Universal Domain Adaptation via Compressive Attention Matching

Didi Zhu [*1], Yinchuan Li [*2], Junkun Yuan[1], Zexi Li[1], Yunfeng Shao[2], Kun Kuang[1], and Chao Wu [†1]

[1]Zhejiang University
[2]Huawei Noah's Ark Lab
[1]{didi_zhu, yuanjk, zexi.li, kunkuang, chao.wu}@zju.edu.cn
[2]{liyinchuan, shaoyunfeng}@huawei.com

Abstract

Universal domain adaptation (UniDA) aims to transfer knowledge from the source domain to the target domain without any prior knowledge about the label set. The challenge lies in how to determine whether the target samples belong to common categories. The mainstream methods make judgments based on the sample features, which overemphasizes global information while ignoring the most crucial local objects in the image, resulting in limited accuracy. To address this issue, we propose a Universal Attention Matching (UniAM) framework by exploiting the self-attention mechanism in vision transformer to capture the crucial object information. The proposed framework introduces a novel Compressive Attention Matching (CAM) approach to explore the core information by compressively representing attentions. Furthermore, CAM incorporates a residual-based measurement to determine the sample commonness. By utilizing the measurement, UniAM achieves domain-wise and category-wise Common Feature Alignment (CFA) and Target Class Separation (TCS). Notably, UniAM is the first method utilizing the attention in vision transformer directly to perform classification tasks. Extensive experiments show that UniAM outperforms the current state-of-the-art methods on various benchmark datasets.

1. Introduction

While deep neural networks have achieved remarkable success on various visual tasks [12, 22, 41, 17], their performance heavily relies on the assumption of independently and identically distributed (iid) training and test data [44]. However, this assumption is frequently violated due to the presence of domain shift in many real-world scenarios. Un-
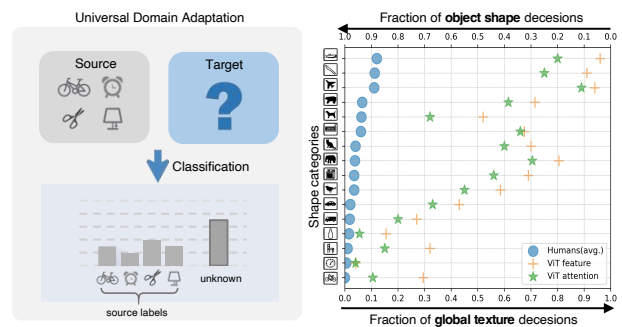

Figure 1: Left: Illustration of Universal Domain Adaptation. Right: Shape-bias Analysis. Plot shows shape-texture tradeoff for attention and feature in ViT and Humans.

supervised Domain adaptation (DA) [1] has emerged as a promising solution to address this limitation by adapting models trained on a source domain to perform well on an unlabeled target domain. Nevertheless, most existing favorable DA approaches assume that the label spaces in the source and target domains are identical, which may not always hold in practical scenarios. Partial Domain Adaptation (PDA) [4] and Open Set Domain Adaptation (OSDA) [34] have been proposed to handle cases where the label spaces in one domain include those in the other, but these still rely on prior knowledge on label set, limiting knowledge generalizing from one scenario to others. Universal domain adaptation (UniDA) [52] considers a more practical and challenging scenario where the relationship of label space between source and target domains is completely unknown i.e. with any number of common, source-private and target-private classes.

In UniDA, the primary objective is to develop a model capable of precisely categorizing target samples as one of the common classes or an "unknown" class as shown in Fig. 1 left. Existing UniDA methods aim to design a transferability criteria to detect common and private
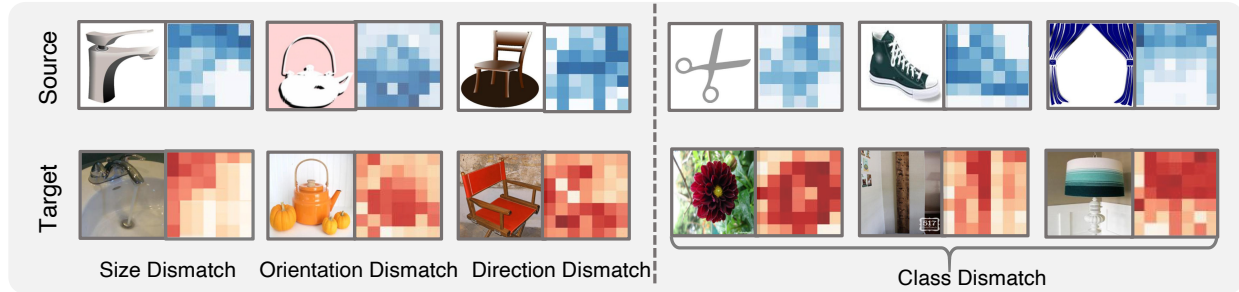

Figure 2: Attention Visualization on different domains.

classes solely based on the discriminability of deep features [6, 7, 8, 9, 13, 24, 26, 38, 39, 40, 52]. However, over-reliance on deep features can impede model adaptation performance, as they have a strong bias towards global information like texture rather than the essential object information like shape [15, 18], which is considered by humans as the most critical cue for recognition [25]. Fortunately, recent studies have demonstrated that vision transformer (ViT) [21] exhibits a stronger shape bias than Convolutional Neural Network (CNN) [32, 43]. As shown in Fig. 1 right, we confirmed that such strong object shape bias is mainly attributed to the self-attention mechanism, verified in a similar way as [15]. Figure 2 demonstrates the attention vectors of samples in different domains. Although we can leverage the attention to focus on more object parts, the *attention mismatch* problem may still exist due to domain shift, which refers to the attention vectors of same-class samples from different domains having some degree of the difference caused by potential variations in object size, orientation, and position across different domains. Attention mismatch can hinder the accurate classification of samples, especially when objects of different classes share similar sizes or positions. For example, in Figure 2, the kettle in the source domain and the flower in the target domain have more similar attention patterns. Therefore, the key challenge in utilizing attention is to effectively explore and leverage the object information embedded in attention while mitigating the negative impact of attention mismatch.

In this paper, we propose a novel Universal Attention Matching (UniAM) framework to address the UniDA problem by leveraging both the feature and attention information in a complementary way. Specifically, UniAM introduces a Compressive Attention Matching (CAM) approach to solve the attention mismatch problem implicitly by sparsely representing target attentions using source attention prototypes. This allows CAM to identify the most relevant attention prototype for each target sample and distinguish irrelevant private labels. Furthermore, a residual-based measurement is proposed in CAM to explicitly distinguish common and private samples across domains. By integrating attention information with features, we can mitigate the interference caused by domain shift and focus on label shift to some ex-

tent. With the guidance of CAM, the UniAM framework achieves domain-wise and category-wise common feature alignment (CFA) and target class separation (TCS). By using an adversarial loss and a source contrastive loss, CFA identifies and aligns the common features across domains, ensuring their consistency and transferability. On the other hand, TCS enhances the compactness of the target clusters, leading to better separation among all target classes. This is accomplished through a target contrastive loss, which encourages samples from the same target class to be closer together and farther apart from samples with other classes.

Main Contributions: (1) We propose the UniAM framework that comprehensively considers both attention and feature information, which allows for more accurate identification of common and private samples. (2) We validate the strong object bias of attention in ViT. To the best of our knowledge, we are the first to directly utilize attention in ViT for classification prediction. (3) We implicitly explore object information by sparsely reconstructing attention, enabling better common feature alignment (CFA) and target class separation (TCS). (4) We conduct extensive experiments to show that UniAM can outperform current state-of-the-art approaches.

