

Divide and Adapt: Active Domain Adaptation via Customized Learning

Duojun Huang^{1,2} Jichang Li³ Weikai Chen⁴ Junshi Huang⁵ Zhenhua Chai⁵ Guanbin Li^{1,2†}

¹School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China

²Research Institute, Sun Yat-sen University, Shenzhen, China

³The University of Hong Kong ⁴Tencent America ⁵Meituan

huangdj9@mail2.sysu.edu.cn, csjcli@connect.hku.hk, chenwk891@gmail.com

{huangjunshi, chaizhenhua}@meituan.com, liguanbin@mail.sysu.edu.cn

Abstract

Active domain adaptation (ADA) aims to improve the model adaptation performance by incorporating active learning (AL) techniques to label a maximally-informative subset of target samples. Conventional AL methods do not consider the existence of domain shift, and hence, fail to identify the truly valuable samples in the context of domain adaptation. To accommodate active learning and domain adaptation, the two naturally different tasks, in a collaborative framework, we advocate that a customized learning strategy for the target data is the key to the success of ADA solutions. We present Divide-and-Adapt (DiaNA), a new ADA framework that partitions the target instances into four categories with stratified transferable properties. With a novel data subdivision protocol based on uncertainty and domainness, DiaNA can accurately recognize the most gainful samples. While sending the informative instances for annotation, DiaNA employs tailored learning strategies for the remaining categories. Furthermore, we propose an informativeness score that unifies the data partitioning criteria. This enables the use of a Gaussian mixture model (GMM) to automatically sample unlabeled data into the proposed four categories. Thanks to the “divide-and-adapt” spirit, DiaNA can handle data with large variations of domain gap. In addition, we show that DiaNA can generalize to different domain adaptation settings, such as unsupervised domain adaptation (UDA), semi-supervised domain adaptation (SSDA), source-free domain adaptation (SFDA), etc.

1. Introduction

Domain adaptation (DA) approaches strive to generalize model trained on a labeled source domain to a target domain with rare annotation [5, 13, 16, 24] by coping with

Source Centroids	Target			
	(1) Confident Consistent	(2) Uncertain Consistent	(3) Uncertain Inconsistent	(4) Confident Inconsistent
Backpack				
Pan				
Fan				
Bike				
TV				
For source model:	Transferable centroids	Transferable margins	Need Annotation	Challenging Samples

Figure 1. The illustration of our proposed Divide-and-Adapt mechanism to divide target samples into different data subsets for customized learning.

the domain disparity. Nevertheless, DA methods are significantly outperformed by their supervised counterparts due to the scarceness of annotation as demonstrated in [4, 14, 27]. In practice, it is cost-effective to get a moderate amount of target samples labeled to boost the performance of domain adaptation. Active learning (AL) approaches seek to select samples with uncertainty [9, 12, 29, 30] and diversity [19, 25] to best benefit the model, which properly matches the demand. However, previous AL methods assume that both the labeled and unlabeled data follow the same distribution, such a strategy may become ineffective to the DA scenarios where the target data suffer from domain shift. The recently proposed active domain adaptation (ADA) [4, 22, 31] aims to resolve this issue by actively selecting the maximally-informative instances such that the performance of the transferred model can be best boosted with a limited annotation budget.

The key to the success of ADA is to strike a good balance between the highly coupled yet inherently different tasks: active learning and domain adaptation. The real-world target data typically exhibit either of the two characteristics:

[†]Corresponding author is Guanbin Li.

source-like or *target-specific*. While the former has similar feature distribution with the source data, the latter tends to be the unique part of target domain and deviates greatly from the source distribution [4, 26, 38]. On one hand, to achieve domain adaptation, applying the same adaptation strategy to all target data equally cannot generalize well to scenarios with varying degrees of domain shift. This is particularly true when the gap between the source-like and target-specific data is unknown. On the other hand, in active learning tasks, the samples with a learning difficulty will be more likely to be selected for labeling. Nonetheless, with large domain gap, incorporating such difficult samples in the adaptation task would hamper the learning of the adaptation model, making the training highly unstable. However, despite the impressive progress that has been made, none of the existing works has fully addressed the above issues.

In this work, we propose Divide-And-Adapt (DiaNA), a novel ADA framework that can scale to large variations of domain gaps while achieving cost-effective data labeling with a significant performance boost for domain adaptation. Our key observation is that customized learning strategies are vital for target data with different characteristics. In particular, DiaNA divides the target data into four subsets with different levels of transferable properties (see Figure 1), each of which is handled with a customized learning strategy. Unlike traditional AL methods that would simply label the most uncertain data [29, 30, 37], we propose to withhold the most challenging samples (Figure 1 category (4)) for training the domain adaptation models. Instead, the selected samples for active annotation would maintain a proper stimulus for the source model, providing informative domain knowledge without jeopardizing the training stability. The subdivision of target data is dynamically updated as the domain disparity is gradually mitigated with more labeled data. Hence, the previous challenging samples could be classified as transferable in the later stage and exploited in the network training.

We introduce a novel protocol for subdividing the target samples for customized learning. In addition to the uncertainty of model prediction, we advocate that the consistency with the learned prior distribution, *i.e.* the domainness, is another key criterion for active domain adaptation [4, 26]. To this end, we divide the target data into four categories as shown in Figure 1 according to the domainness and uncertainty of the instances. We further propose that the samples with 1) being *uncertain* to the model and 2) having *inconsistent* prediction with the label of its closest category prototype in the learned feature space (*i.e.* high domainness) are the most “profitable” instances for bringing informative knowledge of target domain if annotated. Thereby, we identify the *uncertain inconsistent* samples for labeling while applying tailored learning strategies for the remaining categories to boost the selectivity of the sampling.

