

HiGDA: Hierarchical Graph of Nodes to Learn Local-to-Global Topology for Semi-Supervised Domain Adaptation

Ba Hung Ngo^{1*}, Doanh C. Bui^{2*}, Nhat-Tuong Do-Tran³, Tae Jong Choi^{1†}

¹Graduate School of Data Science, Chonnam National University, South Korea

²School of Electrical Engineering, Korea University, South Korea

³Department of Computer Science, National Yang Ming Chiao Tung University, Taiwan

ngohung@chonnam.ac.kr, doanbhc@korea.ac.kr, tuongdotn.cs11@nycu.edu.tw, ctj17@jnu.ac.kr

Abstract

The enhanced representational power and broad applicability of deep learning models have attracted significant interest from the research community in recent years. However, these models often struggle to perform effectively under domain shift conditions, where the training data (the *source domain*) is related to but exhibits different distributions from the testing data (the *target domain*). To address this challenge, previous studies have attempted to reduce the domain gap between source and target data by incorporating a few labeled target samples during training—a technique known as semi-supervised domain adaptation (SSDA). While this strategy has demonstrated notable improvements in classification performance, the network architectures used in these approaches primarily focus on exploiting the *features* of individual images, leaving room for improvement in capturing rich representations. In this study, we introduce a Hierarchical Graph of Nodes designed to simultaneously present representations at both *feature* and *category* levels. At the *feature level*, we introduce a local graph to identify the most relevant patches within an image, facilitating adaptability to defined main object representations. At the *category level*, we employ a global graph to aggregate the features from samples within the same category, thereby enriching overall representations. Extensive experiments on widely used SSDA benchmark datasets, including *Office-Home*, *DomainNet*, and *VisDA2017*, demonstrate that both quantitative and qualitative results substantiate the effectiveness of HiGDA, establishing it as a new state-of-the-art method.

Introduction

Over the past decade, deep learning models have achieved remarkable success in various vision-based tasks, largely based on the assumption that the training and testing datasets share identical distributions. However, in real-world applications, the training samples (source domain) are often collected from a domain that differs from the testing samples (target domain). As a result, models trained on the source domain frequently underperform when applied to the target domain due to the domain shift problem (Long et al. 2017). To address this issue, domain adaptation (DA) methods (Ganin




Figure 1: Overview of the hierarchical graph of nodes. Input images are considered global nodes, each comprising multiple sub-patches that are considered local nodes.

et al. 2016; Saito et al. 2018; Li et al. 2021b; Yan et al. 2022; Yu and Lin 2023) have been proposed to bridge the domain gap between source and target domains at both the *feature* and *category* levels.

At the *feature level*, early DA approaches (Ganin et al. 2016; Saito et al. 2018) employed convolutional neural networks (CNNs) as the standard backbone to encode the semantic representations of input images. CNNs (Kayhan and van Gemert 2020) use sliding convolutions with learnable kernels to introduce shift-invariance. As a result, the CNN architecture tends to preserve local spatial relationships, extracting discriminative features that cover only a specific region of pixels at a time. Consequently, the CNN backbone is not flexible enough to capture complex objects, such as images consisting of multiple similar contexts located at different positions. In contrast, recent DA methods (Zhu, Bai, and Wang 2023; Yang et al. 2023) leverage the cross-attention mechanism in the vision transformer (ViT) (Dosovitskiy et al. 2020) to better comprehend image features. The ViT architecture treats an image as a sequence of patches and explores their relationships, effectively capturing global relations at the *feature level*. However, it struggles to recognize representations of the main object when the input images contain irrelevant objects or a dominant background, as it assumes all patches in an image are correlated. To address this problem, (Ngo et al. 2024) introduces a hybrid approach that leverages the advantages of both ViT and CNN models. These models complement each other’s information via a co-training strategy, alleviating their weaknesses

*These authors contributed equally.

†Corresponding author.

and thereby enhancing image classification performance in CNN during inference. Although the hybrid model significantly improves classification performance, two concerns remain: 1) In the case of ViT, all patches are assumed to have the same level of correlation with each other, which allows various types of noise to be attended to, thus hindering comprehensive representation exploration. 2) As for the hybrid approach, it results in a large number of parameters that need to be optimized. Based on these observations, we pose the following question: **Q1: How can we explicitly present the semantic information of an image at the feature level with a compact model?**

At the *category level*, alignment across source and target domains can be effectively achieved through pseudo-labeling, where generated pseudo-labels from the unlabeled target samples are used to update the parameters of the trained model. Additionally, several semi-supervised domain adaptation (SSDA) methods (Huang, Zhu, and Chen 2023; Yu and Lin 2023; He, Liu, and Yin 2024) have demonstrated that incorporating one or three labeled target samples per class during the training phase, alongside labeled source data, can substantially enhance classification accuracy. The primary objective of the pseudo-labeling process or the inclusion of a few labeled target samples is to increase the representation of the target domain in the learning model. These studies reveal that the explicit presentation of features is a crucial process in domain adaptation tasks. However, it is important to note that previous DA studies often overlook the structural information of labeled samples during training. As a result, samples within the same class may consist of similar patterns that are not effectively integrated to enrich features. Based on this observation, another question arises: **Q2: How can connections be established among samples belonging to the same category for feature aggregation to achieve better representations?**

In this study, we propose a **Hierarchical Graph of nodes for Domain Adaptation (HiGDA)**, as illustrated in Fig. 1, to address the two aforementioned questions. HiGDA is designed with two levels: the *local graph* and the *global graph*. The *local graph* is constructed to exploit the *feature level* of each individual image. Similar to ViT, we partition an image into patches and arrange these patches into a sequence, with each patch conceptualized as a *local node*. However, unlike ViT, our local graph is constructed with adaptable connections, where local nodes establish pairwise connections solely with their closest neighbors for feature aggregation. In this approach, features within an image are leveraged through a graph-based methodology. By doing so, the impact of noise is mitigated since the representations of the main object are kept at a distance from irrelevant objects and background noise in the embedding space. The output features of the *local graph* are then used to build the *global graph*. Here, each image is conceptualized as a *global node*. Next, we construct *global graphs* to explore the category-level representations. Specifically, representations of *global nodes* belonging to the same category are concatenated to enrich the context of visual recognition, thereby addressing the concern raised in **Q2**.

Furthermore, to enhance the alignment of representations

across domains at the *category level*, we propose a Graph Active Learning (GAL) algorithm. In GAL, the trained HiGDA model generates pseudo-labels from unlabeled target samples in each training episode. These pseudo-labels are then combined with given labeled samples to construct the *global graphs*. As a result, the domain gap between the source and target domains at the *category level* is minimized.

