

Learn from the Learnt: Source-Free Active Domain Adaptation via Contrastive Sampling and Visual Persistence

Mengyao Lyu^{1,2*}, Tianxiang Hao^{1,2*}, Xinhao Xu^{1,2}, Hui Chen^{1,2✉}, Zijia Lin¹,
Jungong Han^{1,2}, and Guiguang Ding^{1,2✉}

¹ Tsinghua University, Beijing, China

² BNRist, Beijing, China

mengyao.lyu@outlook.com, beyondhtx@gmail.com, {xxh22@mails,
huichen@mail}.tsinghua.edu.cn, linzijia07@tsinghua.org.cn,
jungonghan77@gmail.com, dinggg@tsinghua.edu.cn

Abstract. Domain Adaptation (DA) facilitates knowledge transfer from a source domain to a related target domain. This paper investigates a practical DA paradigm, namely Source data-Free Active Domain Adaptation (SFADA), where source data becomes inaccessible during adaptation, and a minimum amount of annotation budget is available in the target domain. Without referencing the source data, new challenges emerge in identifying the most informative target samples for labeling, establishing cross-domain alignment during adaptation, and ensuring continuous performance improvements through the iterative query-and-adaptation process. In response, we present *learn from the learnt (LFTL)*, a novel paradigm for SFADA to leverage the learnt knowledge from the source pretrained model and actively iterated models without extra overhead. We propose *Contrastive Active Sampling* to learn from the hypotheses of the preceding model, thereby querying target samples that are both informative to the current model and persistently challenging throughout active learning. During adaptation, we learn from features of actively selected anchors obtained from previous intermediate models, so that the *Visual Persistence-guided Adaptation* can facilitate feature distribution alignment and active sample exploitation. Extensive experiments on three widely-used benchmarks show that our LFTL achieves state-of-the-art performance, superior computational efficiency and continuous improvements as the annotation budget increases. Our code is available at <https://github.com/lyumengyao/lftl>.

Keywords: Transfer learning · Domain adaptation · Active learning

1 Introduction

Deep neural networks have thrived in computer vision but struggle when real-world data deviates from the ideal independent and identical distribution of training data, causing performance drops in well-trained models. This is where domain adaptation (DA) comes into play, which enables knowledge transfer from a source domain to

* Equal Contribution ✉ Corresponding Author

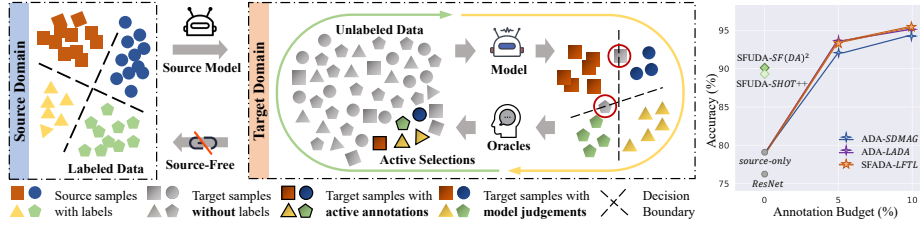


Fig. 1: (L) Source data-Free Active Domain Adaptation (SFADA) paradigm. Note that the ratio of labeled/unlabeled targets is much lower ($\leq 1\%$ or $\leq 5\%$ in our experiments) than in the illustration. (R) Performance comparison among DA SOTAs of different settings on Office-31.

a target domain of different data distribution. Most DA studies assume concurrent access to data from both domains to leverage the inter-domain relationship. Nonetheless, the reliance on labeled source data during adaptation could impede its widespread usage in real-world scenarios, considering stringent data protection regulations and resource deficiency in storage and computation. Therefore, source-free unsupervised DA (SFUDA) [16, 17, 25, 28, 38, 57, 63, 66] advocates for encapsulating knowledge within a source-trained model. They typically employ parameter sharing and hypothesis-driven learning for model adaptation. Yet, the absence of direct supervision from any domain can intensify the ill-posed nature of unsupervised DA [48], leading to the performance bottlenecks despite considerable efforts made.

In fact, the presence of the target sample pool during adaptation implies that minimal annotation effort is possible, yielding substantial performance gains. Hence we study a more practical DA setting, dubbed source-free active DA (SFADA). As shown in Fig 1 (L), the adaptation process is freed from source data, and meantime a minimum amount of annotation budget is available for iteratively querying labels in the target domain. Limited yet definitive supervisory signals intrinsically alleviate the ill-posedness phenomenon, and additionally, SFADA offers an appealing trade-off between labeling costs and adaptation performance regarding accuracy and efficiency.