To avoid heuristic thresholding for data subdivision, we propose an automatic data sampling mechanism based on Gaussian mixture model (GMM). In particular, we propose an *informativeness function* that incorporates the domainness and uncertainty in a unified scoring system. The computed informativeness score of the labeled data is used to train a four-component GMM model, which is then applied to sample the unlabeled target data into four categories.

We evaluate DiaNA over a large variety of domain shift scenarios on DomainNet [21], Office-Home [28] and CIFAR-10 [11]. Furthermore, the proposed sampling strategy of DiaNA can be generalized to various domain adaptation problems with different supervision settings, including unsupervised domain adaptation (UDA), semi-supervised domain adaptation (SSDA), and source-free domain adaptation (SFDA).

We summarize our contributions as follows:

- A general “divide-and-adapt” framework, coded DiaNA, for active domain adaptation that can handle diversified degrees of domain gaps while being able to generalize to different domain adaptation problems, including UDA, SSDA, and SFDA.
- A new target data partition strategy based on domainness and uncertainty that enables stratified learning to achieve more stable training, superior adaptation performance, and better generality.
- A novel informativeness scoring system and the corresponding sampling paradigm based on GMM model for automatic data partitioning.
- New state-of-the-art performance over the mainstream public datasets in the task of active domain adaptation.

2. Related Work

Domain adaptation (DA). Most of the prior methods involving domain adaptation follow the unsupervised domain adaptation (UDA) paradigm that only unlabeled data is accessible in the target domain. Early UDA approaches [16, 18, 32] utilized maximum mean discrepancy (MMD) to minimize the discrepancy of features from different domains to address the domain shift, while recent adversarial learning based methods [5, 17] have become popular solutions for UDA. Despite effectiveness, UDA is still significantly outperformed by its supervised variant as indicated in [4, 14, 27]. Semi-supervised domain adaptation (SSDA) [13, 24, 35] allows a few target instances to be labeled and generally achieves improved performance. However, SSDA assumes that the labeled instances are passively provided in advance, depriving the choice of informative samples beneficial to the model. In this work, we propose an active domain adaptation method to select the most informative samples, which can be incorporated into other DA methods to further boost their performance.

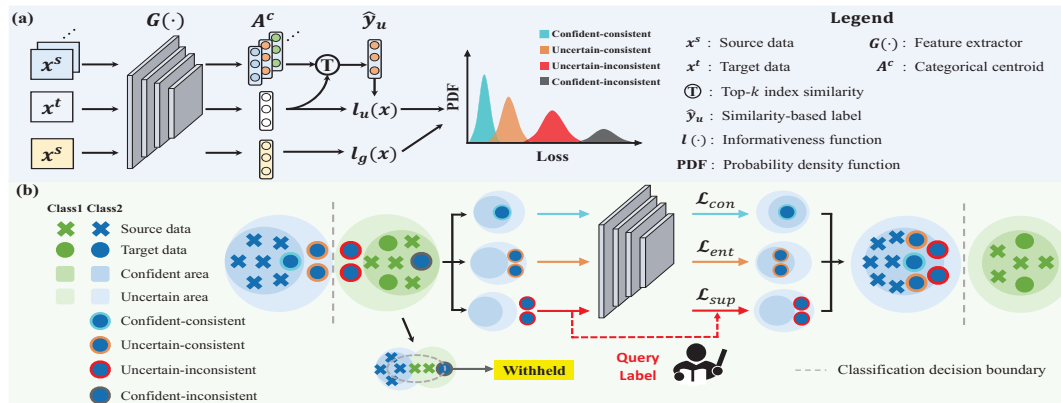


Figure 2. We propose a general “divide-and-adapt” framework, named DiaNA, for active domain adaptation. (a) A new target data partition strategy is presented to divide target data into four categories via building an informativeness scoring function that incorporates both domainness and uncertainty metrics. Besides, an automatic data partitioning function based on a learned four-component GMM model is applied to adaptively separate four sample categories from unlabeled target data. (b) To achieve customized learning, different auxiliary techniques are integrated into model fine-tuning w.r.t the four categories of unlabeled target data.

Active Learning (AL). Active learning is proposed to select a maximally-informative subset of unlabeled data for annotation via a query function, in an effort to best improve the model performance with a limited annotation budget. There exist several efforts to tackle the challenges involving active learning. For example, the uncertainty-based AL methods mainly focus on designing query functions based on prediction confidence [30], entropy [9, 29], or margin [10, 23] of the posterior probabilities. On the other hand, the diversity-based approaches [6, 19, 25] concentrate on selecting samples that can well represent the entire dataset. Meanwhile, expected model change-based methods [3, 15] aim to query the instances that would lead to a significant change to the current model. Recently proposed active learning approaches attain better performance by utilizing hybrid strategies which taking multiple metrics into consideration [1, 20]. Notwithstanding, the efficacy of previous AL resources may falter in DA scenarios as their selection strategy fails to meticulously consider the implications of the domain gap.

Active Domain Adaptation (ADA). Early ADA researches [4, 26] propose to measure the uncertainty and domainness of each target instance via a domain discriminator with adversarial training. CLUE [22] presents an entropy-weighted clustering algorithm to query uncertain and diverse samples in the target domain. Recently, SDM [33] is proposed to optimize a margin loss function for exploring target instances similar to potential hard samples in the source domain. Nevertheless, the majority of existing ADA methods deliberately design hand-crafted query functions to evaluate sample’s annotation value and apply the same adaptation strategy to all target data equally [4, 26, 31]. The rigid criteria make them easily overfit to certain domain adaptation scenarios, which limits the generalization of the

method. In contrast, we propose to integrate uncertainty and domainness metrics within a unified scoring function to select informative samples. Moreover, we apply tailored learning strategies for target data with different characteristics, constructing an effective framework capable of generalizing across large variations of domain shifts.

