Reconfigurability-Aware Selection for Contrastive Active Domain Adaptation

Zeyu Zhang[1,2] , **Chun Shen**[1,2,3] , **Shuai Lü**[1,2,3,∗] and **Shaojie Zhang**[1,3]
[1]Key Laboratory of Symbolic Computation and Knowledge Engineering (Jilin University), Ministry of Education, China
[2]College of Computer Science and Technology, Jilin University, China
[3]College of Software, Jilin University, China
{lus, shenchun}@jlu.edu.cn, {zeyuz22, sjzhang22}@mails.jlu.edu.cn

Abstract

Active domain adaptation (ADA) aims to label a small portion of target samples to drastically improve the adaptation performance. The existing ADA methods mostly rely on the output of domain discriminator or the original prediction probability to design sample selection strategies and do not fully explore the semantic information of source and target domain features, which may lead to selecting the valueless target samples. Moreover, most of them require complex network structures (such as introducing additional domain discriminator, multiple classifiers, or loss predictors) and multiple query functions. In this work, we propose a concise but effective ADA method called Reconfigurability-Aware Selection for Contrastive active domain adaptation (RASC). With the reconfigurability-aware sample selection strategy, RASC can select the most valuable target samples for annotation in the presence of domain shift. To better utilize the selected target samples, we further design a contrastive learning-based gradual active domain adaptation framework. In addition, we propose a variant of RASC called RASC-Ob, which uses a simpler sample annotation method and supplements the learning of misclassified samples. Extensive experimental results on multiple benchmarks demonstrate the superiority of RASC.

1 Introduction

Unsupervised domain adaptation (UDA) is one of the most researched paradigms, which aims to generalize models trained on labeled source domain to unlabeled target domain, alleviating the dependence of training on a large amount of labeled data. However, the performance of UDA methods is still significantly worse than that of corresponding supervised learning methods, especially when there is a significant domain shift [Fu *et al.*, 2021; Xie *et al.*, 2022a; Huang *et al.*, 2023a]. In practical scenarios, labeling a small portion of target samples is feasible. However, as shown in

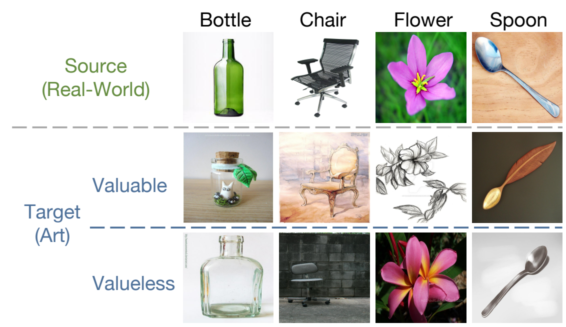

Figure 1: Target domain contains both valuable and valueless samples. It is difficult for the model trained through supervised learning on the source domain to correctly classify valuable target samples because they differ significantly from the source domain. The opposite is true for valueless target samples.

Figure 1, not all target samples are beneficial for further improving performance. It is not necessary to label valueless samples that are similar to the source domain.

Active learning (AL) aims to select the most valuable samples from a large number of unlabeled samples for annotation [Tang and Huang, 2021; Liu *et al.*, 2021; Zhan *et al.*, 2021; Ren *et al.*, 2022]. However, AL assumes that unlabeled and labeled samples subject to the same distribution. Therefore, in the presence of domain shift, AL methods cannot select the most valuable target samples for annotation. Recently, a new domain adaptation paradigm called active domain adaptation (ADA) has been proposed to address the above-mentioned issue. ADA aims to select the most valuable target samples for annotation in the presence of domain shift and use these samples for training to further improve performance. However, current ADA methods [Su *et al.*, 2020; Fu *et al.*, 2021; Xie *et al.*, 2022b; Han *et al.*, 2023] rely on the original prediction probability, auxiliary modules, or multiple query functions for sample selection, and do not fully explore the semantic information of source and target domain features.

With the consideration above, in this paper, we propose the Reconfigurability-Aware Selection for Contrastive active domain adaptation (RASC). Different from previous methods, we utilize the relationship between semantic information of source and target domain features for sample selection. We not only consider the sample-level value but also

[∗]Corresponding author.

design a class prediction reliability-based sample rejection strategy to jointly consider the class-level value. To better utilize the selected target samples, inspired by [He *et al.*, 2020] and [Singh, 2021], we further design a contrastive learning-based gradual active domain adaptation framework to reduce the distribution discrepancy between labeled target domain and source domain, as well as between unlabeled target domain and labeled domains. In addition, considering that the current ADA methods require experts to label the selected target samples correctly, however, it is difficult for experts to remember and distinguish all categories when there are thousands of categories. Inspired by [Zhang *et al.*, 2022], we also propose a variant of RASC called RASC-Ob, where experts only need to judge whether the model predictions are correct, without having to label the samples correctly.

Our main contributions can be summarized as follows:

- We propose a reconfigurability-aware sample selection strategy, which utilizes the semantic information of source and target domain features to select the most valuable target samples for annotation.
- We propose a contrastive learning-based gradual active domain adaptation framework to drastically improve the adaptation performance.
- We propose a variant of RASC called RASC-Ob that experts only need to judge whether the prediction is correct and supplements the learning of misclassified samples.
- RASC attains excellent results on multiple benchmarks. Numerous experiments and analyses demonstrate the effectiveness of RASC.

2 Related Work

2.1 Domain Adaptation

Domain adaptation (DA) aims to transfer the knowledge learned from the source domain to the target domain. Unsupervised domain adaptation assumes that the target domain has no labels at all. Previous methods [Long *et al.*, 2015; Long *et al.*, 2017; Li *et al.*, 2020] are based on discrepancy metric minimization. Recently, adversarial learning-based methods [Ganin and Lempitsky, 2015; Long *et al.*, 2018; Wang and Zhang, 2020; Li *et al.*, 2021] inspired by GAN further improve the performance. Although UDA methods have achieved excellent results, the corresponding supervised learning methods still significantly outperform UDA methods. Semi-supervised domain adaptation (SSDA) [Saito *et al.*, 2019; Jiang *et al.*, 2020; Singh, 2021; Yan *et al.*, 2022; Huang *et al.*, 2023b] assumes a small number of labeled target domain samples are available and uses contrastive learning or adversarial learning to alleviate the domain shift. However, the labeled target samples in SSDA are passively provided before the training begins and not all of these samples are beneficial for training. In addition, contrastive learning-based SSDA methods [Singh, 2021; Huang *et al.*, 2023b] use noisy target pseudo labels to construct prototypes, and do not consider the distribution discrepancy between the source and the labeled target domains.