In summary, our contributions are as follows:

- We propose a hierarchical graph of nodes (HiGDA) to explore the local-to-global structure in representations associated at *feature* and *category* levels.
- We demonstrate that HiGDA, structured as a graph, is more compact and efficient in model size than existing methods while achieving impressive results.
- We introduce the GAL strategy, which efficiently uses pseudo-labels generated from unlabeled target samples to enhance cross-domain representations at the *category level*.
- We validate HiGDA on various domain adaptation datasets under the SSDA setting. Extensive experimental results demonstrate HiGDA's effectiveness, achieving state-of-the-art performance.

Related Work

Network Architectures in SSDA

Early SSDA methods (Saito et al. 2019; Li et al. 2021b; Kim and Kim 2020; Li et al. 2021a) follow a similar methodology when designing network architectures, typically comprising two main components: convolutional neural networks (CNNs) as feature extractors and multilayer perceptrons (MLPs) as single classifiers. The various network architectures are evolving rapidly by flexibly incorporating multiple components to improve classification accuracy. The representative works include DECOTA (Yang et al. 2021), UODA (Qin et al. 2021), and ASDA (Qin et al. 2022). DECOTA introduces a decomposing framework consisting of two distinct branches that share the same network architecture (CNN+MLP). Each branch is trained on a different labeled dataset, and they exchange knowledge using a co-training strategy to achieve consistency on the unlabeled samples. Similarly, UODA and ASDA use a CNN feature extractor followed by two MLP classifiers to provide two different views. They demonstrate that the data bias in SSDA can be alleviated by using this type of network architecture. However, CNNs are only effective in capturing local spatial hierarchical representations. Therefore, using CNNs as the feature extractor faces challenges in achieving a comprehensive global visual context. Furthermore, these network architectures solely focus on exploiting the semantic representations of each individual image (*feature-level*).

Graph Neural Networks

Various Graph Neural Network (GNN) models have been developed to address graph learning tasks on graph-structured data, such as biochemical, social, or citation datasets. A prominent example of GNNs is the Graph Convolutional Network (GCN) (Kipf and Welling 2017), which

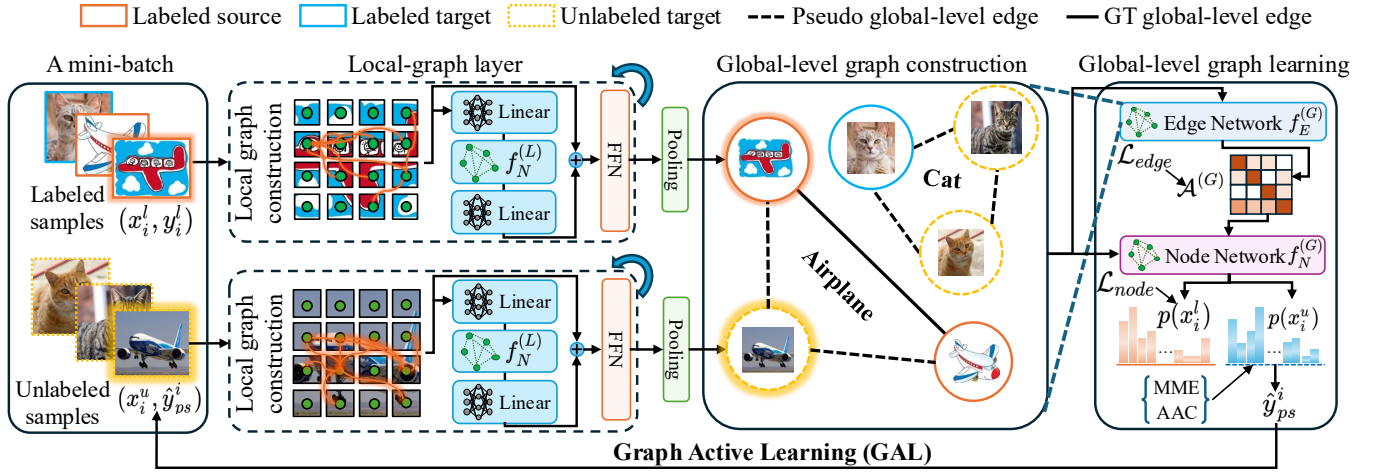


Figure 2: Overview of HiGDA. Each image forms a local graph $\mathcal{G}^{(L)}$, and a mini-batch forms a global graph $\mathcal{G}^{(G)}$ to explore feature to *category level* representations. Then, GAL minimizes bias from the source dataset.

employs the adjacency matrix in conjunction with convolution operations. These graph-based models have consistently demonstrated impressive performance in numerous graph-learning tasks. Consequently, they have gained widespread adoption in domain adaptation tasks (Ma, Zhang, and Xu 2019; Zhu et al. 2023; Li, Li, and Yu 2023; Xiao et al. 2023). For instance, (Ma, Zhang, and Xu 2019) employs the triplet loss to construct the adjacency matrix, incorporating similarity information of source and target samples. In contrast, (Zhu et al. 2023) and (Li, Li, and Yu 2023) build graphs using pairwise label similarity, establishing cross-domain connectivity if the ground-truth labels of source samples share the same information with pseudo labels generated from the unlabeled target samples. However, the aforementioned approaches primarily focus on constructing graphs at the higher level (*category level*), neglecting the lower level (*feature level*). Consequently, it becomes challenging to effectively present features at the *category level*.

Methodology

Problem Definition

In semi-supervised domain adaptation (SSDA), we are given substantial labeled set from the source domain, $S = \{x_i^s, y_i^s\}_{i=1}^{N_s}$, alongside a much smaller labeled set from the target domain, $T = \{x_i^t, y_i^t\}_{i=1}^{N_t}$, to form a combined labeled set $D_l = S \cup T = \{x_i^l, y_i^l\}_{i=1}^{N_s+N_t}$, where $N_s \gg N_t$. The final goal is to minimize the classification error on the unlabeled target set, $D_u = \{x_i^u\}_{i=1}^{N_u}$.

