

TOWARDS REALIZING THE VALUE OF LABELED TARGET SAMPLES: A TWO-STAGE APPROACH FOR SEMI-SUPERVISED DOMAIN ADAPTATION

Mengqun Jin¹, Kai Li², Shuyan Li¹, Chunming He¹, Xiu Li^{1‡}

¹Tsinghua Shenzhen International Graduate School, Tsinghua University, ShenZhen, China

² NEC Labs, America

ABSTRACT

Semi-Supervised Domain Adaptation (SSDA) is a recently emerging research topic that extends from the widely-investigated Unsupervised Domain Adaptation (UDA) by further having a few target samples labeled, i.e., the model is trained with labeled source samples, unlabeled target samples as well as *a few labeled* target samples. Compared with UDA, the key to SSDA lies how to most effectively utilize the few labeled target samples. Existing SSDA approaches simply merge the few precious labeled target samples into vast labeled source samples or further align them, which dilutes the value of labeled target samples and thus still obtains a biased model. To remedy this, in this paper, we propose to decouple SSDA as an UDA problem and a semi-supervised learning problem where we first learn an UDA model using labeled source and unlabeled target samples and then adapt the learned UDA model in a semi-supervised way using labeled and unlabeled target samples. By utilizing the labeled source samples and target samples separately, the bias problem can be well mitigated. We further propose a consistency learning based mean teacher model to effectively adapt the learned UDA model using labeled and unlabeled target samples. Experiments show our approach outperforms existing methods.

Index Terms— Semi-supervised domain adaptation, domain adaptation, semi-supervised learning

1. INTRODUCTION

Domain adaptation (DA) studies the performance degradation problem of a model trained on a source dataset when tested on an out-of-distribution target dataset. Most existing methods tackle this problem by learning an adaptive model using labeled source samples and unlabeled target samples. These methods are called Unsupervised (target samples are not labeled) Domain Adaptation (UDA) methods. Recently some works [1, 2, 3, 4, 5] started investigating the effect of having a few target samples labeled, and formulated a new research topic called Semi-Supervised Domain Adaptation (SSDA), which aims to learn a better adaptive model using labeled

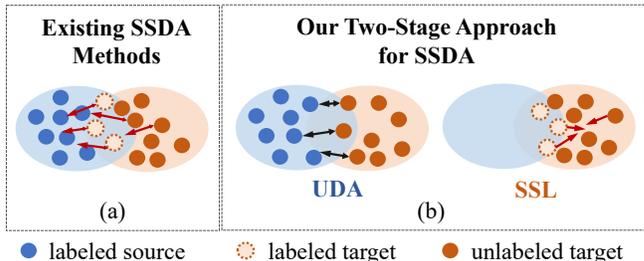


Fig. 1. Schematic difference of the proposed method from existing methods.

source samples, unlabeled target samples, as well as *a few* (e.g., one sample per class) labeled target samples.

A naive way to leverage the few labeled target samples is to merge them into the labeled source samples and train a model in an UDA fashion [2, 6]. However, due to the overwhelming dominance of the labeled source samples, the value of the few labeled target samples cannot be fully realized. Some methods strive to mitigate this problem: [7] performed explicitly class-wise domain alignment using labeled samples from both domains, aiming to push close samples from the same class across domains. Similarly, [3] use contrastive learning to achieve class-wise alignment. [1] instead used adversarial adaptive clustering [1] to encourage samples from the same classes to agglomerate around the labeled target samples. Albeit achieving improved performance, these methods shall still suffer from the imbalance problem as the model could easily learn how to slightly adjust the classes boundaries to include several more labeled samples, resulting in a model still biased towards the source domain.

In this paper, we approach SSDA in a new perspective which fundamentally resolves the imbalance problem by utilizing labeled source samples and labeled target samples in two separated stages. In this first stage, we use labeled source samples and unlabeled target samples to train an UDA model. Next, we aim to adapt this UDA method using only target samples, both labeled and unlabeled, in a Semi-Supervised Learning (SSL) fashion. As in the second stage, only the labeled target samples provide supervision signals, their values are more likely to be realized, without risking being diluted by vast label source samples, facilitating to produce a model more favorable to the target domain. Fig. 1 shows the contrast

