

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

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

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

Abstract. Although unsupervised domain adaptation methods have been widely adopted across several computer vision tasks, it is more desirable if we can exploit a few labeled data from new domains encountered in a real application. The novel setting of the semi-supervised domain adaptation (SSDA) problem shares the challenges with the domain adaptation problem and the semi-supervised learning problem. However, a recent study shows that conventional domain adaptation and semi-supervised learning methods often result in less effective or negative transfer in the SSDA problem. In order to interpret the observation and address the SSDA problem, in this paper, we raise the intra-domain discrepancy issue within the target domain, which has never been discussed so far. Then, we demonstrate that addressing the intra-domain discrepancy leads to the ultimate goal of the SSDA problem. We propose an SSDA framework that aims to align features via alleviation of the intra-domain discrepancy. Our framework mainly consists of three schemes, i.e., attraction, perturbation, and exploration. First, the attraction scheme globally minimizes the intra-domain discrepancy within the target domain. Second, we demonstrate the incompatibility of the conventional adversarial perturbation methods with SSDA. Then, we present a domain adaptive adversarial perturbation scheme, which perturbs the given target samples in a way that reduces the intra-domain discrepancy. Finally, the exploration scheme locally aligns features in a class-wise manner complementary to the attraction scheme by selectively aligning unlabeled target features complementary to the perturbation scheme. We conduct extensive experiments on domain adaptation benchmark datasets such as DomainNet, Office-Home, and Office. Our method achieves state-of-the-art performances on all datasets.

Keywords: Domain Adaptation · Semi-supervised Learning

1 Introduction

Despite the promising success of deep neural networks in several computer vision tasks, these networks often show performance degradation when tested beyond the training environment. One way to mitigate this problem is to collect

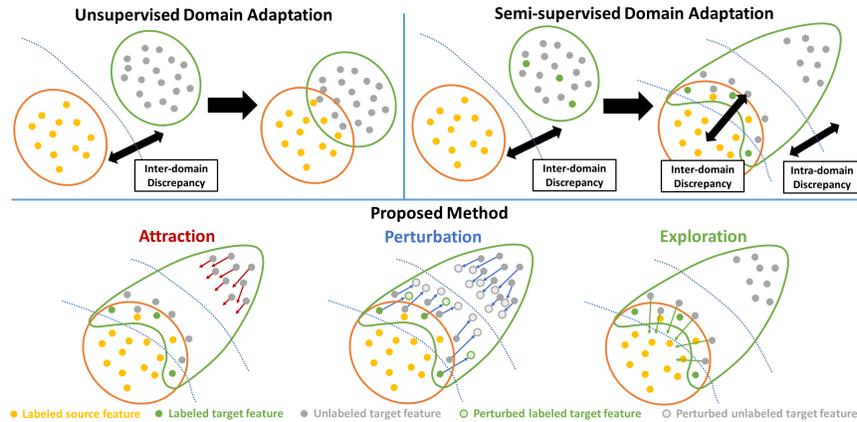


Fig. 1: Conceptual descriptions of the feature alignment approaches. The top row describes the different feature alignment behaviors between the UDA and SSDA problem. Supervision on labeled target samples attracts the corresponding features and their neighborhood toward the source feature cluster, which causes the intra-domain discrepancy. The bottom row describes the proposed attraction, perturbation, and exploration schemes, which are explained in Section 4 in detail.

large amounts of data from the new domain and train the network. Such heavy demands on data annotation cause great interest in domain adaptation and semi-supervised learning on deep neural networks. However, most recent studies on deep domain adaptation are focused on unsupervised approaches, and deep semi-supervised learning is still concentrated on addressing the identical domain problem. Though these methods can be directly applied to the semi-supervised domain adaptation (SSDA) problem only with an additional supervision on the extra labeled samples, a recent study [26] reveals that unsupervised domain adaptation (UDA) methods and semi-supervised learning (SSL) methods often show less effective or even worse performances than just training on the labeled source and target samples in the SSDA problem.

In this paper, we introduce a new concept called *intra-domain discrepancy* to analyze the failure of the UDA and SSL methods and address the SSDA problems. Intra-domain discrepancy is a chronic issue in the SSDA problem that occurs during labeled sample supervision, but has never been discussed so far. In the UDA problem, supervision on the labeled source samples does not critically affect the target domain distribution in general but implicitly attracts some alignable target features similar to the source features. Thus, aligning the source and target domains by reducing their inter-domain discrepancy is reasonable. However, in the SSDA problem, supervision on the labeled target samples enforces the corresponding features and their neighborhood to be attracted toward source feature clusters, which guarantees partial alignment between two domain distributions. Besides, unlabeled target samples that less correlate with the la-

beled target samples are less affected by the supervision and eventually remain unaligned (Top row in Fig. 1). Thus, the target domain distribution is separated into an aligned target subdistribution and an unaligned target subdistribution, causing the intra-domain discrepancy within the target domain. The failure of the UDA and SSL methods will be discussed in Section 3 in detail.