Overview of Hierarchical Graph Design

Figure 1 illustrates our methodology in designing a hierarchical graph of nodes for SSDA. In this approach, we introduce a local-to-global topology to present a hierarchical graph of nodes, $\mathcal{G} = \{\mathcal{G}^{(L)}, \mathcal{G}^{(G)}\}$, where $\mathcal{G}^{(L)}$ and $\mathcal{G}^{(G)}$ are the local and global graphs, respectively, for representing features across domains. At the *feature level*, an image is split into sub-patches, with each patch conceptualized as

a *local node*. Then, $\mathcal{G}^{(L)}$ is designed to explicitly produce the semantic representations of each image by aggregating the embeddings of local nodes with their neighbors. At the *category level*, output features of each input image are considered as a *global node*, which is handled by $\mathcal{G}^{(G)}$. Specifically, $\mathcal{G}^{(G)}$ is created to aggregate global nodes sharing to the same category.

Local Graph. Given an image sample $x_i \in \mathbb{R}^{H \times W \times 3}$, we conceptualize it as a local graph $\mathcal{G}^{(L)} = \{\mathcal{V}^{(L)}, \mathcal{E}^{(L)}\}$. In detail, we define a set of local nodes $\mathcal{V}^{(L)} = \{\mathbf{v}_i^{(L)}\}_{i=1}^{N^{(L)}}$, where $N^{(L)}$ is the number of local nodes, which is fixed for all samples, and each local node $\mathbf{v}_i^{(L)}$ represents a feature vector of a sub-patch image tiled from x_i . Regarding the construction of the local edge set $\mathcal{E}^{(L)}$, an edge $e_{i,j}^{(L)} \in \mathcal{E}^{(L)}$ connecting $\mathbf{v}_i^{(L)}$ and $\mathbf{v}_j^{(L)}$ exists if and only if $\mathbf{v}_j^{(L)} \in \mathcal{N}(\mathbf{v}_i^{(L)})$, where $\mathcal{N}(\mathbf{v}_i^{(L)})$ is the set of k -nearest neighbors of $\mathbf{v}_i^{(L)}$.

Global Graph. Similarly the local graph, we construct a global graph, $\mathcal{G}^{(G)} = \{\mathcal{V}^{(G)}, \mathcal{E}^{(G)}\}$. A set of global nodes is defined as $\mathcal{V}^{(G)} = \{\mathbf{v}_i^{(G)}\}_{i=1}^{N^{(G)}}$, where each global node $\mathbf{v}_i^{(G)}$ is output feature vector of the image sample x_i handled by the local graph. $N^{(G)}$ denotes the number of global nodes. To determine the edge set $\mathcal{E}^{(G)}$, an edge $e_{i,j}^{(G)} \in \mathcal{E}^{(G)}$ between two global nodes exists if and only if they have the same categories, *i.e.*, $y_i = y_j$.

HiGDA Learning Process

In this section, we present the HiGDA learning process including local-graph (*LoG*) and global-graph (*GoG*) networks, as illustrated in Fig. 2.

Local-Graph Network. To learn the local graph $\mathcal{G}^{(L)}$, we design a local-graph network *LoG*. This network incorporates an embedding network for image patching to construct

nodes and a series of local-graph layers. The core of these layers is the local-node network $f_N^{(L)}$, which is primarily responsible for performing message passing, i.e., exchanging information between nodes.

Given $x_i \in \mathbb{R}^{H \times W \times 3}$, it is divided into sub-patches and their corresponding embeddings as a set of local nodes $\mathcal{V}^{(L)} = \{\mathbf{v}_i^{(L)}\}_{i=1}^{N^{(L)}}$ by the embedding network. Each local node $\mathbf{v}_i^{(L)}$ represents a $h \times w$ sub-patch. Hence, there is $N^{(L)} = \frac{H}{h} \cdot \frac{W}{w}$ local nodes. To define the set of neighbors for each local node $\mathbf{v}_i^{(L)}$, the Euclidean distances between each local node $\mathbf{v}_i^{(L)}$ and all remaining nodes are calculated. Then, the top K nodes most similar to $\mathbf{v}_i^{(L)}$ are considered its neighbors $\mathcal{N}(\mathbf{v}_i^{(L)})$. After the set of neighbors for each local node is specified, the local edge set $\mathcal{E}^{(L)}$ is constructed, where an edge $e_{i,j}$ exists if and only if $\mathbf{v}_j^{(L)} \in \mathcal{N}(\mathbf{v}_i^{(L)})$. Then, the local-node network $f_N^{(L)}$ learns to aggregate these neighbors to the local node $\mathbf{v}_i^{(L)}$ as follow:

$$\mathbf{v}_i^{(L)'} = \sigma\left(f_N^{(L)}(\mathbf{v}_i^{(L)}, \mathcal{E}^{(L)})\right), \quad (1)$$

where σ denotes nonlinear activation. To be specific, Max-relative graph convolution (Li et al. 2019) is adopted, followed by a projection layer to design $f_N^{(L)}$:

$$f_N^{(L)}(\cdot) = F\left([\mathbf{v}_i^{(L)}; \max(\{\mathbf{v}_i^{(L)} - \mathbf{v}_j^{(L)} | e_{i,j} \in \mathcal{E}^{(L)}\})\right], \quad (2)$$

where F is a linear projection layer, and $[\cdot; \cdot]$ denotes the concatenation operation. In this manner, differences between a local node $\mathbf{v}_i^{(L)}$ and its neighbors are calculated and pooled by taking the maximum values per channel dimension. Then, it is concatenated with $\mathbf{v}_i^{(L)}$ and processed by a projection layer F as an information-exchanging operation. To enhance feature diversity, two linear projections are applied to each local node: one before and one after fitting to $f_N^{(L)}$. Additionally, we implement a skip connection to prevent gradient vanishing. To strengthen feature transformation, a feed-forward network is placed between local-graph layers. Finally, the set of local nodes $\mathcal{V}^{(L)}$ undergoes max pooling, selecting the highest value in each channel to form the global node $\mathbf{v}^{(G)}$.