Taking both target annotation availability and source data inaccessibility into consideration, SFADA presents new challenges. During query selection, general active learning criteria [3, 9, 19, 45, 46, 62] tend to fail under the domain shift phenomenon as observed in both previous Active DA (ADA) studies [8, 37, 39, 50, 51, 58] and our experiments 4.2. On top of that, ADA typically employs source data as a reference to identify distinctive target samples, which violates our source-free setting and thus is not compatible here either. Therefore, SFADA necessitates a new active learning strategy to *pinpoint the most innovative and informative target samples utilizing solely a source-trained model*. Another challenge comes at the transfer stage. Without the source data, manifold alignment, typically employed via minimizing distribution divergence to boost performance, is impeded. Yet the labeling costs have increased the expectations for performance enhancement. This prompts us to ask *how to exploit the newly acquired knowledge while consolidating the learnt domain-invariant information during adaptation?* And following these, *how to guarantee continuous performance improvement during the iterative query-and-adaptation process?*

To fill in this gap, this paper formulates a novel SFADA paradigm to *learn from the learnt* (LFTL), as illustrated in Fig. 2. We employ the hypothesis and visual representations of target samples obtained from the learnt source model and actively iterated models to ascertain what to supplement and where to adapt in the target domain. We first conduct *active learning from the learnt model*. Although it is impossible to directly identify target samples the most divergent from the source, the hypothesis of learnt model can be utilized as an indication of how well the target domain knowledge is understood. Furthermore, as the model progressively adapts towards the target domain, the hypothesis of the current model becomes increasingly important to query newer labels. This has not been addressed by the one-round query of MHPL [55]. Inspired from contrastive decoding [6, 11, 23, 27, 61], we propose *Contrastive Active Sampling (CAS)* to leverage the hypothesis of the model from the previous active round. It emphasizes samples that are both informative to the current model and persistently challenging throughout the iterative process, while deprioritizing samples that are redundant or deliver knowledge already acquired during active selection. Upon the contrastively decoded hypothesis, we take the difference between the best-versus-second best (BvSB) guesses [19], integrate a class-balancing factor, so that temporally-stagnant, sample-uncertain and class-minor targets can be effectively queried. Secondly, we perform *transfer learning from the learnt visual representations*. Since the alignment of cross-domain distributions cannot be directly achieved, we establish *Visual Persistence-guided Adaptation (VPA)* to maintain feature representations of active target samples throughout the whole process, where the understanding initially derived from the source domain and subsequently obtained from previous active rounds are effectively conserved via momentum updates. Then our learning objective encourages unlabeled samples to be centralized around these memory-preserving representatives, approximating a source-similar distribution while exploiting target-specific knowledge obtained from queried labels.

Gaining insights from the intermediate results already computed during iteration makes our method simple in architecture, effective in active learning and knowledge transfer, superior in terms of both accuracy and efficiency, and flexible in the trade-off between annotation plus time budget and adaptation accuracy. Take the performance on the VisDA-C [36] benchmark for example, the entire query-and-adaptation process obtains 87.4% with merely 1% labeling budget and 780 training iterations (0.3h). In contrast, SHOT++ operating under the SFUDA setting requires 21.6K iterations (5.8h) to achieve 87.3%. Even when factoring in the estimated annotation time (1.83h), our method is significantly more time-saving. Compared with LADA [51], the ADA state-of-the-art (SOTA), we achieve comparable performance without accessing the source data, while benefiting from a 25% deduction in active sampling time and an approximately 17-fold increase in adaptation speed. Furthermore, it exceeds prior SFADA work MHPL [55] by a clear margin of 1.5%. Our main contributions are:

- a novel SFADA paradigm to learn from the learnt source model and actively iterated models, which frees it from specialized architecture and sophisticated learning schemes with superior accuracy and efficiency;
- a CAS strategy to learn from previous model hypothesis, thereby prioritizing targets that are confusing for the current model, less transferable in class membership and consistently challenging in previous active rounds;

- a VPA design to learn from previous feature representations, so that intrinsic distribution alignment and active sample exploitation can be achieved;
- SOTA adaptation accuracy, superior computational efficiency, and continual improvements extensively validated on varied-scaled benchmarks.

2 Related Work

Domain Adaptation (DA) is a specialized case of transfer learning that enables knowledge generalization from a source domain to a related target domain. Varying in the assumption of the label-set relationship between source and target domains, namely semantic shift, DA settings can be mainly divided into close-set DA, open-set DA [33, 44], partial DA [4, 26] and universal DA [43, 64]. According to the assumption of target data accessibility, DA settings fall into unsupervised DA (UDA) [10, 59], semi-supervised DA (SSDA) [13, 42], weakly-supervised DA (WSDA) [52], zero-shot DA [35, 56], one-shot DA [30], few-shot DA [12, 32, 60] as well as active DA (ADA) [8, 37, 39, 50, 51, 58].