Motivated by the insight, we propose an SSDA framework that aligns cross-domain features by addressing the intra-domain discrepancy within the target domain. Our framework focuses on enhancing the discriminability on the unaligned target samples and modulating the class prototypes, the representative features of each class. It consists of three schemes, i.e., attraction, perturbation, and exploration, as shown in Fig. 1. First, the attraction scheme aligns the unaligned target subdistribution to the aligned target subdistribution through the intra-domain discrepancy minimization. Second, we discuss why conventional adversarial perturbation methods are ineffective in the SSDA problem. Unlike these approaches, our perturbation scheme perturbs target subdistributions into their intermediate region to propagate labels to the unaligned target subdistribution. Note that our perturbation scheme does not ruin the already aligned target features since it additionally generates perturbed features temporarily for regularization. Finally, the exploration scheme locally modulates the prototypes in a class-aware manner complementary to the attraction and perturbation schemes. We perform extensive experiments to evaluate the proposed method on domain adaptation datasets such as DomainNet, Office-Home, Office, and achieved state-of-the-art performances. We also deeply analyze our methods in detail.

Our contributions can be summarized as follows:

- We introduce the intra-domain discrepancy issue within the target domain in the SSDA problem.
- We propose an SSDA framework that addresses the intra-domain discrepancy issues via three schemes, i.e., attraction, perturbation, and exploration.
 - The attraction scheme aligns the unaligned target subdistribution to the aligned target subdistribution through the intra-domain discrepancy minimization.
 - The perturbation scheme perturb target subdistributions into their intermediate region to propagate labels to the unaligned target subdistribution.
 - The exploration scheme locally modulate the prototypes in a class-aware manner complementary to the attraction and perturbation schemes.
- We conduct extensive experiments on DomainNet, Office-Home, and Office. We achieve state-of-the-art performances among various methods, including vanilla deep neural networks, UDA, SSL, and SSDA methods.

2 Related Work

2.1 Unsupervised Domain Adaptation

The recent success of deep learning-based approaches and the following enormous demand for massive amounts of data attract great interest in domain adaptation

(DA). Even in the midst of significant interest, most recent works are focused on unsupervised domain adaptation (UDA). Recent UDA methods can be categorized into three approaches. The first approach is to reduce the cross-domain divergence. This can be achieved by minimizing the estimated domain divergence such as MMD [9] or assimilating feature distributions through adversarial confusion using a domain classifier [7, 17–19]. The second approach is to translate the appearance of one domain into the opposite domain so that the translated data can be regarded as sampled from the opposite domain [13, 12, 10]. The last approach is to consider the source domain as partially labeled data and utilize the semi-supervised learning schemes. For example, Drop-to-Adapt [16] exploits a virtual adversarial perturbation scheme [21]. Recently, these approaches are widely adopted across several computer vision tasks beyond image classification such as object detection [4, 28, 14], semantic segmentation [30, 11], person re-identification [35], and even in depth estimation [34].

2.2 Semi-supervised Learning

Similar to domain adaptation (DA), semi-supervised learning (SSL) has also attracted great attention as a way to overcome the shortages of the labeled data. The difference between DA and SSL is that domain adaptation assumes to deal with data sampled from two distributions with significant domain discrepancy, while SSL assumes to deal with the labeled and unlabeled data sampled from the identical distribution. With the rise of the deep learning approaches, several methods have been recently proposed for deep SSL. Some works add data augmentation and regularize the model by enforcing a consistency between the given and the augmented data [15, 29]. Miyato *et al.* [21] extend this scheme by adversarially searching the bounded and small perturbation which leads the model to the most unstable state. Laine and Aila [15] ensemble the prediction of the model by averaging them throughout the training phase, while Targainen and Valpola [29] ensemble the parameter of the model itself. Other few works use self-training schemes with a memory module or a regularization through convergence speed [3, 5]. Recently, Wang *et al.* [32] propose an augmentation distribution alignment approach to explicitly address the empirical distribution mismatch problem in semi-supervised learning.

2.3 Semi-supervised Domain Adaptation

Semi-supervised domain adaptation (SSDA) is an important task which bridges the well-organized source distribution toward target distribution via partially labeled target samples, while a few works have explored so far [1, 6, 33, 26]. Donahue *et al.* [6] address the domain discrepancy by optimizing the auxiliary constraints on the labeled data. Yao *et al.* [33] learn a subspace that can reduce the data distribution mismatch. Ao *et al.* [1] estimate the soft label of the given labeled target sample with the source model and interpolated with the hard label for target model supervision. Saito *et al.* [26] minimize the distance between the unlabeled target samples and the class prototypes through minimax training on

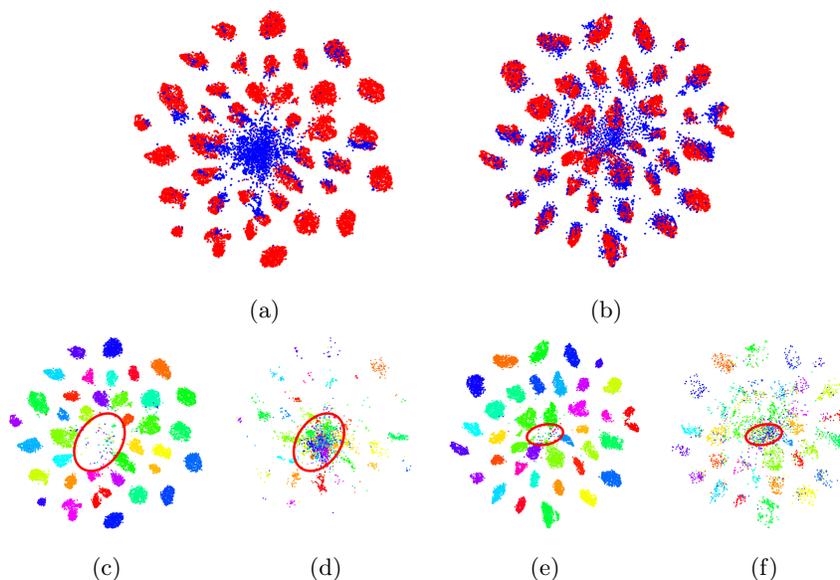


Fig. 2: (a)-(f) The t-SNE visualization of the source and target features in (a) the UDA problem and (b) the SSSDA problem with three target labels for each class. We adopted the Real to Sketch scenario of the DomainNet dataset on the AlexNet backbone and visualized for partial classes. (c) and (e) visualize the source distribution of the UDA and SSSDA problems and (d) and (f) visualize the target distribution of the UDA and SSSDA problems, respectively. Even only three labeled target samples per class can attract their neighborhoods in this degree and separate the target domain into the aligned and unaligned subdistributions.

