

Adversarial Domain Adaptation With Prototype-Based Normalized Output Conditioner

Dapeng Hu¹, Jian Liang², Qibin Hou³, Hanshu Yan¹, *Graduate Student Member, IEEE*, and Yunpeng Chen

Abstract—Domain adversarial training has become a prevailing and effective paradigm for unsupervised domain adaptation (UDA). To successfully align the multi-modal data structures across domains, the following works exploit discriminative information in the adversarial training process, e.g., using multiple class-wise discriminators and involving conditional information in the input or output of the domain discriminator. However, these methods either require non-trivial model designs or are inefficient for UDA tasks. In this work, we attempt to address this dilemma by devising simple and compact conditional domain adversarial training methods. We first revisit the simple concatenation conditioning strategy where features are concatenated with output predictions as the input of the discriminator. We find the concatenation strategy suffers from the weak conditioning strength. We further demonstrate that enlarging the norm of concatenated predictions can effectively energize the conditional domain alignment. Thus we improve concatenation conditioning by normalizing the output predictions to have the same norm of features, and term the derived method as Normalized Output conditioner (NOUN). However, conditioning on raw output predictions for domain alignment, NOUN suffers from inaccurate predictions of the target domain. To this end, we propose to condition the cross-domain feature alignment in the prototype space rather than in the output space. Combining the novel prototype-based conditioning with NOUN, we term the enhanced method as PROTOTYPE-based Normalized Output conditioner (PRONOUN). Experiments on both object recognition and semantic segmentation show that NOUN can effectively align the multi-modal structures across domains and even outperform state-of-the-art domain adversarial training methods. Together with prototype-based conditioning, PRONOUN further improves the adaptation performance over NOUN on multiple object recognition benchmarks for UDA. Code is available at <https://github.com/tim-learn/NOUN>.

Index Terms—Domain adaptation, adversarial learning, prototype, semantic structures, pseudo-labels.

I. INTRODUCTION

DEEP neural networks (DNNs) have achieved impressive success in many fields, relying on large-scale labeled datasets. However, DNNs models have difficulty generalizing

Manuscript received February 8, 2021; revised September 8, 2021; accepted October 11, 2021. Date of publication November 10, 2021; date of current version November 16, 2021. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Andrea Fusiello. (Corresponding author: Jian Liang.)

Dapeng Hu and Hanshu Yan are with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore 119077.

Jian Liang is with the Center for Research on Intelligent Perception and Computing (CRIPAC), National Laboratory of Pattern Recognition (NLPR), Chinese Academy of Sciences (CASIA), Beijing 100190, China (e-mail: liangjian92@gmail.com).

Qibin Hou is with the Department of Computer Science, Nankai University Jinnan Campus, Tianjin 300350, China.

Yunpeng Chen is with Meitu Inc., Beijing 100083, China.
Digital Object Identifier 10.1109/TIP.2021.3124674

to new data with a large distribution shift. As a result, given an unlabeled dataset, the typical learning paradigm of DNNs is to perform laborious manual labeling work for the supervised training. In contrast, unsupervised domain adaptation (UDA) provides an efficient and effective approach to improving the performance of DNNs on unlabeled target data. UDA aims to leverage the knowledge of a labeled dataset (source domain) to help train a predictive model for another unlabeled dataset (target domain). Recently, deep UDA methods have significantly improved the performance of many tasks, including image classification [1], [2] and semantic segmentation [3]–[7], by exploiting supervisions from heterogeneous sources.

The closed-set UDA problem discussed in this paper assumes the shared label space but different data distributions across two domains. Previous methods exploit maximum mean discrepancy (MMD) [1], [8] or other distribution statistics like central moments [9]–[11] for domain adaptation. Recently, adversarial learning [12] provides a promising alternative solution by introducing an extra binary domain discriminator. The task model in domain adversarial training aims to learn domain-invariant features through promoting domain confusion. However, the confusion of the discriminator does not always guarantee the discriminative alignment of multi-class distributions between domains. The probable case is that the multi-modal structures are misaligned between the source and target domains.

Following adversarial methods [13]–[17] are proposed to pursue discriminative domain alignment. However, concerning these methods, [13], [15]–[17] require complex designs on either model architecture or objective function, and [14] involves the inefficient Cartesian product between features and predictions. In contrast, in this work, we pursue simple and compact conditional domain adversarial training methods. We first revisit a simple conditional domain adversarial training strategy, i.e., the concatenation conditioning strategy [14], where cross-domain feature alignment is conditioned on output predictions by concatenating features with the output of the model. As demonstrated in [14], the concatenation conditioning strategy is heavily inferior to the informative multi-linear conditioning strategy. Similarly, we also have such a consistent observation across various object recognition UDA tasks.

[14] ascribes the failure of the naive concatenation conditioning to the independence between the concatenated features and predictions. Instead, we argue that the failure is because the weak conditioning strength renders conditional domain alignment ineffective. We observe that the Euclidean norm of the feature is usually tens of times that of the concatenated

output prediction. The effect is that the contribution of concatenated predictions to the domain confusion objective is heavily ignored. Thus the discriminator would be dominantly optimized with the domain confusion loss between marginal features of two domains, making the conditional domain adversarial training ineffective. As expected, we observe obvious performance improvement over the naive concatenation conditioning by enlarging the norm of concatenated output predictions. Though effective, there is no guarantee to find a suitable range for the norm of conditioned output. Instead of directly tuning the norm of output, we first normalize the output predictions to have the same norm of features. We then tune the norm of the normalized output with a norm control factor if necessary. This method is called Normalized OutpUt coNditioner (NOUN).

Similar to previous conditional methods [13], [14], [16], [18], NOUN applies raw output predictions for the conditional feature alignment across domains. However, due to the domain shift, target predictions generated by the source-trained classifier may be inaccurate, leading to possible misalignment between multi-modal structures across domains. We thus propose to condition the domain alignment in the prototype space rather than in the prediction space. Specifically, we leverage accurate source prototypes to project predictions into the same prototype space. The projected outputs are discriminative, simultaneously considering network predictions and semantic structures shared across domains. Incorporating the structure-aware conditioning into NOUN, we derive a simple yet strong UDA method called PROtotype-based Normalized OutpUt coNditioner (PRONOUN).

Experimental results on both object recognition and semantic segmentation tasks demonstrate the advantages of the proposed methods over previous state-of-the-arts [5], [6], [13], [14], [19] of domain adversarial training methods. The main contributions of this work are summarized as follows:

- We revisit the concatenation conditioning strategy for conditional domain alignment and provide a novel explanation to its previous failure in [14], i.e., the smaller norm of conditioned output predictions.
- We propose a novel and effective conditional domain adversarial training method NOUN. NOUN enlarges the norm of conditioned output by normalizing the output to have the same norm with features and involving an extra norm control factor for better flexibility.
- To deal with the possible misalignment due to inaccurate pseudo-labels, we further propose a novel prototype-based conditioning strategy and combine it with NOUN to be a strong UDA method called PRO-NOUN.
- Both NOUN and PRONOUN are simple and compact. The effectiveness of our methods is further validated on multiple cross-domain object recognition benchmarks. In addition, NOUN is shown to be generic and competitive for synthetic-to-real semantic segmentation tasks.

The rest of this paper is organized as follows. In Section II, we review related domain adaptation works. In Section III, we first introduce the settings of UDA and the vanilla

domain adversarial training solution [20]. Then we revisit the naive concatenation conditioning strategy and explain why it fails in conditional domain alignment. Finally, we introduce NOUN and PRONOUN. In Section IV, we conduct extensive experiments on cross-domain object recognition benchmarks and synthetic-to-real semantic segmentation benchmarks to evaluate NOUN and PRONOUN. Besides, we provide extensive ablations and analyses on our methods. In Section V, we briefly summarize our work and discuss the future work.

II. RELATED WORK

A. Unsupervised Domain Adaptation

Unsupervised domain adaptation (UDA) is first modeled as the covariate shift problem [22] where different domains have different marginal distributions but identical conditional distributions. At early times, researchers address UDA by instance re-weighting [23], [24], feature space alignment [25]–[27] or domain-invariant feature transformation [28]–[31]. In the past years, various deep UDA works [1], [4], [14], [20], [32], [33] have been developed and have significantly boosted the UDA performance in visual tasks. Generally, they can be classified into discrepancy-based methods and adversary-based methods. Discrepancy-based methods address the dataset shift by mitigating specific discrepancies defined on different layers of a shared model between domains, e.g., [1], [34], [35] adopt MMD measure and [10], [11] leverage central moments. Recently, adversarial learning [12] has become an increasingly popular solution to domain adaptation problems. It aims to learn domain-invariant features by introducing an extra domain discriminator to promote domain confusion. [20] optimizes a minimax objective via inserting a gradient reversal layer between the classification network and the discriminator. Differently, [32] adopts two independent optimization objectives with inverted labels following GAN [12].

B. Pseudo-Labeling

UDA can be also regarded as a semi-supervised learning (SSL) problem where unlabeled data from the same domain are replaced by data from another target domain. As a simple yet effective technique for SSL, pseudo-labeling [36] has been widely used in UDA methods. For example, [2], [37], [38] exploit intermediate pseudo-labels with tri-training and self-training, respectively. Recently, curriculum learning [39], self-paced learning [40] and re-weighting schemes [14], [41] are further leveraged to protect the domain adaptation from inaccurate pseudo-labels.