3. Divide-and-Adapt Framework

Problem formulation. Active domain adaptation aims at seeking an optimal model for a target domain when given labeled source data $\mathcal{S} = \{(x_s, y_s)\}$ as well as unlabeled target data $\mathcal{U} = \{x_u\}$, assisted by the iteratively queried labeled data $\mathcal{T} = \{(x_t, y_t)\}$ of the target domain. Here, \mathcal{S} , \mathcal{U} , and \mathcal{T} share the same label space over C categories. Initially, \mathcal{T} is an empty set. The model is iteratively trained using \mathcal{S} , \mathcal{T} , and \mathcal{U} for R active learning loops until a given annotation budget B is reached. In each active learning loop, $b = B/R$ samples from \mathcal{U} would be selected, labeled by human experts, and then moved into \mathcal{T} . The proposed Divide-and-Adapt framework consists of a feature extractor and a classifier. Our goal is to design a query function to identify informative samples to annotate and utilize limited amount of labeled samples to best enhance the classification performance for the target domain.

Overview. In this paper, we propose Divide-And-Adapt, coded DiaNA, to tackle the problem of active domain adaptation. In detail, we first propose to design an informativeness function for subdividing the target samples into four categories, one of which is the informative target samples used for actively annotating. Such a function is a unified scoring system that incorporates the domainness and the uncertainty. To avoid heuristic thresholding for data subdivision, we propose an automatic data sampling mechanism to divide the unlabeled target data into four

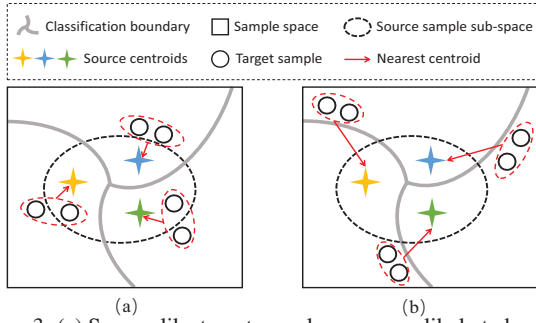


Figure 3. (a) Source-like target samples are more likely to be close to the source centroid in the feature space with respect to its predicted category. (b) Target-specific target samples are prone to be inconsistent between model predictions and similarity-based labels.

sample categories. This is achieved by learning a four-component Gaussian mixture model using semi-supervised Expectation-Maximization algorithm. Finally, we design tailored learning strategies for different categories of target samples during model fine-tuning, thus further improving the model performance. The overview of the proposed algorithm has been summarized in Figure 2.

3.1. Informativeness Scoring Mechanism

Following previous ADA works [4, 26, 31], we aim to search for the target instances with high domainness and uncertainty. To achieve this goal, we propose to first estimate category-wise centroids based on the labeled samples, and then define a similarity-based label for each unlabeled target sample by computing its top- k feature similarity with the categorical centroids. Finally, we integrate both the domainness and the uncertainty metrics into an informativeness function, which is formulated as the consistency between model prediction and similarity-based label.

Categorical centroids. We propose to average the features of labeled samples to estimate the category-specific data distribution, i.e., category prototype. Here, we formulate the centroid with respect to the c -th category as follows,

$$A^c = \frac{\sum_{(x,y) \in \mathcal{S}} \mathbb{1}\{y = c\} \cdot G(x)}{\sum_{(x,y) \in \mathcal{S}} \mathbb{1}\{y = c\}}, \quad (1)$$

where $\mathbb{1}\{\cdot\}$ is an indicator function, and $G(x)$ is a feature extractor that outputs the sample feature of x . Note that, as the active learning loops proceed, labeled target samples would be incorporated into the calculation of categorical centroids to mitigate the impact of distribution change.

Similarity-based data label. As introduced in the works [7, 13], samples share more top- k indices in their respective lists of rank-ordered feature elements would have a higher probability of being the same category. To obtain the

similarity-based label for each target instance, we first measure pairwise similarity of the indices of the sorted feature elements based on the magnitude. Then, we can formulate a similarity-based label for a sample x from \mathcal{U} as follows,

$$\hat{y}(x) = \operatorname{argmax}_{1 \leq c \leq C} \mathbf{IoU}(\operatorname{top}k(G(x)), \operatorname{top}k(A^c)), \quad (2)$$

where $\operatorname{top}k(\cdot)$ computes the top- k feature element indices ranked according to their magnitudes. In addition, $\mathbf{IoU}(\cdot, \cdot)$ denotes an Intersection-over-Union function, which returns the extent of overlap of the two respective index lists.

We set the value of k to be significantly smaller than the full dimension of the feature vector. In this way, the function $\operatorname{top}k(\cdot)$ only extracts the principle components of the feature of sample. Therefore, the IoU function in Eq. 2 is equivalent to measuring pairwise image sample similarity under a low-resolution condition. In ADA tasks, the training of model is inevitably dominated by the source domain especially in the early stage [33]. Therefore, source-like target samples that are similar to the source domain would naturally have more accurate representation than the target-specific ones that are distinct from the source. Under the low-resolution condition, i.e., k set to be small, source-like samples tend to have the same similarity-based label as the model predicted class thanks to their reliable and discriminative features extracted by the model. In contrast, target-specific samples are prone to produce inconsistent results since they are typically underfitted by the current model as shown in Figure 3. Based on the analysis above, we utilize the consistency of the model predicted class and similarity-based label to evaluate the domainness of each target sample. Detailed theoretical and experimental supports are demonstrated in our supplementary.