entropy. However, none of these methods discuss and mitigate the intra-domain discrepancy issue in the SSSDA problem. Different from previous works, we address the SSSDA problem with a new perspective of the intra-domain discrepancy.

3 Intra-domain Discrepancy

Intra-domain discrepancy of a domain is an internal distribution gap among subdistributions within the domain. Though we demonstrate the intra-domain discrepancy issue in the semi-supervised domain adaptation (SSDA) problem, such subdistributions can also appear in the unsupervised domain adaptation (UDA) problem since there usually exist target samples alignable to the source clusters. However, since each domain generally has a unique correlation among domain samples, the target domain distribution is not easily separated into distinctive subdistributions, which eventually causes the insufficient intra-domain discrepancy. Thus, the conventional inter-domain discrepancy minimization approaches have been effectively applied to the UDA problem. In contrast, in the

SSDA problem, supervision on the labeled target samples enforces the target domain to be separated into the aligned subdistribution and the unaligned subdistribution deterministically. More specifically, as shown in the top row of Fig. 1, the presence of the label pulls the target samples and its neighborhoods toward source feature clusters of each corresponding labels. Besides, the unlabeled target samples which less correlate with the given labeled target samples are still located distant from source feature clusters, producing inaccurate and even meaningless inference results. Figure 2 demonstrates the existence of the intra-domain discrepancy within the target domain. Though only three target labels per class are given, significant number of target samples are aligned while wrongly predicted target samples (red circle in Fig. 2 (f)) are still located far from the source domain.

The presence of the intra-domain discrepancy makes the conventional domain adaptation methods less suitable for the SSDA problem. The ultimate goal of domain adaptation is to enhance the discriminability on the target domain, and most of the error occurs on the unaligned target subdistribution in this case. Thus, solving SSDA problems depends on how far the unaligned subdistribution is aligned. However, common domain adaptation methods focus on reducing the inter-domain discrepancy between the source and target domains regardless of the intra-domain discrepancy within the target domain. Since the existence of the aligned target subdistribution cause underestimation of the inter-domain discrepancy, the inter-domain discrepancy reduction approaches work less effectively in the SSDA problems. Moreover, since the aligned target subdistribution is aligned in a class-aware manner, such approaches can even negatively affect.

Similarly, conventional semi-supervised learning (SSL) methods also suffer from the intra-domain discrepancy issue in the SSDA problem. It stems from the different assumptions between the SSDA and SSL problems. Since the SSL problem assumes to sample labeled and unlabeled data from the identical distribution, SSL methods mainly focus on propagating the correct labels to their neighbors. In contrast, SSDA problems assume that there is a significant distribution divergence between the source and target domains, and that labeled samples are dominated by the source domain. Since correctly predicted target samples are highly aligned with the source distribution, whereas incorrectly predicted target samples are located far from them, we can no longer assume that these target samples share the same distribution. Thus, the SSL methods only propagate errors within the wrongly predicted subdistribution, and the propagation is also meaningless in the correctly predicted subdistribution due to the rich distribution of the source domain. Motivated by the interpretation, we propose a framework that addresses the intra-domain discrepancy.

4 Method

4.1 Problem Formulation

Let us denote the set of source domain samples by $\mathcal{D}_s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{m_s}$. For the target domain, $\mathcal{D}_t = \{(\mathbf{x}_i^t, y_i^t)\}_{i=1}^{m_t}$ and $\mathcal{D}_u = \{(\mathbf{x}_i^u)\}_{i=1}^{m_u}$ denote the sets of

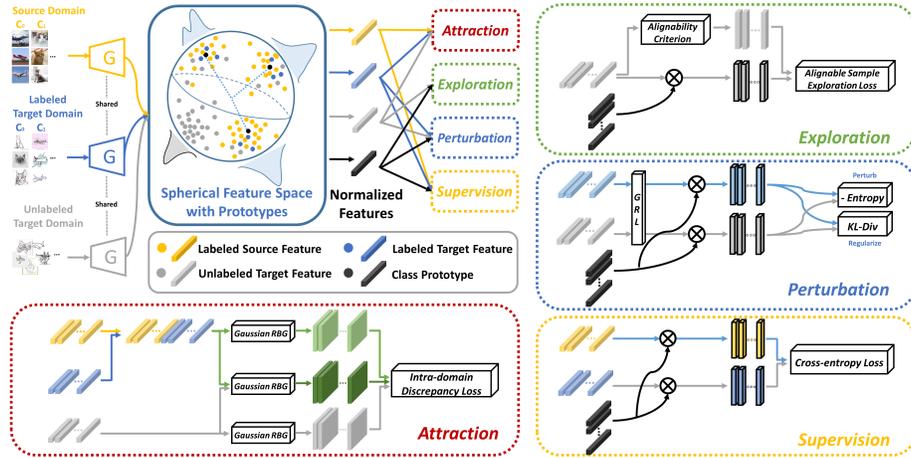


Fig. 3: An overall framework of the proposed method. Our framework consists of the feature extractor, trainable class prototypes, supervision module, and each module for the proposed schemes. The class prototypes and all normalized features of the samples are embedded in the same spherical feature space.

labeled and unlabeled target samples, respectively. SSDA aims to enhance the target domain discriminability through training on \mathcal{D}_s , \mathcal{D}_t , and \mathcal{D}_u .