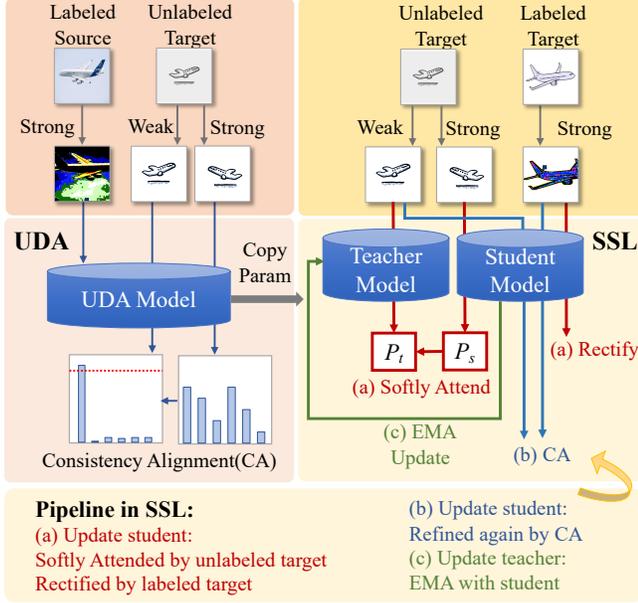


Fig. 2. Framework of our proposed method.

of our method to existing SSDA methods.

Decoupling SSDA as an UDA problem and an SSL problem brings another benefit: In the second stage, we no longer need access to the labeled source samples which could be under protection, privacy sensitive, or large-scale such that it is less likely to be shared for adaptation [8, 9]. Our two-stage solution would thus be favored in the scenario where solution provider only offers models which are trained with protected labeled source data and cheap unlabeled target data and the scenario with resource-contracted devices.

Another imbalance problem arises immediately with the proposed two-stage solution - the imbalance between labeled target samples and unlabeled target samples. We propose a novel Consistency Learning based Mean Teacher (CLMT) model to handle this. For each unlabeled target image, we generates one *weakly-augmented* image and one *strongly-augmented* image. We use a teacher model to produce a pseudo label for the image and use the pseudo label to train the student model in supervised way by taking the strongly-augmented image as input. In this way, we gradually turn unlabeled target samples into labeled ones and thus alleviate this newly-arising imbalance problem.

The contributions of this paper are as follows: (1) We propose to address the SSDA problem in a new perspective by decoupling it as a UDA problem and a SSL problem, which not only resolves the imbalance problem between labeled source samples and labeled target samples, but also offers a solution to protecting labeled source samples. (2) We propose the consistency learning based mean teacher model which is able to mitigate the imbalance problem between unlabeled and labeled target samples. (3) Our method outperforms existing SSDA methods on the common evaluation datasets.

2. METHOD

In the setting of SSDA, we have a large bunch of source data $S = \{(\mathbf{s}_i, y_i^s)\}_{i=1}^{N_s}$, abundant unlabeled target domain data $T_u = \{\mathbf{u}_i\}_{i=1}^{N_u}$ and only a few labeled target data $T_l = \{(\mathbf{t}_i, y_i^t)\}_{i=1}^{N_t}$. Here, the labeled target data is usually 1-shot or 3-shot, which is extremely scarce compared to N_u . We name the labeled target samples as ‘‘Anchors’’. We use T to denote the target domain data, where $T = T_u \cup T_l$ and $T_u \cap T_l = \emptyset$. S and T have different distribution but share the same label space $Y = \{1, 2, \dots, C\}$. C is the number of classes. We aim to learn a domain adaptive model which performs well on the target domain. Let the model be $\theta = f \circ g$, where f extracts features and feeds them into the classifier g that outputs the predictions. Our method includes two stage, the UDA stage and the SSL stage.

2.1. Stage I: Unsupervised Domain Adaptation

We train a UDA model using the following learning objective:

$$\mathcal{L}_{uda} = \mathcal{L}_s(S) + \alpha \mathcal{L}_u(S, T_u), \quad (1)$$

where $\mathcal{L}_s(S)$ is the supervised loss with labeled source samples, and $\mathcal{L}_u(S, T_u)$ is the domain alignment loss with labeled source and unlabeled target samples. α is a hyper-parameter.

Recent studies demonstrate that data augmentation [10, 11] can significantly improve the generalization of supervised learning, especially in image classification and object detection. To extract robust feature, we process label-rich source data S with RandAugment [11], a random augmentation technique, including color, brightness, rotation, sharpness, etc., and output $S' = \{(\mathbf{s}'_i, y_i^s)\}_{i=1}^{N_s}$.

$$\mathcal{L}_s = \frac{1}{N_s} \sum_{(\mathbf{s}'_i, y_i^s) \sim S'} L(\theta(\mathbf{s}'_i), y_i^s) \quad (2)$$

where $L(p, y)$ is the cross-entropy loss over labeled samples.