2. Related Works

2.1. Universal Domain Adaptation

UniDA [52] does not require prior knowledge of label set relationship. To address this problem, UAN [52] proposes a criterion based on entropy and domain similarity to quantify sample transferability. CMU [13] follows this paradigm to detect open classes by setting the mean of three uncertain scores including entropy, consistency and confidence as a new measurement. Afterward, [24] proposes a real-time adaptive source-free UniDA method. In [38] and [26], clustering is developed to solve this problem. [28]. OVANet [39] employs a One-vs-All classifier for each class and decides known or unknown by using the output. Recent works [7, 8, 9, 6] Recent works have shifted their focus towards finding mutually nearest neighbor samples of target samples [7, 8, 9] or constructing relationships between target samples and source domain prototypes [9, 23].
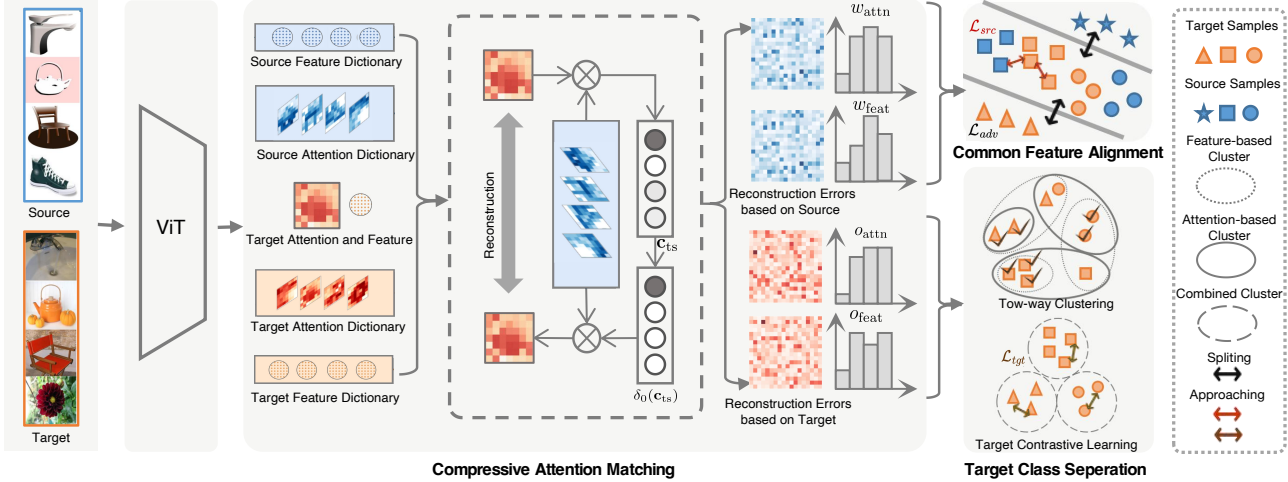
Figure 3: Illustration of the proposed UniAM framework. The framework consists of three integral components: Compressive Attention Matching (CAM), Common Feature Alignment (CFA) and Target Class Separation (TCS). At its core, CAM reconstructs all target attentions and features based on the source dictionary (with feature reconstruction omitted in the figure for simplicity), and attention and feature commonness scores w_{attn} and w_{feat} are computed from residual vectors. Then, domain- and category-wise CFA is achieved by minimizing $\mathcal{L}_{adv}$ and $\mathcal{L}_{src}$ guided by w_{attn} and w_{feat}. Similarly, o_{attn} and o_{feat} are obtained by reconstructing target attentions and features based on the target dictionary in CAM. TCS performs two-way clustering from both the attention and feature views and minimizes $\mathcal{L}_{tgt}$ to achieve effective separation of target classes.

2.2. Vision Transformer

Inspired by the success of Transformer [45] in the NLP field, many researchers have attempted to exploit it for solving computer vision tasks. One of the most pioneering works is Vision Transformer (ViT)[21], which decomposes input images into a sequence of fixed-size patches. Different from CNNs that rely on image-specific inductive bias, ViT takes the advantage of large-scale pre-training data and global context modeling on the entire images. Due to the outstanding performance of ViT, many approaches have been proposed based on it [42, 31, 47, 19, 50], such as Touvron et al. [42] propose DeiT, which introduces a distillation strategy specific to transformers to reduce computational costs. In general, ViT and its variants have achieved excellent results on many computer vision tasks, such as object detection [5, 55, 47], image segmentation [54, 48], and video understanding [16, 33], etc.

2.3. Sparse Representation Classification

Sparse Representation Classification (SRC) [49] and Collaborative Representation Classification (CRC)[53], along with their numerous extensions [30, 51, 11, 10], have been extensively investigated in the field of face recognition using single images and videos. These methods have demonstrated promising performance in the presence of occlusions and variations in illumination. By modeling the test data in terms of a sparse linear combination of a dictionary, SRC can capture non-linear relationships between features. Our UniAM is inspired by them but uses a novel measurement instead of a sparsity concentration index.

3. Problem Formulation and Preliminary

3.1. Problem Formulation

Denoting $\mathbb{X}$, $\mathbb{Y}$, $\mathbb{Z}$ as the input space, label space and latent space, respectively. Elements of $\mathbb{X}$, $\mathbb{Y}$, $\mathbb{Z}$ are noted as $\boldsymbol{x}$, y and $\boldsymbol{z}$. Let P_s and P_t be the source distribution and target distribution, respectively. We are given a labeled source domain $\mathbb{D}_s = \{\boldsymbol{x}_i, y_i\}_{i=1}^m$ and an unlabeled target domain $\mathbb{D}_t = \{\boldsymbol{x}_i\}_{i=1}^n$ are respectively sampled from P_s and P_t, where m and n denote the number of samples of source and target domains, respectively. Denote $\mathbb{L}_s$ and $\mathbb{L}_t$ as the label sets of the source and target domains, respectively. Let $\mathbb{L} = \mathbb{L}_s \cap \mathbb{L}_t$ be the common label set shared by both domains, while $\overline{\mathbb{L}}_s = \mathbb{L}_s \backslash \mathbb{L}$ and $\overline{\mathbb{L}}_t = \mathbb{L}_t \backslash \mathbb{L}$ be the label sets private to source and target domains, respectively. Denote $M = |\mathbb{L}_s|$ as the number of source labels. Universal domain adaptation aims to predict labels of target data in $\mathbb{L}$ while rejecting the target data in $\overline{\mathbb{L}}^t$ based on $\mathbb{D}_s$ and $\mathbb{D}_t$.

Our overall architecture consists of a ViT-based feature extractor, an adversarial domain classifier, and a label classifier. Suppose the function for learning embedding features is $G_f : \mathbb{X} \to \mathbb{Z} \in \mathbb{R}^{d_z}$ where d_z is the length of each feature vector, the discrimination function of the label classifier is $G_c : \mathbb{Z} \to \mathbb{Y} \in \mathbb{R}^M$, and the function of the domain classifier is $G_d : \mathbb{Z} \to \mathbb{R}^1$.

3

3.2. Preliminary

To start with, we provide an overview of the self-attention mechanism used in ViT. First, the input image $\boldsymbol{x}$ is divided into N fixed-size patches, which are linearly embedded into a sequence of vectors. Next, a special token called the class token is prepended to the sequence of image patches for classification. The resulting sequence of length $N + 1$ is then projected into three matrices: queries $\boldsymbol{Q} \in \mathbb{R}^{(N+1) \times d_k}$, keys $\boldsymbol{K} \in \mathbb{R}^{(N+1) \times d_k}$ and values $\boldsymbol{V} \in \mathbb{R}^{(N+1) \times d_v}$ with d_k and d_v being the length of each query and value vector, respectively. Then, $\boldsymbol{Q}$ and $\boldsymbol{K}$ are passed to the self-attention layer to compute the patch-to-patch similarity matrix $\boldsymbol{A}^{(N+1) \times (N+1)}$, which is given by

$$\boldsymbol{A} = \frac{\boldsymbol{Q}\boldsymbol{K}^\top}{\sqrt{d_k}}, \tag{1}$$

For ease of further processing, we flatten $\boldsymbol{A}$ into a vector $\boldsymbol{a} \in \mathbb{R}^{(N+1)^2 \times 1}$. It is worth noting that multiple attention heads are utilized in the self-attention mechanism. Each head outputs a separate attention, and the final attention is obtained by concatenating the vectors from all heads. As a result, the dimensionality of $\boldsymbol{a} \in \mathbb{R}^{d_a \times 1}$, where $d_a = N_H \times (N + 1)^2$ and N_H is the number of attention heads. The utilization of multiple heads allows the model to jointly attend to information from different feature subspaces at different positions.

Once the attention vector $\boldsymbol{a}$ is available, the corresponding k-th attention prototype $\boldsymbol{p}_k$ is calculated by averaging all attention vectors of samples in class k, which will be used in the subsequent matching process.

4. Proposed Methodology

4.1. Compressive Attention Matching

Since the attention mismatch problem exists due to domain shift mentioned in Section 1, how to effectively utilize the core object information and avoid interference from redundant information poses a challenge in applying attention to UniDA. To address this challenge, compressive attention matching (CAM) is proposed to capture the most informative object structures by sparsely representing target attentions. Define the attention dictionary in CAM as the collection of source attention prototypes for efficient matching, i.e., $\boldsymbol{P}_s = [\boldsymbol{p}_1^s, \boldsymbol{p}_2^s, \cdots, \boldsymbol{p}_M^s] \in \mathbb{R}^{d_a \times M}$. Definition 1 gives the definition of CAM.

Definition 1 (**Compressive Attention Matching**). *Given an attention vector $\boldsymbol{a}_t \in \mathbb{R}^{d_a \times 1}$ of the target sample $\boldsymbol{x}_t$ and a source attention dictionary $\boldsymbol{P}_s$, Compressive Attention Matching aims to match $\boldsymbol{a}_t$ with one prototype in $\boldsymbol{P}_s$ to determine its commonness, which is achieved by assuming that $\boldsymbol{a}_t$ can be approximated by a linear combination of $\boldsymbol{P}_s$:*

$$\boldsymbol{a}_t = \boldsymbol{P}_s \boldsymbol{c}_{ts}, \tag{2}$$

*where the coefficient vector $\boldsymbol{c}_{ts} \in \mathbb{R}^{M \times 1}$ satisfies a **sparsity constraint** in order to achieve a compressive representation. Based on $\boldsymbol{c}_{ts}$, $\boldsymbol{x}_t$ is regarded as belonging to common classes from an attention perspective when the following inequality is satisfied:*

$$w_{attn}(\boldsymbol{x}_t) < \delta.$$

$w_{attn}(\cdot)$ *indicates a measurement to evaluate the commonness of $\boldsymbol{x}_t$ which is defined later and δ is a threshold.*

Why Compressive Attention Matching is desirable? By enforcing sparsity on the coefficients in CAM, we can obtain a compressive representation of the attention vectors, which facilitates the extraction and utilization of low-dimensional structures embedded in high-dimensional attention vectors. In the context of UniDA, this compressive representation enables us to identify the most relevant attention prototype for each target sample and distinguish irrelevant private labels, which is crucial for achieving effective common and private class detection. Therefore, CAM with sparse coefficients plays a vital role in solving UniDA.