Informativeness function. To explore the most informative target samples, we integrate the domainness and uncertainty into a unified informativeness function (InfoF). The informativeness score for a sample x from unlabeled target domain can be formulated as follows,

$$\ell_u(x) = - \sum_{c=1}^C \mathbb{1}\{c = \hat{y}(x)\} \cdot \log P_c(x), \quad (3)$$

where $P(x)$ denotes the probabilistic prediction of a sample x from the network model and $P_c(x)$ is the c -th element of the vector $P(x)$. According to Eq. 3, we can divide unlabeled target samples into two subsets depending on whether the model predicted class and the similarity-based label are identical: consistent and inconsistent subsets, while both sets have confident samples and uncertain samples depending on whether the predicted probability is greater than a given confidence threshold. We name these four categories of unlabeled target data as *confident-consistent* (CC), *uncertain-consistent* (UC), *uncertain-inconsistent* (UI) and

Algorithm 1: The training procedure of four-component Gaussian Mixture Model (GMM) using semi-supervised Expectation-Maximization

Input: $\mathcal{D}_L; \mathcal{D}_U; \alpha$

Output: Optimal GMM parameters \mathcal{W}

```

1 Initialize  $\mathcal{W} = \{(\pi_k, \mu_k, \sigma_k) | 1 \leq k \leq K\}$  using  $\mathcal{D}_L$ 
  via Bayes estimation.
2 while  $\mathcal{W}$  is not converged do
3    $\{(\ell_g^i, p_g^i)\}_{i=1}^{|\mathcal{D}_L|} \leftarrow \mathcal{D}_L$ 
4    $\{(\ell_u^j, )\}_{j=1}^{|\mathcal{D}_U|} \leftarrow \mathcal{D}_U$ 
5   // E step:
6    $\gamma_{gk}(\ell_g^i) = \mathbb{1}\{p_g^i = k\},$ 
7    $\gamma_{uk}(\ell_u^j) = \frac{\pi_k \mathcal{N}(\ell_u^j; \mu_k, \sigma_k)}{\sum_{k'=1}^K \pi_{k'} \mathcal{N}(\ell_u^j; \mu_{k'}, \sigma_{k'})},$ 
8   // M step:
9    $\pi_k = \frac{\alpha \sum_i \gamma_{gk}(\ell_g^i) + (1-\alpha) \sum_j \gamma_{uk}(\ell_u^j)}{\alpha |\mathcal{D}_L| + (1-\alpha) |\mathcal{D}_U|};$ 
10   $\mu_k = \frac{\alpha \sum_i \gamma_{gk}(\ell_g^i) \cdot \ell_g^i + (1-\alpha) \sum_j \gamma_{uk}(\ell_u^j) \cdot \ell_u^j}{\alpha \sum_i \gamma_{gk}(\ell_g^i) + (1-\alpha) \sum_j \gamma_{uk}(\ell_u^j)};$ 
11   $\sigma_k = \frac{\alpha \sum_i \gamma_{gk}(\ell_g^i) \cdot (\ell_g^i - \mu_k)^2 + (1-\alpha) \sum_j \gamma_{uk}(\ell_u^j) \cdot (\ell_u^j - \mu_k)^2}{\alpha \sum_i \gamma_{gk}(\ell_g^i) + (1-\alpha) \sum_j \gamma_{uk}(\ell_u^j)};$ 

```

confident-inconsistent (CI), respectively. In this way, samples that are both uncertain to the model and contain target-specific knowledge can be extracted based on the InfoF score. In next section, we devise a strategy for informative sampling to adaptively distinguish those subsets from unlabeled target samples.

3.2. Informative Sampling Function

During training, the varying scale of InfoF scores obtained above would hinder accurate data categorization using manual data division, leading to a biased sampling strategy and degraded performance. Therefore, our work here formulates an Informative Sampling Function to achieve the goal of adaptively selecting informative target samples. In general, it can be observed that the four categories of unlabeled target data as stated above are significantly separable in the distributions of the InfoF scores, making it possible to learn mixture models to separate four sample partitioning from unlabeled target data. Thereafter, by learning the distribution of InfoF scores using a four-component Gaussian mixture model (GMM), the target data can be adaptively mapped to the probability distribution of four categories and then be separated, thereby avoiding sampling bias that may be created by manual division.

In specific, we first take inspiration from [34] that labeled samples can be used to construct supervision information, therefore building a more trustworthy GMM model with K Gaussian components (namely $K = 4$ in this work) via a semi-supervised Expectation-Maximization al-

gorithm. To this end, we then calculate the InfoF score for a labeled sample with (x, y) from \mathcal{S} or \mathcal{T} as follows,

$$\ell_g(x, y) = - \sum_{c=1}^C \mathbb{1}\{c = y\} \cdot \log P_c(x). \quad (4)$$

The similarity-based label is determined using categorical centroids of the labeled data, resulting in the totally identical similarity-based label and model prediction for labeled samples. In order to obtain reliable supervision and build a better GMM model, we have replaced the similarity-based label with the ground-truth class label here.

Afterwards, to obtain the supervision information for training GMM, we here define a piecewise function to obtain the observation labels that assign all labeled samples across domains to the corresponding Gaussian components. Thus, the observation label of a candidate instance x can be computed as,

$$q(x, y) = \begin{cases} 1, & \max P(x) \geq \tau \text{ and } y = \arg \max P(x), \\ 2, & \max P(x) < \tau \text{ and } y = \arg \max P(x), \\ 3, & \max P(x) < \tau \text{ and } y \neq \arg \max P(x), \\ 4, & \max P(x) \geq \tau \text{ and } y \neq \arg \max P(x), \end{cases} \quad (5)$$

where the numbers, i.e., $k = 1, \dots, 4$, indicate the indices of different Gaussian components corresponding to the CC, UC, UI and CI sample subsets, respectively. As well, τ is a confidence threshold to classify labeled samples into the confident subset and the uncertain subset. Furthermore, the comparison between the ground-truth and predicted class labels is used to divide them into the consistent subset and the inconsistent subset of labeled training samples.