2.2 Active Learning

Active learning (AL) aims to select the most valuable samples for annotation and these samples will be used for training to further improve performance. AL methods require defining a query function to measure the value of samples and finally select the samples with the highest value for annotation [Zhan *et al.*, 2021; Wu *et al.*, 2022]. The uncertainty-based methods [Wang and Shang, 2014; Yan and Huang, 2018] select the difficult-to-distinguish samples near the decision boundary for annotation. The diversity-based methods [Sener and Savarese, 2018; Sinha *et al.*, 2019] aim to select samples to more comprehensively represent the entire dataset. However, AL assumes that unlabeled and labeled samples are subject to the same distribution, therefore, the AL methods are ineffective in the presence of domain shift.

2.3 Active Domain Adaptation

AADA [Su *et al.*, 2020] is the first work to introduce active domain adaptation into vision tasks, utilizing the outputs of domain discriminator and the entropy to measure the domainness and uncertainty of samples, and combining with the adversarial-based DA method. TQS [Fu *et al.*, 2021] proposes a transferable query selection strategy that combines multiple mechanisms. EADA [Xie *et al.*, 2022a] utilizes the free energy biases across domains for sample selection, and proposes a regularization term to reduce domain shift. TL-ADA [Han *et al.*, 2023] proposes a transferable loss-based sample selection strategy. DiaNA [Huang *et al.*, 2023a] proposes a data subdivision selection strategy and designs customized learning for samples with different values.

However, these methods mostly rely on the output of domain discriminator or the original prediction probability to design selection strategies, and most of them require complex networks or multiple query functions. More importantly, they have not fully explored the semantic information of source and target domain features, and the DA strategies are also relatively rough. Thus motivated, we propose Reconfigurability-Aware Selection for Contrastive active domain adaptation (RASC), which utilizes the relationship between semantic information of source and target domain features for sample selection, and design a more refined contrastive learning-based gradual active domain adaptation framework.

3 Methodology

In this section, we first formalize ADA and then introduce the details of the proposed Reconfigurability-Aware Selection for Contrastive active domain adaptation (RASC) and its variant.

3.1 Problem Formulation

In ADA, we can access a source domain $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$ with n_s labeled data and a target domain $\mathcal{D}_{ut} = \{x_i^{ut}\}_{i=1}^{n_{ut}}$ with n_{ut} unlabeled data. $\mathcal{D}_s$ and $\mathcal{D}_{ut}$ are sampled from $P(x^s, y^s)$ and $Q(x^t, y^t)$ with $P \neq Q$, respectively. We conduct R rounds sample selection and select $n_{lt} = B \cdot |\mathcal{D}_{ut}|$ target samples in total for annotation, where B is the labeling budget. These labeled samples will be removed from $\mathcal{D}_{ut}$ and form $\mathcal{D}_{lt} = \{(x_i^{lt}, y_i^{lt})\}_{i=1}^{n_{lt}}$. For convenience, we denote $\mathcal{D}_l = \mathcal{D}_s \cup \mathcal{D}_{lt}$ as all labeled data across domains. Our model

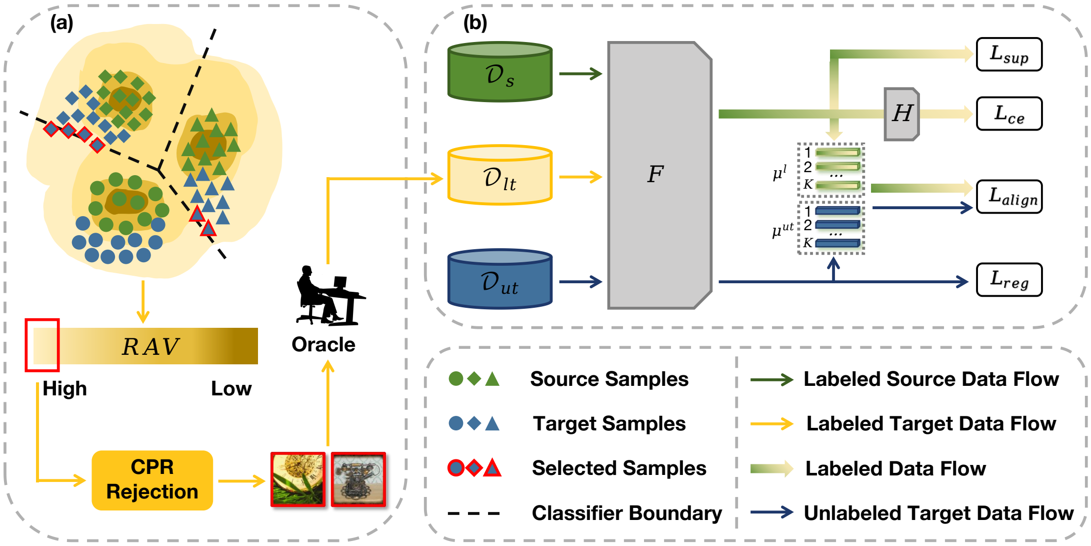

Figure 2: The framework of RASC. (a) For the reconfigurability-aware selection process, we utilize reconfigurability-aware value (RAV) and class prediction reliability-based sample rejection strategy (CPR) to select the most valuable target samples for annotation in the presence of domain shift. (b) For the training process, we utilize contrastive learning-based gradual active domain adaptation to reduce the distribution discrepancy between labeled target domain and source domain, as well as between unlabeled target domain and labeled domains.

consists of a feature extractor F to extract feature $f = F(x)$ and a classifier H to categorize a feature into a logit vector over all K classes, without adding any auxiliary modules. The framework of RASC is illustrated in Figure 2.