For each unlabeled target sample, we apply both weak augmentations φ and strong augmentation ψ .

$$\mathbf{u}_i^w = \varphi(\mathbf{u}_i), \quad \mathbf{u}_i^s = \psi(\mathbf{u}_i). \quad (3)$$

We feed \mathbf{u}_i^w and \mathbf{u}_i^s into the model θ and output the predictions \mathbf{p}_i^w and \mathbf{p}_i^s . If the max prediction $\max(\mathbf{p}_i^w)$ is greater than the threshold μ , it indicates that the prediction has strong discrimination and tendency, thus regarded as a pseudo label $\tilde{\mathbf{p}}^w = \operatorname{argmax}(\mathbf{p}_i^w)$, which returns a one-hot label to supervise the strong augmented unlabeled prediction \mathbf{p}_i^s .

$$\mathcal{L}_u = \sum_{\mathbf{u}_i \sim U} [\mathbb{1}(\max(\mathbf{p}^w) \geq \mu) H(\tilde{\mathbf{p}}^w, \mathbf{p}^s)], \quad (4)$$

$H(\cdot, \cdot)$ is the cross-entropy of two possibility distributions. $\mathbb{1}(\cdot)$ is an indicator function that filters and returns the satisfied \mathbf{p}_i^w statistically.

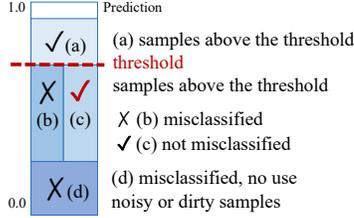


Fig. 3. CLMT rescues samples(c) whose $\max(\mathbf{p}^w)$ below the threshold but not misclassified.

2.2. Stage II: Semi-Supervised Learning

After obtaining the UDA model, we start to adapt it to the target domain with proposed consistency learning based mean teacher (CLMT) model, by taking labeled and unlabeled target samples as input. Following existing SSL methods, we optimize our model with the following learning objective:

$$\mathcal{L}_{ssl} = \gamma \mathcal{L}_t(T_l) + \eta \mathcal{L}_d(U) \quad (5)$$

where $\mathcal{L}_t(T_l)$ is the supervised loss with labeled target samples T_l , and $\mathcal{L}_d(U)$ is the unsupervised loss with unlabeled target samples U .

To fully utilize the discriminative ‘‘Anchors’’, we extract strong augmented features $T_l' = \{(\mathbf{t}'_i, y_i^t)\}_{i=1}^{N_{T_l}}$ and obtain explicit decision boundaries using \mathcal{L}_t . The soft distillation loss \mathcal{L}_d and target supervised loss \mathcal{L}_t together optimize the student model θ_s shown in Figure 2(a).

$$\mathcal{L}_t = \frac{1}{N_t} \sum_{(\mathbf{t}'_i, y_i^t) \sim T_l} L(\theta(\mathbf{t}'_i), y_i^t) \quad (6)$$

In most UDA and SSDA methods [12, 13, 5, 14, 2], as the source domain provides sufficient supervised information, the good model drives the unlabeled target samples to hurdle $\max(\mathbf{p}_i^w)$ above the threshold, the latter ones then join the training as supervisors. However, as shown in Figure 3, a bunch of correctly classified samples are ignored due to the threshold, which is expected to rescue. Inspired by SimCLR v2 [15], we copy the UDA model θ as teacher model θ_t and student model θ_s . We feed the strong augmented view \mathbf{u}^s into student model and the weak augmented view \mathbf{u}^w into the teacher model, and output the prediction \mathbf{p}^s and \mathbf{p}^w , respectively. We assume that the teacher computes relatively precise soft labels for the student model to learn from.

$$\mathcal{L}_d = - \sum_{\mathbf{u}_i \sim U} \theta_t(\mathbf{u}^w) \log \theta_s(\mathbf{u}^s) = - \sum \mathbf{p}^w \log \mathbf{p}^s \quad (7)$$

After one round of rectification, we assume that the θ_s has preciser decision boundaries and can stride across more reliable unlabeled samples. We align the target domain implicitly using pseudo labeling and consistency alignment \mathcal{L}_u again, shown in Figure 2(b).