Once we obtain $\{\ell_g(x, y) | (x, y) \in \mathcal{S} \cup \mathcal{T}\}$, $\{\ell_u(x) | x \in \mathcal{U}\}$, and $q(x, y)$, we incorporate them into building training instances for learning the GMM model as follows,

$$\mathcal{D}_L \leftarrow \{(\ell_g(x, y), q(x, y)) | (x, y) \in \mathcal{S} \cup \mathcal{T}\}, \quad (6)$$

$$\mathcal{D}_U \leftarrow \{(\ell_u(x),) | x \in \mathcal{U}\}. \quad (7)$$

For convenience, we denote the GMM model by $\mathcal{W} = \{(\pi_k, \mu_k, \sigma_k) | 1 \leq k \leq K\}$. Afterwards, given a training instance ℓ_x from \mathcal{D}_L or \mathcal{D}_U , the output of the probability density functions of the GMM model regarding ℓ_x can be computed as

$$p(\ell_x) = \sum_{k=1}^K \pi_k \mathcal{N}(\ell_x; \mu_k, \sigma_k), \quad (8)$$

where π_k represents the weight of the k -th Gaussian component subject to $\sum_{k=1}^K \pi_k = 1$ and $0 \leq \pi_k \leq 1$ for $k = 1, \dots, K$, while μ_k and σ_k are vectors denoting the mean and variance of such a component. Here, $\mathcal{N}(\cdot; \mu_k, \sigma_k)$ is used to model the Gaussian distribution of the k -th component. Note that a given coefficient $\alpha = |\mathcal{D}_L|/|\mathcal{D}_U|$ is

used for the weighted integration of the labeled and unlabeled training instances. Thus, we can summarize the training procedure of the GMM model using a semi-supervised EM algorithm in Algorithm 1 according to [34].

After obtaining the trained GMM parameters \mathcal{W} , the probability of a sample x from \mathcal{U} being classified into the k -th component can be computed as follows,

$$\Pr(z = k|x, \mathcal{W}) = \frac{\pi_k \mathcal{N}(\ell_u(x); \mu_k, \sigma_k)}{\sum_{k'=1}^K \pi_{k'} \mathcal{N}(\ell_u(x); \mu_{k'}, \sigma_{k'})}, \quad (9)$$

where z is a discrete variable to infer which component the sample x belongs to.

3.3. Candidate Selection and Training Objectives

After acquiring the informativeness sampling function, it is utilized for carrying out data partitioning. For ADA, the main focus should be on annotating samples that are uncertain to the model and can effectively represent the target dataset, which is termed as *uncertain-inconsistent* samples within our subdivision framework. Therefore, during each sampling step, we first rank the unlabeled target samples in descending order according to the scores predicted by $\Pr(z = 3|\cdot, \mathcal{W})$, and then the top b examples are selected to annotate and moved into labeled target data. Following previous ADA works, we apply the standard cross-entropy loss to the labeled data as follows:

$$\mathcal{L}_{sup} = \mathbb{E}_{(x,y) \sim \mathcal{S} \cup \mathcal{T}} \left[- \sum_{c=1}^C \mathbb{1}\{c = y\} \cdot \log P_c(x) \right]. \quad (10)$$

To boost the selectivity of the sampling strategy, we further design tailored training techniques for the other data subsets with varied transferable properties. We perform data partitioning on the remaining unlabeled target samples, i.e., $\mathcal{U} \leftarrow \mathcal{U} \setminus \mathcal{T}$, according to the posterior probabilities regarding the rest of Gaussian components.

$$\hat{\mathcal{U}}_k = \{x | \arg \max \mathbf{Pr}(x) = k, \forall (x, \cdot) \in \mathcal{U}\}, \text{ for } k = 1, 2, 4, \quad (11)$$

where $\hat{\mathcal{U}}_1$, $\hat{\mathcal{U}}_2$ and $\hat{\mathcal{U}}_4$ refer to $\hat{\mathcal{U}}_{CC}$, $\hat{\mathcal{U}}_{UC}$, $\hat{\mathcal{U}}_{CI}$. Here, $\mathbf{Pr}(x)$ is the concatenated probability vector corresponding to different components, i.e., $\mathbf{Pr}(x) = [\Pr(z = 1|x, \mathcal{W}), \Pr(z = 2|x, \mathcal{W}), \Pr(z = 3|x, \mathcal{W}), \Pr(z = 4|x, \mathcal{W})]$ w.r.t the sample x .

As shown in Figure 1(4), the source-oriented model could produce confident but unreliable predictions for some target samples having large domain gap with the source. Samples from $\hat{\mathcal{U}}_{CI}$ have extremely controversial results between model prediction and similarity-based label, which are most likely to be some excessively challenging instances. Therefore, we withhold them in the current stage to avoid jeopardizing the training stability.

The samples in $\hat{\mathcal{U}}_{CC}$ have confident and reliable predictions, which tend to be some class centroid in the target feature space. Therefore, we apply consistency regularization to them to enforce the prediction consistency under different perturbations, which can be formulated as follows:

$$\mathcal{L}_{con} = \mathbb{E}_{(x, \cdot) \sim \hat{\mathcal{U}}_{CC}} \left[- \sum_{c=1}^C \mathbb{1}\{c = \hat{y}(x)\} \cdot \log P_c(\mathbf{Aug}(x)) \right], \quad (12)$$

where $\mathbf{Aug}(\cdot)$ is a function to create perturbations for the sample x using classic data augmentation techniques such as AutoAugment [2].