4.2 Spherical Feature Space with Prototypes

When aligning feature distributions, it is crucial to determine which feature space to adapt. Even if the same method is used, performance may not be improved depending on the feature space applied. Thus, we adopt similarity-based prototypical classifier in [2] to prepare suitable feature space for better adaptation. Briefly, the prototypical classifier inputs normalized feature and compare the similarities among all class-wise prototypes, which reduces intra-class variations as results. For the classifier training, we use a cross-entropy loss as our classification loss to train an embedding function $f_\theta(\cdot)$ with parameters θ and the prototypes \mathbf{p}_k ($k = 1, \dots, K$) on the source domain samples and the labeled target samples:

$$\begin{aligned} \mathcal{L}_{cls} &= \mathbb{E}_{(\mathbf{x}, y) \in \mathcal{D}_s \cup \mathcal{D}_t} [-\log p(y|\mathbf{x}, \mathbf{p})] \\ &= \mathbb{E}_{(\mathbf{x}, y) \in \mathcal{D}_s \cup \mathcal{D}_t} \left[-\log \left(\frac{\exp(\mathbf{p}_y \cdot f_\theta(\mathbf{x})/T)}{\sum_{i=1}^K \exp(\mathbf{p}_i \cdot f_\theta(\mathbf{x})/T)} \right) \right]. \end{aligned} \quad (1)$$

While the prototypical classifier is trying to reduce the intra-class variation of the labeled sample features, the proposed schemes focus on aligning distributions of the normalized features on the spherical feature space.

4.3 Attraction Scheme

The attraction scheme aims to globally align the unaligned target subdistribution to the aligned target subdistribution in a subdistribution-level through the estimated intra-domain discrepancy minimization. The scheme measures the feature distribution divergence between the target subdistributions to estimate the intra-domain discrepancy within the target domain. However, the limited number of the labeled target samples would not be sufficient to represent the feature distribution of the aligned target subdistribution. Thus, motivated by the observation that the features of the aligned target subdistribution are highly aligned with that of the source domain in a class-aware manner, we instead use the complex distribution of the labeled source and target data. For the empirical estimation of the intra-domain discrepancy, we adopt Maximum Mean Discrepancy (MMD) [9], a kernel two-sample test that measures the distribution difference. We exploit a mixture $k(\cdot, \cdot)$ of Gaussian Radial Basis Function (RBF) kernels with multiple kernel widths σ_i ($i=1, \dots, N$). Thus, the estimated intra-domain discrepancy on the spherical feature space can be written as:

$$\begin{aligned} d(\mathcal{D}_s \cup \mathcal{D}_t, \mathcal{D}_u) &= \mathbb{E}_{(\mathbf{x}, y), (\mathbf{x}', y') \in \mathcal{D}_s \cup \mathcal{D}_t} [k(f_\theta(\mathbf{x}), f_\theta(\mathbf{x}'))] \\ &+ \mathbb{E}_{(\mathbf{z}, w), (\mathbf{z}', w') \in \mathcal{D}_u} [k(f_\theta(\mathbf{z}), f_\theta(\mathbf{z}'))] \\ &- 2\mathbb{E}_{(\mathbf{x}, y) \in \mathcal{D}_s \cup \mathcal{D}_t, (\mathbf{z}, w) \in \mathcal{D}_u} [k(f_\theta(\mathbf{x}), f_\theta(\mathbf{z}))], \end{aligned} \quad (2)$$

where \mathbf{x}' , \mathbf{z} , and \mathbf{z}' represent samples, y' , w , and w' represent the corresponding labels. Since our attraction scheme directly minimizes the intra-domain discrepancy, the attraction loss can be written by:

$$L_a = d(\mathcal{D}_s \cup \mathcal{D}_t, \mathcal{D}_u). \quad (3)$$

4.4 Perturbation scheme

Conventional adversarial perturbation, one of the semi-supervised learning (SSL) approaches, turns out to be ineffective or even cause negative transfer in the SSDA problem. In the same context as discussed in Section 3, the labeled target samples and its neighborhoods are aligned to the source domain separated from the inaccurate target samples, causing the intra-domain discrepancy. Then, the aligned features already guaranteed its confidence by rich information of the source domain, while the unaligned features can only propagate inaccurate predictions. Thus, the perturbation on both the aligned and unaligned target subdistribution are less meaningful.