3.2 Reconfigurability-Aware Selection

To select the most valuable target samples in the presence of domain shift, we propose the reconfigurability-aware selection strategy (RAS), which utilizes the relationship between semantic information of source and target domain features. RAS not only considers the sample-level value but also jointly considers the class-level value by a class prediction reliability-based sample rejection strategy (CPR).

Reconfigurability-aware value. Existing ADA methods do not fully explore the semantic information of source and target domain features, which may lead to the selection of uninformative target samples. In contrast to the previous methods, we propose a reconfigurability-aware selection strategy that utilizes the relationship between semantic information of source and target domain features for sample selection and integrates the value of the target domain sample into a reconfigurability-aware value. First, we attain the source domain global prototypes for each class as follows:

$$\mu^s = [\mu_c^s]_{c=1}^K, \mu_c^s = \frac{\sum_{(x,y)\in\mathcal{D}_s} \mathbb{1}_{\{y=c\}} F(x)}{\sum_{(x,y)\in\mathcal{D}_s} \mathbb{1}_{\{y=c\}}} \tag{1}$$

where μ^s contains semantic information for each class feature in the source domain. We utilize μ^s and combine it with the prediction probability of x^{ut} to reconstruct the original feature of the unlabeled target sample x^{ut}:

$$\widetilde{f}^{ut} = \sigma(H(F(x^{ut})))\mu^s \tag{2}$$

where $\sigma(\cdot)$ is the softmax operation.

Suppose x^{ut} does not contain target domain-specific feature, i.e. x^{ut} has a high similarity to the source domain (x^{ut} is valueless). Since the model is obtained by supervised training on the source domain, it can classify x^{ut} confidently and correctly. Therefore, the feature reconstructed using the prediction probability of x^{ut} and μ^s can effectively express the semantic information of x^{ut}. In other words, the semantic information of the original feature is the same as that contained in the reconstructed feature $\widetilde{f}^{ut}$. Thus, the discrepancy between the prediction probabilities of the original feature and $\widetilde{f}^{ut}$ is small. On the contrary, for valuable target samples, their semantic information cannot be reconstructed using the source domain prototypes, therefore the discrepancy between the prediction probabilities of the original feature and the reconstructed feature is significant.

Motivated by this, we integrate the value of the target sample into a simple value function: reconfigurability-aware value (RAV). We can obtain $\mathrm{RAV}(x^{ut})$ by using only the prediction probability for the original feature of x^{ut}: $p^{ut} = \sigma(H(F(x^{ut})))$ and the prediction probability of the reconstructed feature $\widetilde{f}^{ut}$: $\widetilde{p}^{ut} = \sigma(H(\widetilde{f}^{ut}))$:

$$\mathrm{RAV}(x^{ut}) = -\sum_{c=1}^{K}(p_c^{ut} + \widetilde{p}_c^{ut})\log(p_c^{ut}\widetilde{p}_c^{ut}) \tag{3}$$

For those valuable target samples, namely those containing a large amount of target domain-specific knowledge, their RAVs are relatively high, while for valueless target samples, their RAVs are relatively low. Additionally, only when the model produces consistent and completely certain predictions for the original and reconstructed features (e.g. [1,0,0] and [1,0,0]), RAV reaches the minimum value of zero. We theoretically explain the role of RAV in measuring the value of target samples in the Appendix.

Class prediction reliability-based sample rejection strategy. RAV reflects the sample-level value. To comprehensively measure the sample value from different perspectives, we propose a class prediction reliability-based sample rejection strategy (CPR) to jointly consider the class-level value. We use the average confidence of c-th class predictions to measure the c-th class prediction reliability:

$$\mathrm{CPR}[c] = \frac{\sum_{x^{ut} \in \mathcal{D}_{ut}} \mathbb{1}_{\{\widehat{y}^{ut}=c\}} p_c^{ut}}{\sum_{x^{ut} \in \mathcal{D}_{ut}} \mathbb{1}_{\{\widehat{y}^{ut}=c\}}} \tag{4}$$

where $\widehat{y}^{ut} = \arg\max p^{ut}$. A high $\mathrm{CPR}[c]$ indicates a high prediction confidence of sample belonging to the c-th class, which means that the predicted label of the c-th class is more likely to be correct. Due to our desire to select more hard samples for annotation, we hope to reject samples with pseudo label c with a high probability. Naturally, we use CPR as the rejection probability for various classes. For a candidate sample with pseudo label c, we first generate a random number ξ from a uniform distribution $U(0, 1)$ and then reject the sample when $\xi < \mathrm{CPR}[c]$.

3.3 Contrastive Learning-Based Gradual Active Domain Adaptation

After sample selection, we obtain the labeled target domain $\mathcal{D}_{lt}$. Previous methods merge $\mathcal{D}_{lt}$ and $\mathcal{D}_s$ as a whole and then use traditional UDA methods for adaptation, which produce suboptimal results. In ADA, sample selection and model training alternate, and a good DA strategy should reduce the discrepancy between the source and target domains while also having a positive impact on the next round of sample selection, prompting it to select more valuable target samples.

With the above consideration, we design a contrastive learning-based gradual active domain adaptation framework. Following [He *et al.*, 2020], we use bank $\mathcal{B}_l$ to store source domain and labeled target domain samples, and bank $\mathcal{B}_{ul}$ to store unlabeled target domain samples. And we use the momentum model $F'(\cdot)$ to encode the samples in the bank. Due to the significant discrepancy between the selected target samples and the source samples, it is suboptimal to directly consider $\mathcal{D}_{lt}$ and $\mathcal{D}_s$ as a whole, as in previous methods. To alleviate the discrepancy between $\mathcal{D}_{lt}$ and $\mathcal{D}_s$ and better learn discriminative features, we hope to pull samples belonging to the same class across $\mathcal{D}_{lt}$ and $\mathcal{D}_s$ closer to each other, while push away samples from different classes. To accomplish it, we design an instance contrastive loss:

$$\mathcal{L}_{sup} = -\mathbb{E}_{x_i \sim \mathcal{D}_l} \log \frac{\sum\limits_{x_j \in \mathcal{P}(x_i)} h(\widehat{x}_i, x_j)}{\sum\limits_{x_j \in \mathcal{B}_l} h(\widehat{x}_i, x_j) + \sum\limits_{x_j \in \mathcal{N}(x_i)} h(\widehat{x}_i, \widehat{x}_j)} \tag{5}$$

where $\widehat{x}_i = \mathrm{Aug}(x_i)$ is the strongly augmented version of x_i, $\mathcal{P}(x_i)$ is the set of samples with the same label as x_i in $\mathcal{B}_l$, $\mathcal{N}(x_i) = \mathcal{B}_l \setminus \mathcal{P}(x_i)$, $h(q, k) = \exp(\mathrm{Sim}(F(q), F'(k))/\tau)$ is the similarity between the features of q and k extracting by F and F' respectively, and τ is the temperature parameter.

