Teacher-Student Competition for Unsupervised Domain Adaption

Ruixin Xiao[1], Zhilei Liu[1*], Baoyuan Wu[2]

[1]College of Intelligence and Computing, Tianjin University, Tianjin, China

[2]School of Data Science, Chinese University of Hong Kong, Shenzhen, China

{ruixinx, zhileiliu}@tju.edu.cn, baoyuanwu1987@gmail.com

Abstract

With the supervision from source domain only in class-level, existing unsupervised domain adaption (UDA) methods mainly learn the domain-invariant representations from a shared feature extractor, which causes the source-bias problem. This paper proposes an unsupervised domain adaption approach with Teacher-Student Competition (TSC). In particular, a student network is introduced to learn the target-specific feature space, and we design a novel competition mechanism to select more credible pseudo-labels for the training of student network. We introduce a teacher network with the structure of existing conventional UDA method, and both teacher and student networks compete to provide target pseudo-labels to constrain every target sample's training in student network. Extensive experiments demonstrate that our proposed TSC framework significantly outperforms the state-of-the-art domain adaption methods on Office-31 and ImageCLEF-DA benchmarks.

Motivation

Domain shift in data distributions across domains is a major obstacle for transferring the model trained on the source domain to the target domain.

To mitigate it, as illustrated in Fig. 1 (a), traditional unsupervised domain adaptation method focuses on aligning source and target features in a common feature space through statistics measures or adversarial learning to learn domain-invariant feature representations.

However, with the supervision from source domain only in class-level, the adapted feature space is source-biased. So target feature distributions are aligned forcibly under such a source-biased criterion, thus making the classification boundaries not fit well for target domain.

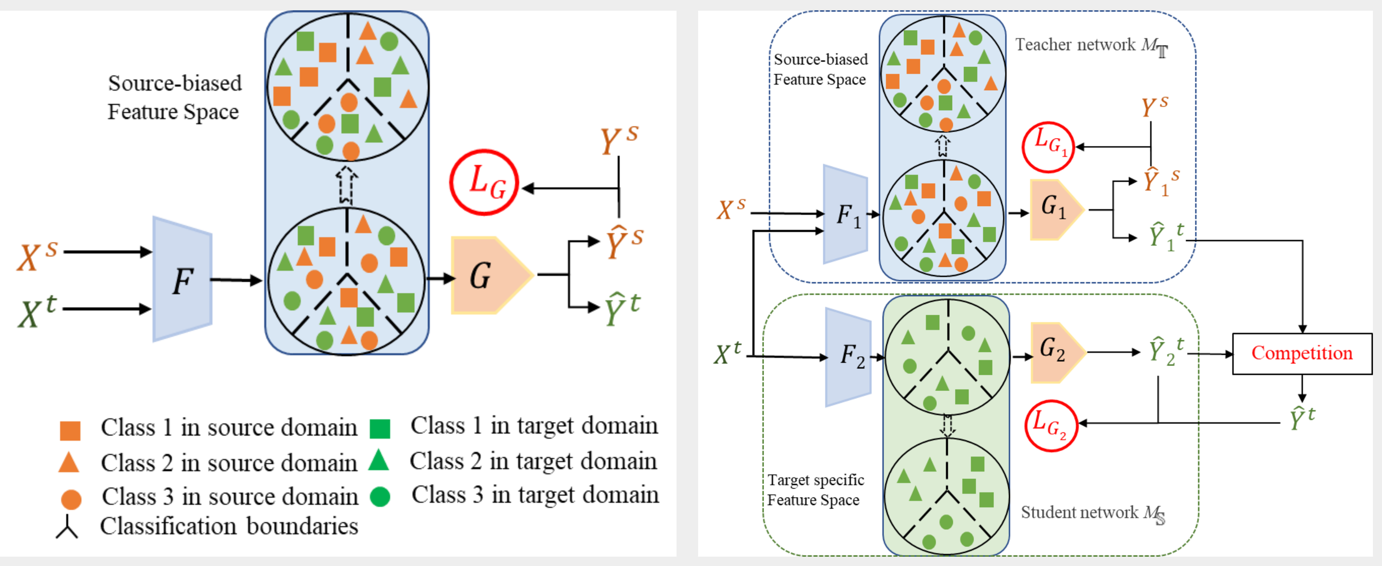

(a) Traditional UDA method (b) Our method with TSC

Fig.1: Comparison between traditional UDA method and our method with TSC.