To solve Eq. 2 in CAM, the coefficient vector $\boldsymbol{c}_{ts}$ is estimated by:

$$\min_{\boldsymbol{c}_{ts}} \|\boldsymbol{a}_t - \boldsymbol{P}_s \boldsymbol{c}_{ts}\|_2^2 + \rho \|\boldsymbol{c}_{ts}\|_1, \tag{3}$$

where $\|\cdot\|_p$ denotes ℓ_p-norm. The ℓ_1-minimization term in Eq. 3 yields a sparse solution, which enforces that $\boldsymbol{c}_{ts}$ has only a small number of non-zero coefficients.

Then we can compute the class reconstruction error vector $\boldsymbol{r}_{ts} \in \mathbb{R}^M$ for each target sample using the sparse matrix $\boldsymbol{c}_{ts}$. The k-th entry of $\boldsymbol{r}_{ts}$ can be represented:

$$\boldsymbol{r}_{ts}(k) = \|\boldsymbol{a}_t - \boldsymbol{P}_s \delta_k(\boldsymbol{c}_{ts})\|_2, \quad k = 1, \dots, M, \tag{4}$$

where $\delta_k(\boldsymbol{c}_{ts})$ is a one-hot vector with the k-th entry in $\boldsymbol{c}_{ts}$ being non-zero while setting all other entries to zero. If $\boldsymbol{x}_t$ corresponds to a common class k, then the reconstruction error corresponding to class k. $\boldsymbol{r}_{ts}(k)$ should be much lower than that corresponding to the other classes. Conversely, if $\boldsymbol{x}_t$ belongs to a private class, the entire reconstruction error vector $\boldsymbol{r}_{ts}$ should be relatively small, without a significant difference between the errors corresponding to different classes.

As a result, the reconstruction error vector $\boldsymbol{r}_{ts}$ is a crucial component in CAM. It serves as the foundation for the design of the measurement $w_{\text{attn}}(\cdot)$ in Definition 1 called Attention Commonness Degree (ACD), defined as belows:

Definition 2 (**Attention Commonness Degree**). *Given the residual vector $\boldsymbol{r}_{ts}$ of $\boldsymbol{x}_t$, the ACD is defined as the difference between the average of non-matched errors and matched errors:*

$$w_{attn}(\boldsymbol{x}_t) = \text{non-match}(\boldsymbol{r}_{ts}) - \text{match}(\boldsymbol{r}_{ts}), \tag{5}$$

where $match(\boldsymbol{r}_{ts}) = \boldsymbol{r}_{ts}(\hat{y})$, $\hat{y} = \arg\min_k \boldsymbol{r}_{ts}(k)$ and non-match($\boldsymbol{r}_{ts}$) is the average of reconstruction errors excepting $\hat{y}$.

Remark 1. *ACD measures the degree of commonness for a target sample $\boldsymbol{x}_t$, which represents the probability of belonging to common classes. A higher ACD value indicates a larger difference between non-matched and matched errors, suggesting the presence of an attention prototype similar to $\boldsymbol{x}_t$, and consequently, a higher degree of sample commonness. Conversely, a smaller ACD value implies a similar reconstruction error between $\boldsymbol{x}_t$ and all source prototypes, indicating a lower degree of sample commonness and a higher degree of privateness.*

To complement the attention information, we retain features that reflect global information. The target feature $\boldsymbol{z}_t$ can be also represented by the linear span of source feature prototypes $\boldsymbol{Q}_s = [\boldsymbol{q}_1^s, \boldsymbol{q}_2^s, \cdots, \boldsymbol{q}_M^s]$, i.e., $\boldsymbol{z}_t = \boldsymbol{Q}_s \boldsymbol{c}_{ts}$. The corresponding residual vector $\boldsymbol{r}'_{ts}(k)$ is computed based on $\boldsymbol{c}_{ts}$. The Feature Commonness Degree (FCD) can be defined as $w_{\text{feat}}(\boldsymbol{x}_t) = \text{non-match}(\boldsymbol{r}'_{ts}) - \text{match}(\boldsymbol{r}'_{ts})$..

It is worth noting that by replacing P_s in Definition 2 with the target dictionary P_t, we can obtain compressive representations of target attentions towards P_t. This leads to a score similar to that in Definition 2, denoted as o_{attn}. The same goes for o_{feat}. These scores can facilitate determining the probability that two target samples belong to the same class, more details will be provided in Section 4.3.

In summary, both attention and feature characteristics are important factors that affect the perception of similarity between different categories. Attention captures the structural properties of objects, while feature captures the appearance properties of the global images. Therefore, we can achieve a more comprehensive and accurate private class detection model that takes into account both object information and global information.

4.2. Common Feature Alignment

To identify and align the common class features across domains, we propose a domain-wise and category-wise Common Feature Alignment (CFA) technique, which considers both attention and feature information.

Domain-wise Alignment. To achieve domain-wise alignment, we first propose a residual-based transferability score d_{t} measuring the probability that the target sample belongs to the common classes, which can be summarized as:

$$w_{\text{t}} = \lambda w_{\text{attn}} + (1 - \lambda) w_{\text{feat}}, \tag{6}$$

where λ is a hyperparameter balancing their contribution. Meanwhile, to measure the probability that the source sample $\boldsymbol{x}_s$ with label j belongs to the common label set, we compute w_j^s with the sum of all target samples' attention

and feature reconstruction errors respectively, i.e.

$$w_s^j = \lambda \sigma(\boldsymbol{r}_{ts})_j + (1 - \lambda) \sigma(\boldsymbol{r}'_{ts})_j \tag{7}$$

where $\boldsymbol{r}_{ts}^i$ indicates the reconstruction error of the i-th target sample. The operator $\sigma[\cdot]$ refers to the normalization sum of all target attention or feature residual vectors. A larger value of w_j^s indicates a higher probability that the source label j belongs to the common label set, while lower values suggest that it is more likely to be a source private label. It is worth noting that samples with the same category label in the source domain are assigned the same weight.

Based on the above two weights, we can derive a domain-wise adversarial loss that aligns the common classes across domains as follows:

$$\begin{aligned}\mathcal{L}_{\text{adv}} = \ &\mathbb{E}_{\boldsymbol{x}_s \in \mathbb{D}_S} \left[w_s \cdot \log\left(1 - G_d\left(\boldsymbol{z}_s\right)\right) \right] \\ &+ \mathbb{E}_{\boldsymbol{x}_t \in \mathbb{D}_T} \left[w_t \cdot \log\left(G_d\left(\boldsymbol{z}_t\right)\right) \right],\end{aligned} \tag{8}$$

In addition, to avoid being interfered with the knowledge of source private samples, we employ an indicator as the weight for the weighted cross-entropy loss $\mathcal{L}_{\text{cls}}$ for the source domain, as shown below:

$$\mathcal{L}_{\text{cls}} = -\mathbb{E}_{(\boldsymbol{x}_s, y_s) \in \mathbb{D}_S} \mathbb{1}_{w_s^y \geqslant \alpha} l_{ce}(y, G_c(G_f(\boldsymbol{x}_s))), \tag{9}$$

where l_{ce} is the standard cross-entropy loss.

Category-wise Alignment. To enhance the source discriminability and align the common features from a category-wise perspective across domains, we propose a contrastive common feature alignment method. In order to quantify the likelihood that the source sample $\boldsymbol{x}_i^s$ and the sample $\boldsymbol{x}_j$ belong to the same category y_i, we design a category-wise target score $w_{i,j}$. If $\boldsymbol{x}_j$ is a source sample, we use its ground truth label y_j to determine if it's a positive or negative example, i.e., $w_{i,j} = \begin{cases} 1, & \text{if } y_j = y_i \\ 0, & \text{if } y_j \neq y_i \end{cases}$. If $\boldsymbol{x}_j$ is a target sample, we estimate the probability that it belongs to y_i based on the residual vectors $\boldsymbol{r}_{ts}$ and $\boldsymbol{r}'_{ts}$. These two vectors can be seen as two vanilla prediction probability vectors and the corresponding pseudo-labels $\hat{y}_j$ and $\hat{y}'_j$ can be obtained by argmin operation. To give a more reliable estimation, the soft label of $\boldsymbol{x}_j$ is determined by these two pseudo-labels together. Specifically, $w_{i,j}$ is set to 1 when both $\hat{y}_j$ and $\hat{y}'_j$ are equal to y_i and set to 0 when both of them are not equal to y_i. The soft label is computed as below when only one of the predictions is y_i:

$$w_{i,j} = \begin{cases} \lambda \cdot w_{\text{attn}}/w_t, & \text{if } \hat{y}'_j \neq \hat{y}_j = y_i, \\ (1 - \lambda) \cdot w_{\text{feat}}/w_t, & \text{if } \hat{y}_j \neq \hat{y}'_j = y_i. \end{cases} \tag{10}$$