After aligning $\mathcal{D}_{lt}$ and $\mathcal{D}_s$, the distribution discrepancy between $\mathcal{D}_{ut}$ and $\mathcal{D}_l$ can be further reduced. Inspired by [Singh, 2021], we mitigate the discrepancy by aligning prototypes of the same class across domains. Unlike prior works, we do not use pseudo labels when calculating prototypes for each class in $\mathcal{D}_{ut}$ because they contain some noise. Instead, we use prediction probability to weight and average the features, which can disperse the features of samples to different classes rather than just the classes corresponding to the pseudo labels to alleviate the impact of noisy pseudo labels. The prototype of the c-th class in $\mathcal{D}_{ut}$ for each training batch is formulated as:

$$\mu_c'^{ut} = \frac{\sum_{i=1}^{n_b} \sigma_c(H(F(x_i^{ut})))F(x_i^{ut})}{\sum_{i=1}^{n_b} \sigma_c(H(F(x_i^{ut})))} \tag{6}$$

where n_b is the size of the batch.

After computing the prototypes for the current batch, we update the global prototypes in a moving average manner:

$$\mu_c^{ut} = \lambda \mu_c^{ut} + (1 - \lambda)\mu_c'^{ut} \tag{7}$$

where λ is the update trade-off parameter. Then, we calculate and update prototypes for each class in $\mathcal{D}_l$ in the same way and use μ_c^l to denote the global prototype of its c-th class. We utilize prototype-based contrastive loss to align prototypes of the same class across domains to reduce the discrepancy between unlabeled target domain and labeled domains:

$$\mathcal{L}_{align} = -\frac{1}{K} \sum_{c=1}^{K} \log \frac{\Omega(\mu_c^{ut}, \mu_c^l)}{\sum_{j=1}^{K} \Omega(\mu_c^{ut}, \mu_j^l)} \tag{8}$$

where $\Omega(q, k) = \exp(\mathrm{Sim}(q, k)/\tau)$.

For unlabeled target domain data, we hope that they can form good clustering, which is not only beneficial for aligning the source and target domains but also has a positive impact on sample selection. We utilize self-supervised contrastive loss to achieve this goal:

$$\mathcal{L}_{reg} = -\mathbb{E}_{x_i \sim \mathcal{D}_{ut}} \log \frac{h(\widehat{x}_i, x_i)}{h(\widehat{x}_i, x_i) + \sum\limits_{x_j \in \mathcal{B}_{ut}} (h(\widehat{x}_i, x_j) + h(\widehat{x}_i, \widehat{x}_j))} \tag{9}$$

Overall, the final objective of RASC can be stated as:

$$\mathcal{L}_{rasc} = \mathcal{L}_{ce} + \alpha(\mathcal{L}_{sup} + \mathcal{L}_{reg}) + \beta \mathcal{L}_{align} \tag{10}$$

where $\mathcal{L}_{ce}$ is the cross entropy loss of labeled data, α and β are the parameters to trade-off different loss.

3.4 Variant: Reconfigurability-Aware Selection with One-bit Annotation

Current ADA methods require experts to label the selected samples correctly. However, when data comes from hundreds or thousands of categories, such annotation is difficult for experts. Inspired by [Zhang *et al.*, 2022], we introduce one-bit annotation into ADA and propose a variant of RASC called RASC-Ob to reduce the workload of experts when labeling samples. Unlike the traditional annotation method, one-bit annotation only requires experts to judge whether the model prediction of the selected sample is correct, that is, experts only need to answer yes or no. From the perspective of information content, one-bit annotation is more efficient. For example, for a K-classes classification problem, labeling the sample correctly requires $\log_2 K$ bits of information, while asking experts to judge whether the model prediction is correct (yes or no) only requires one bit. Therefore, using one-bit annotation allows us to query more samples under the same labeling budget and effectively reduce the burden on experts.

Unlike [Zhang *et al.*, 2022], which requires the initial labeled data used for training, our RASC-Ob does not require labeled initial data at all, meaning that all labeling budgets are used in a one-bit annotation manner. In RASC-Ob, for samples judged by experts to be correctly predicted, the utilization method is the same as in RASC. For samples judged by experts to be incorrectly predicted, we use them to form $\mathcal{D}_{err}$. Inspired by [Zhang *et al.*, 2022], we design the negative contrastive loss:

$$\mathcal{L}_{neg} = -\mathbb{E}_{x_i \sim \mathcal{D}_{err}} \log \frac{h(\widehat{x}_i, x_i)}{h(\widehat{x}_i, x_i) + \sum_{x_j \in \mathcal{N}_{err}(x_i)} (h(\widehat{x}_i, x_j) + h(\widehat{x}_i, \widehat{x}_j))} \tag{11}$$

where $\mathcal{N}_{err}(x_i)$ contains samples with the same real label as the incorrect label of x_i in $\mathcal{B}_l$ and all samples in $\mathcal{B}_{ut}$.

In addition, to keep the model prediction of incorrect samples away from known incorrect category, we also utilize negative loss [Kim *et al.*, 2019]:

$$L_{nl} = -\mathbb{E}_{(x_i, y_i^-) \in \mathcal{D}_{err}} \sum_{i=1}^{K} \log(1 - \sigma_{y_i^-}(H(F(x_i)))) \tag{12}$$

where y_i^- is the known incorrect category of x_i.

Overall, the final objective of RASC-Ob can be stated as:

$$\mathcal{L}_{rasc-ob} = \mathcal{L}_{ce} + \mathcal{L}_{nl} + \alpha(\mathcal{L}_{sup} + \mathcal{L}_{reg} + \mathcal{L}_{neg}) + \beta \mathcal{L}_{align} \tag{13}$$

where α and β are the same as Eq. (10).