The teacher model θ_t is then updated by the EMA [16] method, shown in Figure 2(c). θ_{t-1} denotes the previous teacher model and θ_{s-1} represents the previous student model. θ_t indicates the updated teacher model. σ is the EMA hyper-parameter for the teacher model to adapt.

$$\theta_t = \sigma \theta_{t-1} + (1 - \sigma) \theta_{s-1} \quad (8)$$

After this overall pipeline, the student model is barely affected by the source domain. Thanks to the ‘‘Anchor’’ and unlabeled target samples, the model is biased toward the target domain.

3. EXPERIMENTS

3.1. Setups

Datasets. We evaluate the effectiveness of our proposed approach on the latest domain adaptation dataset DomainNet [17]. It is a dataset of common objects, 6 different domains with 126 categories. Similar to MME [2], we select 4 domains, Real(R), Clipart(C), Painting(P), and Sketch(S). We perform experiments and evaluate classification accuracy on 7 pairs, R→C, R→P, P→C, C→S, S→P, R→S, and P→R. For each pair, we evaluate 1-shot SSDA and 3-shot SSDA settings. The labeled target samples are randomly selected.

Implementation details. Similar to previous SSDA work, we choose Alexnet as our backbone networks and the MME method [2] as our baseline. For fair comparisons, experimental settings in our proposed method, i.e., feature extractor, the linear classification layer, optimizer, learning rate, is set the same as MME. We train for 50K iterations. Threshold is fixed $\mu = 0.95$. The hyper-parameters of different losses are $\alpha = 1.0$, $\gamma = 0.2$, $\eta = 0.8$. σ is set as 0.99.

Compared methods. We compare with the following methods including DANN[14], ADR[18], CDAN[19], FAN[20], BiAT[21], MME[2], Meta-MME[22], PAC[23], CLDA[3], CDAC[1], and ECACL [7].

3.2. Comparative Results

Comparing with the existing SSDA methods, we can see from the Table 1 that our method attains 10.2(1-shot) and 9.6(3-shot) point gains over the baseline MME and surpasses the existing best performing three approaches, i.e., CLDA, CDAC and ECACL-P. CLMT achieves the best accuracy in 6 cases, including R→C(1&3-shot), C→S(1-shot), R→S(1&3-shot), P→R(3-shot), shown below. Our method with 6 best pairs significantly exceeds other methods, e.g. R→S(+4.3, +3.7). In contrast, the most competitive method ECACL achieves the best accuracy in only 4 cases, including R→P(3-shot), P→C(1&3-shot), S→P(1-shot), shown below. For the rest 8 cases, we still achieve competitive results.

An interesting observation from Table 1 is that, for different pairs, the gain varies within [0.9, 5.4] when adding 2-shot per-class. We suspect that a tug-of-war happens on the feature extractor to grab the data distribution between the two domains. The learning difficulty and dragging ability are diverse for each domain. The larger amount or better ability, the better performance as a target domain, etc Real. $\frac{N_{T_u}}{N_{T_l}}$ varies between 60:1 and 200:1 and reveal that ‘Anchors’ are tiny, comparing to the enormous(N_{T_u}) unlabeled target samples. It is reasonable to leave the ‘Anchors’ for the second stage, which amplifies the contribution and helps pin down the decision boundary.

Table 1. Results on the *DomainNet* dataset on AlexNet. Best results are in **bold**.