Thus the category-wise common feature alignment can be improved by minimizing the source contrastive loss $\mathcal{L}_{\text{src}}$:

$$\mathcal{L}_{\text{src}} = -\mathbb{E}_{\boldsymbol{x}_i \in \mathbb{D}_S, \boldsymbol{x}_j \in \mathbb{D}_{S \cup T}} w_{i,j} l(\boldsymbol{z}_i, \boldsymbol{z}_j) \tag{11}$$

with

$$l(\boldsymbol{z}_i, \boldsymbol{z}_j) = \frac{\exp\left(\boldsymbol{z}_i \boldsymbol{z}_j / \tau\right)}{\sum_{k=1}^{m+n} \exp\left(\boldsymbol{z}_i \boldsymbol{z}_k / \tau\right)}.$$

5

Table 1: H-score (%) on **Office-31** and **DomainNet**

Method		**Office-31**							**DomanNet**						
		A2W	D2W	W2D	A2D	D2A	W2A	Avg	P2R	R2P	P2S	S2P	R2S	S2R	Avg
ResNet [17]	ResNet50	47.92	54.94	55.60	49.78	48.48	48.96	50.94	30.06	28.34	26.95	26.95	26.89	29.74	28.15
DANN [14]		48.82	52.73	54.87	50.18	47.69	49.33	50.60	31.18	29.33	27.84	27.84	27.77	30.84	29.13
OSBP [40]		50.23	55.53	57.20	51.14	49.75	50.16	52.34	33.60	33.03	30.55	30.53	30.61	33.65	32.00
UAN [52]		58.61	70.62	71.42	59.68	60.11	60.34	63.46	41.85	43.59	39.06	38.95	38.73	43.69	40.98
CMU [13]		67.33	79.32	80.42	68.11	71.42	72.23	50.78	52.16	45.12	44.82	45.64	50.97	48.25	73.14
DCC [26]		78.54	79.29	88.58	88.50	70.18	75.87	80.16	56.90	50.25	43.66	44.92	43.31	56.15	49.20
OVANet [39]		79.45	95.43	94.35	85.67	80.43	84.23	86.59	56.0	51.7	47.1	47.4	44.9	57.2	50.7
UniOT [6]		89.16	98.93	96.87	86.35	89.85	88.08	91.54	59.30	47.79	51.79	46.81	48.32	58.25	52.04
OVANet*	ViT	87.75	93.14	85.72	82.96	**92.67**	91.25	88.92	71.24	61.14	51.28	55.30	47.51	66.48	58.83
UniOT*		96.35	99.13	99.43	88.40	89.67	**93.81**	94.47	72.40	59.47	49.30	56.86	47.38	69.43	59.14
Ours		**95.46**	**99.62**	**99.81**	**95.28**	92.35	93.23	**95.95**	**73.87**	**60.89**	**52.31**	**59.98**	**51.41**	**70.68**	**61.52**

Table 2: H-score (%) on **Office-Home** and **VisDA2017**

Method		**Office-Home**													**VisDA**
		Ar2Cl	Ar2Pr	Ar2Rw	Cl2Ar	Cl2Pr	Cl2Rw	Pr2Ar	Pr2Cl	Pr2Rw	Rw2Ar	Rw2Cl	Rw2Pr	Avg	S2R
ResNet [17]	ResNet50	44.65	48.04	50.13	46.64	46.91	48.96	47.47	43.17	50.23	48.45	44.76	48.43	47.32	25.44
DANN [14]		42.36	48.02	48.87	45.48	46.47	48.37	45.75	42.55	48.70	47.61	42.67	47.40	46.19	25.65
OSBP [40]		39.59	45.09	46.17	45.70	45.24	46.75	45.26	40.54	45.75	45.08	41.64	46.90	44.48	27.31
UAN [52]		51.64	51.70	54.30	61.74	57.63	61.86	50.38	47.62	61.46	62.87	52.61	65.19	56.58	30.47
CMU [13]		56.02	56.93	59.15	66.95	64.27	67.82	54.72	51.09	66.39	68.24	57.89	69.73	61.60	34.64
DCC [26]		57.97	54.05	58.01	74.64	70.62	77.52	64.34	**73.60**	74.94	80.96	**75.12**	80.38	70.18	43.02
OVANet [39]		62.81	75.54	78.59	70.72	68.78	75.03	71.27	58.64	80.52	76.09	64.13	78.91	71.75	53.10
UniOT [6]		67.27	80.54	86.03	73.51	77.33	84.28	75.54	63.33	85.99	77.77	65.37	81.92	76.57	57.32
OVANet*	ViT	58.09	86.06	89.38	81.86	81.03	86.22	84.49	57.06	88.54	83.67	57.32	86.67	77.45	56.98
UniOT*		63.77	**88.19**	90.23	74.99	81.02	84.55	78.91	61.29	87.60	82.38	63.70	88.30	78.40	63.25
Ours		**72.04**	87.07	**90.67**	**80.30**	**82.39**	79.81	**85.02**	68.35	**88.98**	**85.44**	72.11	86.12	**81.68**	**65.18**

4.3. Target Class Seperation

To better distinguish the common and private classes in the target domain, we propose a Target Class Separation (TCS) technique. As mentioned in Section 4.1, we can construct a CAM problem based on the target attention dictionary P_t. As the target labels are unknown, $P_t = [p_1^t, p_2^t, \cdots, p_K^t] \in \mathbb{R}^{d_a \times K}$ is initialized by performing traditional K-means algorithm on target attentions $\{a_t^i\}_{i=1}^n$, where K is pre-defined. In the subsequent attention clustering process, we calculate the residual vector $r_{tt} \in \mathbb{R}^K$ and use it as a metric for measuring the distance between samples, which allows for dynamically update P_t. Iteratively refining P_t makes it more reliable and discriminative. Meanwhile, the feature clustering based on the target feature dictionary $Q_t = [q_1^t, q_2^t, \cdots, q_K^t] \in \mathbb{R}^{d_z \times K}$ is also performed. After the two-way clustering, each target sample x_i^t is assigned two cluster indexes $\hat{c}_i$ and $\hat{c}_i'$ from attention and feature view like Fig. **??** depicted. The final soft pseudo label $o_{c,i}$ determining whether x_i belong to the c-th cluster is obtained based on these two cluster indexes, similar to $w_{i,j}$. Based on $o_{i,j}$, the target contrastive loss $\mathcal{L}_{\text{tgt}}$ is computed as follows:

$$\mathcal{L}_{\text{tgt}} = -\mathbb{E}_{x_i \in \mathbb{D}_T, x_j \in \mathbb{D}_T} o_{i,j} l(z_i, z_j), \quad (12)$$

with

$$l(z_i, z_j) = \frac{\exp(z_i z_j / \tau)}{\sum_{k=1}^n \exp(z_i z_k / \tau)}.$$

where $o_{i,j} = o_{c,i} \cdot o_{c,j}$ is the probability weight determining whether x_i and x_j belong to the same cluster c. By minimizing $\mathcal{L}_{\text{tgt}}$, we can enhance the compactness of target clusters making a better separation among target classes.

4.4. Overall Framework

Overall, our framework is jointly optimized with four terms, i.e., cross-entropy loss $\mathcal{L}_{\text{cls}}$, adversarial loss $\mathcal{L}_{\text{adv}}$, source and target contrastive loss $\mathcal{L}_{\text{src}}$ and $\mathcal{L}_{\text{tgt}}$ as shown in Fig. 3,

$$\max_{G_d} \min_{G_f, G_c} \mathcal{L}_{\text{cls}} + \eta_1 \mathcal{L}_{\text{src}} + \eta_2 \mathcal{L}_{\text{tgt}} - \mathcal{L}_{\text{adv}}, \quad (13)$$

where η_1 and η_2 are set as 0.5 to balance each loss component. In the testing phase, given each input target sample x_t, we compute w_t in (6). For those samples that satisfy $w_t < \beta$ are assigned with the predicted source class, where β is a validated threshold. Otherwise, the samples are marked as unknown.
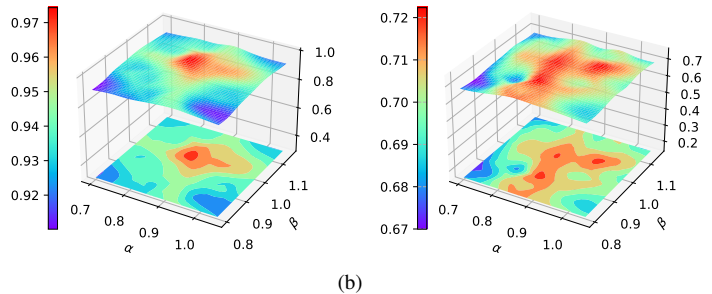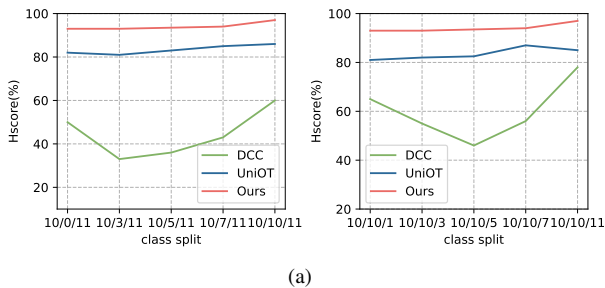
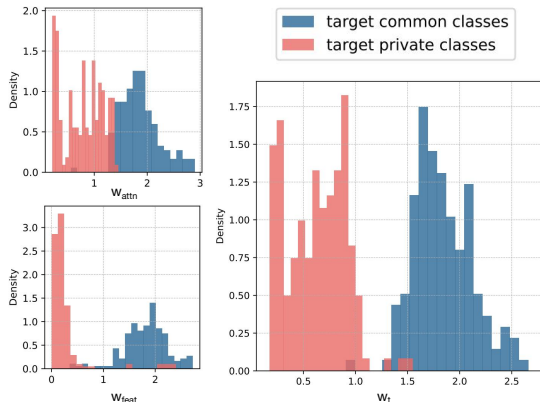Figure 4: (a) Effectiveness on different label set relationships. (b) Effectiveness of varying decision threshold α and β.