Among the subfields, the one most related to our setting is **ADA**, where the goal is to query the most informative target samples for annotation to best benefit classification on the target domain. Saha *et al.* [41] either infer pseudo labels on target samples or query them for oracle annotation, which is based on inter-domain similarity between source and target domains. Su *et al.* [50] train a domain discriminative model for domain alignment and target uncertainty estimation. While most of them utilize or are orthogonal to off-the-shelf active learning mechanisms, Fu *et al.* [8] argue that constructing a committee [31], i.e., ensemble to measure target sample uncertainty based on consensus between multiple predictions is more robust under domain shift. CLUE [37] is a k-means clustering-weighted uncertainty-based active learning strategy, and it optimizes with cross-domain minmax entropy [42]. Xie *et al.* [58] apply the margin loss to the source domain training to exploit hard source samples and a less domain-biased decision boundary, and they further supplement margin sampling with expected error reduction consistent with the training objective. LADA [51] bases its sampling criterion on the predictive purity of a local structure, and exploits both data sources during adaptation. Despite previous efforts, ADA methods typically necessitate the simultaneous presence of data from both the source and target domains to exploit the relationship between domains. However, when this assumption fails in practical scenarios, we must tackle the challenges associated with the unavailability of source data.

On the other hand, **SFUDA** posits a transfer setting where neither the source data nor the target annotations are available. Kundu *et al.* [22] adopt a two-stage paradigm for universal DA, where the procurement stage trains a source model in consideration of future semantic shifts and domain gaps, so that the deployment stage is capable of operating without source data. Liang *et al.* [28] harness information maximization [15, 48] to align different domains, and utilize self-supervision and curriculum learning techniques via pseudo label clustering. Xia *et al.* [57] introduce adversarial training to distinguish between source-similar and -dissimilar target samples for cross-domain alignment, facilitated with self-supervision and clustering to improve performance. Based on prediction confidence, DaC [66] also divides the target samples based on source similarity, employing a global structure for source-like targets while implementing a

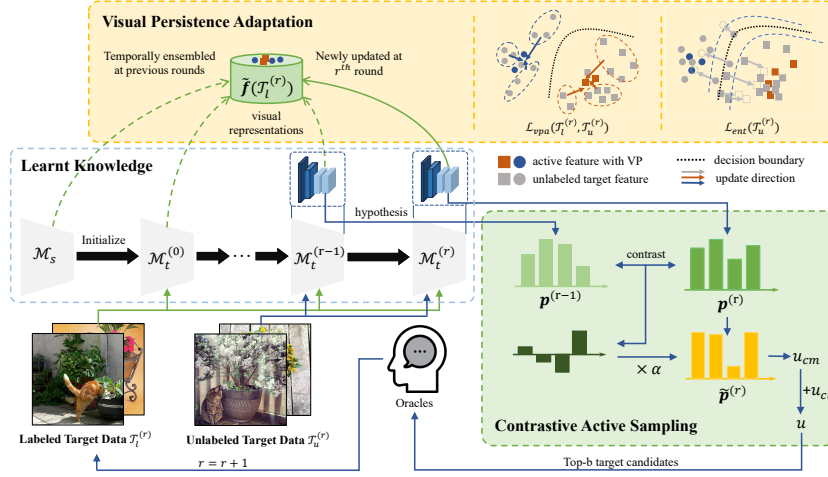


Fig. 2: The proposed LFTL framework for SFADA. Contrastive Active Sampling emphasizes freshly acquired knowledge in the posterior distribution so that novel samples are more likely to be queried for annotations. Then during adaptation the persistence vault retains previous domain-invariant knowledge to facilitate alignment in the target domain via \mathcal{L}_{vpa} and \mathcal{L}_{ent} . As the iterative process continues, it yields more informative targets and an improved target model.

local structure for the distinctive samples. Without definitive supervision signal from either domain, SFUDA makes the alignment more challenging and exacerbates the ill-posedness of UDA. Consequently, it frequently demands intricate design and increased computations to attain desirable outcomes.

Given the accessibility of target samples, this paper focuses on the more practical **SFADA** paradigm, leveraging minimal resources for better performance. Prior work MHPL [55] integrates three active learning strategies to sample targets that are uncertain, diverse and different from the source domain, and then utilizes a neighbor focal loss to emphasize them in adaptation. However, the dependency on the source dissimilarity metric confines it to a one-round sampling method, failing to guarantee sustained performance improvement in real-world applications. SALAD [21] samples target with a binary weighted function between an entropy score and the expected model change. However the adaptation procedure incorporates an additional Guided Attention Transfer Network, leading to undesirable computational overhead and, as evidenced by the results, inefficient annotation utilization. In contrast, our LFTL offers a straightforward, broadly applicable, and growth-oriented solution that can make the best from previous knowledge and a limited budget to guarantee accuracy, efficiency, and a flexible trade-off between annotation resources and performance.

3 Learn from the Learnt

In this section, we present *LFTL* for SFADA. As illustrated in Fig. 2, it alternates between Contrastive Active Sampling (CAS) and Visual Persistence-guided Adaptation (VPA).