C. Conditional Domain Adaptation

Apart from being explicitly used to re-train the classification network, pseudo-labels can also be exploited to enhance the feature-level domain alignment. [29], [42] show that pseudo-labels can help mitigate the joint distribution discrepancy via minimizing multiple class-wise MMD measures. [35] proposes to align joint distributions of multiple domain-specific layers across domains based on a joint MMD

TABLE I

COMPARISON BETWEEN RELATED CONDITIONAL DOMAIN ADVERSARIAL TRAINING METHODS AND OUR METHODS, I.E., NOUN AND PRONOUN, WHERE DANN [20], [21] IS AS A COMMON BASELINE. D MEANS THE DOMAIN CLASSIFIER, I.E., THE DISCRIMINATOR. WE DENOTE THE FEATURE DIMENSION AS d AND THE CLASS NUMBER AS c . IN THE FIRST COLUMN, ‘INPUT DIMENSION OF D ’: THE INPUT FEATURE DIMENSION OF THE DISCRIMINATOR; ‘OUTPUT DIMENSION OF D ’: THE OUTPUT PREDICTION DIMENSION OF THE DISCRIMINATOR; ‘NUMBER OF D ’: THE NUMBER OF DISCRIMINATORS INVOLVED IN THE METHOD; ‘SENSITIVITY TO PSEUDO-LABELS’: THE SENSITIVITY OF THE METHOD TO PSEUDO-LABELS. WE CLASSIFY THIS SENSITIVITY INTO THREE DEGREES: LOW, MEDIUM, AND HIGH. METHODS NOT USING PSEUDO LABELS ARE CLASSIFIED AS THE ‘LOW’ SENSITIVITY. METHODS WITH ‘HIGH’ SENSITIVITY DIRECTLY RELY ON PSEUDO LABELS TO PERFORM CLASS-WISE FEATURE ALIGNMENT ACROSS DOMAINS. METHODS WITH THE ‘MEDIUM’ SENSITIVITY IMPLICITLY LEVERAGE PSEUDO LABELS FOR CONDITIONAL DOMAIN ALIGNMENT

	DANN [20]	MADA [13]	CDAN [14]	DANN-CA [17]	RCA [16]	IDDA [15]	NOUN	PRONOUN
Input dimension of D	d	d	cd	d	d	d	$c + d$	$2d$
Output dimension of D	1	1	1	$c + 1$	$2c$	$c + 1$	1	1
Amount of D	1	c	1	0	1	1	1	1
Sensitivity to pseudo-labels	low	high	medium	low	high	low	medium	medium

criterion. Recent methods resort to conditional domain adversarial training for better alignment of multi-modal structures across domains. [13], [18] leverage pseudo-labels and multiple class-wise domain discriminators to enable fine-grained alignment between domains. In contrast, [14] conditions the adversarial domain alignment on discriminative information via the outer product between features and predictions. [15], [16] introduce a classification-aware domain discriminator. Especially, [16] imposes the fine-grained adversarial loss with pseudo-labels and [15] performs the adversarial training asymmetrically between semantic labels and the domain labels. Different from [15], [17] replaces the original discriminator with a domain-aware classifier and achieves the asymmetrical adversarial training within the exact classifier. [43] improves [17] by introducing discriminative interaction between domain predictions and semantic predictions.

Concerning the above works, the conditional domain adversarial training methods [13]–[18], [20] are the most related to ours. We thus compare NOUN and PRONOUN with these works in Table I in terms of the implementation of the discriminator and the sensitivity to pseudo-labels. The comparison shows that NOUN and PRONOUN are compact, easy to implement, and reliable conditional domain adversarial training methods.

III. METHOD

In this section, we begin with the basic settings of UDA as well as the introduction of domain adversarial training [20]. Then we revisit a naive conditional domain adversarial training method, i.e., the concatenation conditioning strategy mentioned in [14] and analyze its failure in domain alignment with multi-modal structures. We then propose NOUN, a simple yet effective concatenation conditioning method. Finally, we propose a novel prototype-based conditioning strategy to mitigate the misalignment induced by inaccurate predictions.

A. Preliminaries

In a vanilla UDA task, we are given label-rich source domain data $\{(\mathbf{x}_s^i, \mathbf{y}_s^i)\}_{i=1}^{n_s}$ sampled from the joint distribution $P_s(\mathbf{x}_s, \mathbf{y}_s)$ and unlabeled target domain data $\{\mathbf{x}_t^i\}_{i=1}^{n_t}$ sampled from the joint distribution $Q_t(\mathbf{x}_t, \mathbf{y}_t)$, where $\mathbf{x}_s^i \in \mathcal{X}_S$ and $\mathbf{y}_s^i \in \mathcal{Y}_S$ denote an image and its corresponding label from the

source domain dataset and $\mathbf{x}_t^i \in \mathcal{X}_T$ denotes an image from the target domain dataset and $P_s \neq Q_t$. The goal of UDA is to learn a discriminative model from $\mathcal{X}_S, \mathcal{Y}_S$, and \mathcal{X}_T to predict labels for unlabeled target samples \mathcal{X}_T .

Domain adversarial training methods aim to reduce the domain discrepancy through learning domain-invariant features. Then the discriminative classifier trained by source labeled data can be freely applied to target unlabeled data. A prevailing domain adversarial training framework proposed in [20], [21] consists of a feature extractor network G , a classifier network F , and a discriminator network D . Given an image \mathbf{x} , we denote the feature vector extracted by G as $\mathbf{f} = G(\mathbf{x}) \in \mathbb{R}^d$ and the probability prediction obtained by F as $\mathbf{p} = F(\mathbf{f}) \in \mathbb{R}^c$, where d means the feature dimension and c means the number of classes. The vanilla domain adversarial training method dubbed DANN in [20], [21] can be formulated as optimizing the following minimax problem:

$$\min_{G, F} \max_D \mathcal{L}_y(G, F) - \lambda_{adv} \mathcal{L}_{dann}(G, D), \quad (1)$$

$$\begin{aligned} \mathcal{L}_{dann}(G, D) = & -\mathbb{E}_{\mathbf{x}_s^i \sim P_s} \log[D(\mathbf{f}_s^i)] \\ & - \mathbb{E}_{\mathbf{x}_t^j \sim Q_t} \log[1 - D(\mathbf{f}_t^j)], \end{aligned} \quad (2)$$

$$\mathcal{L}_y(G, F) = -\mathbb{E}_{(\mathbf{x}_s^i, \mathbf{y}_s^i) \sim P_s} \mathbf{y}_s^{i,T} \log(\mathbf{p}_s^i), \quad \mathbf{p}_s^i = F(G(\mathbf{x}_s^i)), \quad (3)$$

where the binary classifier $D : \mathbb{R}^d \rightarrow [0, 1]$ predicts the domain assignment probability for the input features \mathbf{f} , $\mathcal{L}_y(G, F)$ is the cross-entropy loss of labeled source data as for the classification task, and λ_{adv} is the coefficient for the adversarial loss.

B. Why Does the Naive Concatenation Conditioning Strategy Fail?

DANN does not explicitly consider discriminative information in cross-domain feature alignment. Therefore, it cannot guarantee the effective alignment of domains with multi-modal structures, even if the domain discriminator is fully confused. Motivated by conditional GANs [14], [44], [45] proposes to leverage the discriminative network prediction \mathbf{p} for conditional domain adversarial training. A simple baseline conditioning strategy named DANN-[f, p] in [14] is to condition the cross-domain feature alignment on network predictions by

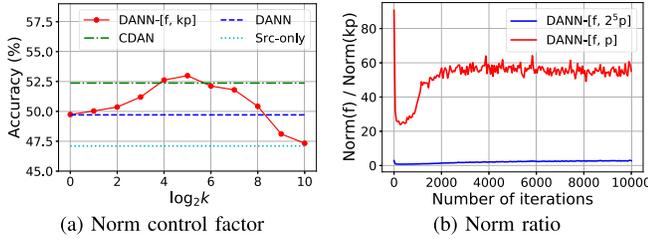


Fig. 1. Results for a 65-way classification UDA task Ar→Cl on Office-Home. ‘ k ’ is the norm control factor to enlarge the norm of conditioned output. ‘Src-only’ means using the source-trained model without domain adaptation. (a) illustrates how the norm of the conditioned output affects the adaptation performance of the concatenation conditioning method. (b) shows the ratio between the Euclidean norm of two concatenated parts during the training.

directly concatenating features with predictions, i.e., $f \oplus p$. Therefore, the adversarial loss of this naive concatenation conditioning method, i.e., DANN-[f, p], can be formulated as the following:

$$\mathcal{L}_{dann-fp}(G, D) = -\mathbb{E}_{x_s^i \sim P_s} \log[D(f_s^i \oplus p_s^i)] - \mathbb{E}_{x_t^j \sim Q_t} \log[1 - D(f_t^j \oplus p_t^j)]. \quad (4)$$

Compared with the adversarial loss of DANN in Eq. (2), the concatenation operation $f \oplus p$ explicitly considers discriminative predictions for conditional feature alignment between domains. However, the naive concatenation conditioning strategy is demonstrated to be unsuccessful for cross-domain alignment with multi-modal structures in [14]. Furthermore, [14] accounts the failure of the naive concatenation conditioning for the independence between f and p in $f \oplus p$. A similar observation for DANN-[f, p] in our experiments is shown in Fig. 1(a) when the norm control factor k is set as 1, i.e., DANN-[f, p] does not outperform DANN.

However, experiments in Fig. 1(a) show that if we simply introduce a norm control factor k to re-weight the network prediction for the concatenation conditioning, the adaptation performance would vary a lot with different k . Since this re-weighting operation only changes the norm of the concatenated prediction p , it is natural to take a closer look at the norm of the two concatenated parts during the model training. We illustrate the norm ratio between f and p in Fig. 1(b) for the naive concatenation conditioning (k equals to 1) and the effective re-weighted concatenation conditioning case (k equals to 32). It is noteworthy that for the naive concatenation strategy, the norm ratio keeps more than 50. While for the re-weighted concatenation strategy, the norm ratio keeps small and close to 1. In addition, as shown in Fig. 1(a), re-weighted with moderate norm factors, the concatenation conditioning can outperform the informative multi-linear conditioning strategy implemented via the outer product between f and p [14].

The above observations inspire us to rethink the reason for the failure of the naive concatenation conditioning strategy. To figure out how the norm control factor ‘ k ’ affects the conditional domain adversarial training, we take a closer look at the optimization process. For concatenation conditioning strategy, denote the output dimension of the first layer of D by d_e and the weight of the first layer of D by $W_1 = W_f \oplus W_g$, where $W_f \in \mathbb{R}^{d \times d_e}$ accepts the input feature f

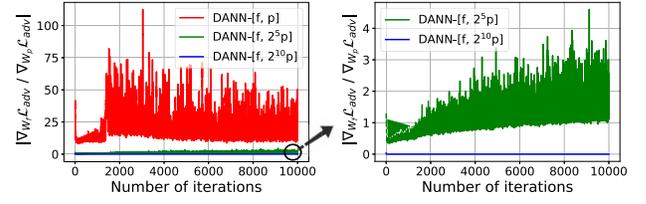


Fig. 2. How does the norm of the conditioned output affect the optimization of D ? $|\nabla_{W_f} \mathcal{L}_{adv}|$ denotes the absolute value of the mean of gradient matrix of W_f , i.e., $\nabla_{W_f} \mathcal{L}_{adv}$. $|\nabla_{W_f} \mathcal{L}_{adv}| / |\nabla_{W_p} \mathcal{L}_{adv}|$ thus measures the ratio between gradient of W_f and gradient of W_p .

and $W_p \in \mathbb{R}^{c \times d_e}$ accepts the conditioned output p . Taking $f \oplus kp$ as the input of D and \mathcal{L}_{adv} as the adversarial loss, we compare the gradients from domain confusion loss to W_f and W_g under different k values in Fig. 2. We analyze the concatenation conditioning with different k values according to the performance in Fig. 1(a) and the optimization process in Fig. 2. For naive concatenation conditioning ($k = 1$), the gradient of W_f is tens of times that of W_p . The marginal features f dominate the optimization of D , and conditioned output predictions p are ignored, making the conditional adversarial learning degenerate to the ‘DANN’ case.