4 Experiments

Datasets. To evaluate the superiority of RASC, we conduct experiments on four benchmark datasets: *Office-31* [Saenko *et al.*, 2010] is a mainstream DA dataset. It contains 3 domains: Amazon (A), Webcam (W), DSLR (D) and 31 categories. *Office-Home* [Venkateswara *et al.*, 2017] is a more challenging dataset. It contains 4 domains: Artistic (A), Clipart (C), Product (P), Real-World (R) and 65 categories. *VisDA* [Peng *et al.*, 2017] is a large-scale dataset containing 150k synthetic images and 55k real images. It contains two domains: Syntatic (S), Real (R) and 12 categories. *MiniDomainNet* [Zhou *et al.*, 2021] is a subset of DomainNet [Peng *et al.*, 2019], consisting of four domains: Clipart (C), Painting (P), Real (R), Sketch (S) and 126 categories. It preserves the complexity of DomainNet and reduces the consumption of computing resources.

Method	A→D	A→W	D→A	D→W	W→A	W→D	Avg
Source Only	81.5	75.0	63.1	95.2	65.7	99.4	80.0
Random	87.1	84.1	75.5	98.1	75.8	99.6	86.7
Entropy	91.0	89.2	76.1	99.7	77.7	**100.0**	88.9
CoreSet	82.5	81.1	70.3	96.5	72.4	99.6	83.7
BADGE	90.8	89.1	79.8	99.6	79.6	**100.0**	89.8
AADA	89.2	87.3	78.2	99.5	78.7	**100.0**	88.8
CLUE	92.0	87.3	79.0	99.2	79.6	<u>99.8</u>	89.5
TQS	92.8	92.2	80.6	**100.0**	80.4	**100.0**	91.1
DBAL	88.2	88.9	75.2	99.4	77.0	**100.0**	88.1
SDM-AG	94.8	93.5	81.9	**100.0**	81.9	**100.0**	92.0
EADA	**97.7**	<u>96.6</u>	82.1	**100.0**	<u>82.8</u>	**100.0**	<u>93.2</u>
TL-ADA	96.6	**96.8**	79.9	<u>99.8</u>	81.7	<u>99.8</u>	92.2
RASC	<u>96.8</u>	<u>96.6</u>	**83.4**	**100.0**	**84.0**	**100.0**	**93.5**
RASC-Ob	96.0	96.2	<u>82.6</u>	99.5	81.4	**100.0**	92.6

Table 1: Comparison results (Accuracy: %) on Office-31 with 5% labeling budget. The best accuracy is indicated in **bold** and the second best one is <u>underlined</u>.

Implementation details. All experiments are implemented using the Pytorch platform. For fair comparison, like [Fu *et al.*, 2021; Xie *et al.*, 2022b; Huang *et al.*, 2023a], we use ResNet-50 [He *et al.*, 2016] pre-trained on ImageNet [Krizhevsky *et al.*, 2017] as the backbone. For training, we use SGD optimizer with a learning rate of 0.01, momentum of 0.9, and weight decay 5e-4. We set the prototype update trade-off parameter λ to 0.9, the temperature τ in the similarity function to 0.1 and the trade-off parameters α and β to 0.5 and 0.7 respectively. The batch size is 36 for VisDA and 32 for other datasets. For sample selection, we set the labeling budget B to 5% and conduct 5 rounds of selection. More implementation details can be seen in the Appendix. Code is available at https://github.com/zeyuz22/RASC.

4.1 Comparative Results

We compare RASC with three types of methods. The first type is two baseline methods: Source Only (ResNet-50 trained with source domain only) and Random (target samples are randomly selected). The second type is several AL methods: Entropy [Wang and Shang, 2014], CoreSet [Sener and Savarese, 2018], and BADGE [Ash *et al.*, 2020]. The third type is the recently proposed ADA methods: AADA [Su *et al.*, 2020], CLUE [Prabhu *et al.*, 2021], TQS [Fu *et al.*, 2021], DBAL [de Mathelin *et al.*, 2022], SDM-AG [Xie *et al.*, 2022b], EADA [Xie *et al.*, 2022a], TL-ADA [Han *et al.*, 2023], and DiaNA [Huang *et al.*, 2023a].

Office-31. Results on Office-31 are shown in Table 1. We can observe that the performance is saturated, but our RASC still outperforms other methods significantly in most tasks, especially in some difficult tasks (e.g. D→A and W→A). In addition, due to the small size of Office-31, one-bit annotation cannot obtain sufficient supervision information, resulting in a decrease in the performance of RASC-Ob, but it still achieves performance second only to EADA.

Method	VisDA	Office-Home												
	Syn→Real	A→C	A→P	A→R	C→A	C→P	C→R	P→A	P→C	P→R	R→A	R→C	R→P	Avg
Source Only	44.7	42.1	66.3	73.3	50.7	59.0	62.6	51.9	37.9	71.2	65.2	42.6	76.6	58.3
Random	78.1	52.5	74.3	77.4	56.3	69.7	68.9	57.7	50.9	75.8	70.0	54.6	81.3	65.8
Entropy	82.7	58.0	78.4	79.1	60.5	73.0	72.6	60.4	54.2	77.9	71.3	58.0	83.6	68.9
CoreSet	81.9	51.8	72.6	75.9	58.3	68.5	70.1	58.8	48.8	75.2	69.0	52.7	80.0	65.1
BADGE	84.3	58.2	79.7	79.9	61.5	74.6	72.9	61.5	56.0	78.3	71.4	60.9	84.2	69.9
AADA	80.8	56.6	78.1	79.0	58.5	73.7	71.0	60.1	53.1	77.0	70.6	57.0	84.5	68.3
CLUE	85.2	58.0	79.3	80.9	68.8	77.5	76.7	66.3	57.9	81.4	75.6	60.8	86.3	72.5
TQS	83.1	58.6	81.1	81.5	61.1	76.1	73.3	61.2	54.7	79.7	73.4	58.9	86.1	70.5
DBAL	82.6	58.7	77.3	79.2	61.7	73.8	73.3	62.6	54.5	78.1	72.4	59.9	84.3	69.6
SDM-AG	80.3	61.2	82.2	82.7	66.1	77.9	76.1	66.1	58.4	81.0	76.0	62.5	87.0	73.1
EADA	88.3	63.6	84.4	83.5	70.7	83.7	80.5	73.0	63.5	85.2	78.4	65.4	88.6	76.7
TL-ADA	86.8	63.7	83.9	82.5	69.7	82.7	81.4	70.3	61.2	84.6	77.4	63.4	85.9	75.6
DiaNA	-	64.5	86.0	84.9	72.3	84.6	82.5	73.3	63.7	85.6	78.5	67.2	89.5	77.7
RASC	**88.7**	**71.6**	**87.3**	86.0	**72.6**	86.7	82.6	73.9	**69.6**	85.4	79.8	**73.3**	91.3	**80.0**
RASC-Ob	87.2	67.7	87.2	**86.8**	72.5	**86.9**	**84.3**	**74.1**	66.8	**87.6**	**80.4**	72.0	**91.6**	79.8

