

Attract, Perturb, and Explore: Learning a Feature Alignment Network for Semi-supervised Domain Adaptation (Supplementary Material)

Taekyung Kim^[0000-0001-7401-098X] and Changick Kim

Korea Advanced Institute of Science and Technology, Daejeon, South Korea
 {tkkim93, changick}@kaist.ac.kr

Table 1: Classification accuracy (%) on the DomainNet dataset with the ResNet-34 backbone network. The performance comparisons were done for 7 scenarios on five-shot and ten-shot settings.

Net	Method	R to C	R to P	P to C	C to S	S to P	R to S	P to R	MEAN
Five-shot									
ResNet	S+T	64.5	63.1	64.2	59.2	60.4	56.2	75.7	63.3
	DANN	63.7	62.9	60.5	55.0	59.5	55.8	72.6	61.4
	CDAN	68.0	65.0	65.5	58.0	62.8	58.4	74.8	64.6
	ENT	77.1	71.0	75.7	61.9	66.2	64.6	81.1	71.1
	MME	75.5	70.4	74.0	65.0	68.2	65.5	79.9	71.2
	Ours	77.7	73.0	76.9	67.0	71.4	68.8	80.5	73.6
Ten-shot									
ResNet	S+T	68.5	66.4	69.2	64.8	64.2	60.7	77.3	67.3
	DANN	70.0	64.5	64.0	56.9	60.7	60.5	75.9	64.6
	CDAN	69.3	65.3	64.6	57.5	61.6	60.2	77.0	65.1
	ENT	79.0	72.9	78.0	68.9	68.4	68.1	82.6	74.0
	MME	77.1	71.9	76.3	67.0	69.7	67.8	81.2	73.0
	Ours	79.8	75.1	78.9	70.5	73.6	70.8	82.9	76.8

1 Additional Results and Analysis

Additional Performance Comparisons with Varying Number of Target Labels. We additionally conducted performance comparisons on five-shot and ten-shot settings for all the scenarios in the DomainNet dataset. We used ResNet-34 as the backbone network for the experiments. As shown in Table 1, our method outperformed all the baselines with a considerable margin on average on both the five-shot and ten-shot settings. Moreover, our method showed superior results on all the scenarios compare to the other baselines except only one scenario, which verifies the superiority of our method on conventional domain

adaptation methods, conventional semi-supervised methods, and the state-of-the-art semi-supervised domain adaptation method. On the other hand, while the conventional domain adaptation methods showed less effective or even worse performances than S+T, ENT showed comparable results on the five-shot setting and outperformed the state-of-the-art SSDA method on the ten-shot setting on average. As discussed in the main paper, it is because the SSDA problem gradually resembles the SSL problem on the target domain as the number of labeled target samples increases. Note that the five-shot and ten-shot settings stand for 1.8% \sim 6.7% and 3.6% \sim 13.5% ratio among all target samples on the target domains of the DomainNet dataset.

References