Global-Graph Network. During training, a mini-batch of image samples is used to construct a global graph $\mathcal{G}^{(G)}$. The number of samples in a mini-batch is the number of global nodes $N^{(G)}$. Herein, we design global-graph network GoG to process $\mathcal{G}^{(G)}$. Following (Luo et al. 2020), we design two networks: the edge network $f_E^{(G)}$ and the node network $f_N^{(G)}$. The edge network $f_E^{(G)}$ is responsible for producing similarity scores between global nodes, while the node network $f_N^{(G)}$ aggregates all global nodes to the considered global node $\mathbf{v}_i^{(G)}$ based on computed similarities. First, the similarity score between $\mathbf{v}_i^{(G)}$ and $\mathbf{v}_j^{(G)}$, is computed:

$$\hat{a}_{i,j}^{(G)} = Sigmoid\left(f_E^{(G)}(\|\mathbf{v}_i^{(G)} - \mathbf{v}_j^{(G)}\|)\right), \quad (3)$$

where $Sigmoid$ denotes the Sigmoid function, and $\hat{a}_{i,j}^{(G)}$ denotes an element of the global-graph unnormalized affinity matrix $\hat{\mathcal{A}}^{(G)} \in \mathbb{R}^{N^{(G)} \times N^{(G)}}$. Then, self-connections of global nodes are added to $\hat{\mathcal{A}}$, and it is normalized:

$$\mathcal{A}^{(G)} = D^{-\frac{1}{2}}(\hat{\mathcal{A}}^{(G)} + I)D^{-\frac{1}{2}}, \quad (4)$$

where D is the degree matrix of $\hat{\mathcal{A}}^{(G)} + I$, I is the identity matrix, and $\mathcal{A}^{(G)}$ is the normalized affinity matrix. The feature aggregation of the global node $\mathbf{v}_i^{(G)}$ is processed as follows:

$$\mathbf{v}_i^{(G)'} = f_N^{(G)}\left([\mathbf{v}_i^{(G)}; \sum_{j \in N^{(G)}} a_{i,j}^{(G)} \cdot \mathbf{v}_j^{(G)}]\right), \quad (5)$$

where $\mathbf{v}_i^{(G)'}$ is the updated feature, which is aggregated from its neighboring node representations. By doing so, all global nodes in $\mathcal{V}^{(G)}$ are multiplied by their similarity scores with the considered global node $\mathbf{v}_i^{(G)}$, followed by computing a weighted sum. This result is then concatenated with the global node $\mathbf{v}_i^{(G)}$ itself and passed through the node network $f_N^{(G)}$ as an update function.

Training Scheme

Supervised Learning. To perform supervised learning on labeled samples, we use the cross-entropy loss function to optimize LoG and GoG :

$$\mathcal{L}_{node} = -\frac{1}{N^{(G)}} \sum_{i=1}^{N^{(G)}} y_i \log p(x_i), \quad (6)$$

where $p(x_i) = \text{softmax}(GoG(LoG(x_i)))$.

Global Edge Supervision. Given the set of similarity scores $\{a_{i,j}^{(G)}\}_{j=1}^{N^{(G)}}$ between $\mathbf{v}_i^{(G)}$ and its neighbor nodes, produced by the global-edge network $f_E^{(G)}$, to guide the edge construction process, the binary cross-entropy loss is used as follows:

$$\mathcal{L}_{edge} = -e_{i,j}^{(G)} \log a_{i,j}^{(G)} - (1 - e_{i,j}^{(G)}) \log(1 - a_{i,j}^{(G)}), \quad (7)$$

where $e_{i,j}^{(G)}$ is the global ground-truth edge, $e_{i,j}^{(G)} = 1$ if and only if $\mathbf{v}_i^{(G)}$ and $\mathbf{v}_j^{(G)}$ belong to the same category; otherwise, $e_{i,j}^{(G)} = 0$.

Final Objective Function. Finally, the loss functions from Eq. 6 and Eq. 7 are merged into a unified objective function for training the whole network, which is formed as:

$$\mathcal{L}_{HiGDA} = \mathcal{L}_{node} + \mathcal{L}_{edge}. \quad (8)$$

Graph Active Learning Strategy

Due to the dominance of source samples, data bias may occur, leading to the suboptimal problem of the HiGDA on the target domain. Hence, we incorporate the concept of active learning to develop the graph active learning algorithm

(GAL). In GAL, the pretraining-finetuning manner is repeatedly conducted in each training episode to select the most informative samples from the unlabeled target data for labeling aiming to optimize the model performance, which is processed as follows:

Step 1. In each episode, the pre-trained model on labeled samples is used to create pseudo labels from the unlabeled set D_u as follows:

$$\hat{y}_{ps}^i = \operatorname{argmax} p(x_i^u) \text{ if only if } \max(p(x_i^u)) \geq \tau, \quad (9)$$

with $p(x_i^u) = \operatorname{softmax}(GoG(LoG(x_i^u)))$, and τ is the predefined threshold. The pseudo label set is denoted as $D_{ps} = \{x_i^u, \hat{y}_{ps}^i\}_{i=1}^{N_{ps}}$, where N_{ps} is the number of pseudo labels.

Step 2. Then, we update the labeled set D_l by combining it with the pseudo label set D_{ps} every episode as follows:

$$\hat{D}_l^{q+1} = D_l \cup D_{ps}^q, \quad (10)$$

where D_{ps}^q is the pseudo label set at the training episode q . The HiGDA model is then fine-tuned with the graph active learning strategy to get the optimal model performance on the updated version \hat{D}_l with the supervised losses in Eq. (8) and Eq. (9). The pairwise connection rule to determine edges among samples in \hat{D}_l for feature aggregation is defined as:

$$e_{i,j}^{(G)} = \begin{cases} 1, & \text{if } y_i^l = y_j^l, \\ 1, & \text{if } y_i^l = \hat{y}_j^{ps}, \\ 1, & \text{if } \hat{y}_i^{ps} = \hat{y}_j^{ps}, \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

In this manner, the performance of HiGDA can be adjusted according to the quality and quantity of generated pseudo labels.

Integrating HiGDA With SSSA Techniques

In this section, we deploy the proposed framework HiGDA with the previous SSSA method.

Minimax Entropy (MME). MME (Saito et al. 2019) is the most popular SSSA method using adversarial learning with a minimax entropy strategy expressed as follows:

$$\begin{aligned} \theta_{LoG}^* &= \operatorname{argmin}_{\theta_{LoG}} \mathcal{L}_{HiGDA} + \lambda \mathcal{L}_{ENT}, \\ \theta_{GoG}^* &= \operatorname{argmin}_{\theta_{GoG}} \mathcal{L}_{HiGDA} - \lambda \mathcal{L}_{ENT}, \end{aligned} \quad (12)$$

where \mathcal{L}_{HiGDA} is the supervised learning loss in Eq. (8) calculated using the labeled samples, \mathcal{L}_{ENT} is the entropy loss estimated by using the unlabeled target samples. $\lambda = 0.1$ is the trade-off hyperparameter.