	R→C		R→P		P→C		C→S		S→P		R→S		P→R		Mean	
	1-shot	3-shot														
ST	43.3	47.1	42.4	45.0	40.1	44.9	33.6	36.4	35.7	38.4	29.1	33.3	55.8	58.7	40.0	43.4
DANN	43.3	46.1	41.6	43.8	39.1	41.0	35.9	36.5	36.9	38.9	32.5	33.4	53.6	57.3	40.4	42.4
ADR	43.1	46.2	41.4	44.4	39.3	43.6	32.8	36.4	33.1	38.9	29.1	32.4	55.9	57.3	39.2	42.7
CDAN	46.3	46.8	45.7	45.0	38.3	42.3	27.5	29.5	30.2	33.7	28.8	31.3	56.7	58.7	39.1	41.0
ENT	37.0	45.5	35.6	42.6	26.8	40.4	18.9	31.1	15.1	29.6	18.0	29.6	52.2	60.0	29.1	39.8
MME	48.9	55.6	48.0	49.0	46.7	51.7	36.3	39.4	39.4	43.0	33.3	37.9	56.8	60.7	44.2	48.2
Meta-MME	-	56.4	-	50.2	-	51.9	-	39.6	-	43.7	-	38.7	-	60.7	-	48.7
BiAT	54.2	58.6	49.2	50.6	44.0	52.0	37.7	41.9	39.6	42.1	37.2	42.0	56.9	58.8	45.5	49.4
FAN	47.7	54.6	49.0	50.5	46.9	52.1	38.5	42.6	38.5	42.2	33.8	38.7	57.5	61.4	44.6	48.9
PAC	55.4	61.7	54.6	56.9	47.0	59.8	46.9	52.9	38.6	43.9	38.7	48.2	56.7	59.7	48.3	54.7
CLDA	56.3	59.9	56.0	57.2	50.8	54.6	42.5	47.3	46.8	51.4	38.0	42.7	64.4	67.0	50.7	54.3
CDAC	56.9	61.4	55.9	57.5	51.6	58.9	44.8	50.7	48.1	51.7	44.1	46.7	63.8	66.8	52.1	56.2
ECACL	55.8	62.6	54.0	59.0	56.1	60.5	46.1	50.6	54.6	50.3	45.0	48.4	62.3	67.4	52.8	57.6
CLMT	58.0	63.1	55.7	57.7	55.6	60.3	46.9	52.3	50.7	51.6	49.3	52.1	63.9	68.2	54.3	57.8

Table 2. Comparison with decomposition in SSDA.

Method	1-shot	3-shot
Co-training SSDA	45.0	48.4
Single Stage1 UDA	43.3	43.3
Single Stage2 SSL	24.2	26.3
CLMT	49.3	52.1

Table 3. Ablation study of the UDA stage in 3-shot.

Method	AlexNet	ResNet34
Source	30.5	45.1
+Anchors	33.3	50.1
+Anchors+Aug	37.6	54.8
+FixMatch	45.6	67.2

3.3. Analysis

We conduct all analysis with the adaptation experiment from the *Real* dataset to the *Sketch* dataset.

Comparison with decomposition in SSDA. To verify the effectiveness of decomposition, we compare our method with ECACL [7], shown in Table 2 as “Co-training SSDA”, which includes data augment and consistency alignment, and prototype alignment. “Single Stage1 UDA” means UDA training without “+Anchors”. “Single Stage2 SSL” applies only FixMatch[24]. The comparison shows that our decomposed two-stage approach outperforms the co-training ones, over 4.3%(1-shot) and 3.7%(3-shot). From “Single Stage1 UDA” to CLMT, 6.0% and 8.8% are gained, reconfirming that UDA model still has domain bias and room for improvement.

Comparison with components in UDA. We perform ablation studies on UDA under the 3-shot setting, to analyze the effectiveness of each loss term, including L_s and L_u . As shown in Table 3, “Source” only uses L_s , serving as baseline. “Source” and “+Anchors” have no augmentation, while “+Anchors+Aug” includes augmentation. It shows that despite labels are vacant in the target domain, the consistency alignment method is still fruitful for unlabeled target samples to pre-train a UDA model.

Table 4. Ablation study of the SSL stage.

Method	1-shot	3-shot
UDA	43.3	43.3
+MT	44.8	44.9
+FixMatch	44.7	44.7
+Anchors	47.1	49.0
+Anchors+FixMatch	47.1	49.3
+Anchors+MT	47.4	49.1
CLMT	49.3	52.1

Comparison with semi-supervised methods. As shown in Table 4, “+Anchors” indicates that the pretrained UDA model is then fine-tuned only on 1-shot or 3-shot labeled target samples using L_t . “+MT” and “+FixMatch” indicates L_d Mean Teacher and L_u FixMatch. “+Anchors+MT” and “+Anchors+FixMatch” are combinations of fine-tuning and semi-supervised methods. CLMT exceeds other semi supervised methods in solving domain adaptation challenge.

4. CONCLUSION

In this paper, we solve domain gap in a task decomposition way. We propose a simple but effective method for SSDA which decomposes SSDA into a two-stage approach, i.e., UDA and SSL. We well-pretrain a UDA model in the first stage. In the second stage, the target samples rectify and shift the model further to the target domain, using our method specially designed for the adaptation challenge. We outperforms existing methods on DomainNet, which demonstrates that the method performs well and alleviates domain bias.