3.1 Overview

Given a labeled source domain \mathcal{S} with restricted access and an unlabeled target domain \mathcal{T} , our goal is to derive a model that can minimize the risk on the target domain through knowledge transfer with a small annotation budget. At the source end, image and label pairs $\mathcal{S} = \{(x_s, y_s)\}$ from the source domain are first utilized to train a model $\mathcal{M}_s : x_s \mapsto y_s$, which consists of a feature extractor $f(\cdot)_s$ and a classifier $g(\cdot)_s$, i.e., $\mathcal{M}_s(x_s) = g_s(f_s(x_s))$. After severing communication with the source domain data, we adapt \mathcal{M}_s , where the prior knowledge is maintained, to the target domain with the assistance of a budget of B labeled instances. The initial target dataset has a total of $n_{tu}^{(0)}$ unlabeled samples $\mathcal{T}_u^{(0)} = \{x_{tu}^i\}_{i=1}^{i=n_{tu}^{(0)}}$ and a shared label space $\mathcal{Y}_{\mathcal{T}}$.

At the beginning of the R -round query-and-adaptation alternation ($r = 0$), the target model $\mathcal{M}_t^{(r)}$, without any architectural adaptation, is initialized from \mathcal{M}_s as $\mathcal{M}_t^{(0)}$. We actively query labels for $b = B/R$ most informative target instances with the criterion of CAS (Sec. 3.2), which initializes the labeled subset of the target pool, denoted as $\mathcal{T}_l^{(1)} = \{x_{tl}^j\}_{j=1}^{j=n_{tl}^{(1)}}$, and thus $\mathcal{T}_u^{(1)} = \mathcal{T}_u^{(0)} \setminus \mathcal{T}_l^{(1)}$. Then the parameters of the target model can be optimized using the newly acquired target data splits, and updated to $\mathcal{M}_t^{(1)}$ via the proposed VPA strategy (Sec. 3.3). Likewise, in subsequent rounds ($r \geq 1$), target data splits $\mathcal{T}_l^{(r)} \leftarrow \mathcal{T}_l^{(r-1)} \cup \{x_{tl}^j\}_{n_{tl}^{(r-1)}}^{n_{tl}^{(r)}}$ and $\mathcal{T}_u^{(r)} = \mathcal{T}_u^{(r-1)} \setminus \mathcal{T}_l^{(r)}$, as well as the task model $\mathcal{M}_t^{(r)}$ are updated, continuously querying novel knowledge and optimizing towards the target domain through the synergistic interaction of sampling and adaptation.

3.2 Contrastive Active Sampling

In each active learning round, the aim is to find the most informative instances from the target domain, especially those diverging from the source, to benefit the adaptation of the model. Without directly referencing the source data, previous source-free approaches [7, 25, 28, 57] often rely on the hypothesis transfer to distinguish between source-like and source-dissimilar samples. However, when applied to the SFADA scenario, such solution [55] is constrained to a single round of active sampling and model updating, which restricts the potential applications of this setting.

Considering the iterative nature of the task context, we tackle this challenge from the active perspective. Just as the source knowledge is encapsulated within the initial model \mathcal{M}_s , the target knowledge learned in previous rounds is manifested in the hypotheses of the preceding model $\mathcal{M}_t^{(r-1)}$. Compared to the hypotheses of the preceding model, samples that obtain higher prediction confidence from the newly updated model reflect the fresh insight just acquired. The larger the gap in predictive confidence, the better the current model has grasped the sample, and thus the less informative for the subsequent phase of sample selection. As illustrated in Fig 2, we first contrastively highlight the less informative samples as follows:

$$\tilde{\mathbf{p}}^{(r)}(\cdot|x_{tu}^i) = \begin{cases} \log \mathbf{p}^{(r)} & \text{if } r = 0, \\ \log \mathbf{p}^{(r)} + \alpha(\log \mathbf{p}^{(r)} - \log \mathbf{p}^{(r-1)}) & \text{if } r > 0, \end{cases} \quad (1)$$

weighted by α , where \mathbf{p} is derived from the softmax of the model logits $\delta(g(f(x_{tu}^i)))$. We omit the superscript of round (r) for simplicity in the following text.

Based on $\tilde{\mathbf{p}}$, various active learning criteria can be employed. Without loss of generality, here we focus on the most confused classes. Specifically, we take the difference between the best-versus-second best (BvSB) guesses as the uncertainty indicator, defined as $u_{cm}(x_{tu}^i) = \tilde{p}(y_a^i|x_{tu}^i) - \tilde{p}(y_b^i|x_{tu}^i)$, where $y_a^i = \arg \max_{y \in \mathcal{Y}} \tilde{p}(y|x_{tu}^i)$, and $y_b^i = \arg \max_{y \in \mathcal{Y} \setminus y_a^i} \tilde{p}(y|x_{tu}^i)$. Then the unlabeled target sample pool is ranked by:

$$\mathbf{r}(x_{tu}^i) = \sum_{k=1}^{n_{tu}} \mathbb{I}(u_{cm}(x_{tu}^k) < u_{cm}(x_{tu}^i)), \quad (2)$$

where \mathbb{I} is the indicator function. A smaller u_{cm} score indicates that the assessed target sample x_{tu}^i furnishes novel insights that diverge from the source domain and, in contrast to prior iterations, has not been grasped. Conversely, when the current stronger model, aided by a marginal increment of active labels, assigns more probability preference to one class for an unlabeled sample, CAS emphasizes it via a larger u_{cm} score, lowering the priority of those familiar candidates. As the query-and-adaptation alternation proceeds, CAS continuously harvests fresh information for transfer by leveraging readily available intermediate results, without additional extra computational burden or memory usage.