For data group $\hat{\mathcal{U}}_{UC}$ whose prediction tend to be reliable but uncertain, we directly minimize the entropy of the predictions to using the following conditional entropy loss:

$$\mathcal{L}_{ent} = \mathbb{E}_{(x, \cdot) \sim \hat{\mathcal{U}}_{UC}} \left[- \sum_{c=1}^C P_c(x) \log P_c(x) \right]. \quad (13)$$

In summary, the overall loss functions used to further optimize the model can be formulated as follows,

$$\mathcal{L} = \mathcal{L}_{sup} + \lambda_c \mathcal{L}_{con} + \lambda_e \mathcal{L}_{ent} \quad (14)$$

where λ_c and λ_e are the weights to trade off different loss terms during the training process. As shown in Figure 2, \mathcal{L}_{sup} can calibrate the potentially underfitting of uncertain and target-specific samples. \mathcal{L}_{con} can guide the model to learn global cross-domain clustering while \mathcal{L}_{ent} serves to minimize the distance of data from the same class. As demonstrated in Sec. 4.3, these training objectives are complementary to each other to best adapt the model to the target domain.

3.4. Compatibility with UDA/SSDA/SFDA

The proposed sampling strategy of DiaNA requires no additional network modules like a domain discriminator or multiple classifiers [4, 26], making it easy to be integrated into existing DA frameworks, such as unsupervised domain adaptation (UDA), semi-supervised domain adaptation (SSDA), and source-free domain adaptation (SFDA); see supplementary document for the implementation details. It should be noted that we further design a variant of DiaNA for the SFDA settings, due to the unavailability of source data. As evidenced by Sec. 4.3, when integrated into diverse DA techniques, the proposed sampling strategy can deliver more performance improvement than random sample selection, thereby demonstrating its compatibility in various DA settings. To the best of our knowledge, we are the first to propose an active selection strategy that can generalize to UDA, SSDA, and SFDA methods.

Method	R → C			C → S			S → P			C → Q			R → S			R → P			AVG		
	1k	2k	5k	1k	2k	5k	1k	2k	5k	1k	2k	5k	1k	2k	5k	1k	2k	5k	1k	2k	5k
Random	50.9	55.2	61.5	42.6	44.8	49.4	41.0	44.4	50.1	26.3	33.2	43.0	38.2	40.7	46.9	46.5	48.8	53.7	40.9	44.5	50.8
CONF	48.6	54.2	61.3	41.9	43.6	48.0	41.6	44.2	48.5	25.5	32.7	42.8	36.9	39.4	45.3	47.1	47.8	51.2	40.3	43.6	49.5
Entropy [29]	48.0	54.2	61.9	42.7	45.2	48.9	41.1	43.7	48.7	24.6	31.0	42.1	37.0	40.4	45.7	46.5	48.4	52.2	40.0	43.8	49.9
Coreset [25]	49.1	52.5	59.8	40.6	42.2	45.8	39.1	41.5	44.8	24.7	30.3	39.7	36.7	38.3	43.3	45.1	46.8	50.3	39.2	41.9	47.3
BADGE [1]	51.7	55.4	62.5	43.6	45.7	50.1	42.3	45.5	49.7	26.0	34.3	43.9	37.8	41.5	47.8	48.0	50.1	54.2	41.6	45.4	51.4
Alpha [20]	52.0	56.3	62.8	44.3	46.8	50.5	41.9	44.9	49.7	27.2	34.6	44.4	38.6	41.3	47.9	47.6	50.4	53.2	41.9	45.7	51.4
AADA [26]	50.8	56.3	64.7	42.8	45.7	51.0	41.8	45.6	50.8	25.2	31.3	40.8	38.5	41.7	49.5	46.5	49.7	55.1	40.9	45.1	52.0
SDM [33]	51.0	55.4	62.3	43.1	45.2	49.6	41.9	44.9	49.3	26.9	35.0	43.8	37.3	40.2	46.5	47.4	49.7	53.2	41.3	45.1	50.8
EADA [31]	51.5	56.1	63.9	44.2	46.3	51.7	43.8	46.1	50.3	27.5	34.8	42.9	39.0	41.5	46.4	48.6	50.4	54.1	42.4	45.9	51.2
CLUE [22]	54.1	59.4	66.0	45.3	48.7	53.7	44.0	48.0	53.2	28.8	35.5	43.2	40.4	44.4	51.1	50.0	52.7	56.4	43.8	48.1	53.9
DiaNA(Ours)	55.6	60.4	68.7	47.0	50.8	56.6	44.2	49.0	55.4	30.2	39.5	52.1	42.4	47.9	55.0	50.5	53.6	59.1	45.0	50.2	57.8

Table 1. Comparison results (Accuracy: %) on DomainNet with 1k, 2k and 5k labeling budgets. “Random” and “CONF” correspond to the classic AL approaches “Random Sampling” and “Least-Confidence Sampling”.

Method	Office-Home												AVG
	A → C	A → P	A → R	C → A	C → P	C → R	P → A	P → C	P → R	R → A	R → C	R → P	
Random	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 [29]	51.3	72.7	76.4	61.7	74.1	72.9	57.6	51.1	76.6	69.8	57.1	82.3	67.0
Coreset [25]	52.8	73.3	75.5	59.8	73.4	70.8	58.1	52.6	75.8	69.3	56.5	82.5	66.7
BADGE [1]	57.2	75.2	76.9	61.5	77.2	71.9	60.4	53.6	78.0	70.8	61.1	84.3	69.0
Alpha [20]	57.0	79.4	78.2	61.6	78.0	74.1	58.9	54.2	78.1	71.7	61.1	84.6	69.7
AADA [26]	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 [22]	58.0	79.1	77.0	61.3	78.0	73.1	60.4	55.9	77.9	70.9	60.3	84.1	69.7
TQS [4]	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
SDM [33]	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 [31]	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
DiaNA(Ours)	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

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

4. Experiments

4.1. Dataset

We validate the effectiveness of the proposed approach DiaNA on three standard benchmark datasets: DomainNet [21], Office-Home [28] and CIFAR-10 [11]. On DomainNet, we select five domains namely Real (R), Clipart (C), Sketch (S), Painting (P), and Quicdraw (Q), while we hire these domains to construct 6 adaptation scenarios. As in [4, 31, 33], we report 12 different adaptation scenarios on Office-Home, constructed from four domains: Real (R), Clipart (C), Art (A), and Product (P). To further evaluate the generalization of the proposed sample selection strategy, we additionally report the classification results on CIFAR-10 with a standard active learning setting.