While for conditioning with extremely large output norm ($k = 2^{10}$), the conditioned output predictions p dominate the discriminator optimization, and features vf are ignored, making the conditional adversarial learning degenerate to the ‘Src-only’ case. Conditioning is beneficial for the domain adaptation only when the norm of conditioned output is moderate, e.g., $k = 2^5$.

To sum up, the success of re-weighted concatenation conditioning indicates that conditioning on the stop-gradient concatenated predictions, the discriminator can capture the multi-modal structures for effective conditional domain alignment. Therefore, it seems that the independence between f and p in the concatenation strategy is not the main obstacle blocking effective conditioning. Instead, the devil is in the norm of the concatenated prediction. In fact, since p means the probability predictions for c categories, the Euclidean norm of p is between $\frac{1}{c}$ and 1. In contrast, the feature f is usually a high-dimension vector. The feature norm of f is probably tens of times that of p . The effect of this overwhelming norm ratio between f and p is that the domain confusion objective neglects the concatenated prediction, leading to ineffective conditioning. We argue this is the main reason for the failure of the naive concatenation conditioning strategy.

C. Normalized Output conditioner (NOUN)

Since the norm of features varies a lot for different tasks, it is non-trivial to find suitable k for a given UDA task, as shown in Fig. 1 (a). Besides, Fig. 1 (b) shows that the norm ratio between two concatenated parts plays an important role in effective conditional feature alignment across domains. Therefore, we normalize p to have the same Euclidean norm of f using Eq. (5). To keep flexible, we retain k as an extra norm control factor. Then k would represent the norm ratio between f and p . Therefore, it will be friendly to decide the value of k within a normal range for different UDA tasks.

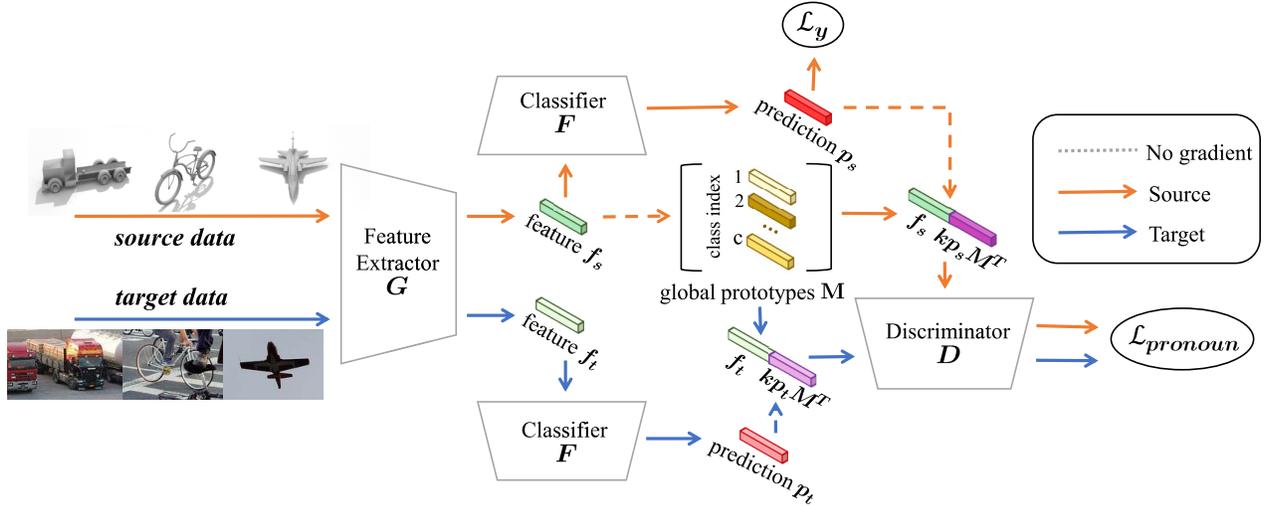


Fig. 3. The framework of PRONOUN. PRONOUN conditions the cross-domain feature alignment on semantic structure-aware predictions via the concatenation conditioning. Note that concatenated output predictions are detached without gradient from the binary discriminator D and the global prototypes M are collected from source domain by Eq. (7) and Eq. (8).

In this way, we obtain an effective concatenation conditioning method called Normalized Output conditioner (NOUN). The adversarial loss of NOUN is formulated as:

$$\hat{p}_a = \frac{\|f_a\|_2}{\|p_a\|_2} p_a, \quad a \in \{s, t\}, \quad (5)$$

$$\begin{aligned} \mathcal{L}_{noun}(G, D) = & -\mathbb{E}_{x_s^i \sim P_s} \log[D(f_s^i \oplus k\hat{p}_s^i)] \\ & -\mathbb{E}_{x_t^j \sim Q_t} \log[1 - D(f_t^j \oplus k\hat{p}_t^j)], \quad (6) \end{aligned}$$

where k is a norm control factor and the default value is 1.

D. Prototype-Based Normalized Output conditioner (PRONOUN)

Effective conditioning on network output is supposed to achieve more discriminative domain alignment than marginal domain feature alignment. However, output predictions of target data are probably inaccurate due to the domain shift. As a result, conditioning the cross-domain feature alignment directly on pseudo-labels may damage the domain adaptation performance. [14] thus proposes the entropy conditioning strategy where entropy is leveraged to explicitly distinguish out easy-to-transfer samples towards safe adaptation. Differently, we propose to alleviate the misalignment of multi-modal structures between domains from the perspective of the conditioning space.

In UDA, a model is trained with labeled source data, making source predictions nearly one-hot vectors. For better illustration, we consider using one-hot target predictions for the conditional domain alignment. Though discriminative, the multiple modes embedded in the output space are generally independent and orthogonal to each other. If conditioning in the output space, concatenated output predictions would directly dominate the domain alignment. Guided by inaccurate pseudo-labels, the cross-domain alignment would be misleading. Note target predictions contain both uncertain predictions and certain but inaccurate ones. The entropy conditioning [14] thus cannot fully mitigate the misalignment. Specifically, imposing domain alignment with confident instances of two

different categories would heavily damage the feature learning in the target domain. Even if the leveraged target predictions are soft probabilities, the output space only encodes implicit and faint information of semantic structures, which helps little in mitigating the misalignment induced by inaccurate pseudo labels.

Therefore, we propose to condition the domain alignment on prototypes, which explicitly encode semantic structures. Semantic correlations embedded in the semantic structures are promising to alleviate the misalignment induced by inaccurate conditioned outputs. We represent the domain-shared semantic structures with source global prototypes. Specifically, we update global prototypes in training by the exponential moving average strategy (ema) over the batch-level prototypes. We formulate the process of obtaining reliable prototypes as follows:

$$M_{\text{batch}} = [m_1 \cdots m_e \cdots m_c], \quad m_e = \frac{\sum_i^b \mathbb{1}(y_i^s = e) f_s^i}{\sum_i^b \mathbb{1}(y_i^s = e)}, \quad (7)$$

$$M = \lambda_{\text{ema}} M + (1 - \lambda_{\text{ema}}) M_{\text{batch}}, \quad (8)$$

where $M_{\text{batch}} \in \mathbb{R}^{d \times c}$ denotes batch-level prototypes, $m_e \in \mathbb{R}^d$ denotes the prototype of class e , $\lambda_{\text{ema}} \in [0, 1]$ is a coefficient and $M \in \mathbb{R}^{d \times c}$ denotes global prototypes.

Taking source prototypes as the basis and projecting output predictions into the same prototype feature space, we would obtain discriminative yet semantic structure-aware predictions for conditioning. Seamlessly combining NOUN with prototype-based conditioning, we obtain the PROTOTYPE-based Normalized Output conditioner (PRONOUN). The domain adversarial training objective of PRONOUN is shown as below:

$$\begin{aligned} \mathcal{L}_{pronoun}(G, D) = & -\mathbb{E}_{x_s^i \sim P_s} \log[D(f_s^i \oplus k\hat{p}_s^i M^T)] \\ & -\mathbb{E}_{x_t^j \sim Q_t} \log[1 - D(f_t^j \oplus k\hat{p}_t^j M^T)]. \quad (9) \end{aligned}$$

Fig. 3 illustrates the framework of PRONOUN. Without complicated network designs or complex optimization objectives, both NOUN and PRONOUN are simple and compact

solutions to cross-domain feature alignment with multi-modal structures.

IV. EXPERIMENTS

In this section, we conduct extensive experiments to evaluate NOUN and PRONOUN. We start with experimental settings, including benchmarks, implementation details, and compared methods. Then we compare our solutions with related works on various UDA benchmarks. At last, we provide detailed analyses on our solutions, including ablation study, compatibility with self-training, stability of training, the sensitivity of hyper-parameters, and qualitative visualizations in the feature space or output space.

A. Experimental Setup

We conduct experiments to verify the effectiveness of NOUN and PRONOUN on cross-domain object recognition benchmarks including *Office-Home* [52], *VisDA2017* [53], *Office31* [54] and *ImageCLEF-DA*.¹ Besides, we validate the generalization of our methods on synthetic-to-real semantic segmentation benchmarks including *GTA5* [55]→*Cityscapes* [56] and *Synthia* [57]→*Cityscapes*.

1) *Datasets*: *Office-Home* is a new challenging dataset that consists of 65 different object categories found typically in 4 different Office and Home settings, i.e., Artistic (**Ar**) images, Clip Art (**Ca**), Product images (**Pr**) and Real-World (**Re**) images. *VisDA2017* is a challenging large-scale benchmark aimed at the synthetic-to-real object recognition task across 12 categories. The source domain consists of 152k synthetic images generated by rendering 3D models. The target domain contains 55k realistic images gathered from *Microsoft COCO* [58]. *Office31* is a popular dataset that includes 31 object categories taken from 3 domains, i.e., Amazon (**A**), DSLR (**D**) and Webcam (**W**). *ImageCLEF-DA* is a dataset built for the ‘ImageCLEF2014:domain-adaptation’ competition. We follow [1] to select 3 subsets, i.e., **C**, **I** and **P**, which share 12 common object classes.

Cityscapes is a realistic dataset of pixel-level annotated urban street scenes. We use its original training split and validation split as training target data and testing target data, respectively. *GTA5* consists of 24,966 densely labeled synthetic road scenes annotated with the same 19 classes as *Cityscapes*. For *Synthia*, we take the SYNTHIA-RAND-CITYSCAPES set as the source domain, which contains 9,400 synthetic images compatible with 16 annotated classes of *Cityscapes* [59].