The contrastively decoded margin evaluates each sample on an individual basis. However, the batch-mode active learning could be susceptible to outliers, since the queried samples are not selected based on the underlying natural density distribution [1, 47]. Particularly in DA tasks, different classes show different sensitivity to the domain shift. Some of them are more robust to certain shifts, i.e., their learnt features present more domain-invariant patterns, whereas others are likely to suffer a steep performance drop [18]. These motivate us to incorporate a broader, semantic view of the target domain into our active criterion, where classes with higher transferability is estimated via the class memberships of reliable model hypotheses:

$$u_{ct}(c|\mathcal{T}_u) = \frac{\sum_{i=1}^{n_{tu}} \mathbb{I}(y_a^i = c) \cdot \mathbb{I}(\mathbf{r}(x_{tu}^i) > (n_{tu} - \kappa))}{\max_c [\sum_{i=1}^{n_{tu}} \mathbb{I}(y_a^i = c) \cdot \mathbb{I}(\mathbf{r}(x_{tu}^i) > (n_{tu} - \kappa))]} \quad (3)$$

Eq. 3 measures the frequency of unlabeled samples being inferred as class c among κ highest contrastive margins, wherein κ controls the range of statistics. Thus the classes with higher u_{ct} scores can be downweighted in favor of less transferable classes.

In this way, we drive a hybrid active sampling criterion that favors target samples with local uncertainty, global intransferability and temporal stagnancy:

$$u^i = u_{cm}^i + \lambda u_{ct}(y_a^i), \quad (4)$$

where λ controls the importance of class-level cross-domain transferability.

3.3 Visual Persistence-guided Adaptation

During each iterative round, the top- b most informative samples are queried for labels towards the adaption to the target domain. They offer reliable signals but also pose a

challenge of making the best use of them. We first warm up the adaptation process via empirical risk minimization on the small amount of target instances with active labels:

$$\mathcal{L}_{ce} = -\mathbb{E}_{x_{tl} \sim \mathcal{T}_l} \mathbf{q}(x_{tl})^T \log[\mathbf{p}(\cdot|x_{tl})], \quad (5)$$

where $\mathbf{q}(x_{tl})$ is the ground truth 1-hot vector for the labeled instance x_{tl} .

To further minimize the risk on the target domain, we are motivated to bridge the distance between actively queried data and unlabeled data. Instead of seeking cluster centroids in the target pool, we directly take active samples as representative anchors, and encourage unlabeled data to form concentrated clusters around the nearest anchor on the visual embedding space via soft similarity minimization:

$$\mathcal{L}_{ac} = -\mathbb{E}_{x_{tu} \sim \mathcal{T}_u} \mathbf{d}_{tu}^T \log(\mathbf{d}_{tu}), \text{ where } \mathbf{d}_{tu} = \delta[\mathcal{D}(f_e(x_{tu}), \mathbf{f}(\mathcal{T}_l))]. \quad (6)$$

In Eq. 6, the feature extractor of the adapted model outputs the visual representation of x_{tu} as $f_t(x_{tu})$. $\mathbf{f}(\mathcal{T}_l)$ is a $n_{tl} \times d$ matrix where each row represents a d -dimensional feature of a labeled anchor x_{tl}^i . \mathcal{D} denotes a distance metric. For simplicity and without loss of generality, we adopt cosine similarity to measure the relationship between different instances, and a normalization function is applied to transform the distances to $[0, 1]$. In decreasing the entropy of distances, each unlabeled sample is encouraged to get close to its nearest labeled anchors and meantime keep other clusters at a considerable distance. Consequently, as shown in Fig. 2, features are aggregated in a class-wise manner, which makes it easier for classifier decision boundaries to be established.

