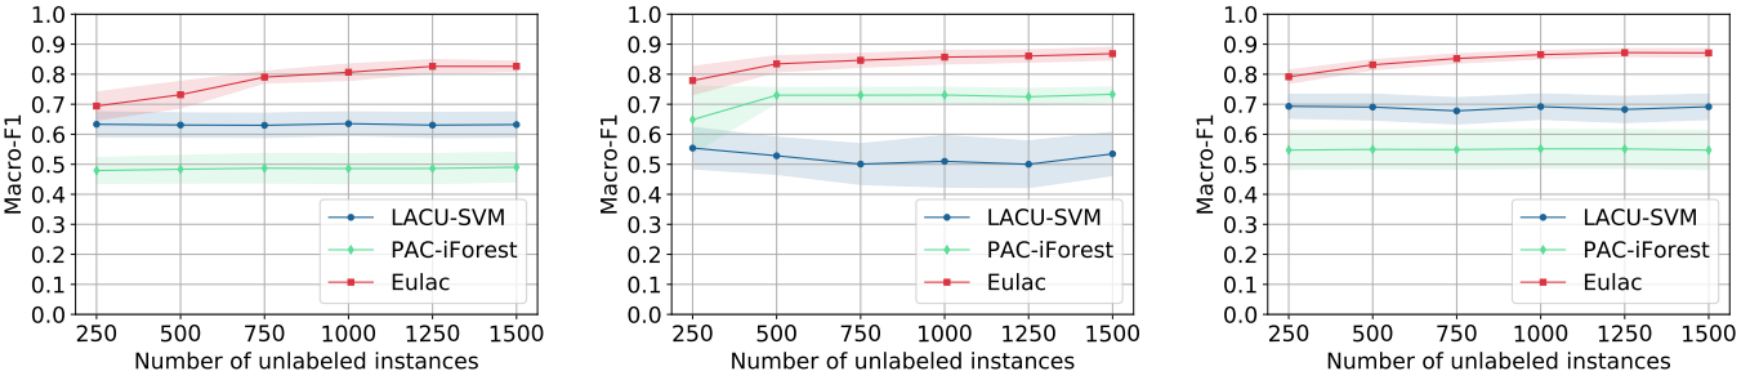

(a) mnist (b) landset (c) usps

Figure 2: Macro-F1 score comparisons when the number of unlabeled data increases.

Experiments: Influence of Mixture Proportion

- Accuracy of estimating mixture prior θ and its influence on EULAC

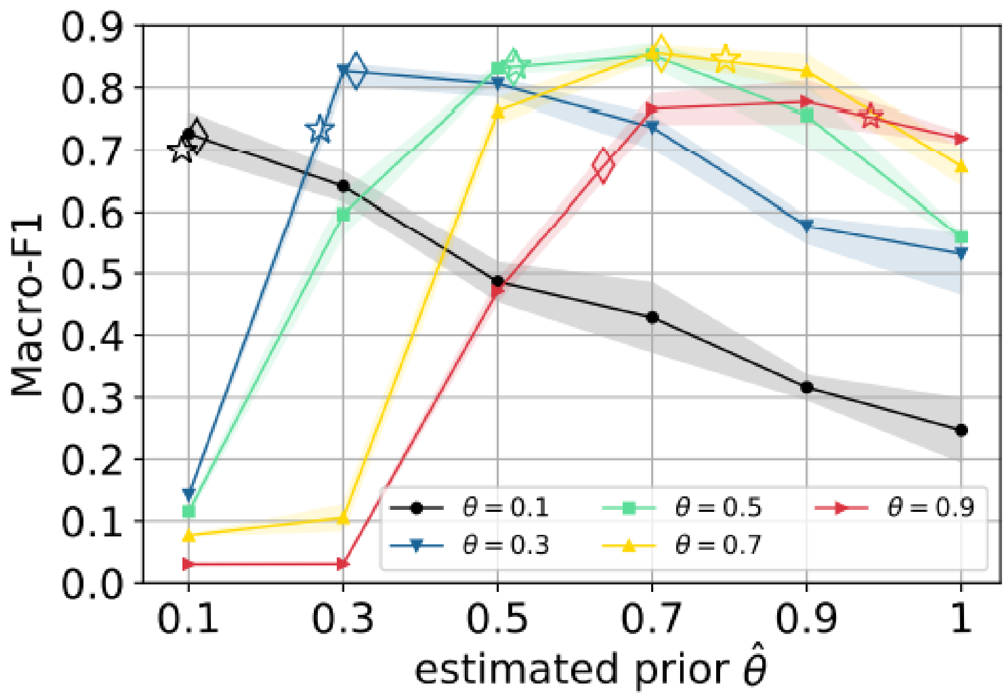

Figure 3: Influence and estimation accuracy of mixture proportion θ.

THANKS