Unlike the common adversarial perturbation approaches, our scheme perturbs target subdistributions toward their intermediate region for 1) accurate prediction propagation from the aligned subdistribution to the unaligned subdistribution and 2) class prototypes modulation toward the region. Such perturbation can be achieved by searching the direction of the anisotropically high entropy of the target features since the element-wise entropy increases as the feature move away from the prototype while the feature far from the prototypes

can be attracted toward the prototypes. Note that the perturbation scheme does not ruin the already aligned subdistribution since it temporally generates additional perturbed features of the aligned feature for regularization. To achieve this, we first perturb the class prototypes in an entropy maximization direction. Then, we optimize a small and bounded perturbation toward the perturbed prototypes. Finally, we regularize the perturbed data and the given data through Kullback–Leibler divergence. To summarize, the perturbation loss can be formulated as follows:

$$\begin{aligned}
H_{\mathbf{p}}(\mathbf{x}) &= - \sum_{i=1}^K p(y = i|\mathbf{x}) \log p(y = i|\mathbf{x}, \mathbf{p}) \\
r_{\mathbf{x}} &= \operatorname{argmin}_{\|r\| < \epsilon} \max_{\mathbf{p}} H_{\mathbf{p}}(\mathbf{x} + r) \\
\mathcal{L}_{\mathbf{p}} &= \mathbb{E}_{\mathbf{x} \in \mathcal{D}_u} \left[\sum_{i=1}^K D_{KL}[p(y = i|\mathbf{x}, \mathbf{p}), p(y = i|\mathbf{x} + \mathbf{r}_{\mathbf{x}}, \mathbf{p})] \right] \\
&\quad + \mathbb{E}_{(\mathbf{z}, w) \in \mathcal{D}_t} \left[\sum_{i=1}^K D_{KL}[p(y = i|\mathbf{z}, \mathbf{p}), p(y = i|\mathbf{z} + \mathbf{r}_{\mathbf{z}}, \mathbf{p})] \right].
\end{aligned} \tag{4}$$

where $H_{\mathbf{p}}(\cdot)$ is an element-wise entropy function defined upon similarities between the given feature and the prototypes, \mathbf{x} and \mathbf{z} represent samples, and y represent the corresponding label.

4.5 Exploration scheme

The exploration scheme aims to locally modulate the prototypes in a class-aware manner complementary to the attraction scheme, while selectively aligns the unlabeled target features via suitable criteria complementary to the perturbation scheme. Though the attraction scheme globally aligns the target subdistributions on the feature space regardless of the prototypes, it does not explicitly enforce the prototypes to be modulated, which can be complemented by local and class-aware alignment. On the other hand, since the perturbation scheme regularizes the perturbed features of the anisotropically high entropy, the entropy of the perturbed feature and its neighborhood gradually became low. The exploration scheme aligns these features so that their entropy became isotropic, and thus the aligned features can be perturbed farther toward the unaligned subdistribution. To practically achieve this, we selectively collect unlabeled target data with its element-wise entropy less than a certain threshold, then apply a cross-entropy loss with the class of the nearest prototype. The objective function of the exploration scheme can be written as follows:

$$\begin{aligned}
M_{\epsilon} &= \{\mathbf{x} \in \mathcal{D}_u | H_{\mathbf{p}}(\mathbf{x}) < \epsilon\} \\
\hat{y}_{\mathbf{x}} &= \operatorname{argmax}_{i \in \{1, \dots, K\}} p(y = i|\mathbf{x}, \mathbf{p}) \\
\mathcal{L}_e &= \mathbb{E}_{\mathcal{D}_u} [-\mathbf{1}_{M_{\epsilon}}(\mathbf{x}) \log p(y = \hat{y}_{\mathbf{x}}|\mathbf{x}, \mathbf{p})].
\end{aligned} \tag{5}$$

Table 1: Classification accuracy (%) on the DomainNet dataset on the AlexNet and ResNet-34 backbone networks. The performance comparisons were done for seven scenarios with one or three labeled target samples for each class.

Net	Method	R to C		R to P		P to C		C to S		S to P		R to S		P to R		MEAN	
		1-shot	3-shot														
AlexNet	S+T	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
	SagNet	45.8	49.1	45.6	46.7	42.7	46.3	36.1	39.4	37.1	39.8	34.2	37.5	54.0	57.0	42.2	45.1
	Ours	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
ResNet	S+T	55.6	60.0	60.6	62.2	56.8	59.4	50.8	55.0	56.0	59.5	46.3	50.1	71.8	73.9	56.9	60.0
	DANN	58.2	59.8	61.4	62.8	56.3	59.6	52.8	55.4	57.4	59.9	52.2	54.9	70.3	72.2	58.4	60.7
	ADR	57.1	60.7	61.3	61.9	57.0	60.7	51.0	54.4	56.0	59.9	49.0	51.1	72.0	74.2	57.6	60.4
	CDAN	65.0	69.0	64.9	67.3	63.7	68.4	53.1	57.8	63.4	65.3	54.5	59.0	73.2	78.5	62.5	66.5
	ENT	65.2	71.0	65.9	69.2	65.4	71.1	54.6	60.0	59.7	62.1	52.1	61.1	75.0	78.6	62.6	67.6
	MME	70.0	72.2	67.7	69.7	69.0	71.7	56.3	61.8	64.8	66.8	61.0	61.9	76.1	78.5	66.4	68.9
	SagNet	59.4	62.0	61.9	62.9	59.1	61.5	54.0	57.1	56.6	59.0	49.7	54.4	72.2	73.4	59.0	61.5
	Ours	70.4	76.6	70.8	72.1	72.9	76.7	56.7	63.1	64.5	66.1	63.0	67.8	76.6	79.4	67.6	71.7

where M_ϵ is a set of unlabeled target data with entropy value less than a hyperparameter ϵ , and $\mathbf{1}_{M_\epsilon}(\cdot)$ is an indicator function that filters out alignable samples from the given unlabeled target samples.