However, in the context of the SFADA framework, as the labeled target subset \mathcal{T}_l and the model \mathcal{M}_t iterate for R rounds, the incremental domain-specific information obtained from the actively queried samples may intensify the cross-domain incompatibility [5, 49, 53]. Furthermore, there is a potential risk of domain-invariant knowledge being forgotten during the process. Without accessing the source domain data, previous source-free methods achieve domain alignment via freezed hypotheses [25, 28, 57] or buffered embeddings [7] from the source classifier. In our multi-round iterative process, we introduce a visual persistence vault to guide the adaptation. It preserves judgements on the active anchors, which have been derived from source and intermediate models, via exponential moving average:

$$\tilde{\mathbf{f}}(x_{tl}) \leftarrow \gamma \mathbf{f}(x_{tl}) + (1 - \gamma) \tilde{\mathbf{f}}(x_{tl}), \quad (7)$$

where $\mathbf{f}(x_{tl})$ is derived from the adapted feature extractor f_t , and γ is the momentum parameter controlling the weighting decrease of previous observations, which is empirically fixed as 0.9. In replacing the most recent feature representation of active anchors $\mathbf{f}(\mathcal{T}_l)$ in Eq. 6 with the temporally ensembled $\tilde{\mathbf{f}}(\mathcal{T}_l)$, we obtain the visual persistence-guided adaption loss \mathcal{L}_{vpa} . The learnt domain-invariant knowledge maintained by the persistence vault effectively supports the alignment in the target domain, and meantime the target-specific information is well exploited.

Additionally, the entropy minimization loss [16, 28, 55] is introduced to approach the ideal adaptation performance and foster discriminative features in the learnt manifold:

$$\mathcal{L}_{ent} = -\mathbb{E}_{x_{tu} \sim \mathcal{T}_u} \mathbf{p}(\cdot|x_{tu})^T \log[\mathbf{p}(\cdot|x_{tu})]. \quad (8)$$

Table 1: Results on the large-scale VisDA-C dataset (ResNet101) in terms of classification accuracy (%). SF represents source inavailability, and AS shows percentage of active annotations (%). Best results are highlighted in bold, and the best in each section are underlined.

Method	SF	AS	plane	bcycl	bus	car	horse	knife	mcycl	person	plant	sktbrd	train	truck	Per-class
SFDA [20]	✓	-	86.9	81.7	84.6	63.9	93.1	91.4	86.6	71.9	84.5	58.2	74.5	42.7	76.7
A ² Net [57]	✓	-	94.0	87.8	85.6	66.8	93.7	95.1	85.8	81.2	91.6	88.2	86.5	56.0	84.3
SHOT [28]	✓	-	95.8	88.2	87.2	73.7	95.2	96.4	87.9	<u>84.5</u>	92.5	89.3	85.7	49.1	85.5
SHOT++ [28]	✓	-	<u>97.7</u>	88.4	<u>90.2</u>	<u>86.3</u>	<u>97.9</u>	98.6	92.9	84.1	<u>97.1</u>	92.2	<u>93.6</u>	28.8	87.3
CPGA [38]	✓	-	94.8	83.6	79.7	65.1	92.5	94.7	90.1	82.4	88.8	88.0	88.9	60.1	84.1
DaC [66]	✓	-	96.6	86.8	86.4	78.4	96.4	96.2	<u>93.6</u>	83.8	96.8	<u>95.1</u>	89.6	50.0	87.3
SF(DA) ² [17]	✓	-	96.8	<u>89.3</u>	82.9	81.4	96.8	95.7	90.4	81.3	95.5	93.7	88.5	<u>64.7</u>	<u>88.1</u>
<hr/>															
AADA [50]	×	1	83.5	64.0	67.2	80.5	87.8	61.4	88.5	79.1	87.9	78.5	84.7	32.6	74.6
TQS [8]	×	1	87.3	77.5	77.1	67.0	90.5	58.4	81.1	82.0	91.1	73.9	68.7	<u>57.6</u>	76.0
CLUE [37]	×	1	77.0	57.8	73.4	76.5	76.9	68.7	87.0	75.5	85.1	66.5	76.1	47.3	72.3
SDM-AG [58]	×	1	84.0	72.3	77.8	82.0	92.2	81.5	83.1	71.4	85.2	74.2	80.2	35.1	76.6
LADA [51]	×	1	<u>97.8</u>	<u>82.3</u>	<u>92.0</u>	<u>86.8</u>	<u>98.0</u>	<u>94.9</u>	<u>94.7</u>	88.8	<u>95.6</u>	<u>95.0</u>	<u>93.9</u>	57.0	<u>89.7</u>
<hr/>															
AADA [50]	×	5	92.3	78.5	87.4	87.4	91.8	90.8	91.2	86.2	92.6	90.3	90.1	61.5	86.7
TQS [8]	×	5	89.6	85.8	82.9	78.5	96.8	82.8	90.2	81.6	93.9	85.2	87.0	<u>69.6</u>	85.3
CLUE [37]	×	5	92.8	81.7	83.0	84.0	93.4	89.1	91.4	89.1	94.2	88.0	85.3	62.0	86.2
SDM-AG [58]	×	5	93.3	85.1	83.7	86.7	94.2	90.2	92.8	85.5	90.2	81.1	80.8	56.2	85.0
LADA [51]	×	5	98.6	<u>87.4</u>	92.9	91.3	98.6	<u>97.0</u>	96.0	<u>91.8</u>	<u>97.4</u>	97.2	95.2	61.6	<u>92.1</u>
<hr/>															
LFTL	✓	1	95.9	84.6	84.6	77.1	95.4	93.6	91.4	87.1	93.2	90.4	87.8	67.6	87.4
MHPL [55]	✓	5	-	-	-	-	-	-	-	-	-	-	-	-	91.3
LFTL	✓	5	<u>98.0</u>	92.5	<u>88.7</u>	<u>89.1</u>	<u>98.0</u>	<u>97.2</u>	<u>94.3</u>	93.5	98.0	<u>96.5</u>	<u>92.6</u>	75.6	92.8