4.2. Implementation

Here, we mainly focus on the detailed descriptions of the settings involving active learning on DomainNet, Office-Home, and CIFAR-10, while more implementation details can be found in the supplementary material. In each adaptation scenario, we report the average accuracy over 3 trials. **DomainNet.** Similar to CLUE [22], we consider ResNet-34 [8] as the network backbone, where such network model is pre-trained on all labeled source samples over 50 training epochs through standard supervision. In all adaptation cases, we set $B = 5000$ and $R = 10$.

Office-Home. To be fair, we follow the experimental settings of existing ADA works [4, 31, 33] to conduct

performance comparison. In particular, we also use ResNet-50 [8] as the backbone model, while we choose a 5% proportion of unlabeled target data to set the labeling budget B , and R is set to 5.

CIFAR-10. On this dataset, we select ResNet-34 [8] as the backbone model. We randomly choose select 10% training samples at random as the initially labeled data set, and we then annotate 4% at each sampling step with a total annotation budget 30%.

4.3. Main Results

Results on DomainNet and Office-Home. We summarize the comparison with baselines in Table 1 and Table 2 on both datasets respectively. Note that we consider varying label budgets, namely i.e., 1k, 2k, and 5k, on DomainNet. We can see that DiaNA achieves the highest average accuracy in all cases on both datasets for solving the ADA tasks. Especially in the most difficult adaptation scenarios, e.g., $C \rightarrow A$ and $R \rightarrow C$, on Office-Home, the proposed method can still exceed all the other methods by more than 1.6% and 1.8% respectively, demonstrating the effectiveness of DiaNA on the commonly used DA benchmark datasets.

Results on CIFAR-10. We also extend the proposed method to standard AL tasks on CIFAR-10, where the domain gap between labeled and unlabeled ones is minor. Figure 4(a) lists the comparisons of DiaNA against classic AL strategies involving [1, 20, 25, 29]. It can be observed that compared to classic AL algorithms, with fewer labeling budgets, the effect of the proposed method is still notice-




Figure 4. (a) Comparison results on CIFAR-10. (b) Combination with UDA/SSDA methods. (c) Combination with SFDA methods.

Method	AL	\mathcal{L}_{con}	\mathcal{L}_{ent}	GMM	AVG
DiaNA	UI	✓	✓	✓	77.7±0.2
DiaNA w/o AL	R	✓	✓	✓	74.8±0.4
DiaNA w/o \mathcal{L}_{con}	UI	✓	✓	✓	76.9±0.3
DiaNA w/o \mathcal{L}_{ent}	UI	✓	✓	✓	75.8±0.3
DiaNA w/o GMM	UI	✓	✓	✓	74.0±0.2
DiaNA-UI	UI			✓	72.1±0.4
DiaNA-UC	UC			✓	68.3±0.6
DiaNA-CI	CI			✓	66.4±0.8
DiaNA-CC	CC			✓	63.0±0.3
Random (Baseline)	R				65.8±0.5

Table 3. Ablation study results of DiaNA. The “AL” column indicates which type of data subset should be selected to be annotated during each sampling step. We evaluated the performance of all adaptation scenarios on Office-Home by averaging their accuracy.

able, resulting from its significant performance improvement. This also indicates that our sampling strategy can be effectively applied to general active learning.

Combination with UDA/SSDA/SFDA methods. To validate the compatibility of the proposed DiaNA with existing UDA, SSDA, and SFDA algorithms, we conduct the experiments for UDA and SSDA in the adaptation $C \rightarrow S$ on DomainNet, and for SFDA in $C \rightarrow A$ and $R \rightarrow C$ on Office-Home. DANN [5], MME [24], and NRC [36] correspond to the SOTA methods towards UDA, SSDA, and SFDA, respectively. As shown in Figure 4(b) and Figure 4(c), combining the selection strategy of DiaNA with other DA methods, i.e., DANN+Ours, MME+Ours, and NRC+Ours, significantly increases the performance of existing DA algorithms that select labeled target samples via random selection, namely DANN+Random, MME+Random and NRC+Random. This suggests that, adaptively considering both the domainness and uncertainty, DiaNA is beneficial for the selection of informative target sample searching, while combined into diverse domain adaptation techniques.

4.4. Ablation study

Efficacy of the targeted training strategies. To illustrate the effect of each component of the proposed DiaNA, we conduct ablation study by removing the corresponding individual components. As displayed in Table 3, the proposed method exceeds all of the model variants by a large accuracy margin on average, demonstrating the effectiveness of each component. By comparing our proposed entire approach with DiaNA w/o GMM, manual data division strategy might considerably cause performance degradation, thereby in turn emphasizing the need for adaptive

data subdivision protocol. Furthermore, as demonstrated, disrupting the correspondence between targeted data subsets provides empirical evidence supporting the validity of our proposed customized learning techniques; see supplementary document for more details.

Efficacy of the sampling strategy. To provide insights into how effectively the domainness-based metric and the uncertainty-based metric work in the sampling step, we use different data subsets to replace the UI subset as active samples for annotation. To eliminate the influence of \mathcal{L}_{con} and \mathcal{L}_{ent} , we remove these two terms but only retain \mathcal{L}_{sup} during training. As shown in Table 3, both DiaNA-UC and DiaNA-CI achieve better classification performance than baseline selection, thereby showing the effect of the proposed uncertainty metric and domainness metric. Moreover, active sampling on UI samples (DiaNA-UI) obtains the best performance among all the variants, further illustrating the superiority of the integration of both metrics.