Table 2: Comparison results (Accuracy: %) on VisDA and Office-Home with 5% labeling budget.

Method	C→P	C→R	C→S	P→C	P→R	P→S	R→C	R→P	R→S	S→C	S→P	S→R	Avg
Source Only	52.1	63.0	49.4	55.9	73.0	51.1	56.8	61.0	50.0	54.0	48.9	60.3	56.3
Random	61.6	78.7	61.6	64.0	78.7	63.7	60.5	64.3	61.1	64.8	58.7	75.2	66.1
AADA	62.4	77.5	61.7	61.9	79.7	61.1	65.6	66.0	60.8	65.1	62.1	80.0	67.0
CLUE	57.6	77.5	58.6	58.9	76.8	65.9	66.3	60.2	60.5	66.2	58.7	76.0	65.3
TQS	67.8	**82.0**	65.4	67.5	84.8	66.1	63.8	67.2	62.5	71.1	64.4	81.6	70.4
DBAL	62.9	79.2	60.8	64.6	78.1	62.5	65.6	65.2	59.2	66.3	61.3	80.3	67.2
EADA	66.0	80.8	63.5	69.4	83.0	65.1	71.1	68.6	65.7	71.0	64.3	81.0	70.8
RASC	68.7	81.9	**73.2**	75.6	**87.3**	68.9	77.5	**72.1**	**72.4**	**77.8**	68.7	80.2	75.4
RASC-Ob	**70.0**	81.3	71.8	**77.0**	85.2	**72.9**	**77.6**	71.8	71.4	76.5	**70.3**	**82.9**	**75.7**

Table 3: Comparison results (Accuracy: %) on MiniDomainNet with 5% labeling budget.

Office-Home. Results on Office-Home are shown in Table 2. We can observe that our method performs best on all tasks. RASC and RASC-Ob also achieve +2.3% and +2.1% improvement in average accuracy compared with the SOTA method DiaNA respectively. It demonstrates the superiority of our sample selection and domain adaptation strategies.

VisDA. As shown in the first column of Table 2, RASC performs significantly better than AL and recent ADA methods. Besides, RASC-Ob also achieves performance second only to EADA. It demonstrates the effectiveness of our method for large-scale datasets with significant domain shift.

MiniDomainNet. We also conduct experiments on a hard dataset: MiniDomainNet, and the results are shown in Table 3. Similarly, our method still achieves the best performance, with RASC and RASC-Ob achieving +4.6% and +4.9% improvement in average accuracy. It is noteworthy that the performance of RASC-Ob is superior to RASC and other methods. This is because MiniDomainNet contains a large number of images from 126 categories, and for large-scale datasets with such a large number of categories, RASC-Ob can use one-bit annotation to obtain enough supervised information to improve performance. It demonstrates that using one-bit annotation in ADA can effectively alleviate the issue of ADA methods being difficult to apply to large-scale datasets containing a large number of categories.

Method	RAS	$\mathcal{L}_{sup}$	$\mathcal{L}_{align}$	$\mathcal{L}_{reg}$	$\mathcal{L}_{neg}$	$\mathcal{L}_{nl}$	VisDA
RASC	✓	✓	✓	✓	-	-	**88.7**
w/o RAS		✓	✓	✓	-	-	80.9
w/o $\mathcal{L}_{sup}$	✓		✓	✓	-	-	87.3
w/o $\mathcal{L}_{align}$	✓	✓		✓	-	-	87.9
w/o $\mathcal{L}_{reg}$	✓	✓	✓		-	-	88.1
RASC-Ob	✓	✓	✓	✓	✓	✓	**87.2**
w/o $\mathcal{L}_{neg}$	✓	✓	✓	✓		✓	86.5
w/o $\mathcal{L}_{nl}$	✓	✓	✓	✓	✓		84.5

Table 4: Ablation studies on VisDA with 5% labeling budget.

4.2 Additional Analysis

Ablation studies. To demonstrate the effect of the main components of RASC and RASC-Ob, we conduct ablation studies on VisDA. The results are shown in Table 4. Specifically, when RAS is removed, it indicates that the target sam-

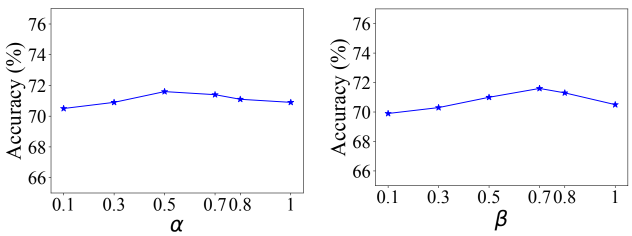

(a) Analysis of α (b) Analysis of β

Figure 3: Hyper-parameter sensitivity analysis of α and β in task A→C on Office-Home.

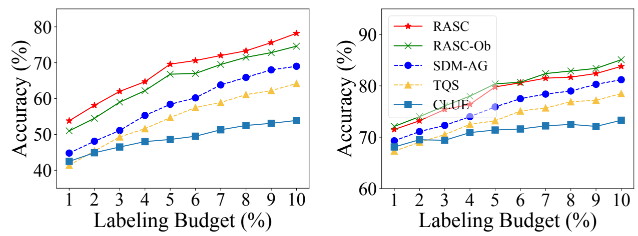

(a) P→C (b) R→A

Figure 4: Comparison results of varying percent of labeling budget (1% to 10%) on Office-Home.