As illustrated in Fig. 1 (b), to tackle the source-bias problem, we build a target-specific student network, which is fed into target data only, to learn a target-specific feature space. But this network can't be constrained without any target ground-truth labels during training.

We introduce existing conventional UDA methods to provide target pseudo-labels as a reference, like a teacher to guide the student. But we are not going to make the student network copy the teacher network. So the reference are not used to train student network directly.

Instead, we design a competition module to select the more reliable pseudo-labels.

Our motivation of competition module comes from teacher-student learning process in human learning, which is a process of chasing and surpassing. At the initial stage of learning, a student learns knowledge with the guidance of his teacher mainly. When student gets further understanding about the task, he makes his own predictions with higher and higher confidence. When student makes prediction with higher credibility than his teacher in the later phase of the learning, he does believe himself even if his prediction is inconsistent with which of teacher, which is key to enable student to breakthrough the inherent error of teacher.

Contributions

To alleviate the source-biased problem, a target-specific student network is built to learn a target-specific feature space.

A novel pseudo-label selection strategy, in which pseudo-labels from teacher network and student network compete to be the final pseudo-label training for student network.

Extensive experimental results demonstrate that our method achieves state of the art performance on common benchmark domain adaption tasks.

Method

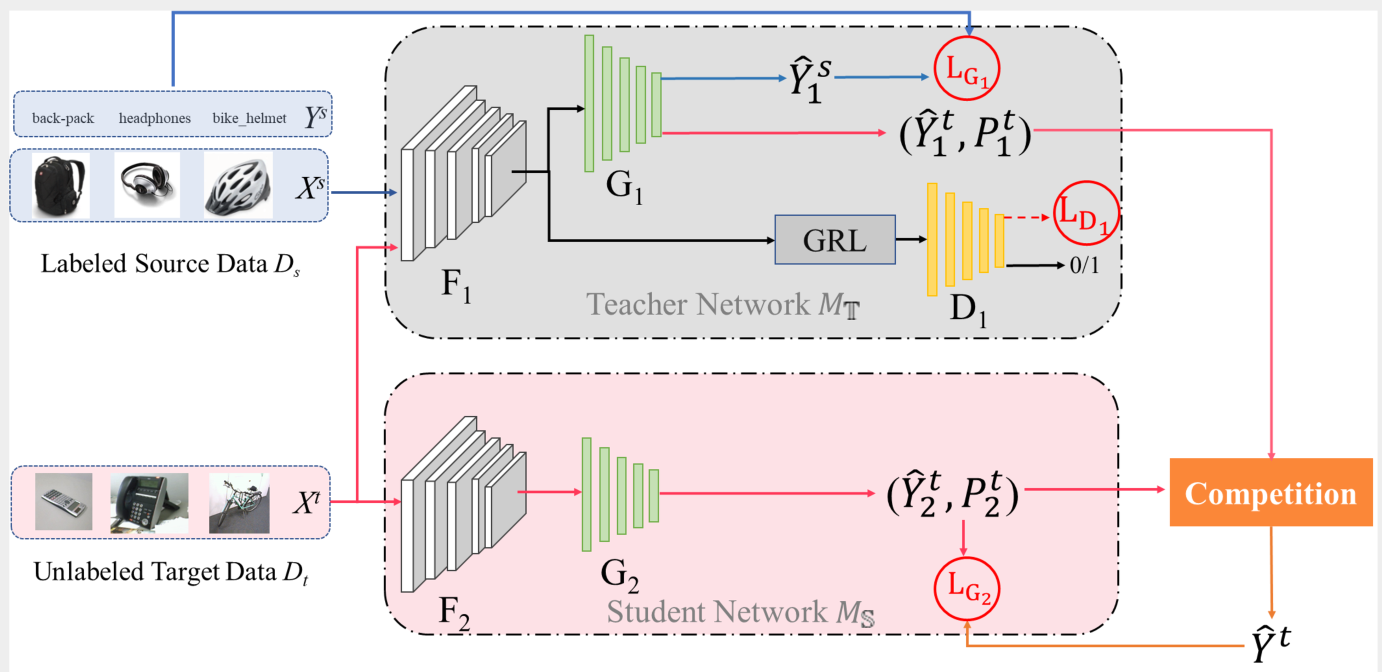

Fig.2: Teacher-Student Competition (TSC) method with DANN as a teacher.

Competition

$$T_p = \frac{1}{1+\exp(-\delta p)}$$

$$\hat{y}_j^t = \begin{cases} \hat{y}_{1,j}^t, & \text{if} \quad p_{1,j}^t > T_p \quad \text{or} \ p_{1,j}^t > p_{2,j}^t \\ \hat{y}_{2,j}^t, & \text{otherwise} \end{cases}$$

δ=10 and p denotes the training process

Results

Table 1 Accuracy (%) on ImageCLEF-DA for unsupervised domain adaptation (ResNet-50).

Methods	I→P	P→I	I→C	C→I	C→P	P→C	AVG
RestNet-50	**74.8**±0.3	91.5±0.3	83.9±0.1	78.0±0.2	65.5±0.3	91.2±0.3	80.7
DAN	74.5±0.4	82.2±0.2	92.8±0.2	86.3±0.4	69.2±0.4	89.8±0.4	82.5
DANN	75.0±0.6	86.0±0.3	96.2±0.4	87.0±0.5	74.3±0.5	91.5±0.6	85.0
JAN	76.8±0.4	88.0±0.2	94.7±0.2	89.5±0.3	74.2±0.3	91.7±0.3	85.8
MADA	75.0±0.3	87.9±0.2	96.0±0.3	88.8±0.3	75.2±0.2	92.2±0.3	85.8
CDAN	76.7±0.3	90.6±0.3	97.0±0.4	90.5±0.4	74.5±0.3	93.5±0.4	87.1
CDAN+E	77.7±0.3	90.7±0.2	**97.7**±0.3	91.3±0.3	74.2±0.2	94.3±0.3	87.7
iCAN	79.5	89.7	94.7	89.9	78.5	92.0	87.4
rDANN+CAT	77.2±0.2	91.0±0.3	95.5±0.3	91.3±0.3	75.3±0.6	93.6±0.5	87.3
TSC+DANN	78.3±0.2	92.8±0.3	96.8±0.2	90.3±0.5	74.5±0.7	96.0±0.2	88.1
TSC+CDAN	**79.0**±0.3	**93.2**±0.5	**97.2**±0.4	**92.7**±0.2	**77.4**±0.4	**96.5**±0.3	**89.3**

Table 2 Accuracy (%) on Office-31 for unsupervised domain adaptation (ResNet-50).