2) *Implementation Details*: For object recognition, we follow the standard protocol [20], i.e., using all labeled source instances and all unlabeled target instances for UDA, and report the average accuracy based on results of three random trials for fair comparisons. Following [14], we experiment with ResNet-50 model pre-trained on ImageNet for all datasets excluding *VisDA2017* where we use ResNet-101 model to compare with other works. We follow [14] to choose the network parameters. The whole model is trained through

back-propagation, and λ_{adv} increases from 0 to 1 with the same strategy as [20]. Regarding the domain discriminator, we design a simple classifier with only one ReLU [60] layer and two linear layers (256→1024→1). Empirically, we fix the batch size to 36 with the initial learning rate as 1e-4.

For semantic segmentation, we adopt DeepLab-V2 [61] based on ResNet-101 [46] as done in [4], [6], [19], [62]. Following DCGAN [63], the discriminator network consists of one LeakyReLU layer [64] with the slope of 0.2 and two 4×4 convolutional layers with stride 2 (512→512→1). In training, we use SGD [65] to optimize the network with momentum (0.9), weight decay (5e-4) and initial learning rate (2.5e-4). We use the same learning rate policy as in [61]. Discriminators are optimized by Adam [66] with momentum ($\beta_1 = 0.9$, $\beta_2 = 0.99$), initial learning rate (1e-4) along with the same decreasing strategy as above. λ_{adv} is set to 1e-3 following [4].

All experiments are implemented via PyTorch on a single 12GB GPU. The total iteration number is set as 10k for object recognition and 100k for semantic segmentation. The momentum value λ_{ema} is set to 0.5 for all tasks without parameter selection. The norm control factor k is determined by reverse validation [20] and is set to 3 in PRONOUN for all tasks. Similar to the entropy-based criterion in [67], the adapted model producing target predictions with the minimum mean entropy is selected for testing in the target domain. Data augmentation skills like random scale or ten-crop ensemble evaluation are not adopted.

3) *Baseline Methods*: The general comparisons between our NOUN as well as PRONOUN and other conditional domain adversarial training methods are shown in Table I.

For cross-domain object recognition tasks, we compare our methods with DANN [20], DANN-[f, p] [14], MADA [13], CDAN [14], IDDA [15], RCA [16], DANN-CA [17] and CDAN + E [14]. Note that we re-implement these most related methods using the same training protocol including batch size and training iterations. As for other hyper-parameters including learning rate and optimizer setting, we strictly follow the released code by authors or the implementation details in original papers. For fair comparisons, we report both results in original papers if any, and the accuracy of our re-implementation. Besides, we also compare with the state-of-the-art works including SAFN [48], SWD [51], ADR [49], DWC-MEC [47], BSP [50], CAT [68] and iCAN [69]. Our reported ‘source-model-only’ means training a model on source data with only the supervised loss in Eq. (3) while using the same training protocol and architecture. In this way, we can make fair comparisons among different methods.

For synthetic-to-real semantic segmentation tasks, we compare our methods with domain adversarial training methods including AdaptSeg [4], SIBAN [5], AdvEnt [62], CLAN [6], AdaptPatch [19] and another generic UDA method SWD [51]. ‘NonAdapt’ shares the similar meaning with ‘source-model-only’ in object recognition benchmarks.

By default, NOUN introduces no hyper-parameter by setting the k in Eq. (6) as 1. For results reported in Table II–Table VII, we report results from original papers in the top rows, the re-implementation results of related methods using the same

¹<https://www.imageclef.org/2014/adaptation>

TABLE II
RECOGNITION ACCURACIES (%) ON *Office-Home* VIA RESNET-50. **BOLD: BEST**

Methods (Source→Target)	Ar→Cl	Ar→Pr	Ar→Re	Cl→Ar	Cl→Pr	Cl→Re	Pr→Ar	Pr→Cl	Pr→Re	Re→Ar	Re→Cl	Re→Pr	Avg.
ResNet-50 [46]	34.9	50.0	58.0	37.4	41.9	46.2	38.5	31.2	60.4	53.9	41.2	59.9	46.1
DANN [20]	45.6	59.3	70.1	47.0	58.5	60.9	46.1	43.7	68.5	63.2	51.8	76.8	57.6
CDAN [14]	49.0	69.3	74.5	54.4	66.0	68.4	55.6	48.3	75.9	68.4	55.4	80.5	63.8
CDAN+E [14]	50.7	70.6	76.0	57.6	70.0	70.0	57.4	50.9	77.3	70.9	56.7	81.6	65.8
DWT-MEC [47]	50.3	72.1	77.0	59.6	69.3	70.2	58.3	48.1	77.3	69.3	53.6	82.0	65.6
SAFN [48]	52.0	73.3	76.3	64.2	69.9	71.9	63.7	51.4	77.1	70.9	57.1	81.5	67.3
DANN [20]	48.1	62.3	74.9	54.2	64.1	65.2	53.6	47.7	74.6	66.9	54.6	79.0	62.1
DANN-[f, p] [14]	49.4	63.4	74.0	56.1	63.9	65.7	54.7	47.4	74.7	67.6	54.8	79.6	62.6
MADA [13]	47.2	69.7	75.1	54.3	68.5	69.8	53.0	42.6	75.0	70.1	51.2	80.8	63.1
CDAN [14]	51.5	69.6	75.2	58.1	69.1	68.9	58.8	49.7	77.7	70.3	55.5	81.4	65.5
IDDA [15]	49.9	66.2	74.2	57.6	66.5	66.9	55.5	49.6	74.9	67.0	53.2	78.6	63.4
RCA [16]	47.8	59.7	72.7	56.6	68.1	68.9	55.2	47.8	75.9	67.9	52.7	79.7	62.7
DANN-CA [17]	45.7	68.4	75.4	58.4	67.1	68.8	55.4	43.7	74.6	66.5	49.2	78.3	62.6
CDAN+E [14]	55.3	72.0	77.8	63.5	71.9	72.7	62.5	54.5	80.2	72.3	57.9	83.7	68.7
source-model-only	46.9 ±0.2	67.5 ±0.7	75.3 ±0.1	56.2 ±0.9	63.8 ±1.1	66.6 ±0.2	55.1 ±0.3	42.3 ±1.1	74.3 ±0.8	66.9 ±0.5	49.3 ±0.2	78.2 ±0.1	61.9 ±0.2
NOUN	53.2 ±0.1	69.2 ±0.2	76.0 ±0.4	60.6 ±0.5	68.5 ±0.5	69.9 ±0.2	62.6 ±0.2	53.3 ±0.0	78.6 ±0.3	71.2 ±0.4	56.0 ±0.4	80.9 ±0.5	66.7 ±0.0
PRONOUN	57.6 ±0.7	75.0 ±0.5	78.4 ±0.1	64.9 ±0.6	74.0 ±0.8	74.8 ±0.0	66.6 ±0.3	58.2 ±0.4	80.4 ±0.0	74.3 ±0.3	60.4 ±1.1	84.3 ±0.1	70.7 ±0.2

TABLE III
ACCURACIES (%) ON *VisDA2017* VIA RESNET-101

Method	plane	bicycl	bus	car	horse	knife	mcycl	person	plant	sktbrd	train	truck	Per-class
ResNet-101 [46]	55.1	53.3	61.9	59.1	80.6	17.9	79.7	31.2	81.0	26.5	73.5	8.5	52.4
DANN [20]	81.9	77.7	82.8	44.3	81.2	29.5	65.1	28.6	51.9	54.6	82.8	7.8	57.4
ADR [49]	94.2	48.5	84.0	72.9	90.1	74.2	92.6	72.5	80.8	61.8	82.2	28.8	73.5
CDAN [14]	85.2	66.9	83.0	50.8	84.2	74.9	88.1	74.5	83.4	76.0	81.9	38.0	73.7
CDAN+BSP [50]	92.4	61.0	81.0	57.5	89.0	80.6	90.1	77.0	84.2	77.9	82.1	38.4	75.9
SWD [51]	90.8	82.5	81.7	70.5	91.7	69.5	86.3	77.5	87.4	63.6	85.6	29.2	76.4
SAFN [48]	93.6	61.3	84.1	70.6	94.1	79.0	91.8	79.6	89.9	55.6	89.0	24.4	76.1
DANN [20]	90.0	58.9	76.9	56.1	80.3	60.9	89.1	72.5	84.3	73.8	89.3	35.8	72.3
DANN-[f, p] [14]	90.7	55.6	76.3	62.9	82.6	64.8	83.9	73.0	84.6	69.6	86.7	35.8	72.2
MADA [13]	84.8	53.3	78.1	66.6	87.4	80.3	85.9	72.4	87.2	69.7	83.3	25.8	72.9
CDAN [14]	91.9	62.4	74.5	72.3	89.0	83.0	88.6	80.9	85.3	72.7	86.6	32.6	76.3
IDDA [15]	76.4	75.8	74.2	58.9	90.4	94.9	80.3	74.5	90.1	69.7	82.5	47.3	76.2
RCA [16]	79.6	55.2	69.4	73.0	87.3	59.3	86.0	61.2	86.7	72.4	84.2	28.3	70.2
DANN-CA [17]	88.9	37.4	78.4	71.3	91.9	63.1	90.9	35.3	84.5	23.9	88.2	20.1	64.5
CDAN+E [14]	90.1	63.4	83.8	71.9	88.5	94.2	91.2	80.9	85.4	77.4	85.0	43.4	79.1
source-model-only	77.5 ±0.2	14.5 ±0.7	46.7 ±0.1	69.3 ±0.9	67.3 ±1.1	4.7 ±0.2	78.3 ±0.3	10.9 ±1.1	66.1 ±0.8	22.1 ±0.5	81.5 ±0.2	5.8 ±0.1	45.4
NOUN	92.5 ±1.3	64.3 ±6.6	78.3 ±5.2	62.9 ±7.3	89.5 ±1.2	88.2 ±2.8	90.8 ±1.8	79.1 ±0.5	87.8 ±2.6	81.2 ±4.5	83.8 ±1.4	48.1 ±11.5	78.9
PRONOUN	92.7 ±2.1	69.3 ±7.6	82.7 ±1.0	70.5 ±0.8	92.8 ±0.7	91.4 ±2.4	91.3 ±1.2	79.9 ±2.1	90.2 ±1.3	82.2 ±1.6	87.8 ±1.6	47.9 ±3.9	81.6

protocol in the middle rows, and results of our methods in the bottom rows.

B. Results on Object Recognition Benchmarks

For cross-domain object recognition, we evaluate our methods on benchmarks including *Office-Home*, *VisDA2017*, *Office31*, and *ImageCLEF-DA*.