Figure 5: Qualitative Analysis.

Table 3: Evaluation of the effectiveness of UniAM

Method	A2W	D2W	W2D	A2D	D2A	W2A	Avg
w/ $\mathcal{L}_{adv}$	92.12	93.37	99.49	93.54	88.32	92.77	93.27
w/ $\mathcal{L}_{src}$	89.61	98.58	99.57	91.65	91.35	92.73	93.91
w/ $\mathcal{L}_{tgt}$	93.26	98.27	99.78	92.10	89.21	89.97	93.76
w/o $\mathcal{L}_{adv}$	94.69	98.10	99.78	**96.12**	92.05	92.35	95.51
w/o $\mathcal{L}_{src}$	93.78	98.43	99.78	94.69	91.94	93.47	95.34
w/o $\mathcal{L}_{tgt}$	90.76	98.20	99.78	92.41	91.69	93.13	94.32
w/o w_{attn}	94.76	97.54	97.26	94.31	91.22	90.69	94.30
Ours	**95.46**	**99.62**	**99.81**	95.28	**92.35**	**93.23**	**95.95**

5. Experiment Results

5.1. Experimental Setup

Datasets. We perform experiments on **Office-31** [37], **Office-Home** [46], **VisDA2017** [36] and **DomainNet** [35] datasets. **Office-31** consists of three domains: Amazon (A), DSLR (D) and Webcam (W). Each domain contains 31 categories. **Office-Home** is a dataset made up of 65 different categories from four domains: Artistic (Ar), Clipart (Cl), Product (Pr) and Real-world images (Rw). **VisDA2017** is a dataset with a single source and target domain testing the ability to perform transfer learning from synthetic images to natural images. The dataset has 12 categories in each domain. **DomainNet** is by far the largest domain adaptation dataset with about 0.6 million images. It consists of six distinct domains: Clipart (C), Infograph (I), Painting (P), Quick-draw (Q), Real (R) and Sketch (S) across 345 classes. We conduct experiments on three subsets from it, i.e., Painting (P), Real (R), and Sketch (S). For a fair comparison, we follow the same dataset split as [52] for the first three dataset and [13] for the last dataset.

Evaluation Protocols. We evaluate all methods using H-score [13]. H-score is the harmonic mean of the accuracy of common classes and the accuracy of the "unknown" classes,

which can make a trade-off between the accuracy of known and unknown classes.

Implementation Details. The method is implemented in Pytorch using a ViT-base model with 16×16 input patch size (or ViT-B/16) [21], pretrained on ImageNet [12]), as the backbone feature extractor. More information about dataset and implementation details are presented in Appendix.

5.2. Comparison Results

The experimental results for the Office-31, Office-Home, VisDA2017, and DomainNet datasets are presented in Table 1 and Table 2, which demonstrate that our proposed UniAM framework outperforms the state-of-the-art approaches in all benchmarks, as evaluated by the H-score metric. Additionally, to ensure a fair comparison, we conducted experiments by replacing the backbone of OVANet and UniOT with ViT, marked as ⋆. The proposed method consistently surpasses these ViT-based methods by a large margin. This indicates that our approach does not solely rely on using ViT as the backbone, but rather it fully exploits the advantage of the attention mechanism in ViT for UniDA tasks to achieve such superior performance.
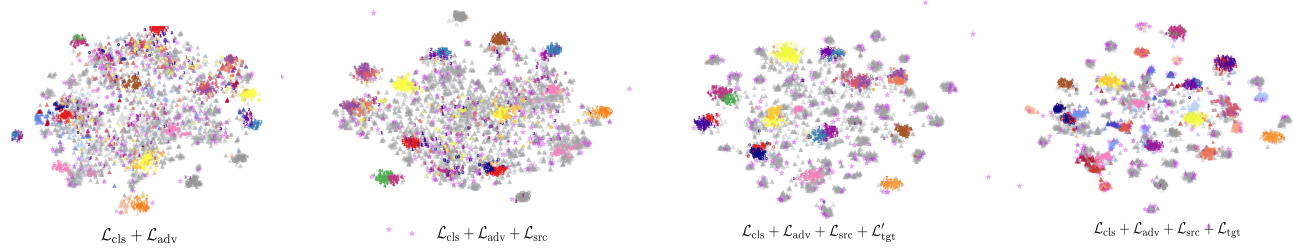
$$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{adv}} \qquad \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{adv}} + \mathcal{L}_{\text{src}} \qquad \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{adv}} + \mathcal{L}_{\text{src}} + \mathcal{L}'_{\text{tgt}} \qquad \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{adv}} + \mathcal{L}_{\text{src}} + \mathcal{L}_{\text{tgt}}$$

Figure 6: Feature visualization of target domain with different losses and whether attention is incorporated into a loss. $\mathcal{L}'_{\text{tgt}}$ and $\mathcal{L}_{\text{tgt}}$ refer to using only w_{feat} to calculate the loss weight $w_{i,j}$ and using both w_{attn} and w_{feat} together to calculate $w_{i,j}$.

5.3. Analysis on Different Label Set Relationships

Varying size of target private label set $\overline{\mathbb{L}}^t$. To explore the performance of our method under different class splitting settings with OVANet, UniOT and UniOT*, we fix $\mathbb{L}^s$, $\mathbb{L}$ and change $\overline{\mathbb{L}}^t$ on task A→W in Office-31 dataset. As shown in Fig. 4 (a) left, our method consistently outperforms all comparison methods under different $\overline{\mathbb{L}}^t$, proving that our method is effective and robust for different $\overline{\mathbb{L}}^t$. As $\overline{\mathbb{L}}^t$ increases, meaning there are many open classes, our method outperforms other methods by a large margin, demonstrating that our method is superior in detecting open classes.

Varying size of common label set $\mathbb{L}$ We fix $\mathbb{L}^s$ and $\mathbb{L}^t$ and varying $\mathbb{L}$ on task A→W in Office-31 dataset. We let $\overline{\mathbb{L}}^s$, $\overline{\mathbb{L}}^t$ to keep 10 and 11 and vary $\mathbb{L}$ from 0 to 10 . In particular, all target data should be marked as "unknown" when the source and target domains do not overlap on label sets. As shown in Fig. 4 (a) right, our method consistently outperforms previous methods on all sizes of $\mathbb{L}$, indicating that our method can detect open classes more effectively.

5.4. Analysis on Our Method

Effectiveness of different losses. As there are three losses excluding classification loss in our method, we conduct another experiment to verify the effectiveness of each loss and any combination of them on Office-31 dataset. As shown in the first six rows of Table. 3, the results indicate that the use of any single loss function or a combination of any two loss functions can lead to a decrease in performance to some extent, with a performance drop of 3%-4% observed when using a single loss. These findings demonstrate the effectiveness of our proposed method.

Effectiveness of attention in ViT. To demonstrate the indispensable roles of attention in ViT in the proposed transferability criteria, we introduce a variant denoted as w/o w_{attn}, which performs sparse reconstruction only on the features. Compared with our method in the last two rows of Table 3, the average performance drop of w/o w_{attn} 2.40%. It indicates that the attention mechanism does play an effec-

tive role in attention enhancement on the basis of features during the process of common category detection.

Qualitative Analysis. As shown in Fig. 5, the three probability density histograms visualize partially w_{attn}, w_{text}, and their weighted sum w_t in A→W task on Office. From Fig. 5, it can be observed that using w_{attn} or w_{text} alone can partially distinguish common samples (colored in blue) and private samples (colored in red), but each has its limitations. w_{attn} is prone to confusion at the boundary, while w_{text} has some outliers, such as private samples with extremely high values and common samples with extremely low values. By combining them together, these two limitations can be effectively alleviated. The weighted sum w_t can result in clearer boundaries between private and common samples, and the outliers are reduced.