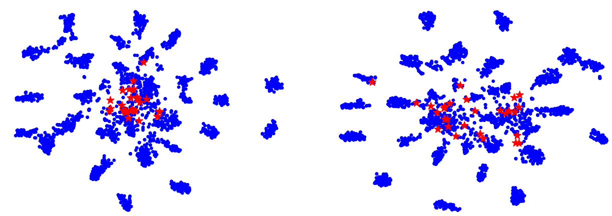

(a) RASC w/o CPR (b) RASC w/ CPR

Figure 5: T-SNE visualization of target features in task W→A on Office-31.

ples are randomly selected for annotation. From Table 4, we can observe that the performance of the full RASC and RASC-Ob significantly outperform other variants. RASC surpassing RASC w/o RAS demonstrates that our sample selection strategy can select valuable target samples for annotation, which can significantly improve adaptation performance when used for training. RASC surpassing RASC (w/o $\mathcal{L}_{sup}$, w/o $\mathcal{L}_{align}$ or w/o $\mathcal{L}_{reg}$) demonstrates that our contrastive learning-based gradual active domain adaptation framework can effectively alleviate domain shift. Furthermore, RASC-Ob surpassing RASC-Ob (w/o $\mathcal{L}_{neg}$ or w/o $\mathcal{L}_{nl}$) demonstrates that it is necessary to supplement the learning of misclassified samples when using one-bit annotation. More experiments and analyses can be seen in the Appendix.

Hyper-parameter sensitivity. To investigate the impact of the trade-off parameters α and β on the performance of RASC, we conduct experiments in task A→C on Office-Home. The results are shown in Figure 3. The larger α emphasizes more on reducing the distribution discrepancy between the labeled target and source domains, while the larger β emphasizes more on reducing the distribution discrepancy between the unlabeled target domain and labeled domains. We first set β to 0.7 while varying α from 0.1 to 1.0. We can observe that when α is very small (0.1), the performance of RASC is poor because the labeled target and source domains are not well aligned. As α increases, the performance gradually improves, demonstrating that it is necessary to consider the inter-domain discrepancy of labeled data in ADA. But when α is too large (> 0.5), the performance will decrease due to the over-alignment of the labeled target and source domains. Then, we set α to 0.5 while varying β from 0.1 to 1.0, and we can observe similar patterns for β.

Varying labeling budget. We show the performance of RASC and RASC-Ob under varying labeling budgets (1% to 10%) in Figure 4. We can observe that RASC and RASC-Ob outperform other ADA methods with varying labeling budgets, demonstrating that our method is suitable for various labeling budgets. More importantly, even if the labeling budget is small (1%), our method can still select the most valuable samples. Moreover, we can observe that as the labeling budget increases, our method can continuously improve the adaptation performance, demonstrating that our sample selection strategy is stable and effective, and can continuously select the most valuable samples for annotation.

How does RAS ensure diversity? While our sample selection strategy RAS mainly focuses on sample value estimation without explicitly introducing constraints on sample diversity, we identify that the class prediction reliability-based sample rejection strategy (CPR) can effectively ensure the diversity of the selected samples, as shown in Figure 5. The blue points are unlabeled target samples and the red stars are the selected samples. We can observe that samples selected with CPR are diverse and more dispersed in the feature space, while without CPR, the selected samples are confined to a small region and many of them are redundant. This is because CPR introduces randomness into the sample selection process, even if the RAVs of samples in a small area are high, these samples still have a probability of being rejected, increasing the probability of selecting samples from other areas.

5 Conclusion

In this paper, we propose a concise but effective method for active domain adaptation problem, termed Reconfigurability-Aware Selection for Contrastive active domain adaptation (RASC). Specifically, we first design a reconfigurability-aware sample selection strategy to select the most valuable target samples in the presence of domain shift. Then, we design a contrastive learning-based gradual active domain adaptation framework to better reduce the distribution discrepancy between labeled target domain and source domain, as well as between unlabeled target domain and labeled domains. Moreover, we propose a variant of RASC called RASC-Ob, which uses one-bit annotation method and supplements the learning of misclassified samples. Extensive experiments and analyses demonstrate the superiority of RASC.

Acknowledgments

We sincerely thank the anonymous reviewers for their careful work and thoughtful suggestions, which have greatly improved this article. This work was supported by the Natural Science Research Foundation of Jilin Province of China under Grant No. 20220101106JC, the Youth Growth Technology Program of Jilin Province Science and Technology Development Plan of China under Grant No. 20240602108RC, and the Fundamental Research Funds for the Central Universities of China under Grant Nos. 2412022ZD018, 2412022Q D040 and 93K172022K10.

References

[Ash *et al.*, 2020] Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *Proceedings of the International Conference on Learning Representations (ICLR 2020)*, 2020.

[de Mathelin *et al.*, 2022] Antoine de Mathelin, François Deheeger, Mathilde Mougeot, and Nicolas Vayatis. Discrepancy-based active learning for domain adaptation. In *Proceedings of the International Conference on Learning Representations (ICLR 2022)*, 2022.

[Fu *et al.*, 2021] Bo Fu, Zhangjie Cao, Jianmin Wang, and Mingsheng Long. Transferable query selection for active domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2021)*, pages 7272–7281, 2021.

[Ganin and Lempitsky, 2015] Yaroslav Ganin and Victor S. Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the International Conference on Machine Learning (ICML 2015)*, pages 1180–1189, 2015.

[Han *et al.*, 2023] Kyeongtak Han, Youngeun Kim, Dongyoon Han, Hojun Lee, and Sungeun Hong. TL-ADA: Transferable loss-based active domain adaptation. *Neural Networks*, 161:670–681, 2023.

[He *et al.*, 2016] 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 (CVPR 2016)*, pages 770–778, 2016.

[He *et al.*, 2020] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2020)*, pages 9726–9735, 2020.

[Huang *et al.*, 2023a] Duojun Huang, Jichang Li, Weikai Chen, Junshi Huang, Zhenhua Chai, and Guanbin Li. Divide and adapt: Active domain adaptation via customized learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2023)*, pages 7651–7660, 2023.