4.6 Overall framework and training objective

The overall training objective of our method is the weighted sum of the supervision loss, the attraction loss, the perturbation loss, and the exploration loss. The optimization problem can be formulated as follows:

$$\min_{\mathbf{p}, \theta} \mathcal{L}_{cls} + \alpha \mathcal{L}_a + \beta \mathcal{L}_e + \gamma \mathcal{L}_p. \quad (6)$$

We integrated all the schemes into one framework, as shown in the Fig. 3.

5 Experiments

5.1 Experimental Setup

Datasets. **DomainNet** [24] is a recently released large-scale domain adaptation benchmark dataset that contains six domains and approximately 0.6 million images with 345 classes. **Office-Home** [31] and **Office** [25] are standard benchmarks for domain adaptation. Office-Home consists of Art, Clipart, Product, and Real-world domain with 65 classes. Office consists of Amazon, Webcam, and DSLR domains with 31 classes.

Evaluation tasks. For a fair comparison with the state-of-the-art SSDA method [26], we performed experiments on 7 adaptation scenarios on the four domains

Table 2: Classification accuracy (%) on the Office-Home dataset with the AlexNet and ResNet-34 backbone networks. The performance comparisons were done for a total of 12 scenarios on three-shot setting.

Net	Method	R to C	R to P	R to A	P to R	P to C	P to A	A to P	A to C	A to R	C to R	C to A	C to P	MEAN
AlexNet	S+T	44.6	66.7	47.7	57.8	44.4	36.1	57.6	38.8	57.0	54.3	37.5	57.9	50.0
	DANN	47.2	66.7	46.6	58.1	44.4	36.1	57.2	39.8	56.6	54.3	38.6	57.9	50.3
	ADR	45.0	66.2	46.9	57.3	38.9	36.3	57.5	40.0	57.8	53.4	37.3	57.7	49.5
	CDAN	41.8	69.9	43.2	53.6	35.8	32.0	56.3	34.5	53.5	49.3	27.9	56.2	46.2
	ENT	44.9	70.4	47.1	60.3	41.2	34.6	60.7	37.8	60.5	58.0	31.8	63.4	50.9
	MME	51.2	73.0	50.3	61.6	47.2	40.7	63.9	43.8	61.4	59.9	44.7	64.7	55.2
	Ours	51.9	74.6	51.2	61.6	47.9	42.1	65.5	44.5	60.9	58.1	44.3	64.8	55.6
ResNet	S+T	55.7	80.8	67.8	73.1	53.8	63.5	73.1	54.0	74.2	68.3	57.6	72.3	66.2
	DANN	57.3	75.5	65.2	69.2	51.8	56.6	68.3	54.7	73.8	67.1	55.1	67.5	63.5
	CDAN	61.4	80.7	67.1	76.8	58.1	61.4	74.1	59.2	74.1	70.7	60.5	74.5	68.2
	ENT	62.6	85.7	70.2	79.9	60.5	63.9	79.5	61.3	79.1	76.4	64.7	79.1	71.9
	MME	64.6	85.5	71.3	80.1	64.6	65.5	79.0	63.6	79.7	76.6	67.2	79.3	73.1
	Ours	66.4	86.2	73.4	82.0	65.2	66.1	81.1	63.9	80.2	76.8	66.6	79.9	74.0

(Real, Clipart, Painting, Sketch) with 126 classes for DomainNet, 12 adaptation scenarios on all the domains for Office-Home, and two challenging adaptation scenarios on Office. One or three labeled target samples are given for each class for these scenarios. Additionally, we compared the performances for 5, 10 and 20 labeled target samples for each class.

Implementation details. We adopted AlexNet and ResNet-34 for the backbone network. Every mini-batch consists of the same number of labeled source and target samples with a doubled number of unlabeled target samples. We prepared 32 and 24 samples for each split of the mini-batch for AlexNet and Resnet-34, respectively. We used the Stochastic Gradient Descent (SGD) optimizer with an initial learning rate of 0.01, a momentum of 0.9, and a weight decay of 0.0005. All implementations were done in PyTorch [23] and on a single GeForce Titan XP GPU.

Baselines. We compared our method with the semi-supervised domain adaptation (SSDA), unsupervised domain adaptation (UDA), semi-supervised learning (SSL), and no adaptation methods. More specifically, the baselines consist of MME [26], **SagNet** [22], **DANN** [7], **ADR** [27], **CDAN** [18], **ENT** [8], and non-adapted model. For the UDA methods (DANN, ADR, and CDAN), the labeled target samples were supervised during the training process. S+T is a vanilla model trained on all labeled samples. DANN confuses the cross-domain distributions through adversarial learning. ADR adopts the dropout scheme to modify the decision boundary for feature alignment. CDAN adversarially aligns the feature by fooling the conditional domain discriminator. ENT is an SSL method that minimizes the entropy of the unlabeled target data. For the fair comparison, all the methods have the same backbone architecture with our method.

Table 3: Classification accuracy (%) on the Office dataset with three-shot setting.

Net	Method	W to A	D to A	MEAN
AlexNet	S+T	61.2	62.4	61.8
	DANN	64.4	65.2	64.8
	ADR	61.2	61.4	61.3
	CDAN	60.3	61.4	60.8
	ENT	64.0	66.2	65.1
	MME	67.3	67.8	67.6
	Ours	67.6	69.0	68.3

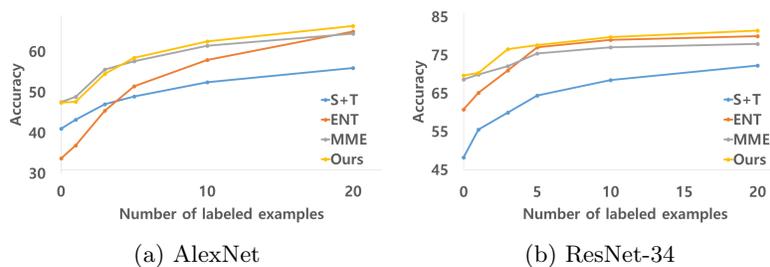


Fig. 4: Trend in classification accuracy (%) with varying number of labeled target samples per class. The experiments are conducted on the Real to Clipart scenario of the DomainNet dataset.