Adversarial Adaptive Clustering (AAC). Similar to MME, AAC (Li et al. 2021a) also uses adversarial learning between a feature extractor and a classifier by proposing the new adversarial adaptive clustering loss to estimate the pairwise feature similarity among unlabeled target data within a mini-batch, as follows:

$$\begin{aligned} \theta_{LoG}^* &= \operatorname{argmin}_{\theta_{LoG}} \mathcal{L}_{HiGDA} + \beta \mathcal{L}_{AAC}, \\ \theta_{GoG}^* &= \operatorname{argmin}_{\theta_{GoG}} \mathcal{L}_{HiGDA} - \beta \mathcal{L}_{AAC}, \end{aligned} \quad (13)$$

where \mathcal{L}_{AAC} is the adversarial adaptive clustering loss with β to be 1.0.

Experiments

We evaluate our framework on various challenging domain adaptation datasets, including *Office-Home* (Venkateswara et al. 2017), *DomainNet* (Peng et al. 2019) and *VisDA2017* (Peng et al. 2018) in Tabs. 1, 2 and 3, respectively. Consistent with previous SSSA methods, we perform 12 DA tasks on *Office-Home* and 7 DA tasks on *DomainNet* under 1-shot and 3-shot settings, while considering only one domain adaptation scenario for *VisDA2017*.

Implemental Details

For fairness to the CNN-based network in terms of the number of parameters, we select the tiny version of Pyramid ViG (Han et al. 2022) for *LoG*, where the number of local nodes $N^{(L)}$ obtained from the last block is 49. The number of neighbors for each local node is set to 9. We are inspired by (Luo et al. 2020) to build *GoG* consisting of the node and edge networks. The number of global nodes $N^{(G)}$ is set to 32. For training, we use the SGD optimizer, the learning rate and weight decay are set to 1×10^{-3} and 5×10^{-5} , respectively. The threshold τ in Eq. (9) is set to 0.95 and the episode in Eq. (10) is set to $q = 50$ consisting of 1,000 training steps. A single NVIDIA GeForce RTX 4090 GPU is used for all experiments.

Comparison Results

As shown in Tabs. 1, 2, and 3, HiGDA-T baseline significantly outperforms the CNN-based baseline model with the same S+T setting. Surprisingly, the baseline can also surpass the existing CNN-based SSSA methods. Specifically, on *Office-Home* in Tab. 1, HiGDA-T exceeds DE-COTA by 5.3% in the 1-shot setting, and overcomes MME (Saito et al. 2019) and CDAC (Li et al. 2021a) by 1.4% and 0.3% in the 3-shot setting, respectively. Furthermore, as in Tab. 2 on *DomainNet*, when HiGDA-T is enhanced with the proposed GAL, referred to as HiGDA-T+GAL, it shows marginal improvements over the second-best method, EFTL (He, Liu, and Yin 2024), with increasing of up to 5.1% and 4.8% under the 1-shot and 3-shot settings, respectively. The proposed method achieves the highest performance with implementation of HiGDA-T+GAL+MME, surpassing EFTL (He, Liu, and Yin 2024) by as much as 15.2% and 15.8% with 1- and 3-shot settings, respectively. Similarly, our method reaches 97.8% on *VisDA2017* with HiGDA-T+GAL+MME under 3-shot setting, as in Tab. 3.

Analyses

To further validate HiGDA, we investigate three key aspects as *ablation studies*: (1) the ability of HiGDA to integrate with existing SSSA methods; (2) the effectiveness of the proposed GAL on cross-domain alignment; and (3) the effectiveness of *LoG* compared to other CNN-based and ViT-based backbones, as well as the performance of *GoG* compared to MLP.

Setting	Method	R→C	R→P	R→A	P→R	P→C	P→A	A→P	A→C	A→R	C→R	C→A	C→P	Avg. (%)	
1-shot	ResNet-34														
	S+T (Baseline)	52.1	78.6	66.2	74.4	48.3	57.2	69.8	50.9	73.8	70.0	56.3	68.1	63.8	
	MME (ICCV'19)	61.9	82.8	71.2	79.2	57.4	64.7	75.5	59.6	77.8	74.8	65.7	74.5	70.4	
	CDAC (CVPR'21)	61.9	83.1	72.7	80.0	59.3	64.6	75.9	61.2	78.5	75.3	64.5	75.1	71.0	
	DECOTA (ICCV'21)	56.0	79.4	71.3	76.9	48.8	60.0	68.5	42.1	72.6	70.7	60.3	70.4	64.8	
	MCL (IJCAI'22)	67.0	85.5	73.8	81.3	61.1	68.0	79.5	64.4	81.2	78.4	68.5	79.3	74.0	
	MME + SLA (CVPR'23)	64.1	83.8	72.9	80.0	59.9	66.7	76.3	62.1	78.6	75.1	67.5	77.1	72.0	
	SPA (NeurIPS'23)	65.2	84.1	71.4	80.7	59.1	65.7	76.7	62.3	79.0	76.4	66.6	77.3	72.0	
	FixMME + EFTL (AAAI'24)	66.6	87.2	74.3	82.6	63.3	68.7	80.5	65.7	80.8	77.5	65.6	79.6	74.4	
	HiGDA-T														
	S+T (Baseline)	57.9	86.9	77.3	81.9	52.4	63.8	76.4	54.7	80.0	76.3	62.6	71.2	70.1	
	Baseline+AAC	67.7	87.3	84.1	87.6	62.9	75.4	78.6	61.8	86.7	83.5	71.3	79.1	77.2	
	Baseline+MME	70.4	92.0	88.1	90.5	65.0	78.9	84.2	66.1	89.5	86.4	76.4	83.2	80.9	
	HiGDA-T+GAL														
	S+T (Baseline)	66.7	87.3	79.9	85.0	59.8	71.1	78.8	62.1	79.1	79.6	72.2	73.8	74.6	
	Baseline+AAC	81.3	91.1	87.7	92.6	77.9	83.2	85.1	76.5	90.0	88.6	82.6	85.4	85.2	
	Baseline+MME	82.5	93.9	90.4	92.7	79.3	84.3	86.8	80.2	94.1	91.6	87.3	86.5	87.5	
	3-shot	ResNet-34													
S+T (Baseline)		55.7	80.8	67.8	73.1	53.8	63.5	73.1	54.0	74.2	68.3	57.6	72.3	66.2	
MME (ICCV'19)		64.6	85.5	71.3	80.1	64.6	65.5	79.0	63.6	79.7	76.6	67.2	79.3	73.1	
CDAC (CVPR'21)		67.8	85.6	72.2	81.9	67.0	67.5	80.3	65.9	80.6	80.2	67.4	81.4	74.2	
DECOTA (ICCV'21)		70.4	87.7	74.0	82.1	68.0	69.9	81.8	64.0	80.5	79.0	68.0	83.2	75.7	
MCL (IJCAI'22)		70.1	88.1	75.3	83.0	68.0	69.9	83.9	67.5	82.4	81.6	71.4	84.3	77.1	
SPA (NeurIPS'23)		67.2	87.0	73.9	82.0	65.2	69.5	81.0	63.1	80.2	77.5	68.5	81.7	74.7	
MME + SLA (CVPR'23)		68.4	87.4	74.7	81.9	67.4	69.7	81.1	65.9	80.5	79.4	69.2	81.9	75.6	
FixMME + EFTL (AAAI'24)		72.8	89.3	77.5	85.4	70.9	72.6	84.8	70.3	83.8	81.5	70.6	84.6	78.7	
HiGDA-T															
S+T (Baseline)		64.7	90.3	78.2	84.8	60.0	69.8	82.6	58.4	83.1	71.9	70.0	80.4	74.5	
Baseline+AAC		73.1	91.0	84.1	89.7	71.1	79.8	85.5	66.9	89.6	85.8	77.2	86.0	81.7	
Baseline+MME		76.8	94.1	88.3	92.6	72.7	82.4	89.8	73.9	92.5	90.1	80.8	88.7	85.2	
HiGDA-T+GAL															
S+T (Baseline)		71.0	92.1	81.5	87.1	66.4	74.8	86.2	66.2	86.5	84.3	74.2	87.8	79.8	
Baseline+AAC		84.6	94.9	88.8	94.1	80.7	85.3	91.2	81.1	94.9	92.7	86.8	91.2	88.9	
Baseline+MME		84.8	95.1	91.8	94.3	84.2	88.4	92.3	81.6	95.8	93.9	87.6	92.5	90.2	