1) *Office-Home*: The results of the challenging 65-way classification benchmark *Office-Home* are reported in Table II. All of the compared conditional domain adversarial training methods beat DANN in terms of the average accuracy,

which verifies the effectiveness of conditional feature alignment for domains with multi-modal structures. Among these re-implemented methods, CDAN [14] models perform the best and achieve markedly higher average accuracy than results in original papers, i.e., 63.8% → 65.5% for CDAN and 65.8% → 68.7% for CDAN + E. The naive concatenation conditioning strategy DANN-[f, p] brings minor improvement of the average accuracy (0.6%) over DANN. While NOUN brings a 4.1% absolute increase of the average accuracy over DANN[f, p] and even outperforms the more informative multi-linear conditioning method CDAN. Suffering from inaccurate pseudo-labels, NOUN still falls behind CDAN + E.

TABLE IV
ACCURACIES (%) ON *Office31* VIA RESNET-50

Method (Source→Target)	A→D	A→W	D→A	D→W	W→A	W→D	Avg.
ResNet-50 [46]	68.9	68.4	62.5	96.7	60.7	99.3	76.1
DANN [20]	79.7	82.0	68.2	96.9	67.4	99.1	82.2
CDAN [14]	89.8	93.1	70.1	98.2	68.0	100.	86.6
CDAN+E [14]	92.9	94.1	71.0	98.6	69.3	100.	87.7
DANN-CA [17]	89.9	91.4	69.6	98.2	68.8	99.5	86.2
MADA [13]	87.8	90.0	70.3	97.4	66.4	99.6	85.2
iCAN [69]	90.1	92.5	72.1	98.8	69.9	100.	87.2
CAT [68]	90.8	94.4	72.2	98.0	70.2	100.	87.6
SAFN [48]	87.7	88.8	69.8	98.4	69.7	99.8	85.7
SAFN+Ent [48]	90.7	90.1	73.0	98.6	70.2	99.8	87.1
<hr/>							
DANN [20]	79.1	85.3	66.0	96.2	69.4	99.8	82.6
DANN-[f, p] [14]	80.3	86.0	67.5	96.3	69.2	99.6	83.1
MADA [13]	91.0	91.5	68.2	98.0	69.3	99.9	86.3
CDAN [14]	89.8	88.6	69.6	98.5	70.4	100.	86.2
IDDA [15]	85.7	88.3	66.0	96.5	67.7	99.8	84.0
RCA [16]	83.3	86.7	67.3	97.9	66.8	100.	83.7
DANN-CA [17]	84.4	84.7	66.1	97.7	65.9	99.8	83.1
CDAN+E [14]	90.2	91.0	73.8	98.8	72.3	99.7	87.6
<hr/>							
source-model-only	80.9 ±1.6	79.5 ±0.9	64.3 ±0.7	97.9 ±0.4	64.2 ±0.5	99.9 ±0.1	81.1 ±0.4
<hr/>							
NOUN	88.5 ±2.2	90.9 ±1.5	73.4 ±0.1	98.5 ±0.7	72.7 ±1.0	100. ±0.0	87.3 ±0.5
<hr/>							
PRONOUN	93.5 ±2.2	93.4 ±0.8	74.6 ±0.5	98.2 ±0.8	73.4 ±0.8	100. ±0.0	88.8 ±0.4

While conditioning on the semantic structure-aware predictions, PRONOUN outperforms CDAN + E for all of the 12 transfer tasks and improves the average accuracy over NOUN by 4%.

2) *VisDA2017*: The results of the large-scale benchmark *VisDA2017* are reported in Table III. DANN-[f, p] achieves almost the same per-class accuracy as DANN, while CDAN models still outperform other compared domain adversarial training methods. NOUN significantly improves the per-class accuracy by 6.7% over DANN-[f, p], comparable with CDAN + E. Both NOUN and CDANE outperform all other comparison methods. PRONOUN notably improves NOUN and achieves the best per-class accuracy among all domain adversarial training methods.

3) *Office31*: The results of the standard benchmark *Office31* are reported in Table IV. DANN-[f, p] improves marginally over DANN. Though CDAN + E outperforms comparison methods markedly, CDAN proves slightly inferior to MADA. The observation is reasonable because tasks in *Office31* such as A→D and A→W are easy and can provide more accurate pseudo-labels, which prioritizes methods involving explicit class-wise cross-domain feature alignment like MADA. Comparable with CDAN + E, NOUN improves the average accuracy by 4.2% over DANN-[f, p]. PRONOUN brings handsome improvement over NOUN on four main transfer tasks, including A→D, A→W, D→A, and W→A, and outperforms CDAN + E.

4) *ImageCLEF-DA*: We report results of the easy benchmark *ImageCLEF-DA* with equal domain size and balanced classes in Table V. Domains of *ImageCLEF-DA* are more visually similar than those of the above benchmarks. Thus the average accuracy of the source-trained model is higher, and various comparison methods exhibit generally comparable

TABLE V
ACCURACIES (%) ON *ImageCLEF-DA* VIA RESNET-50

Method (Source→Target)	C→I	C→P	I→C	I→P	P→C	P→I	Avg.
ResNet-50 [46]	78.0	65.5	91.5	74.8	91.2	83.9	80.7
DANN [20]	87.0	74.3	96.2	75.0	91.5	86.0	85.0
CDAN [14]	90.5	74.5	97.0	76.7	93.5	90.6	87.1
CDAN+E [14]	91.3	74.2	97.7	77.7	94.3	90.7	87.7
MADA [13]	88.8	75.2	96.0	75.0	92.2	87.9	85.8
iCAN [69]	89.9	78.5	94.7	79.5	92.0	89.7	87.4
CAT [68]	91.3	75.3	95.5	77.2	93.6	91.0	87.3
SAFN [48]	91.1	77.0	96.2	78.0	94.7	91.7	88.1
SAFN+Ent [48]	91.7	77.6	96.3	79.3	95.3	93.3	88.9
<hr/>							
DANN [20]	91.6	76.5	95.3	76.8	92.3	90.1	87.1
DANN-[f, p] [14]	91.3	77.3	95.0	76.6	92.9	89.9	87.2
MADA [13]	91.6	77.1	97.0	77.1	94.6	90.7	88.0
CDAN [14]	91.2	76.5	95.9	77.0	93.6	91.0	87.5
IDDA [15]	90.7	75.8	95.6	74.1	94.2	88.6	86.5
RCA [16]	90.1	75.5	95.5	76.1	89.4	88.9	85.9
DANN-CA [17]	90.2	75.3	96.2	76.4	94.1	90.7	87.1
CDAN+E [14]	91.4	77.1	95.9	77.7	94.3	91.9	88.1
<hr/>							
source-model-only	85.3 ±0.3	72.6 ±0.8	93.2 ±0.3	75.7 ±0.4	91.8 ±0.9	88.2 ±0.6	84.5 ±0.1
<hr/>							
NOUN	92.3 ±1.1	78.1 ±0.5	96.8 ±0.2	77.6 ±0.3	94.6 ±0.2	91.8 ±1.0	88.5 ±0.3
<hr/>							
PRONOUN	92.3 ±0.8	78.3 ±0.9	97.4 ±0.3	78.0 ±0.4	95.2 ±0.8	92.6 ±0.9	89.0 ±0.4

performance. Specifically, DANN-[f, p] shows few improvements over DANN. Beyond the observation on *Office31*, MADA is shown to be comparable with CDAN + E and outperforms CDAN. While NOUN slightly beats CDAN + E and PRONOUN further improves the average accuracy over NOUN.

Across the above results on four different cross-domain object recognition benchmarks, there are two consistent observations about our methods. The first one is that NOUN always brings a significant improvement over DANN-[f, p] and NOUN beats all of the other domain adversarial training methods compared in Table I. This observation strongly demonstrates the effectiveness of NOUN. It also indicates that the failure of the naive concatenation conditioning method mainly lies in the weak conditioning strength, since NOUN only normalizes the norm of concatenated output in DANN-[f, p] with the norm of the feature. The second consistent observation is that PRONOUN always further improves NOUN and outperforms all of the compared baselines, which verifies the superiority of our prototype-based conditioning strategy.

C. Results on Semantic Segmentation Benchmarks

NOUN can be seamlessly extended to tasks with dense prediction. As for PRONOUN, we obtain prototypes by simply averaging features of corresponding pixels. We evaluate both methods on synthetic-to-real semantic segmentation benchmarks. Following the VisDA challenge [53], we adopt the mean intersection-over-union (mIoU) as the evaluation metric. The results of *GTA5*→*Cityscapes* are reported in Table VI and the results of *Synthia*→*Cityscapes* are reported in Table VII. For both tasks, NOUN outperforms all of the compared

TABLE VI
COMPARISON RESULTS OF *GTA5* [55]→*Cityscapes* [56] SEMANTIC SEGMENTATION USING RESNET-101 AS THE BACKBONE

Methods	road	sdwk	blndg	wall	fence	pole	light	sign	veg.	ter.	sky	per.	rider	car	truck	bus	train	mbike	bike	mIoU
NonAdapt [4]	75.8	16.8	77.2	12.5	21.0	25.5	30.1	20.1	81.3	24.6	70.3	53.8	26.4	49.9	17.2	25.9	6.5	25.3	36.0	36.6
AdaptSeg [4]	86.5	25.9	79.8	22.1	20.0	23.6	33.1	21.8	81.8	25.9	75.9	57.3	26.2	76.3	29.8	32.1	7.2	29.5	32.5	41.4
AdvEnt [62]	89.9	36.5	81.6	29.2	25.2	28.5	32.3	22.4	83.9	34.0	77.1	57.4	27.9	83.7	29.4	39.1	1.5	28.4	23.3	43.8
SIBAN [5]	88.5	35.4	79.5	26.3	24.3	28.5	32.5	18.3	81.2	40.0	76.5	58.1	25.8	82.6	30.3	34.4	3.4	21.6	21.5	42.6
CLAN [6]	87.0	27.1	79.6	27.3	23.3	28.3	35.5	24.2	83.6	27.4	74.2	58.6	28.0	76.2	33.1	36.7	6.7	31.9	31.4	43.2
AdaptPatch [19]	89.2	38.4	80.4	24.4	21.0	27.7	32.9	16.1	83.1	34.1	77.8	57.4	27.6	78.6	31.2	40.2	4.7	27.6	27.6	43.2
SWD (PSP) [51]	92.0	46.4	82.4	24.8	24.0	35.1	33.4	34.2	83.6	30.4	80.9	56.9	21.9	82.0	24.4	28.7	6.1	25.0	33.6	44.5
NonAdapt	83.7	10.2	73.9	20.7	15.2	23.4	32.9	22.4	76.7	29.6	66.7	55.0	28.5	69.4	29.8	40.1	0.9	16.2	33.6	38.4
NOUN	88.5	20.9	81.6	26.7	23.6	31.1	37.5	29.2	84.0	37.2	76.4	59.8	26.6	85.4	36.8	43.9	3.4	26.5	42.2	45.3
PRONOUN	88.7	25.8	81.7	25.8	23.1	30.5	36.3	25.7	84.1	34.8	76.6	59.4	28.9	84.1	34.6	42.7	6.4	25.3	34.1	44.7