Methods	A→W	D→W	W→D	A→D	D→A	W→A	AVG
RestNet-50	68.4±0.2	96.7±0.1	99.3±0.1	68.9±0.2	62.5±0.3	60.7±0.3	76.1
DAN	80.5±0.4	97.1±0.2	99.6±0.1	78.6±0.2	63.6±0.3	62.8±0.2	80.4
DANN	82.0±0.4	96.9±0.2	99.1±0.1	79.7±0.4	68.2±0.4	67.4±0.5	82.2
JAN	85.4±0.3	97.4±0.2	99.8±0.2	84.7±0.3	68.6±0.3	70.0±0.4	84.3
MADA	90.0±0.1	97.4±0.1	99.6±0.1	87.8±0.2	70.3±0.3	66.4±0.3	85.2
GTA	89.5±0.5	97.9±0.3	99.8±0.4	87.7±0.5	72.8±0.3	71.4±0.4	86.5
CDAN	93.1±0.2	98.2±0.2	**100.0**±0.0	89.8±0.3	70.1±0.4	68.0±0.4	86.6
CDAN+E	94.1±0.1	98.6±0.1	**100.0**±0.0	92.9±0.2	71.0±0.3	69.3±0.3	87.7
iCAN	92.5	**98.8**	**100.0**	90.1	72.1	68.9	87.2
rDANN+CAT	94.4±0.1	98.0±0.2	**100.0**±0.0	90.8±1.8	72.2±0.6	70.2±0.1	87.6
TSC+DANN	85.0±0.3	98.0±0.1	**100.0**±0.0	80.3±0.4	69.3±0.2	67.7±0.1	83.4
TSC+CDAN	**94.6**±0.3	98.2±0.2	**100.0**±0.0	**94.7**±0.1	**74.0**±0.2	**71.6**±0.7	**88.9**

Table 1 and Table 2 show that our TSC+DANN and TSC+CDAN outperform state of the art results on most DA tasks and bring excellent increments over DANN and CDAN.

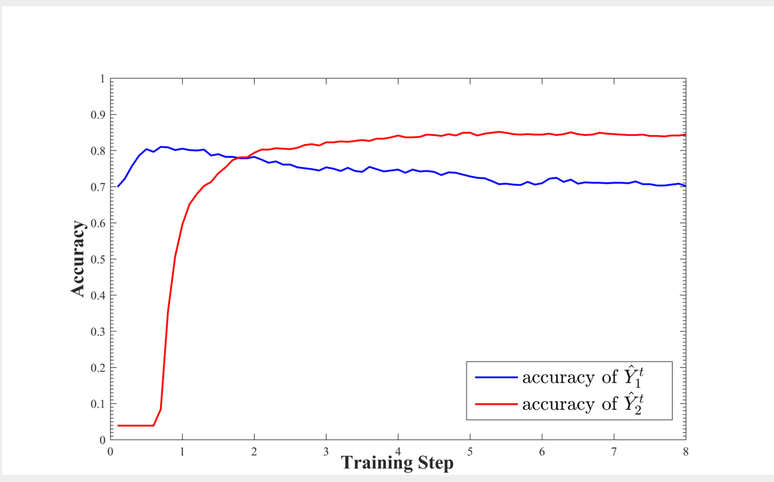

Fig. 3:Accuracy of TSC+DANN on A→W in office-31

To analyze the poor performances when TSC+DANN in office-31, Fig. 3 trace accuracy fluctuations of teacher and student as the training goes on A→W.

It can be observed that the accuracy of teacher can't maintain its peak stably. The student outperforms teacher at the later stage with a lower performance than other methods.