Table 1: Classification accuracy (%) on *Office-Home* for 1-shot and 3-shot settings.

Method	rel→clp		rel→pnt		pnt→clp		clp→skt		skt→pnt		rel→skt		pnt→rel		Avg. (%)	
	1 _{shot}	3 _{shot}	1 _{shot}	3 _{shot}	1 _{shot}	3 _{shot}	1 _{shot}	3 _{shot}	1 _{shot}	3 _{shot}	1 _{shot}	3 _{shot}	1 _{shot}	3 _{shot}	1 _{shot}	3 _{shot}
ResNet-34																
S+T (Baseline)	55.6	60.0	60.6	62.2	56.8	59.4	50.8	55.0	56.0	59.5	46.3	50.1	71.8	73.9	56.8	60.0
MME (ICCV'19)	70.0	72.2	67.7	69.7	69.0	71.7	56.3	61.8	64.8	66.8	61.0	61.9	76.1	78.5	66.4	68.9
DECOTA (ICCV'21)	79.1	80.4	74.9	75.2	76.9	78.7	65.1	68.6	72.0	72.7	69.7	71.9	79.6	81.5	73.9	75.6
CDAC (CVPR'21)	77.4	79.6	74.2	75.1	75.5	79.3	67.6	69.9	71.0	73.4	69.2	72.5	80.4	81.9	73.6	76.0
MCL (IJCAI'22)	77.4	79.4	74.6	76.3	75.5	78.8	66.4	70.9	74.0	74.7	70.7	72.3	82.0	83.3	74.4	76.5
SPA (NeurIPS'23)	75.3	76.0	71.8	72.2	74.8	76.5	65.9	67.0	69.7	71.1	65.8	67.2	81.1	82.3	72.1	73.2
CDAC + SLA (CVPR'23)	79.8	81.6	75.6	76.0	77.4	90.3	68.1	71.3	71.7	73.5	71.7	73.5	80.4	82.5	75.0	76.9
FixMME + EFTL (AAAI'24)	79.6	81.2	74.9	77.1	78.2	81.8	69.3	72.8	71.8	74.4	69.9	71.5	83.1	84.4	75.3	77.6
HiGDA-T																
S+T (Baseline)	68.8	75.3	74.3	78.3	69.4	75.6	61.0	71.1	68.6	73.9	60.2	65.6	82.2	87.0	69.2	75.3
Baseline+AAC	80.2	84.7	86.4	87.6	80.5	84.8	76.2	79.8	84.2	86.2	75.1	77.9	90.1	92.5	81.8	84.8
Baseline+MME	81.5	85.7	88.0	89.7	84.7	88.4	77.9	80.4	85.4	87.8	75.6	79.1	90.9	93.0	83.5	86.4
HiGDA-T + GAL																
S+T (Baseline)	78.5	85.1	80.9	84.6	73.1	83.7	64.5	76.0	77.6	80.4	71.9	78.3	87.8	88.6	76.3	82.4
Baseline+AAC	90.9	92.7	92.4	94.2	88.4	93.5	85.2	89.3	90.6	93.8	89.7	90.9	90.2	94.3	89.6	92.7
Baseline+MME	92.1	93.9	93.8	95.0	89.6	95.4	85.7	89.7	91.2	94.0	90.1	91.5	91.0	94.6	90.5	93.4

Table 2: Classification accuracy (%) on *DomainNet* for 1-shot and 3-shot settings.

Effectiveness of HiGDA on SSDA Methods (1). Results reported in Tabs. 1 and 2, highlighted under yellow cells,

demonstrate that HiGDA can be flexibly coupled with existing SSDA methods. Under the 3-shot setting, integrating

Method	1-shot	3-shot
ResNet-34		
S+T (Baseline)	60.2	64.6
MME (ICCV'19)	68.7	70.9
APE (ECCV'20)	78.9	81
CDAC (CVPR'21)	69.9	80.6
DECOTA (ICCV'21)	64.9	80.7
ECACL (ICCV'21)	81.1	83.3
MCL (IJCAI'22)	86.3	87.3
ProML (IJCAI'23)	87.6	88.4
HiGDA-T + GAL		
S+T (Baseline)	80.9	93.6
Baseline+AAC	91.4	96.3
Baseline+MME	88.9	97.8

Table 3: Mean class-wise accuracy (%) on *VisDA2017* for 1-shot and 3-shot settings.

HiGDA-T with AAC and MME improves performance by 7.2% and 10.7% compared to only the HiGDA-T baseline on *Office-Home*, respectively. For *DomainNet*, these improvements are 9.5% and 11.1%, respectively.

Effectiveness of Graph Active Learning (2). We report the impact of GAL (HiGDA-T+GAL)—highlighted under pink cells, by comparing it to the HiGDA baseline model—highlighted under yellow cells. As listed in Tabs 1 and 2, the results show significant performance gaps when GAL is removed. Under the S+T setting with 3-shot, there is a performance drop of 5.3% on *Office-Home* and 7.1% on *DomainNet*. These results reveal that data bias still exists when only a few target labeled samples are available, and the proposed GAL can effectively address this issue by providing additional pseudo labels from unlabeled target samples.