Overall, the combination of supervised cross-entropy, visual persistence guidance and entropy minimization losses yields:

$$\mathcal{L} = \mathcal{L}_{ce} + \beta_1 \mathcal{L}_{vpa} + \beta_2 \mathcal{L}_{ent}. \quad (9)$$

We emphasize that no explicit supervision (e.g., pseudo labels) is enforced for unlabeled samples during the optimization, so that the error accumulation problem is greatly alleviated in the source data-free and target label-scarce setting.

4 Experiments

4.1 Setup

Datasets. We conduct experiments on three widely-used DA benchmarks: a) **VisDA-C** [36] aims to solve the simulation-to-reality shift, in which the source domain contains 152K synthetic images with varying angles and lightning conditions, and the target domain is composed of 55K images cropped from the Microsoft COCO dataset [29], each with the same 12 categories. b) **Office-Home** [54] is a medium-sized dataset consisting of 4 domains: Artistic (Ar), Clip Art (Cl), Product (Pr) and Real-World (Rw) images, each with the same 65 classes. c) **Office-31** [40] is small-sized, which consists of 31 classes in 3 different domains: Amazon (A), DSLR (D) and webcam (W).

Active Setting. Taking the dataset scales into consideration, we budget 1% and 5% of VisDA-C for annotation. For the medium- and small-scaled office datasets, 5% and

Table 2: Results on the Office-Home dataset (ResNet50) in terms of classification accuracy (%). SF represents source inavailability, and AS shows percentage of active annotations (%). Best results are highlighted in bold, and the best in each section are underlined.

Method	SF	AS	Ar→Cl	Ar→Pr	Ar→Rw	Cl→Ar	Cl→Pr	Cl→Rw	Pr→Ar	Pr→Cl	Pr→Rw	Rw→Ar	Rw→Cl	Rw→Pr	Avg
SFDA [20]	✓	-	48.4	73.4	76.9	64.3	69.8	71.7	62.7	45.3	76.6	69.8	50.5	79.0	65.7
A ² Net [57]	✓	-	58.4	79.0	82.4	67.5	79.3	78.9	68.0	56.2	82.9	<u>74.1</u>	60.5	<u>85.0</u>	72.8
SHOT [28]	✓	-	57.7	79.1	81.5	67.6	77.9	77.8	68.1	55.8	82.0	72.8	59.7	84.4	72.0
SHOT++ [28]	✓	-	57.9	<u>79.7</u>	<u>82.5</u>	68.5	<u>79.6</u>	<u>79.3</u>	<u>68.5</u>	57.0	<u>83.0</u>	73.7	60.7	84.9	<u>73.0</u>
CPGA [38]	✓	-	<u>59.3</u>	78.1	79.8	65.4	75.5	76.4	65.7	<u>58.0</u>	81.0	72.0	<u>64.4</u>	83.3	71.6
DaC [66]	✓	-	59.1	79.5	81.2	<u>69.3</u>	78.9	79.2	67.4	56.4	82.4	74.0	61.4	84.4	72.8
AADA [50]	×	5	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
TQS [8]	×	5	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
CLUE [37]	×	5	50.7	76.1	78.1	62.6	75.2	71.0	64.0	52.3	79.7	72.7	56.9	83.8	68.6
SDM-AG [58]	×	5	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
LADA [51]	×	5	<u>71.2</u>	<u>87.4</u>	<u>84.6</u>	<u>72.1</u>	<u>87.0</u>	<u>83.6</u>	<u>71.5</u>	<u>71.6</u>	<u>85.3</u>	<u>79.3</u>	<u>75.5</u>	90.4	<u>80.0</u>
AADA [50]	×	10	65.8	84.5	82.2	64.1	80.6	76.1	67.6	62.6	80.1	73.7	66.1	88.6	74.3
TQS [8]	×	10	68.0	87.7	85.7	67.0	83.0	78.7	69.3	64.5	83.9	77.8	68.9	90.6	77.1
CLUE [37]	×	10	62.1	79.1	80.3	64.1	77.1	76.7	65.4	63.9	84.2	74.2	68.6	84.7	73.4
SDM-AG [58]	×	10	68.5	87.6	86.4	69.5	84.8	81.2	71.0	66.2	84.8	79.3	69.9	90.8	78.3
LADA [51]	×	10	<u>77.2</u>	<u>91.9</u>	<u>88.1</u>	<u>76.9</u>	<u>91.1</u>	<u>86.8</u>	<u>76.6</u>	<u>78.1</u>	<u>88.3</u>	<u>82.0</u>	<u>79.0</u>	<u>93.8</u>	<u>84.2</u>
MHPL [55]	✓	5	<u>69.0</u>	85.7	<u>86.4</u>	72.6	<u>87.4</u>	84.2	<u>73.3</u>	<u>67.4</u>	<u>86.4</u>	<u>80.1</u>	69.6	89.8	<u>79.3</u>
LFTL	✓	5	66.9	<u>86.6</u>	85.5	<u>73.1</u>	86.3	<u>84.5</u>	72.2	65.7	85.9	79.2	69.0	<u>90.2</u>	78.8
LFTL	✓	10	76.6	<u>92.2</u>	<u>89.7</u>	<u>78.9</u>	<u>93.0</u>	<u>89.2</u>	<u>78.6</u>	77.1	<u>90.0</u>	<u>83.4</u>	77.8	<u>94.6</u>	<u>85.1</u>