Results

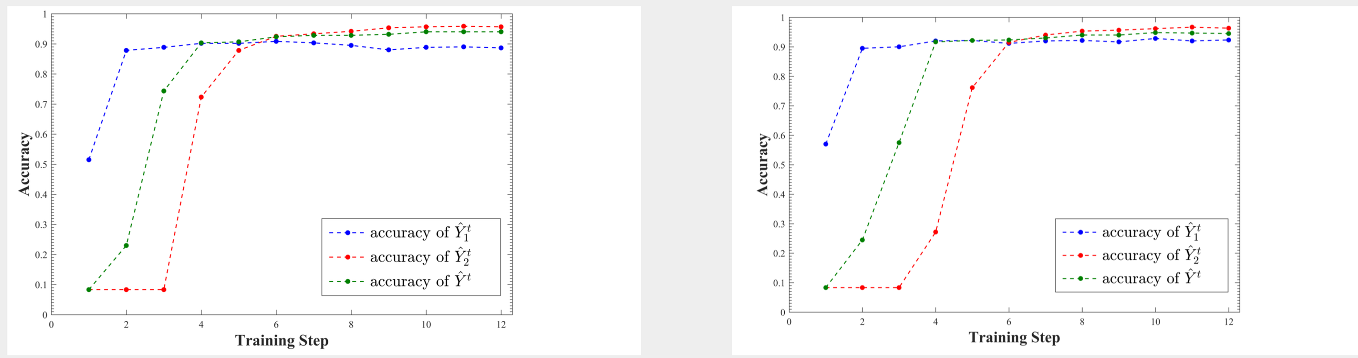

(a) TSC+DANN (b) TSC+CDAN

Fig. 4: Pseudo-labeling accuracy comparison on task P→C in ImageCLEF-DA.

Fig. 4 verifies the effectiveness of the competition module during training. Blue curves show the performance fluctuation of DANN and CDAN as the training goes. Green curves for the competition module hoist with the help of teacher and competition strategy and has higher accuracy than teacher at the later phase. Student raises under the constraint of pseudo-labels competition module from and climbs to the top finally.

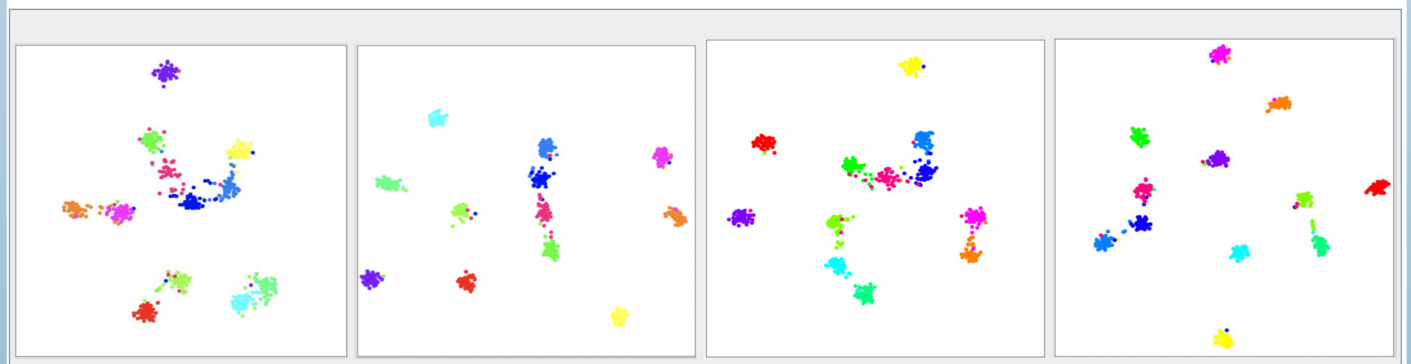

(a) M_T in TSC+DANN (b) M_S in TSC+DANN (c) M_T in TSC+CDAN (d) M_S in TSC+CDAN

Fig.5: Visualization of target feature

Fig. 5 shows no matter when teacher is DANN or CDAN, our student network learns more discriminative representations, has more clear inter-class boundaries and less classified falsely samples.

Conclusion

Our proposed TSC significantly outperforms the state-of-the-art domain adaption methods.

More separable target feature space can be achieved by introducing our competition model to tackle the source-bias problem.

Future Work

Consider to set the hyper parameter T_p dynamically and smoothly to improve the performance further

Acknowledgment

This work is supported by the National Natural Science Foundation of China (No. 61503277). We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan V GPU used for this research.

Main references

[1] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in International Conference on Machine Learning, 2015, pp. 1180–1189.

[2] M. Long, Z. Cao, J. Wang, and M. I. Jordan, "Conditional adversarial domain adaptation," in Advances in Neural Information Processing Systems, 2018, pp. 1640–1650.

25th International Conference on Pattern Recognition (ICPR 2020), Milan, Italy, January 10-15, 2021