Effectiveness of Local and Global Graphs (3). In the context of the *feature-level*, *LoG* significantly outperforms ResNet-34 (He et al. 2016) and PVT-Tiny (Wang et al. 2021) across all evaluated head classifiers, as detailed in Table 4. Specifically, *LoG* paired with the MLP head classifier demonstrates an increase in accuracy of 2.9% and 7.0% compared to ResNet-34 and PVT-Tiny, respectively, while utilizing fewer parameters. Similarly, when employing the *GoG* as the classifier, *LoG* also exceeds ResNet-34 and PVT-Tiny by 2.3% and 11.1%, respectively. Regarding the *category level*, *GoG* consistently demonstrates enhanced performance compared to MLP, independent of the various backbones utilized. Notably, ResNet-34 equipped with *GoG* yields the most significant improvement over the MLP, with an increase of 11.4%. Finally, the highest overall performance of 75.3% is attained by integrating *LoG* with *GoG* (HiGDA-T), thereby proving the effectiveness and robustness of HiGDA design.

Qualitative Results

We present qualitative results of HiGDA, including GradCAM visualizations with the constructed local graph $\mathcal{G}^{(L)}$ in Fig. 3 and t-SNE comparisons in Fig. 4, to provide visible

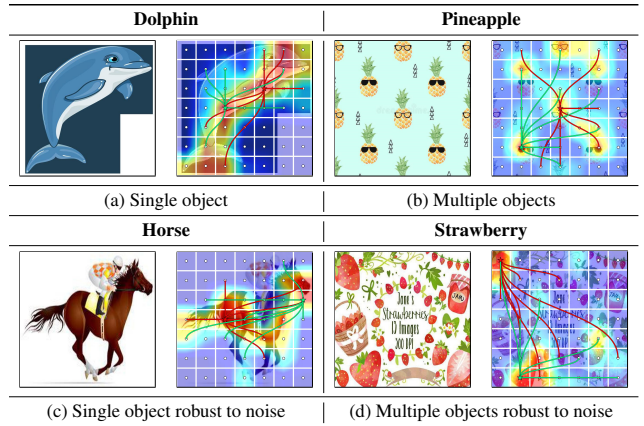


Figure 3: GradCAM (Selvaraju et al. 2017) results extracted by the local graph. Please zoom in for viewing ease.

and explainable insights of the model operation that further demonstrates the superiority of HiGDA.

GradCAM With *LoG*. Among the local nodes in $\mathcal{V}^{(L)}$, representing patches, two anchor nodes are selected where the model focuses on (based on GradCAM). Edges to their neighbors are visualized, with anchor nodes shown as stars (\star and \star). Under HiGDA, the local graph $\mathcal{G}^{(L)}$ effectively connects highly correlated patches to each other. For a single-object image in Fig. 3a, $\mathcal{G}^{(L)}$ easily covers the entire *dolphin* starting from two anchor nodes in the middle of the body. In the multiple-object image in Fig. 3b, it links objects with similar patterns such as *pineapples*. Furthermore, $\mathcal{G}^{(L)}$ demonstrates robustness to noise by disregarding non-essential objects in Fig. 3c, focusing on a *horse* while excluding the rider. Even in complex backgrounds, as illustrated in Fig. 3d, it accurately links relevant objects such as *strawberries*, while omitting irrelevant ones.

t-SNE Comparisons. As illustrated in Fig. 4a, the result indicates that the HiGDA-T baseline model with the S+T setting struggles to distinguish representations of different classes. In contrast, GAL proves to be effective by leading to better-distinguished representations in Fig. 4b. However, the differences between samples of several categories, such as *anvil*, *axe*, and *bear*, remain unclear. Consistent with the quantitative results, integrating HiGDA with MME produces higher quality representations than AAC, with samples from different categories well clustered, as shown in Fig. 4c and Fig. 4d, respectively.

Conclusion

In this study, we introduced a hierarchical graph of nodes to explicitly represent objects from the *feature level* to the *category level*. At the *feature level*, an input image is divided into multiple patches, each conceptualized as a local node. Patches with strong relationships are connected to form a local graph, which better represents complex objects while filtering out irrelevant elements. At the *category level*, each image is treated as a global node. We construct a global graph

Feature Extractor	Classifier	Parameter (M)	rel→clp	rel→pnt	pnt→clp	clp→skt	skt→pnt	rel→skt	pnt→rel	Avg. (%)
ResNet-34	MLP	21.6	60.6	63.0	62.8	54.9	62.0	51.5	76.5	61.6
PVT-Tiny		13.8	56.9	60.0	57.2	49.6	57.1	45.3	76.1	57.5
LoG-Tiny		9.5	64.7	66.9	64.5	56.5	64.9	54.3	79.8	64.5
ResNet-34	GoG	22.6	73.1	75.3	74.1	67.1	73.4	64.0	84.1	73.0
PVT-Tiny		14.8	65.0	67.4	64.7	54.9	62.5	52.7	82.0	64.2
LoG-Tiny		10.5	75.3	78.3	75.6	71.1	73.9	65.6	87.0	75.3

Table 4: Ablation study on *DomainNet* under the 3-shot setting to evaluate the effectiveness of different classifiers.

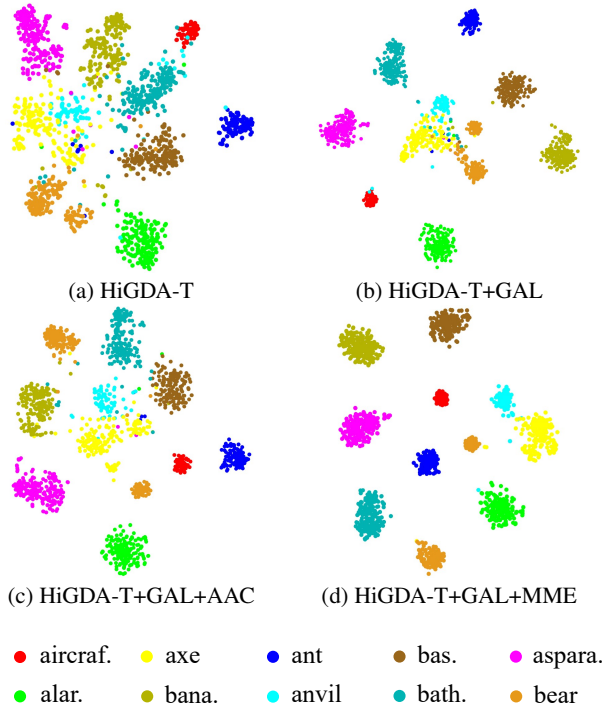


Figure 4: t-SNE (Van der Maaten and Hinton 2008) visualization on *DomainNet* of 10 classes in the *real to sketch* (rel→skt) task. Please zoom in for viewing ease.

to aggregate the features of global nodes that share the same label information, thereby enriching the overall representations. Extensive experiments on various semi-supervised domain adaptation datasets, along with qualitative and quantitative analyses, demonstrate that the proposed method is more effective than previous approaches.