Feature visualization. We use t-SNE to visualize the learned target features for Pr→Rw of Office-Home. As shown in Fig. 6, the gray dots represent private samples, while the non-gray dots represent common samples, and their colors indicate their ground-truth classes. Fig. 6 (a)-(c) shows that $\mathcal{L}_{\text{src}}$ increased the distance between common and private categories while all target-private samples are treated as a single class, and $\mathcal{L}_{\text{tgt}}$ improved the discriminability of the target private classes. Especially, Fig. 6 (d) validates that UniOT learns a better target representation introducing attention as a guide for attention enhancement can further improve the discriminability in the target domain by bringing same-class samples closer and pushing different-class samples farther away.

Sensitivity to decision threshold. We investigate the sensitivity of thresholds α and β, which are used to determine whether source and target samples belong to common classes respectively. The analysis was done in A→D on Office-31 and Ar→Cl on Office-Home. As shown in Fig. 4 (b), the H-score has little variance. α varies in a reasonable range [0.7, 1.0] while β varies in a range [0.8.1.1]. These demonstrate that our method is robust to α and β.

6. Conclusions

In conclusion, we proposed UniAM, a Compressive Attention Matching framework that leverages the self-attention mechanism in vision transformers to capture the most useful information for universal domain adaptation. We introduced a compressive reconstruction module and a residual-based transferability criterion to achieve domain alignment. UniAM is the first method to utilize attention in vision transformers directly for classification tasks, and our experiments show that it outperforms the current state-of-the-art UniDA methods on various benchmark datasets.

References

[1] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010. 1

[2] Silvia Bucci, Mohammad Reza Loghmani, and Tatiana Tommasi. On the effectiveness of image rotation for open set domain adaptation. In *European Conference on Computer Vision*, pages 422–438. Springer, 2020. 12

[3] Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Michael I Jordan. Partial transfer learning with selective adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2724–2732, 2018. 12

[4] Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 135–150, 2018. 1, 12

[5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 3

[6] Wanxing Chang, Ye Shi, Hoang Duong Tuan, and Jingya Wang. Unified optimal transport framework for universal domain adaptation. *arXiv preprint arXiv:2210.17067*, 2022. 2, 6, 12

[7] Liang Chen, Qianjin Du, Yihang Lou, Jianzhong He, Tao Bai, and Minghua Deng. Mutual nearest neighbor contrast and hybrid prototype self-training for universal domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6248–6257, 2022. 2

[8] Liang Chen, Yihang Lou, Jianzhong He, Tao Bai, and Minghua Deng. Evidential neighborhood contrastive learning for universal domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6258–6267, 2022. 2

[9] Liang Chen, Yihang Lou, Jianzhong He, Tao Bai, and Minghua Deng. Geometric anchor correspondence mining with uncertainty modeling for universal domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16134–16143, 2022. 2, 12

[10] Yi-Chen Chen, Vishal M Patel, P Jonathon Phillips, and Rama Chellappa. Dictionary-based face recognition from video. In *Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part VI 12*, pages 766–779. Springer, 2012. 3

[11] Yuejie Chi and Fatih Porikli. Classification and boosting with multiple collaborative representations. *IEEE transactions on pattern analysis and machine intelligence*, 36(8):1519–1531, 2013. 3

[12] Jia Deng. A large-scale hierarchical image database. *Proc. of IEEE Computer Vision and Pattern Recognition, 2009*, 2009. 1, 7, 12

[13] Bo Fu, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Learning to detect open classes for universal domain adaptation. In *European Conference on Computer Vision*, pages 567–583. Springer, 2020. 2, 6, 7, 12

[14] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. 6, 12

[15] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*, 2018. 2

[16] Rohit Girdhar, Joao Carreira, Carl Doersch, and Andrew Zisserman. Video action transformer network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 244–253, 2019. 3

[17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 6, 12

[18] Katherine Hermann, Ting Chen, and Simon Kornblith. The origins and prevalence of texture bias in convolutional neural networks. *Advances in Neural Information Processing Systems*, 33:19000–19015, 2020. 2

[19] Zilong Huang, Youcheng Ben, Guozhong Luo, Pei Cheng, Gang Yu, and Bin Fu. Shuffle transformer: Rethinking spatial shuffle for vision transformer. *arXiv preprint arXiv:2106.03650*, 2021. 3

[20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 12

[21] Alexander Kolesnikov, Alexey Dosovitskiy, Dirk Weissenborn, Georg Heigold, Jakob Uszkoreit, Lucas Beyer, Matthias Minderer, Mostafa Dehghani, Neil Houlsby, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. 2021. 2, 3, 7, 12

[22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 1

[23] Jogendra Nath Kundu, Suvaansh Bhambri, Akshay Kulkarni, Hiran Sarkar, Varun Jampani, and R Venkatesh Babu. Subsidiary prototype alignment for universal domain adaptation. *arXiv preprint arXiv:2210.15909*, 2022. 2

[24] Jogendra Nath Kundu, Naveen Venkat, R Venkatesh Babu, et al. Universal source-free domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4544–4553, 2020. 2

[25] Barbara Landau, Linda B Smith, and Susan S Jones. The importance of shape in early lexical learning. *Cognitive development*, 3(3):299–321, 1988. 2

[26] Guangrui Li, Guoliang Kang, Yi Zhu, Yunchao Wei, and Yi Yang. Domain consensus clustering for universal domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9757–9766, 2021. 2, 6, 12

[27] Jian Liang, Yunbo Wang, Dapeng Hu, Ran He, and Jiashi Feng. A balanced and uncertainty-aware approach for partial domain adaptation. In *European Conference on Computer Vision*, pages 123–140. Springer, 2020. 12

[28] Omri Lifshitz and Lior Wolf. A sample selection approach for universal domain adaptation. *arXiv preprint arXiv:2001.05071*, 2020. 2

[29] Hong Liu, Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Qiang Yang. Separate to adapt: Open set domain adaptation via progressive separation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2927–2936, 2019. 12

[30] Xiaofeng Liu, Zhaofeng Li, Lingsheng Kong, Zhihui Diao, Junliang Yan, Yang Zou, Chao Yang, Ping Jia, and Jane You. A joint optimization framework of low-dimensional projection and collaborative representation for discriminative classification. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 1493–1498. IEEE, 2018. 3

[31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 3

[32] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *Advances in Neural Information Processing Systems*, 34:23296–23308, 2021. 2

[33] Daniel Neimark, Omri Bar, Maya Zohar, and Dotan Asselmann. Video transformer network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3163–3172, 2021. 3

[34] Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *Proceedings of the IEEE international conference on computer vision*, pages 754–763, 2017. 1

[35] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415, 2019. 7

[36] Xingchao Peng, Ben Usman, Neela Kaushik, Dequan Wang, Judy Hoffman, and Kate Saenko. Visda: A synthetic-to-real benchmark for visual domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 2021–2026, 2018. 7

[37] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European conference on computer vision*, pages 213–226. Springer, 2010. 7

[38] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, and Kate Saenko. Universal domain adaptation through self supervision. *Advances in neural information processing systems*, 33:16282–16292, 2020. 2, 12

[39] Kuniaki Saito and Kate Saenko. Ovanet: One-vs-all network for universal domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9000–9009, 2021. 2, 6, 12

[40] Kuniaki Saito, Shohei Yamamoto, Yoshitaka Ushiku, and Tatsuya Harada. Open set domain adaptation by backpropagation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 153–168, 2018. 2, 6, 12

[41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1

[42] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. 3

[43] Shikhar Tuli, Ishita Dasgupta, Erin Grant, and Thomas L Griffiths. Are convolutional neural networks or transformers more like human vision? *arXiv preprint arXiv:2105.07197*, 2021. 2

[44] Vladimir Vapnik. Principles of risk minimization for learning theory. *Advances in neural information processing systems*, 4, 1991. 1

[45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3

[46] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5018–5027, 2017. 7

[47] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 568–578, 2021. 3

[48] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-end video instance segmentation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8741–8750, 2021. 3

[49] John Wright, Allen Y Yang, Arvind Ganesh, S Shankar Sastry, and Yi Ma. Robust face recognition via sparse representation. *IEEE transactions on pattern analysis and machine intelligence*, 31(2):210–227, 2008. 3

[50] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing

convolutions to vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22–31, 2021. 3

[51] Yang Wu, Vansteenberge Jarich, Masayuki Mukunoki, and Michihiko Minoh. Collaborative representation for classification, sparse or non-sparse? *arXiv preprint arXiv:1403.1353*, 2014. 3

[52] Kaichao You, Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Universal domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2720–2729, 2019. 1, 2, 6, 7, 12

[53] Lei Zhang, Meng Yang, and Xiangchu Feng. Sparse representation or collaborative representation: Which helps face recognition? In *2011 International conference on computer vision*, pages 471–478. IEEE, 2011. 3

[54] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6881–6890, 2021. 3

[55] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 3

G. More Experiment Details

To demonstrate the effectiveness and generalization ability of our proposed UniAM framework, we further evaluate it in the sub-cases of UniDA (i.e., PDA [4] and OSDA [40]) on **Office-31** and **Office-Home** dataset and compare it with other existing methods. In this section, we will supplement the main text in four aspects: dataset split, evaluation protocols, comparison baselines and implementation details.