TABLE VII

COMPARISON RESULTS OF *Synthia* [57]→*Cityscapes* SEMANTIC SEGMENTATION USING RESNET-101 AS THE BACKBONE. mIoU* DENOTES THE MEAN IOU OF 16 CLASSES, INCLUDING THE CLASSES WITH *

Methods	road	sdwk	blndg	wall*	fence*	pole*	light	sign	veg.	sky	per.	rider	car	bus	mbike	bike	mIoU	mIoU*
NonAdapt [4]	55.6	23.8	74.6	-	-	-	6.1	12.1	74.8	79.0	55.3	19.1	39.6	23.3	13.7	25.0	38.6	-
AdaptSeg [4]	79.2	37.2	78.8	-	-	-	9.9	10.5	78.2	80.5	53.5	19.6	67.0	29.5	21.6	31.3	45.9	-
AdvEnt [62]	87.0	44.1	79.7	9.6	0.6	24.3	4.8	7.2	80.1	83.6	56.4	23.7	72.7	32.6	12.8	33.7	47.6	40.8
SIBAN [5]	82.5	24.0	79.4	-	-	-	16.5	12.7	79.2	82.8	58.3	18.0	79.3	25.3	17.6	25.9	46.3	-
CLAN [6]	81.3	37.0	80.1	-	-	-	16.1	13.7	78.2	81.5	53.4	21.2	73.0	32.9	22.6	30.7	47.8	-
AdaptPatch [19]	82.4	38.0	78.6	8.7	0.6	26.0	3.9	11.1	75.5	84.6	53.5	21.6	71.4	32.6	19.3	31.7	46.5	40.0
SWD (PSP) [51]	82.4	33.2	82.5	-	-	-	22.6	19.7	83.7	78.8	44.0	17.9	75.4	30.2	14.4	39.9	48.1	-
NonAdapt	38.2	17.5	73.8	4.6	0.0	25.0	5.8	10.9	73.3	80.7	55.7	17.3	40.3	16.4	12.3	21.7	35.7	30.9
NOUN	89.3	45.2	79.4	4.0	0.4	27.0	8.0	13.2	78.9	82.3	57.0	20.4	80.9	30.8	15.5	38.0	49.2	41.9
PRONOUN	89.2	44.4	78.8	4.5	0.2	27.0	6.5	11.0	78.0	81.3	57.0	22.0	81.5	30.3	18.4	39.4	49.1	41.8

domain adversarial training methods, including those specially designed for semantic segmentation, which verifies the generalization and effectiveness of NOUN. Inconsistent with observations in cross-domain object recognition tasks, PRONOUN does not further improve NOUN on both synthetic-to-real semantic segmentation tasks. A possible reason is that the estimated prototypes in semantic segmentation are inaccurate. Even so, PRONOUN achieves the second-best results on both segmentation benchmarks.

D. Model Analysis and Discussions

1) *Ablation Study*: The benefit of enlarging the norm of concatenated predictions introduced in NOUN has been demonstrated by comparing DANN-[f, p] and NOUN. To validate the effectiveness of the prototype-based conditioning in PRONOUN, we further conduct ablation on the four classification benchmarks and show the respective average accuracy in Table VIII. ‘NOUN + E’ means combining the entropy conditioning [14] with NOUN. Increasing the norm control factor k from 1 to 3, NOUN models show notable performance drop on *Office-Home* and *Office31*, while both NOUN + E and PRONOUN models achieve performance improvement on all benchmarks. This comparison indicates that NOUN suffers from inaccurate predictions with too strong conditioning.

TABLE VIII

ABLATION STUDY: THE AVERAGE ACCURACY (%) OF DIFFERENT VARIANTS OF NOUN ON OBJECT RECOGNITION BENCHMARKS

Ablation	<i>Home</i>	<i>VisDA</i>	<i>Office31</i>	<i>CLEF</i>
NOUN ($k = 1$)	66.7	78.9	87.3	88.5
NOUN+E ($k = 1$)	68.9(↑)	79.8(↑)	88.1(↑)	88.5(-)
PRONOUN ($k = 1$)	68.1(↑)	80.0(↑)	88.0(↑)	88.6(↑)
NOUN ($k = 3$)	66.2(↓)	79.5(↑)	86.9(↓)	88.5(-)
NOUN+E ($k = 3$)	69.1(↑)	80.6(↑)	88.2(↑)	88.6(↑)
PRONOUN ($k = 3$)	70.7(↑)	81.6(↑)	88.8(↑)	89.0(↑)

Both the entropy conditioning [14] and our prototype-based conditioning are able to alleviate the misalignment induced by inaccurate predictions. In addition, the consistent improvement of NOUN + E and PRONOUN by simply increasing the norm control factor k , in turn, demonstrates that a larger norm of concatenated predictions can contribute to more effective conditional domain alignment. With stronger conditioning, i.e., $k = 3$, PRONOUN evidently surpasses the competitive counterpart of NOUN + E on all benchmarks, which explicitly verifies the superiority of our prototype-based conditioning.

2) *Training Stability*: We investigate the convergence of our methods by testing the model every 50 iterations during the training and illustrate the empirical accuracy curves of

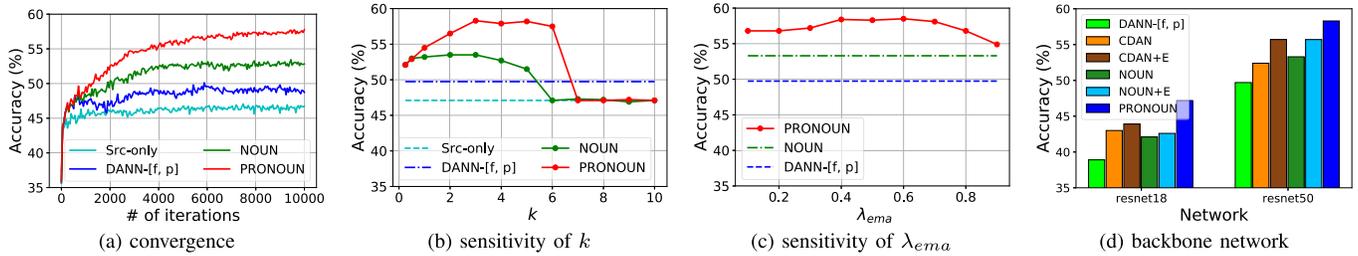


Fig. 4. Further analysis of NOUN and PRONOUN on the task Ar→Cl on *office-home*.

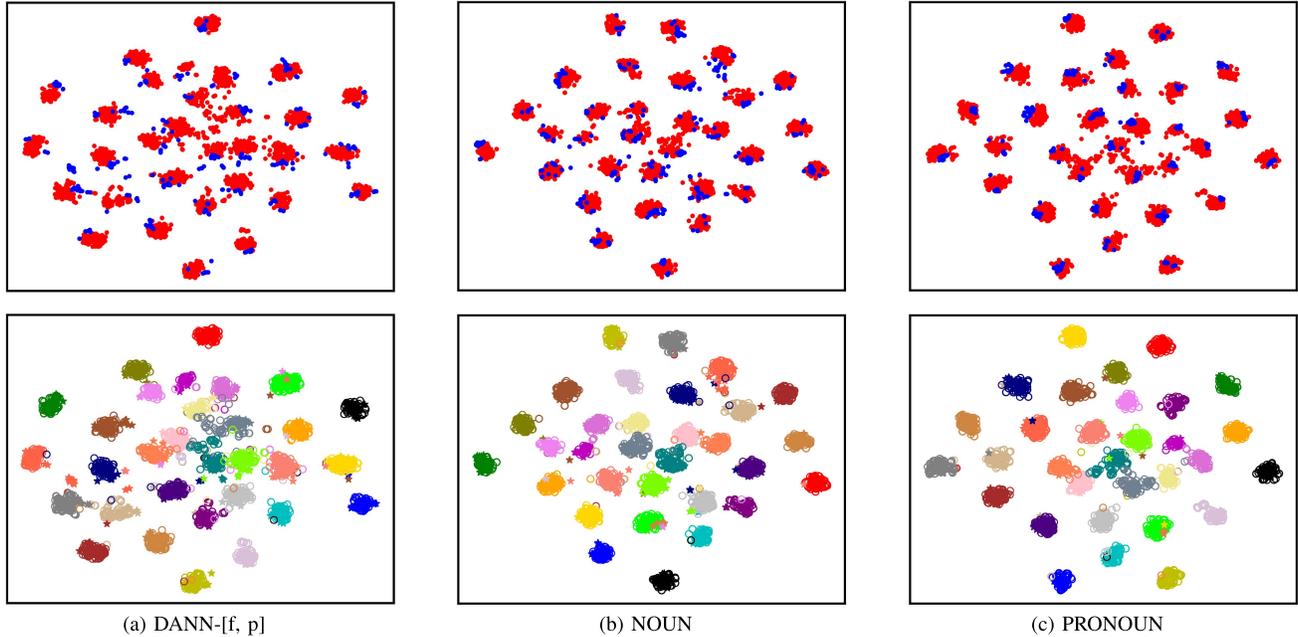


Fig. 5. t-SNE [70] embedding visualizations of UDA methods for the A→D task on *Office31*. In the upper row, colors denote different domains (red: source, blue: target). In the bottom row, colors denote different classes, and shapes denote the domain information (source in \circ and target in \star).

Ar→Cl on *Office-Home* in Fig. 4(a). It is obvious that for both NOUN and PRONOUN, the model training is stable and finally converges.