Limitations

We found that *GoG* is sensitive to noise, such as image samples that do not align well with the annotated labels, leading to cumulative errors and performance degradation. Controlling the contribution of node and edge loss to find the optimal solution is crucial to addressing this issue, which we plan to explore in future research.

Acknowledgments

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea Government (MSIT) (No. RS-2023-00214326 and No. RS-2023-00242528).

References

- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *J. Mach. Learn. Res.*, 17(1): 2096–2030.
- Han, K.; Wang, Y.; Guo, J.; Tang, Y.; and Wu, E. 2022. Vision GNN: An Image is Worth Graph of Nodes. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 8291–8303. Curran Associates, Inc.
- He, J.; Liu, B.; and Yin, G. 2024. Enhancing Semi-supervised Domain Adaptation via Effective Target Labeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12385–12393.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Huang, X.; Zhu, C.; and Chen, W. 2023. Semi-supervised Domain Adaptation via Prototype-based Multi-level Learning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Kayhan, O. S.; and van Gemert, J. C. 2020. On translation invariance in cnns: Convolutional layers can exploit absolute spatial location. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 14274–14285.
- Kim, T.; and Kim, C. 2020. Attract, perturb, and explore: Learning a feature alignment network for semi-supervised domain adaptation. In *Eur. Conf. Comput. Vis. (ECCV)*, 591–607.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*.

- Li, G.; Müller, M.; Thabet, A.; and Ghanem, B. 2019. DeepGCNs: Can GCNs Go As Deep As CNNs? In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9266–9275.
- Li, J.; Li, G.; Shi, Y.; and Yu, Y. 2021a. Cross-Domain Adaptive Clustering for Semi-Supervised Domain Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2505–2514.
- Li, J.; Li, G.; and Yu, Y. 2023. Adaptive Betweenness Clustering for Semi-Supervised Domain Adaptation. *IEEE Transactions on Image Processing*, 32: 5580–5594.
- Li, K.; Liu, C.; Zhao, H.; Zhang, Y.; and Fu, Y. 2021b. ECACL: A Holistic Framework for Semi-Supervised Domain Adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 8578–8587.
- Long, M.; Zhu, H.; Wang, J.; and Jordan, M. I. 2017. Deep transfer learning with joint adaptation networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, 2208–2217.
- Luo, Y.; Wang, Z.; Huang, Z.; and Baktashmotlagh, M. 2020. Progressive Graph Learning for Open-Set Domain Adaptation. In *Proceedings of the 37th International Conference on Machine Learning*, 6468–6478.
- Ma, X.; Zhang, T.; and Xu, C. 2019. GCAN: Graph Convolutional Adversarial Network for Unsupervised Domain Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ngo, B. H.; Do-Tran, N.-T.; Nguyen, T.-N.; Jeon, H.-G.; and Choi, T. J. 2024. Learning CNN on ViT: A Hybrid Model to Explicitly Class-specific Boundaries for Domain Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 28545–28554.
- Peng, X.; Bai, Q.; Xia, X.; Huang, Z.; Saenko, K.; and Wang, B. 2019. Moment Matching for Multi-Source Domain Adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Peng, X.; Usman, B.; Kaushik, N.; Wang, D.; Hoffman, J.; and Saenko, K. 2018. VisDA: A Synthetic-to-Real Benchmark for Visual Domain Adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Qin, C.; Wang, L.; Ma, Q.; Yin, Y.; Wang, H.; and Fu, Y. 2021. Contradictory Structure Learning for Semi-supervised Domain Adaptation. In *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*, 576–584.
- Qin, C.; Wang, L.; Ma, Q.; Yin, Y.; Wang, H.; and Fu, Y. 2022. Semi-Supervised Domain Adaptive Structure Learning. *IEEE Transactions on Image Processing*, 31: 7179–7190.
- Saito, K.; Kim, D.; Sclaroff, S.; Darrell, T.; and Saenko, K. 2019. Semi-Supervised Domain Adaptation via Minimax Entropy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Saito, K.; Watanabe, K.; Ushiku, Y.; and Harada, T. 2018. Maximum Classifier Discrepancy for Unsupervised Domain Adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 618–626.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep Hashing Network for Unsupervised Domain Adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wang, W.; Xie, E.; Li, X.; Fan, D.-P.; Song, K.; Liang, D.; Lu, T.; Luo, P.; and Shao, L. 2021. Pyramid Vision Transformer: A Versatile Backbone for Dense Prediction Without Convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 568–578.
- Xiao, Z.; Wang, H.; Jin, Y.; Feng, L.; Chen, G.; Huang, F.; and Zhao, J. 2023. SPA: a graph spectral alignment perspective for domain adaptation. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*.
- Yan, Z.; Wu, Y.; Li, G.; Qin, Y.; Han, X.; and Cui, S. 2022. Multi-level Consistency Learning for Semi-supervised Domain Adaptation. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Yang, J.; Liu, J.; Xu, N.; and Huang, J. 2023. TVT: Transferable Vision Transformer for Unsupervised Domain Adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 520–530.
- Yang, L.; Wang, Y.; Gao, M.; Shrivastava, A.; Weinberger, K. Q.; Chao, W.-L.; and Lim, S.-N. 2021. Deep Co-Training With Task Decomposition for Semi-Supervised Domain Adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 8906–8916.
- Yu, Y.-C.; and Lin, H.-T. 2023. Semi-Supervised Domain Adaptation With Source Label Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 24100–24109.
- Zhu, J.; Bai, H.; and Wang, L. 2023. Patch-Mix Transformer for Unsupervised Domain Adaptation: A Game Perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3561–3571.
- Zhu, R.; Jiang, X.; Lu, J.; and Li, S. 2023. Cross-Domain Graph Convolutions for Adversarial Unsupervised Domain Adaptation. *IEEE Transactions on Neural Networks and Learning Systems*, 34(8): 3847–3858.