5.2 Results

Performance Comparison on DomainNet. We summarized the classification accuracies of 7 scenarios on the DomainNet dataset in Table 1. On average, our method outperformed the best-performed baseline by 2.8% in the three-shot setting and 1.2% in the one-shot setting on ResNet-34, and by 0.7% in the three-shot setting and 0.4% in the one-shot setting on AlexNet. Moreover, our method outperformed most of the cases except for a few adaptation tasks. On the other hand, though UDA methods like DANN and ADR performed slightly better than S+T when only one labeled target per class is given, these methods become less effective or even cause negative transfer as the number of the labeled target samples increases. It verifies our statement that conventional domain adaptation methods are often less beneficial than the partial alignment effect from the given target labels. ENT showed significant improvement on ResNet-34, while it shows degenerative performance on AlexNet. Moreover, the performance enhancement gap increased as the number of labeled target samples increases, which will be discussed in more detail in Section 5.3.

Performance Comparison on Office-Home and Office. The comparison results of our method with the baselines on the Office-Home dataset are reported in Table 2. Our method outperformed all the baselines regardless of the backbone network on average. Similar to DomainNet, our method showed the best

Table 4: Ablation study results of the proposed schemes on the Real to Sketch task of DomainNet with three-shot setting.

Net	Method	Attract	Explore	Perturb	R to C	R to P	P to C	C to S	S to P	R to S	P to R	MEAN
AlexNet	S+T				47.1	45.0	44.9	36.4	38.4	33.3	58.7	43.4
	DANN				46.1	43.8	41.0	36.5	38.9	33.4	57.3	42.4
	MMD				47.9	45.5	44.6	38.1	38.4	35.5	56.6	43.8
	VAT				46.1	43.8	44.3	35.6	38.2	31.8	57.7	42.5
		✓			50.2	46.2	47.5	40.8	41.3	37.2	59.8	46.1
		✓	✓		53.9	49.8	50.5	42.0	41.9	38.0	60.7	48.3
	Ours			✓	57.2	47.5	54.1	38.8	39.7	38.5	59.2	47.9
		✓	✓	✓	54.6	50.5	52.1	42.6	42.2	38.7	61.4	48.9

performance in most of the scenarios. While DANN performed at least similar to S+T on the AlexNet backbone, it showed degenerative performance on the ResNet-34 backbone. It demonstrates that capacity difference of the backbone network causes the difference in the degree of the target label exploitation, and DANN performed less effective than the exploitation of ResNet-34. The considerable results of ENT are reasonable since the three-shot setting provides approximately 5 ~ 10% of the target labels for training, and such ratio is quite rich in a perspective of the SSL problem. Table 3 showed the performance comparison on the Office dataset and our method also outperformed other baseline on the dataset.

5.3 Analysis

Performance comparison with varying number of target labels. We compared the behavior of the methods by varying the number of labeled target samples from 0 to 20 for each class. As shown in Fig. 4, our methods showed superior performance for a large number of target labels even on the scenario where our method worked less effectively on a one-shot or three-shot setting. Moreover, it outperformed the other baselines throughout all the cases on the ResNet-34 backbone. On the other hand, ENT also significantly enhanced the accuracy for a large number of target labels, and it even outperformed the state-of-the-art SSDA methods when more than twenty and five target labels are given per class for the AlexNet and ResNet-34 backbone networks, respectively. It is reasonable since the increase of the labeled target sample ratio assimilates the SSDA problem to the SSL problem, which is suitable for SSL methods.

Ablation Study on the proposed schemes. We conducted an ablation study on our schemes. To verify the effectiveness, we additionally evaluated DANN, MMD, and VAT [21] on the DomainNet dataset. As shown in Table 4, DANN and MMD rarely worked or even caused negative transfer while our attraction scheme showed meaningful improvement on average. It verifies that conventional UDA methods that focus on reducing the inter-domain discrepancy suffer from the intra-domain discrepancy issue, and it can be addressed by the intra-domain

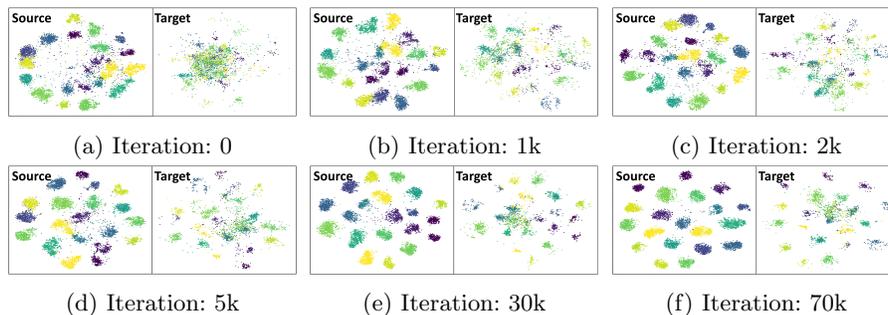


Fig. 5: (a)-(f) The t-SNE visualization of the feature alignment progress through our method during the training phase.

discrepancy minimization. Moreover, VAT also caused degenerative effect while our perturbation scheme significantly enhanced the performance, which demonstrates that conventional adversarial perturbation methods are not suitable for the SSDA problem, and our perturbation scheme can address it by modulating the perturbation direction toward the intermediate region of the target subdistributions. The exploration scheme also worked complementary to other schemes.