3) *Parameter Sensitivity*: The introduced hyper-parameters are k for both NOUN and PRONOUN and λ_{ema} only for PRONOUN. We find NOUN generally performs well with k of 1, i.e., NOUN can work without extra hyper-parameters. Thus we mainly investigate the sensitivity of k and λ_{ema} for PRONOUN on the UDA task Ar→Cl on *Office-Home* and show results in Fig. 4(b) and Fig. 4(c), respectively. For k , we observe that within a moderate range, i.e., when k is less than 3, increasing the norm of concatenated output brings improvement. However, an especially large norm of output leads to performance degradation. The moderate range is 7 for PRONOUN and 6 for NOUN, larger values of k lead to the worst performance close to the non-adapted source model. Because with an extreme value of k , the detached concatenation part will dominate the discriminator optimization. Then the contribution of features to the adversarial loss would be ignored, making domain adversarial training ineffective. For λ_{ema} , we set it as 0.5 for all tasks without selection. As shown in Fig. 4(c), 0.4–0.7 is a suitable range, and an extremely small or large value of λ_{ema} would weaken the performance. The reason is that moderate λ_{ema} can provide a better representation of semantic structures by keeping the prototypes both

consistent and up-to-date, which is beneficial to PRONOUN. Generally, the newly involved hyper-parameters in PRONOUN are not sensitive within the normal range.

4) *Network Sensitivity*: Pseudo-labels play an important role in conditional domain adversarial training. The network architecture has a great influence on the quality of pseudo-labels. Thus we evaluate the sensitivity of our methods to the network architecture. In Fig. 4(d) we report the respective accuracy of DANN-[f, p], DANN, CDAN, CDAN + E, NOUN, NOUN + E and PRONOUN for Ar→Cl on *Office-Home* with two backbones of different capacity, i.e., ResNet-18 and ResNet-50, where ‘NOUN + E’ denotes only replacing our prototype-based conditioning with the entropy conditioning [14]. For the stronger model ResNet-50, NOUN slightly beats CDAN and PRONOUN outperforms both CDAN + E and NOUN + E. For ResNet-18 with worse pseudo-labels, NOUN slightly falls behind CDAN but still significantly exceeds DANN-[f, p], and PRONOUN outperforms other methods by a large margin. Generally, both of our methods show desirable robustness to the network change. NOUN always significantly outperforms DANN-[f, p], and PRONOUN always performs the best.

5) *Compatibility With Self-Training*: To evaluate the compatibility of NOUN and PRONOUN with other types of domain adaptation methods, we choose a state-of-the-art

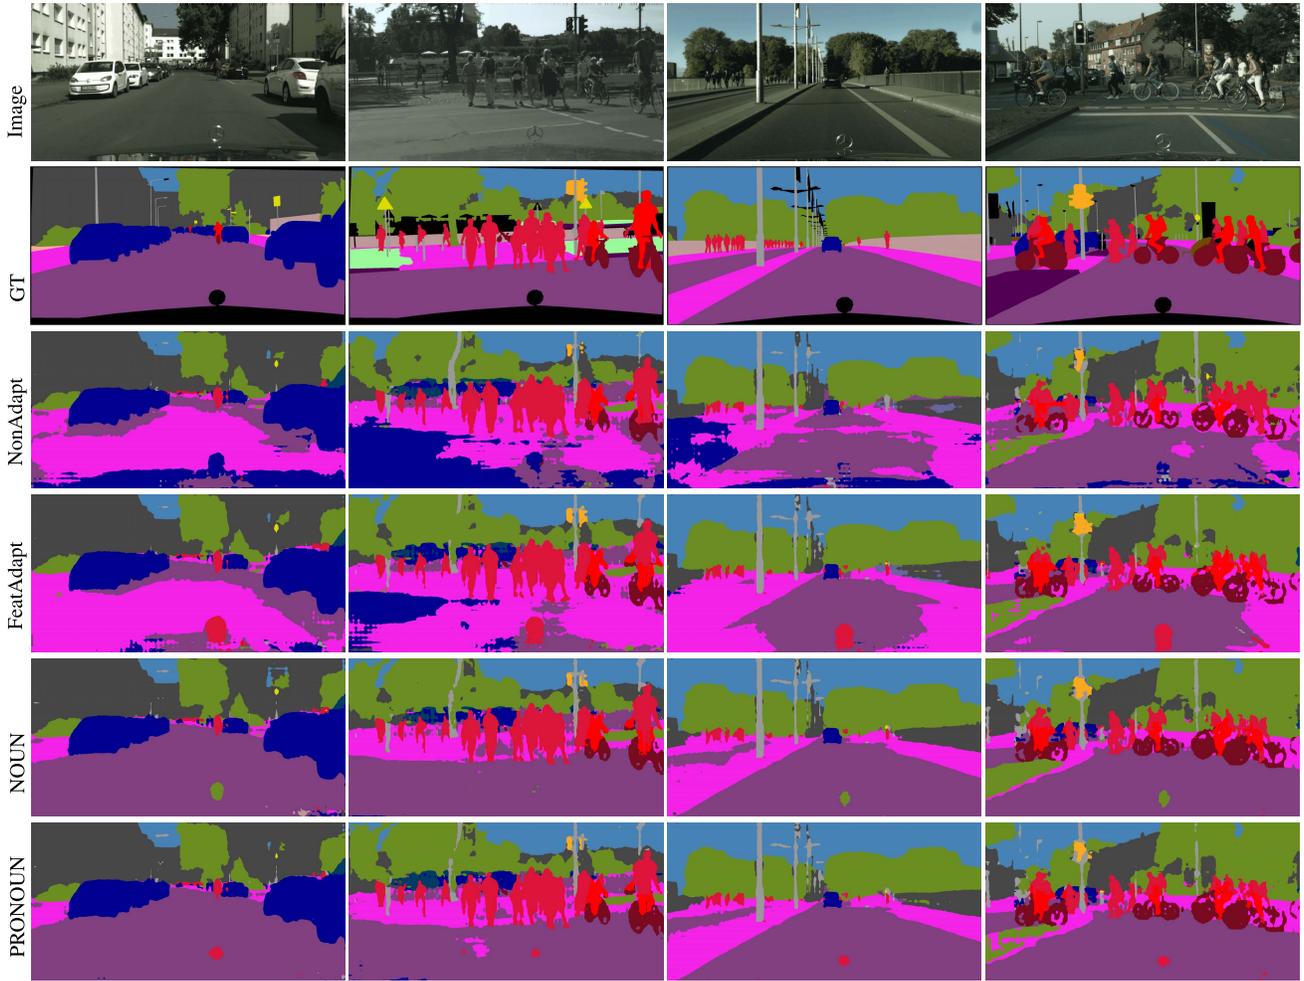


Fig. 6. Qualitative results of UDA methods on the synthetic-to-real semantic segmentation task of *Synthia*→*Cityscapes*.

TABLE IX

COMPATIBILITY WITH THE SELF-TRAINING METHOD: THE AVERAGE ACCURACY (%) ON OBJECT RECOGNITION BENCHMARKS

Methods	<i>Home</i>	<i>VisDA</i>	<i>Office31</i>
source-model-only [38]	59.4	49.1	76.5
ATDOC [38]	72.2	80.3	89.7
ATDOC+NOUN	73.0	80.6	90.0
ATDOC+PRONOUN	74.1	81.9	90.4

self-training method ATDOC [38] for experiments. We report results on three object recognition benchmarks in Table IX. We find both NOUN and PRONOUN consistently bring evident performance improvement when combined with ATDOC [38]. Especially, PRONOUN helps ATDOC [38] achieve the new-state-of-the-art average accuracy of 74.1% on *Office-Home*.

6) *Visualization*: For object recognition, we follow the *de facto* practice to study the t-SNE [70] visualization of aligned features generated by different UDA methods in Fig. 5. As expected, DANN-[f, p] fails to align multi-class distributions between domains. With effective conditioning, NOUN better aligns target samples to source clusters and learns more discriminative target features. With the prototype-based

conditioning strategy, PRONOUN alleviates the misalignment induced by inaccurate pseudo-labels. For semantic segmentation, we present some qualitative results in Fig. 6. Compared with the domain adversarial training involving only marginal features (FeatAdapt), our NOUN produces cleaner and more discriminative adapted segmentation results, which verifies the benefit of the effective conditioning for domain adaptation.

V. CONCLUSION

In this work, we target the UDA problem with multi-modal structures. We aim to develop simple also compact domain adversarial training solutions to address UDA. We first rethink the previous failure of the naive concatenation conditioning strategy. We find that the small norm of the concatenated output prediction makes the conditioning ineffective. We thus propose NOUN to ensure effective conditional domain alignment by enlarging the vector norm of the concatenated prediction. We provide a strong domain adversarial training method PRONOUN, which involves the novel prototype-based conditioning strategy to protect NOUN from inaccurate pseudo-labels. Extensive evaluations on various object recognition benchmarks justify the effectiveness of NOUN and PRONOUN over well-established UDA baselines. NOUN is also verified to be generic and competitive for

semantic segmentation UDA tasks. In the future, we aim to explore applications of PRONOUN to more challenging UDA tasks like semantic segmentation and object detection.

ACKNOWLEDGMENT

The authors would like to thank the reviewers and the associate editor for their valuable comments. The authors also would like to appreciate Jiashi Feng and Shuicheng Yan for their insightful discussions throughout this project and comments on the writing of this paper.