5. Conclusions

In this paper, we introduce a “divide-and-adapt” framework, named DiaNA, to address the problem of active domain adaptation. Specifically, we first design an informativeness function that jointly captures sample’s uncertainty and domainness. Moreover, we devise an automatic data subdivision protocol to partition the target instances into four categories with different characteristics. While selecting the most informative samples to annotate, we also design tailored learning strategies for the other target data subsets with stratified transferable properties. Extensive experimental results, as well as ablation studies, have confirmed the superiority of the proposed approach.

Acknowledgments

This work was supported in part by the Guangdong Basic and Applied Basic Research Foundation (NO. 2020B1515020048), in part by the National Natural Science Foundation of China (NO. 61976250), in part by the Shenzhen Science and Technology Program (NO. JCYJ20220530141211024) and in part by the Fundamental Research Funds for the Central Universities under Grant 22lqgb25.

References

- [1] Jordan T Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *International Conference on Learning Representations*, 2019.
- [2] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 113–123, 2019.
- [3] Alexander Freytag, Erik Rodner, and Joachim Denzler. Selecting influential examples: Active learning with expected model output changes. In *European conference on computer vision*, pages 562–577. Springer, 2014.
- [4] Bo Fu, Zhangjie Cao, Jianmin Wang, and Mingsheng Long. Transferable query selection for active domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7272–7281, 2021.
- [5] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- [6] Yuhong Guo. Active instance sampling via matrix partition. *Advances in Neural Information Processing Systems*, 23, 2010.
- [7] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Automatically discovering and learning new visual categories with ranking statistics. In *International Conference on Learning Representations*, 2020.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [9] Sheng-Jun Huang, Jia-Wei Zhao, and Zhao-Yang Liu. Cost-effective training of deep cnns with active model adaptation. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1580–1588, 2018.
- [10] Ajay J Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for image classification. In *2009 IEEE conference on computer vision and pattern recognition*, pages 2372–2379. IEEE, 2009.
- [11] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [12] David D Lewis and Jason Catlett. Heterogeneous uncertainty sampling for supervised learning. In *Machine learning proceedings 1994*, pages 148–156. Elsevier, 1994.
- [13] Jichang Li, Guanbin Li, Yemin Shi, and Yizhou Yu. Cross-domain adaptive clustering for semi-supervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2505–2514, 2021.
- [14] Jian Liang, Dapeng Hu, and Jiashi Feng. Domain adaptation with auxiliary target domain-oriented classifier. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16632–16642, 2021.
- [15] Zhuoming Liu, Hao Ding, Huaping Zhong, Weijia Li, Jifeng Dai, and Conghui He. Influence selection for active learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9274–9283, 2021.
- [16] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International conference on machine learning*, pages 97–105. PMLR, 2015.
- [17] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. *Advances in neural information processing systems*, 31, 2018.
- [18] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *International conference on machine learning*, pages 2208–2217. PMLR, 2017.
- [19] Hieu T Nguyen and Arnold Smeulders. Active learning using pre-clustering. In *Proceedings of the twenty-first international conference on Machine learning*, page 79, 2004.
- [20] Amin Parvaneh, Ehsan Abbasnejad, Damien Teney, Gholamreza Reza Haffari, Anton van den Hengel, and Javen Qinfeng Shi. Active learning by feature mixing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12237–12246, 2022.
- [21] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415, 2019.
- [22] Viraj Prabhu, Arjun Chandrasekaran, Kate Saenko, and Judy Hoffman. Active domain adaptation via clustering uncertainty-weighted embeddings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8505–8514, 2021.
- [23] Dan Roth and Kevin Small. Margin-based active learning for structured output spaces. In *European Conference on Machine Learning*, pages 413–424. Springer, 2006.
- [24] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8050–8058, 2019.
- [25] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018.
- [26] Jong-Chyi Su, Yi-Hsuan Tsai, Kihyuk Sohn, Buyu Liu, Subhansu Maji, and Manmohan Chandraker. Active adversarial domain adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 739–748, 2020.
- [27] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7472–7481, 2018.

- [28] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5018–5027, 2017.
- [29] Dan Wang and Yi Shang. A new active labeling method for deep learning. In *2014 International joint conference on neural networks (IJCNN)*, pages 112–119. IEEE, 2014.
- [30] Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. Cost-effective active learning for deep image classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(12):2591–2600, 2016.
- [31] 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*, volume 36, pages 8708–8716, 2022.
- [32] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*, pages 478–487. PMLR, 2016.
- [33] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 7993–8002, 2022.
- [34] Heng-Chao Yan, Jun-Hong Zhou, and Chee Khiang Pang. Gaussian mixture model using semisupervised learning for probabilistic fault diagnosis under new data categories. *IEEE Transactions on Instrumentation and Measurement*, 66(4):723–733, 2017.
- [35] Zizheng Yan, Yushuang Wu, Guanbin Li, Yipeng Qin, Xiaoguang Han, and Shuguang Cui. Multi-level consistency learning for semi-supervised domain adaptation. *arXiv preprint arXiv:2205.04066*, 2022.
- [36] Shiqi Yang, Joost van de Weijer, Luis Herranz, Shangling Jui, et al. Exploiting the intrinsic neighborhood structure for source-free domain adaptation. *Advances in Neural Information Processing Systems*, 34:29393–29405, 2021.
- [37] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 93–102, 2019.
- [38] Ziyi Zhang, Weikai Chen, Hui Cheng, Zhen Li, Siyuan Li, Liang Lin, and Guanbin Li. Divide and contrast: Source-free domain adaptation via adaptive contrastive learning. *arXiv preprint arXiv:2211.06612*, 2022.