5. ACKNOWLEDGEMENT

This research was partly supported by Shenzhen Stable Supporting Program (WDZC20200820200 655001) & the National Key RD Program of China (Grant No. 2020AAA01083 03), and by Shenzhen Science and Technology Project (Grant No. JCYJ20200109143041798) & Shenzhen Key Laboratory of next generation interactive media innovative technology-Grant No. ZDSYS20210623092001004).

6. REFERENCES

- [1] Jichang Li, Guanbin Li, Yemin Shi, and Yizhou Yu, “Cross-domain adaptive clustering for semi-supervised domain adaptation,” in *CVPR*, 2021.
- [2] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko, “Semi-supervised domain adaptation via minimax entropy,” in *ICCV*, 2019.
- [3] Ankit Singh, “Clda: Contrastive learning for semi-supervised domain adaptation,” *NeurIPS*, 2021.
- [4] Luyu Yang, Yan Wang, Mingfei Gao, Abhinav Shrivastava, Kilian Q Weinberger, Wei-Lun Chao, and Ser-Nam Lim, “Deep co-training with task decomposition for semi-supervised domain adaptation,” in *ICCV*, 2021.
- [5] Rui Wang, Zuxuan Wu, Zejia Weng, Jingjing Chen, Guo-Jun Qi, and Yu-Gang Jiang, “Cross-domain contrastive learning for unsupervised domain adaptation,” *IEEE T Multimedia*, 2022.
- [6] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le, “Unsupervised data augmentation for consistency training,” *NeurIPS*, vol. 33, pp. 6256–6268, 2020.
- [7] Kai Li, Chang Liu, Handong Zhao, Yulun Zhang, and Yun Fu, “Ecacl: A holistic framework for semi-supervised domain adaptation,” in *ICCV*, 2021.
- [8] Jian Liang, Dapeng Hu, and Jiashi Feng, “Do we really need to access the source data? Source hypothesis transfer for unsupervised domain adaptation,” in *ICML*, 2020, pp. 6028–6039.
- [9] Shiqi Yang, Joost van de Weijer, Luis Herranz, Shangling Jui, et al., “Exploiting the intrinsic neighborhood structure for source-free domain adaptation,” *NeurIPS*, 2021.
- [10] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le, “Autoaugment: Learning augmentation strategies from data,” in *CVPR*, 2019, pp. 113–123.
- [11] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le, “Randaugment: Practical automated data augmentation with a reduced search space,” in *CVPR workshops*, 2020.
- [12] Kendrick Shen, Robbie M Jones, Ananya Kumar, Sang Michael Xie, Jeff Z HaoChen, Tengyu Ma, and Percy Liang, “Connect, not collapse: Explaining contrastive learning for unsupervised domain adaptation,” in *ICML*, 2022, pp. 19847–19878.
- [13] Yaroslav Ganin and Victor Lempitsky, “Unsupervised domain adaptation by backpropagation,” in *ICML*, 2015.
- [14] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky, “Domain-adversarial training of neural networks,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [15] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton, “Big self-supervised models are strong semi-supervised learners,” *NeurIPS*, 2020.
- [16] Antti Tarvainen and Harri Valpola, “Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results,” *NeurIPS*, 2017.
- [17] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang, “Moment matching for multi-source domain adaptation,” in *ICCV*, 2019.
- [18] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko, “Adversarial dropout regularization,” *arXiv preprint arXiv:1711.01575*, 2017.
- [19] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan, “Conditional adversarial domain adaptation,” *NeurIPS*, 2018.
- [20] Taekyung Kim and Changick Kim, “Attract, perturb, and explore: Learning a feature alignment network for semi-supervised domain adaptation,” in *ECCV*, 2020.
- [21] Pin Jiang, Aming Wu, Yahong Han, Yunfeng Shao, Meiyu Qi, and Bingshuai Li, “Bidirectional adversarial training for semi-supervised domain adaptation,” in *IJCAI*, 2020, pp. 934–940.
- [22] Da Li and Timothy Hospedales, “Online meta-learning for multi-source and semi-supervised domain adaptation,” in *ECCV*, 2020.
- [23] Samarth Mishra, Kate Saenko, and Venkatesh Saligrama, “Surprisingly simple semi-supervised domain adaptation with pretraining and consistency,” *arXiv preprint arXiv:2101.12727*, 2021.
- [24] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li, “Fixmatch: Simplifying semi-supervised learning with consistency and confidence,” *NeurIPS*, 2020.