[Huang *et al.*, 2023b] Xinyang Huang, Chuang Zhu, and Wenkai Chen. Semi-supervised domain adaptation via prototype-based multi-level learning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2023)*, pages 884–892, 2023.

[Jiang *et al.*, 2020] Pin Jiang, Aming Wu, Yahong Han, Yunfeng Shao, Meiyu Qi, and Bingshuai Li. Bidirectional adversarial training for semi-supervised domain adaptation. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2020)*, pages 934–940, 2020.

[Kim *et al.*, 2019] Youngdong Kim, Junho Yim, Juseung Yun, and Junmo Kim. NLNL: Negative learning for noisy labels. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV 2019)*, pages 101–110, 2019.

[Krizhevsky *et al.*, 2017] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.

[Li *et al.*, 2020] Shuang Li, Chi Harold Liu, Qiuxia Lin, Binhui Xie, Zhengming Ding, Gao Huang, and Jian Tang. Domain conditioned adaptation network. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2020)*, pages 11386–11393, 2020.

[Li *et al.*, 2021] Shuang Li, Fangrui Lv, Binhui Xie, Chi Harold Liu, Jian Liang, and Chen Qin. Bi-classifier determinacy maximization for unsupervised domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2021)*, pages 8455–8464, 2021.

[Liu *et al.*, 2021] Zhuoming Liu, Hao Ding, Huaping Zhong, Weijia Li, Jifeng Dai, and Conghui He. Influence selection for active learning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV 2021)*, pages 9254–9263, 2021.

[Long *et al.*, 2015] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I. Jordan. Learning transferable features with deep adaptation networks. In *Proceedings of the International Conference on Machine Learning (ICML 2015)*, pages 97–105, 2015.

[Long *et al.*, 2017] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. Deep transfer learning with joint adaptation networks. In *Proceedings of the International Conference on Machine Learning (ICML 2017)*, pages 2208–2217, 2017.

[Long *et al.*, 2018] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I. Jordan. Conditional adversarial domain adaptation. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS 2018)*, pages 1647–1657, 2018.

[Peng *et al.*, 2017] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.

[Peng *et al.*, 2019] 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 International Conference on Computer Vision (ICCV 2019)*, pages 1406–1415, 2019.

[Prabhu *et al.*, 2021] Viraj Prabhu, Arjun Chandrasekaran, Kate Saenko, and Judy Hoffman. Active domain adaptation via clustering uncertainty-weighted embeddings. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV 2021)*, pages 8485–8494, 2021.

[Ren *et al.*, 2022] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B. Gupta, Xiaojiang Chen, and Xin Wang. A survey of deep active learning. *ACM Computing Surveys*, 54(9):180:1–180:40, 2022.

[Saenko *et al.*, 2010] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European Conference on Computer Vision (ECCV 2010)*, pages 213–226, 2010.

[Saito *et al.*, 2019] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV 2019)*, pages 8049–8057, 2019.

[Sener and Savarese, 2018] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *Proceedings of the International Conference on Learning Representations (ICLR 2018)*, 2018.

[Singh, 2021] Ankit Singh. CLDA: Contrastive learning for semi-supervised domain adaptation. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS 2021)*, pages 5089–5101, 2021.

[Sinha *et al.*, 2019] Samarth Sinha, Sayna Ebrahimi, and Trevor Darrell. Variational adversarial active learning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV 2019)*, pages 5971–5980, 2019.

[Su *et al.*, 2020] Jong-Chyi Su, Yi-Hsuan Tsai, Kihyuk Sohn, Buyu Liu, Subhransu Maji, and Manmohan Chandraker. Active adversarial domain adaptation. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV 2020)*, pages 728–737, 2020.

[Tang and Huang, 2021] Ying-Peng Tang and Sheng-Jun Huang. Dual active learning for both model and data selection. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2021)*, pages 3052–3058, 2021.

[Venkateswara *et al.*, 2017] 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 (CVPR 2017)*, pages 5385–5394, 2017.

[Wang and Shang, 2014] Dan Wang and Yi Shang. A new active labeling method for deep learning. In *International Joint Conference on Neural Networks (IJCNN 2014)*, pages 112–119, 2014.

[Wang and Zhang, 2020] Shanshan Wang and Lei Zhang. Self-adaptive re-weighted adversarial domain adaptation. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2020)*, pages 3181–3187, 2020.

[Wu *et al.*, 2022] Mingfei Wu, Chen Li, and Zehuan Yao. Deep active learning for computer vision tasks: Methodologies, applications, and challenges. *Applied Sciences*, 12(16):8103, 2022.

[Xie *et al.*, 2022a] Binhui Xie, Longhui Yuan, Shuang Li, Chi Harold Liu, Xinjing Cheng, and Guoren Wang. Active learning for domain adaptation: An energy-based approach. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2022)*, pages 8708–8716, 2022.

[Xie *et al.*, 2022b] Ming Xie, Yuxi Li, Yabiao Wang, Zekun Luo, Zhenye Gan, Zhongyi Sun, Mingmin Chi, Chengjie Wang, and Pei Wang. Learning distinctive margin toward active domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2022)*, pages 7983–7992, 2022.

[Yan and Huang, 2018] Yifan Yan and Sheng-Jun Huang. Cost-effective active learning for hierarchical multi-label classification. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2018)*, pages 2962–2968, 2018.

[Yan *et al.*, 2022] Zizheng Yan, Yushuang Wu, Guanbin Li, Yipeng Qin, Xiaoguang Han, and Shuguang Cui. Multi-level consistency learning for semi-supervised domain adaptation. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2022)*, pages 1530–1536, 2022.

[Zhan *et al.*, 2021] Xueying Zhan, Huan Liu, Qing Li, and Antoni B. Chan. A comparative survey: Benchmarking for pool-based active learning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2021)*, pages 4679–4686, 2021.

[Zhang *et al.*, 2022] Yuhang Zhang, Xiaopeng Zhang, Lingxi Xie, Jie Li, Robert C. Qiu, Hengtong Hu, and Qi Tian. One-bit active query with contrastive pairs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2022)*, pages 9687–9695, 2022.

[Zhou *et al.*, 2021] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain adaptive ensemble learning. *IEEE Transactions on Image Processing*, 30:8008–8018, 2021.