Convergence Analysis. To analyze the convergence of our method, we depicted the t-SNE visualization [20] of the cross-domain features over the training progress in Fig. 5. We conducted the experiment on the Real to Sketch scenario of the DomainNet. All 126 classes were used for the experiment, but we choose 20 classes for better visualization. Note that we did not specifically pick classes of top-20 accuracies. Figure 5 (a) clearly shows the initial domain divergence between the source and the target domain. Moreover, the feature depiction of the early stages often showed many unaligned source and target clusters. As the training goes on, our method aligned the corresponding source and target clusters and finally obtained well-accumulated target clusters, as shown in Fig. 5 (f).

6 Conclusions

In this work, we demonstrated the intra-domain discrepancy issue of the target domain in the SSDA problem. Motivated by this, we proposed an SSDA framework that aligns the cross-domain feature distributions by addressing the intra-domain discrepancy through the attraction, exploration, and perturbation schemes. The attraction scheme directly minimized the estimated intra-domain discrepancy within the target domain. The perturbation scheme perturbed the well-aligned and unaligned target features into the intermediate region of the target subdistributions. The exploration scheme locally aligned features in a selective and class-wise manner complementary to the attraction and perturbation schemes. The experiments conducted on DomainNet, Office-Home, and Office datasets validate the effectiveness of our method, and it outperformed the conventional UDA and SSL methods on all the datasets.

References

1. Ao, S., Li, X., Ling, C.X.: Fast generalized distillation for semi-supervised domain adaptation. In: AAAI (2017)
2. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C., Huang, J.B.: A closer look at few-shot classification. In: International Conference on Learning Representations (ICLR) (2019)
3. Chen, Y., Zhu, X., Gong, S.: Semi-supervised deep learning with memory. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018)
4. Chen, Y., Li, W., Sakaridis, C., Dai, D., Van Gool, L.: Domain adaptive faster r-cnn for object detection in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
5. Cicek, S., Fawzi, A., Soatto, S.: SaaS: Speed as a supervisor for semi-supervised learning. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018)
6. Donahue, J., Hoffman, J., Rodner, E., Saenko, K., Darrell, T.: Semi-supervised domain adaptation with instance constraints. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2013)
7. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: Proceedings of the International Conference on Machine Learning (ICML) (2015)
8. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. In: Advances in Neural Information Processing Systems (NeurIPS) (2005)
9. Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A.: A kernel two-sample test. *Journal of Machine Learning Research (JMLR)* (2012)
10. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A.A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: Proceedings of the International conference on machine learning (ICML) (2018)
11. Hong, W., Wang, Z., Yang, M., Yuan, J.: Conditional generative adversarial network for structured domain adaptation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
12. Hu, L., Kan, M., Shan, S., Chen, X.: Duplex generative adversarial network for unsupervised domain adaptation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
13. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
14. Kim, T., Jeong, M., Kim, S., Choi, S., Kim, C.: Diversify and match: A domain adaptive representation learning paradigm for object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
15. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. *International Conference on Learning Representations (ICLR)* (2017)
16. Lee, S., Kim, D., Kim, N., Jeong, S.G.: Drop to adapt: Learning discriminative features for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
17. Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: Proceedings of the International Conference on Machine Learning (ICML) (2015)

18. Long, M., Cao, Z., Wang, J., Jordan, M.I.: Conditional adversarial domain adaptation. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2018)
19. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Unsupervised domain adaptation with residual transfer networks. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2016)
20. van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *Journal of Machine Learning Research (JMLR)* (2008)
21. Miyato, T., Maeda, S.i., Ishii, S., Koyama, M.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (2018)
22. Nam, H., Lee, H., Park, J., Yoon, W., Yoo, D.: Reducing domain gap via style-agnostic networks (2019)
23. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
24. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2019)
25. Saenko, K., Kulis, B., Fritz, M., Darrell, T.: Adapting visual category models to new domains. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2010)
26. Saito, K., Kim, D., Sclaroff, S., Darrell, T., Saenko, K.: Semi-supervised domain adaptation via minimax entropy. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2019)
27. Saito, K., Ushiku, Y., Harada, T., Saenko, K.: Adversarial dropout regularization. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2018)
28. Saito, K., Ushiku, Y., Harada, T., Saenko, K.: Strong-weak distribution alignment for adaptive object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
29. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: *Advances in neural information processing systems (NeurIPS)* (2017)
30. Tsai, Y.H., Hung, W.C., Schuler, S., Sohn, K., Yang, M.H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
31. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
32. Wang, Q., Li, W., Gool, L.V.: Semi-supervised learning by augmented distribution alignment. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2019)
33. Yao, T., Yingwei Pan, Ngo, C., Houqiang Li, Tao Mei: Semi-supervised domain adaptation with subspace learning for visual recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)
34. Zheng, C., Cham, T.J., Cai, J.: T2Net: Synthetic-to-realistic translation for solving single-image depth estimation tasks. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2018)
35. Zheng, Z., Yang, X., Yu, Z., Zheng, L., Yang, Y., Kautz, J.: Joint discriminative and generative learning for person re-identification. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)