REFERENCES

- [1] M. Long, Y. Cao, J. Wang, and M. I. Jordan, "Learning transferable features with deep adaptation networks," in *Proc. ICML*, 2015, pp. 97–105.
- [2] K. Saito, Y. Ushiku, and T. Harada, "Asymmetric tri-training for unsupervised domain adaptation," in *Proc. ICML*, 2017, pp. 2988–2997.
- [3] J. Hoffman, D. Wang, F. Yu, and T. Darrell, "FCNs in the wild: Pixel-level adversarial and constraint-based adaptation," 2016, *arXiv:1612.02649*.
- [4] Y.-H. Tsai, W.-C. Hung, S. Schuler, K. Sohn, M.-H. Yang, and M. Chandraker, "Learning to adapt structured output space for semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7472–7481.
- [5] Y. Luo, P. Liu, T. Guan, J. Yu, and Y. Yang, "Significance-aware information bottleneck for domain adaptive semantic segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6778–6787.
- [6] Y. Luo, L. Zheng, T. Guan, J. Yu, and Y. Yang, "Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2507–2516.
- [7] Y. Luo, P. Liu, T. Guan, J. Yu, and Y. Yang, "Adversarial style mining for one-shot unsupervised domain adaptation," in *Proc. NeurIPS*, 2020, pp. 20612–20623.
- [8] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel method for the two-sample problem," *J. Mach. Learn. Res.*, vol. 1, pp. 1–10, May 2008.
- [9] B. Sun and K. Saenko, "Deep coral: Correlation alignment for deep domain adaptation," in *Proc. ECCV*, 2016, pp. 443–450.
- [10] W. Zellinger, T. Grubinger, E. Lughofer, T. Natschläger, and S. Saminger-Platz, "Central moment discrepancy (CMD) for domain-invariant representation learning," in *Proc. ICLR*, 2017.
- [11] P. Koniusz, Y. Tas, and F. Porikli, "Domain adaptation by mixture of alignments of second-order scatter tensors," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 7139–7148.
- [12] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. NeurIPS*, 2014, pp. 2672–2680.
- [13] Z. Pei, Z. Cao, M. Long, and J. Wang, "Multi-adversarial domain adaptation," in *Proc. AAAI*, 2018, pp. 3934–3941.
- [14] M. Long, Z. Cao, J. Wang, and M. I. Jordan, "Conditional adversarial domain adaptation," in *Proc. NeurIPS*, 2018, pp. 1647–1657.
- [15] V. K. Kurmi and V. P. Nambodiri, "Looking back at labels: A class based domain adaptation technique," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2019, pp. 1–8.
- [16] S. Cicek and S. Soatto, "Unsupervised domain adaptation via regularized conditional alignment," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 1416–1425.
- [17] L. Tran, K. Sohn, X. Yu, X. Liu, and M. Chandraker, "Gotta adapt'em all: Joint pixel and feature-level domain adaptation for recognition in the wild," in *Proc. CVPR*, 2019, pp. 2672–2681.
- [18] Y.-H. Chen, W.-Y. Chen, Y.-T. Chen, B.-C. Tsai, Y.-C.-F. Wang, and M. Sun, "No more discrimination: Cross city adaptation of road scene segmenters," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 1992–2001.
- [19] Y.-H. Tsai, K. Sohn, S. Schuler, and M. Chandraker, "Domain adaptation for structured output via discriminative patch representations," 2019, *arXiv:1901.05427*.
- [20] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *Proc. ICML*, 2015, pp. 1180–1189.
- [21] Y. Ganin *et al.*, "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2030–2096, 2016.
- [22] H. Shimodaira, "Improving predictive inference under covariate shift by weighting the log-likelihood function," *J. Statist. Planning Inference*, vol. 90, no. 2, pp. 227–244, 2000.
- [23] M. Dudík, S. J. Phillips, and R. E. Schapire, "Correcting sample selection bias in maximum entropy density estimation," in *Proc. NeurIPS*, 2006, pp. 323–330.
- [24] J. Huang, A. Gretton, K. Borgwardt, B. Schölkopf, and A. J. Smola, "Correcting sample selection bias by unlabeled data," in *Proc. NeurIPS*, 2007, pp. 601–608.
- [25] B. Gong, Y. Shi, F. Sha, and K. Grauman, "Geodesic flow kernel for unsupervised domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2012, pp. 2066–2073.
- [26] B. Fernando, A. Habrard, M. Sebban, and T. Tuytelaars, "Unsupervised visual domain adaptation using subspace adaptation," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 2960–2967.
- [27] B. Sun, J. Feng, and K. Saenko, "Return of frustratingly easy domain adaptation," in *Proc. AAAI*, 2016, pp. 2058–2065.
- [28] S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang, "Domain adaptation via transfer component analysis," *IEEE Trans. Neural Netw.*, vol. 22, no. 2, pp. 199–210, Feb. 2010.
- [29] M. Long, J. Wang, G. Ding, J. Sun, and P. S. Yu, "Transfer feature learning with joint distribution adaptation," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 2200–2207.
- [30] S. Herath, M. Harandi, and F. Porikli, "Learning an invariant Hilbert space for domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 3956–3965.
- [31] J. Liang, R. He, Z. Sun, and T. Tan, "Aggregating randomized clustering-promoting invariant projections for domain adaptation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 5, pp. 1027–1042, May 2018.
- [32] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2962–2971.
- [33] K. Saito, K. Watanabe, Y. Ushiku, and T. Harada, "Maximum classifier discrepancy for unsupervised domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3723–3732.
- [34] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell, "Deep domain confusion: Maximizing for domain invariance," 2014, *arXiv:1412.3474*.
- [35] M. Long, J. Wang, and M. I. Jordan, "Deep transfer learning with joint adaptation networks," in *Proc. ICML*, 2017, pp. 2208–2217.
- [36] D.-H. Lee *et al.*, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks," in *Proc. Workshop Challenges Represent. Learn. (ICML)*, 2013, p. 896.
- [37] Y. Li, L. Yuan, and N. Vasconcelos, "Bidirectional learning for domain adaptation of semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 6936–6945.
- [38] J. Liang, D. Hu, and J. Feng, "Domain adaptation with auxiliary target domain-oriented classifier," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 16632–16642.
- [39] J. Choi, M. Jeong, T. Kim, and C. Kim, "Pseudo-labeling curriculum for unsupervised domain adaptation," in *Proc. BMVC*, 2019, p. 67.
- [40] Y. Zou, Z. Yu, B. V. Kumar, and J. Wang, "Unsupervised domain adaptation for semantic segmentation via class-balanced self-training," in *Proc. ECCV*, 2018, pp. 297–313.
- [41] J. Liang, R. He, Z. Sun, and T. Tan, "Exploring uncertainty in pseudo-label guided unsupervised domain adaptation," *Pattern Recognit.*, vol. 96, Dec. 2019, Art. no. 106996.
- [42] J. Zhang, W. Li, and P. Ogunbona, "Joint geometrical and statistical alignment for visual domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1859–1867.
- [43] H. Tang and K. Jia, "Discriminative adversarial domain adaptation," in *Proc. AAAI*, 2020, pp. 5940–5947.
- [44] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1125–1134.
- [45] A. Odena, C. Olah, and J. Shlens, "Conditional image synthesis with auxiliary classifier gans," in *Proc. ICML*, 2017, pp. 2642–2651.
- [46] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [47] S. Roy, A. Siarohin, E. Sangineto, S. R. Buló, N. Sebe, and E. Ricci, "Unsupervised domain adaptation using feature-whitening and consensus loss," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 9471–9480.
- [48] R. Xu, G. Li, J. Yang, and L. Lin, "Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 1426–1435.

- [49] K. Saito, Y. Ushiku, T. Harada, and K. Saenko, "Adversarial dropout regularization," in *Proc. ICLR*, 2018.
- [50] X. Chen, S. Wang, M. Long, and J. Wang, "Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation," in *Proc. ICML*, 2019, pp. 1081–1090.
- [51] C.-Y. Lee, T. Batra, M. H. Baig, and D. Ulbricht, "Sliced wasserstein discrepancy for unsupervised domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 10285–10295.
- [52] H. Venkateswara, J. Eusebio, S. Chakraborty, and S. Panchanathan, "Deep hashing network for unsupervised domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 5018–5027.
- [53] X. Peng, B. Usman, N. Kaushik, J. Hoffman, D. Wang, and K. Saenko, "VisDA: The visual domain adaptation challenge," 2017, *arXiv:1710.06924*.
- [54] K. Saenko, B. Kulis, M. Fritz, and T. Darrell, "Adapting visual category models to new domains," in *Proc. ECCV*, 2010, pp. 213–226.
- [55] S. R. Richter, V. Vineet, S. Roth, and V. Koltun, "Playing for data: Ground truth from computer games," in *Proc. ECCV*, 2016, pp. 102–118.
- [56] M. Cordts *et al.*, "The cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3213–3223.
- [57] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez, "The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3234–3243.
- [58] T.-Y. Lin *et al.*, "Microsoft COCO: Common objects in context," in *Proc. ECCV*, 2014, pp. 740–755.
- [59] Y. Zhang, P. David, and B. Gong, "Curriculum domain adaptation for semantic segmentation of urban scenes," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2020–2030.
- [60] V. Nair and G. Hinton, "Rectified linear units improve restricted Boltzmann machines," in *Proc. 27th Int. Conf. Mach. Learn. (ICML)*, 2010, pp. 807–814.
- [61] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 834–848, Apr. 2017.
- [62] T.-H. Vu, H. Jain, M. Bucher, M. Cord, and P. Perez, "ADVENT: Adversarial entropy minimization for domain adaptation in semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2517–2526.
- [63] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," 2015, *arXiv:1511.06434*.
- [64] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *Proc. Workshop Deep Learn. Audio, Speech, Lang. Process. (ICML)*, 2013.
- [65] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *Proc. COMPSTAT*. Heidelberg, Germany: Physica-Verlag, 2010, pp. 177–186.
- [66] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.
- [67] P. Morerio, J. Cavazza, and V. Murino, "Minimal-entropy correlation alignment for unsupervised deep domain adaptation," in *Proc. ICLR*, 2018.
- [68] Z. Deng, Y. Luo, and J. Zhu, "Cluster alignment with a teacher for unsupervised domain adaptation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9944–9953.
- [69] W. Zhang, W. Ouyang, W. Li, and D. Xu, "Collaborative and adversarial network for unsupervised domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3801–3809.
- [70] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, Nov. 2008.



Dapeng Hu received the B.Sc. degree in electronic information science and technology from Nanjing University in 2017. He is currently pursuing the Ph.D. degree with the National University of Singapore. His research interests focus on domain adaptation, self-supervised learning, and computer vision.



Jian Liang received the B.Eng. degree in electronic information and technology from Xi'an Jiaotong University, China, in July 2013, and the Ph.D. degree in pattern recognition and intelligent systems from NLPR, CASIA, in January 2019. He was a Research Fellow with the National University of Singapore, Singapore, from June 2019 to April 2021. Currently, he joins NLPR as an Associate Professor. His research interests focus on transfer learning, pattern recognition, and computer vision.



Qibin Hou received the Ph.D. degree from the School of Computer Science, Nankai University, under Prof. Ming-Ming Cheng's supervision. Then, he worked at the National University of Singapore as a Research Fellow. He is currently an Associate Professor at Nankai University. His research interests include deep learning, image processing, and computer vision.



Hanshu Yan (Graduate Student Member, IEEE) received the B.Eng. and B.Sc. degrees from Beihang University (BUAA) in 2015 and the M.Sc. degree from the National University of Singapore (NUS) in 2017, where he is currently pursuing the Ph.D. degree with the ECE Department, working with Prof. Vincent Y. F. Tan and Dr. Jiashi Feng. His current research interests include robust and efficient machine learning algorithms.



Yunpeng Chen received the Ph.D. degree from the National University of Singapore (NUS), advised by an Assistant Professor Jiashi Feng and an Associate Professor Shuicheng Yan. He is currently the Algorithm Director at Meitu Inc., where he is working on machine learning and computer vision, including video understanding, image representation learning, and deep learning architecture engineering. Till now, he has published over 30 papers in top international journals and conferences, such as CVPR, NeurIPS, ECCV, ICCV, and IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, with over 2,700 citations. He also received the PREMIA Best Student Conference Paper Award. He is the winner of ILSVRC 2017 and ILSVRC 2015 and HIE 2020.