G.1. Dataset Split

The experimental procedure was carried out following the dataset split protocols introduced by previous works [13, 52]. As detailed in Section 3.1 in the main paper, we consider $\mathbb{L} = \mathbb{L}_s \cap \mathbb{L}_t$ as the shared set of labels between the domains while defining $\overline{\mathbb{L}}_s = \mathbb{L}_s \backslash \mathbb{L}$ and $\overline{\mathbb{L}}_t = \mathbb{L}_t \backslash \mathbb{L}$ as the label sets of the source and target private classes. Following existing studies [38, 39], we show the category split ($|\mathbb{L}|/|\overline{\mathbb{L}}_s|/|\overline{\mathbb{L}}_t|$) of each experimental setting in the corresponding result tables.

The dataset split in the UniDA setting is described as below: **Office-31:** For Office-31, the 10 common classes existing between Office-31 and Caltech-256 are selected as $\mathbb{L}$, whereas the subsequent 10 classes in alphabetical order form the set $\overline{\mathbb{L}}_s$ and the remaining 11 classes constitute $\overline{\mathbb{L}}_t$. **Office-Home:** For Office-Home, we use the first 10 classes in alphabetical order as the shared label set $\mathbb{L}$, the next 5 as the source private set $\overline{\mathbb{L}}_s$ and the rest are the target private set $\overline{\mathbb{L}}_t$. **VisDA2017:** For VisDA2017, the first 6 classes are chosen for $\mathbb{L}$, along with the next 3 classes designated to form the set $\overline{\mathbb{L}}_s$ and all other classes assigned to $\overline{\mathbb{L}}_t$. **DomainNet:** For DomainNet, the first 150 classes are selected as $\mathbb{L}$, then the next 50 classes were designated for $\overline{\mathbb{L}}_s$ and all other classes were dedicated to $\overline{\mathbb{L}}_t$.

The dataset split in the PDA setting is described as below: **Office-31:** For Office-31, the same 10 common classes existing between Office-31 and Caltech-256 are used as $\mathbb{L}$, the remaining 21 categories are used as source private label set $\overline{\mathbb{L}}_s$. **Office-Home:** For Office-Home, the first 25 classes in alphabetical order are used as the shared label set $\mathbb{L}$, the remaining 40 classes as the source private set $\overline{\mathbb{L}}_s$.

The dataset split in the OSDA setting is described below: **Office-31:** For Office-31, we use the same 10 common classes existing between Office-31 and Caltech-256 as $\mathbb{L}$, the selected 11 classes ("tape dispenser", "ring binder", "stapler", "scissors", "punchers", "speaker", "pen", "trash can", "phone", "ruler" and "printer") as the target private label set $\overline{\mathbb{L}}_t$. This setting is the same as [40]. **Office-Home:** For Office-Home, the first 15 classes in alphabetical order are selected as the common label set $\mathbb{L}$ and the rest are the target private set $\overline{\mathbb{L}}_t$.

G.2. Evaluation Metrics

In the PDA setting, we only calculate the classification accuracy across all target samples. In the OSDA setting, as target private samples are merged into a single class called "unknown" similar to UniDA, we still leverage H-score mentioned in the main paper to evaluate the effectiveness of methods. In the UniDA setting, we introduce another metric called average class accuracy used in [13, 52] to measure the superiority of our method in addition to the H-score. Average class accuracy reports the average of per-class accuracy over $|\mathbb{L}| + 1$ classes, where the last label contains all target private labels. The H-score reflects the balance between the unknown and known class accuracies, while the average class accuracy prioritizes overall classification performance, making it more practical.

G.3. Comparison Baselines

Follow the previous existing works [13], we compare our method with (1) ResNet [17], (2) close-set domain adaptation: DANN [14], (3) partial domain adaptation: PADA [4], ETN [3], BA³US [27] (4) open set domain adaptation: OSBP [40], STA [29], ROS [2]. (5)universal domain adaptation: UAN [52], CMU [13], DANCE [38], DCC [26], OVANet [39], UniOT [6], GATE [9]. We use some results from [9]. In all experiments, we assume that none of the UniDA methods have prior knowledge of category shift, while baselines tailored for each setting consider this prior.

G.4. Implementation Details

The method is implemented in Pytorch using a ViT-base model with 16×16 input patch size (or ViT-B/16) [21], pre-trained on ImageNet [12]), as the backbone feature extractor. The transformer encoder of ViT-B/16 contains a total of 12 Transformer layers. The label classifier consists of a fully connected network with BatchNorm [20]. The domain discriminator is a three-layer MLP with ReLU activations. We train all models using a minibatch Stochastic Gradient Descent (SGD) optimizer with a momentum of 0.9 and a weight decay of 5×10^{-4}. The learning rate decays by a factor of $(1 + \alpha i/N)^{-\beta}$, where i and N respectively denote the current iteration and the global iteration. The batch size is set to 36. We initialize the initial learning rate to 0.01 for Office-31 and Office-Home, while set 0.001 for VisDA2017 and DomainNet. For the regularization hyperparameters, we set $\gamma = 100$ and $\lambda = 0.3$ for all dataset. For the decision threshold, we set $\alpha = 0.85$ and $\beta = 1.0$ for all dataset in UniDA and OSDA. In PDA, we set $\alpha = 0.8$ for all dataset except $alpha = 0.85$ in the Office-31 W2A task. For the pre-defined number of target prototypes, a larger size of the target domain indicates a larger K. Therefore, we empirically set $K = 50$ for Office-31, $K = 150$ for Office-Home, $K = 500$ for VisDA, $K = 1000$ for DomainNet.

Table D: Average Accuracy (%) on Universal Domain Adaptation.

Method	Office-31(10 / 10 / 11)							Office-Home(10 / 5 / 50)													VisDA
	A2W	D2W	W2D	A2D	D2A	W2A	**Avg**	A2C	A2P	A2R	C2A	C2P	C2R	P2A	P2C	P2R	R2A	R2C	R2P	**Avg**	(6/3/3)
ResNet	75.9	89.6	90.9	80.5	78.8	81.4	82.9	59.4	76.6	87.5	68.9	71.1	81.7	73.3	56.3	86.1	78.7	59.2	78.6	73.2	48.2
DANN	80.7	80.9	88.1	82.7	74.8	83.5	81.8	56.2	81.7	85.9	68.7	73.4	83.8	69.9	56.8	85.8	79.4	57.3	78.3	73.2	52.3
OSBP	66.1	73.6	85.6	73.6	47.4	60.5	67.7	47.8	60.9	76.8	59.2	61.6	74.3	61.7	44.5	79.3	70.6	54.9	75.2	63.9	**66.5**
UAN	85.6	94.8	98.0	86.5	85.5	85.1	89.2	63.0	82.8	87.9	76.9	78.7	85.4	78.2	58.6	86.8	83.4	63.2	79.4	77.0	51.9
CMU	86.9	95.7	98.0	89.1	88.4	88.6	91.1	63.5	83.8	88.9	77.7	79.4	86.9	78.6	59.3	88.3	84.1	64.6	81.4	**78.0**	54.2
DCC	91.7	94.5	96.2	93.7	90.4	92.0	93.1	63.1	80.9	92.1	69.3	75.8	87.1	81.5	55.8	92.1	82.4	62.1	87.3	77.5	59.6
OVANet	76.8	94.3	97.4	84.8	79.9	81.1	85.7	62.0	77.6	86.3	70.0	70.4	78.2	71.7	60.4	83.2	76.9	63.2	81.7	73.5	60.7
UniOT	91.5	98.9	96.3	88.4	89.8	90.9	92.7	-	-	-	-	-	-	-	-	-	-	-	-	-	-
OVANet⋆	86.8	95.3	94.9	87.1	84.0	89.5	89.6	-	-	-	-	-	-	-	-	-	-	-	-	-	-
UniOT⋆	96.0	99.3	98.9	82.0	91.4	91.4	93.2	56.0	91.2	93.1	68.7	82.3	83.8	74.3	53.1	88.0	80.1	55.2	89.7	76.3	56.6
Ours	97.8	98.1	97.9	94.0	92.9	94.3	**95.8**	61.7	88.1	90.8	76.5	63.4	68.7	83.6	64.6	90.5	87.8	71.1	85.6	77.7	61.8

Table E: H-score (%) on Open-Set Domain Adaptation.

Method	Office-31 (10 / 0 / 11)							Office-Home (25 / 0 / 40)													VisDA
	A2W	D2W	W2D	A2D	D2A	W2A	**Avg**	A2C	A2P	A2R	C2A	C2P	C2R	P2A	P2C	P2R	R2A	R2C	R2P	**Avg**	(6/0/6)
STA	75.9	75.0	69.8	75.2	73.2	66.1	72.5	35.0	55.2	59.7	37.5	48.4	53.5	36.0	32.2	59.9	54.3	38.5	64.6	47.9	48.2
OSBP	82.7	82.4	97.2	91.1	75.1	73.7	83.7	55.1	65.2	72.9	64.3	64.7	70.6	63.2	53.2	73.9	66.7	54.5	72.3	64.7	52.3
ROS	82.1	82.4	96.0	99.7	77.9	77.2	85.9	60.1	69.3	76.5	58.9	65.2	68.6	60.6	56.3	74.4	68.8	60.4	75.7	66.2	66.5
UAN	46.8	38.9	68.8	53.0	68.0	54.9	55.1	40.3	41.5	46.1	53.2	48.0	53.7	40.6	39.8	52.5	53.6	43.7	56.9	47.5	51.9
CMU	55.7	52.6	75.9	64.7	76.5	65.8	65.2	45.1	48.3	51.7	58.9	55.4	61.2	46.5	43.8	58.0	58.6	50.1	61.8	53.3	54.2
DANCE	78.8	84.9	78.8	88.9	79.1	68.3	79.8	61.9	61.3	63.7	64.2	58.6	62.6	67.4	61.0	65.5	65.9	61.3	64.2	63.0	67.5
DCC	54.8	58.3	89.4	80.9	67.2	85.3	72.6	56.1	67.5	66.7	49.6	66.5	64.0	55.8	53.0	70.5	61.6	57.2	71.9	61.7	59.6
OVANet	88.3	90.5	98.2	98.4	86.7	88.3	91.7	58.9	66.0	70.4	62.2	65.7	67.8	60.0	52.6	69.7	68.2	59.1	67.6	64.0	66.1
OVANet⋆	95.2	93.0	100.0	94.1	92.6	93.6	94.8	63.5	63.4	74.5	64.7	73.8	76.3	70.8	71.1	77.8	69.9	65.6	78.8	70.9	-
Ours	96.9	99.3	100.0	97.9	92.2	92.2	**96.4**	73.0	84.6	80.3	80.0	88.1	83.6	79.4	68.1	82.9	82.4	76.3	74.5	**79.4**	70.2

H. Additional Results

UniDA. Table D illustrates the average class accuracy of our proposed method compared to other methods in the UniDA setting. Combining Tables 1 and 2 in the main text, it is evident that our approach outperforms other methods not only in terms of H-score but also in terms of average class accuracy on Office-31, demonstrating its superior ability to accurately identify both known and unknown classes. On Office-home and VisDA datasets, the average class accuracy of our method is comparable to CMU and OSBP, respectively. However, CMU and OSBP only focus on average class accuracy, while neglecting the balance between identifying unknown and known categories, which leads to their particularly low H-scores. These results are indicative of the effectiveness of our method in addressing the challenges posed by category shift.

OSDA. As shown in Table E, our method achieves the highest H-score compared to all baselines for the OSDA setting, including methods tailored for this scenario such as OSBP and ROS. In particular, our method achieves 79.4% on the Office-Home dataset under the OSDA setting, which is noticeable 9% higher than the best-performing baseline (OVANet⋆) with 70.9%. This demonstrates the effectiveness and robustness of our method in the face of different label shift scenarios.

PDA. As shown in Table F, our method still outperforms all other methods, including those specifically designed for PDA, in terms of average class accuracy on Office-31 dataset. This indicates the robustness and effectiveness of our approach even in PDA scenario. On the Office-Home and VisDA datasets, our method's performance is comparable with GATE. The PDA task mainly reflects the ability to identify source private categories, while the OSDA task reflects our method's ability to explore target private categories. The latter is more challenging as the target labels are not available. Overall, our method has achieved a good balance between these two capabilities, resulting in an impressive performance in both aspects.

I. Why Compressive Attention Matching

To demonstrate the indispensability of Compressive Attention Matching, we qualitatively illustrate the similarity matrix obtained by computing the cosine similarity between target attentions and source attention prototypes (Figure G

Table F: Average Accuracy (%) on Partial Domain Adaptation.

| Method	Office-31 (10 / 21 / 0)							Office-Home (15 / 50 / 0)													VisDA
	A2W	D2W	W2D	A2D	D2A	W2A	Avg	A2C	A2P	A2R	C2A	C2P	C2R	P2A	P2C	P2R	R2A	R2C	R2P	Avg	(6/6/0)
PADA	82.2	86.5	92.7	99.3	95.4	**100.0**	92.7	52.0	67.0	78.7	52.2	53.8	59.1	52.6	43.2	78.8	73.7	56.6	77.1	62.1	53.5
ETN	94.5	95.0	**100.0**	100.0	96.2	94.6	96.7	59.2	77.0	79.5	62.9	65.7	75.0	68.3	55.4	84.4	75.7	57.7	84.5	70.5	59.8
BA^{3}US	98.9	99.4	**100.0**	98.7	94.8	95.0	97.8	60.6	83.2	88.4	71.8	72.8	83.4	75.5	61.6	86.5	79.3	62.8	86.1	76.0	54.9
UAN	76.8	79.7	93.4	98.3	82.7	83.7	85.8	24.5	35.0	41.5	34.7	32.3	32.7	32.7	21.1	43.0	39.7	26.6	46.0	34.2	39.7
CMU	84.2	84.1	97.2	98.8	69.2	66.8	83.4	50.9	74.2	78.4	62.2	64.1	72.5	63.5	47.9	78.3	72.4	54.7	78.9	66.5	65.5
DANCE	71.2	77.1	94.6	96.8	83.7	92.6	86.0	53.6	73.2	84.9	70.8	67.3	82.6	70.0	50.9	84.8	77.0	55.9	81.8	71.1	73.7
DCC	81.3	87.3	100.0	100.0	95.4	95.5	93.3	54.2	47.5	57.5	83.8	71.6	86.2	63.7	65.0	75.2	85.5	78.2	82.6	70.9	72.4
OVANet	61.7	69.4	90.2	98.7	61.4	66.4	74.6	34.1	54.6	72.1	42.4	47.3	55.9	38.2	26.2	61.7	56.7	35.8	68.9	49.5	34.3
GATE	86.2	89.5	**100.0**	98.6	93.5	94.4	93.7	55.8	75.9	85.3	73.6	70.2	83.0	72.1	59.5	84.7	79.6	63.9	83.8	**73.9**	**75.6**
OVANet*	68.7	75.7	96.0	100.0	76.6	80.6	82.9	43.2	60.6	81.1	53.5	56.3	67.2	50.4	32.3	66.8	63.2	40.5	73.4	57.4	40.6
Ours	90.0	100.0	99.6	98.1	87.6	90.6	**94.3**	67.8	81.0	78.8	68.9	55.2	64.9	75.0	78.3	74.7	80.4	65.8	77.8	72.4	74.3
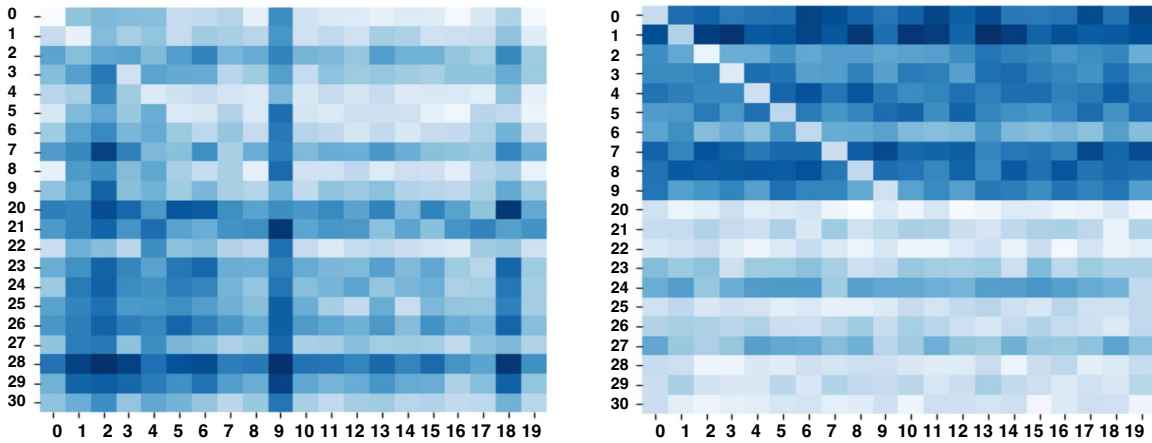
Figure G: Qualitative illustration of Attention Cosin Similarity and Compressive Attention Matching.

left) and the residual matrix obtained through Eq.4 in Compressive Attention Matching in the main paper (Figure G right). The two matrixes both have 21 rows representing 21 target samples corresponding to all target classes (0-9 are common classes and 20-30 are target private classes), and 20 columns representing all source domain prototypes (0-9 are common classes and 10-19 are source private classes). In the matrices, lighter colors indicate higher probabilities for the sample belonging to the corresponding category. As shown in Figure G, it can be seen that the residual matrix is more accurate in classifying attention than the similarity matrix. In particular, in the first ten rows of the residual matrix corresponding to common classes, the diagonal blocks are much lighter than other blocks in the same row, indicating that they belong to one of the source domain prototypes. However, in the following twenty rows corresponding to target private classes, the colors between the blocks do not show any significant difference, indicating that they do not belong to any of the source domain prototypes.

The phenomenon indicates that our designed Attention Commonness Degree (Definition 2 in the main paper) is very effective in distinguishing the samples belonging to the common category. The reason for the significant difference between the attention similarity matrix and the attention residual matrix is that, by enforcing sparsity on the coefficients of attention reconstruction, we can implicitly capture and amplify the differences between attentions, thus making the final residual matrix more discriminative. On the other hand, an attention similarity matrix is easier to be confused different categories with similar objects due to without sparse constraints.