10% data is actively sampled. In each round, 0.1% annotation increment is made. In all experiments, the average performance of 3 repeated trials is reported.

Implementation. Following SHOT [28], A²Net [57], CPGA [38] and MHPL [55], our main results are reported on ResNet101 [14] for VisDA-C and ResNet50 for Office, on top of which is a bottleneck layer that compresses the 2048 high dimensional features into 256 units, followed by a task-dependent FC layer. In the source domain, we randomly split source data into 90% training and 10% validation to obtain a well-trained model. Considering that the number of active selections is quite small, we harness Mixup [65] regularization to exploit the potential of data during target adaptation. Component verification and more training details can be found in supplementary.

4.2 Comparison with SOTAs

We consider the following competitors: SFUDA methods including SFDA [20], A²Net [57], SHOT [28], SHOT++ [28], CPGA [38], DaC [66] and SF(DA)² [17], ADA methods including AADA [50], TQS [8], CLUE [37], SDM-AG [58] and LADA [51], and SFADA prior work MHPL [55]. For *fair inter- and intra-setting comparison*, consistent experimental conditions provided by SHOT are utilized for ADA methods based on the official codes of TQS [8], CLUE [37] and SDM-AG [58], CLUE [51] as well as the re-implementation of AADA [50] provided by CLUE [37]. Results not reported by MHPL and not reproducible due to unavailability of source code are indicated by a dash.

We summarize the adaption accuracies on VisDA-C, Office-Home and Office-31 in Tab. 1, 2 and 3, respectively. On the large-scale dataset VisDA-C, given merely 1% label annotations, LFTL is already capable of outperforming most of ADA methods. When the budget increased to 5%, it exceeded the SFADA SOTA by 1.5% and also outdid the source-accessible ADA SOTA under the same conditions by 0.7%. On the medium-sized Office-Home dataset, LFTL elevates the adaptation performance to a level of 78.8% with

Table 3: Results on the Office-31 dataset (ResNet50) in terms of classification accuracy (%). SF represents source inavailability, and AS shows percentage of active annotations (%). Best results are highlighted in bold, and the best in each section are underlined.

Method	SF	AS	A→D	A→W	D→A	D→W	W→A	W→D	Avg
SFDA [20]	✓	-	92.2	91.1	71.0	98.2	71.2	99.5	87.2
A ² Net [57]	✓	-	94.5	94.0	<u>76.7</u>	<u>99.2</u>	76.1	100.0	<u>90.1</u>
SHOT [28]	✓	-	93.9	90.1	75.3	98.7	75.0	99.9	88.8
SHOT++ [28]	✓	-	94.3	90.4	76.2	98.7	75.8	99.9	89.2
CPGA [38]	✓	-	94.4	<u>94.1</u>	76.0	98.4	76.6	99.8	89.9
SF(DA) ² [17]	✓	-	<u>95.8</u>	92.1	75.7	99.0	<u>76.8</u>	99.8	89.9
<hr/>									
AADA [50]	×	5	89.2	87.3	78.2	99.5	78.7	100.0	88.8
TQS [8]	×	5	92.8	92.2	80.6	100.0	80.4	100.0	91.1
CLUE [37]	×	5	90.7	94.3	78.7	99.1	76.1	99.8	89.8
SDM-AG [58]	×	5	93.5	94.8	81.9	100.0	81.9	100.0	92.0
LADA [51]	×	5	<u>96.3</u>	<u>97.7</u>	<u>83.1</u>	99.6	<u>85.0</u>	99.7	<u>93.6</u>
<hr/>									
AADA [50]	×	10	93.5	93.1	83.2	99.7	84.2	100.0	92.3
TQS [8]	×	10	96.4	96.4	86.4	100.0	87.1	100.0	94.4
CLUE [37]	×	10	96.2	94.7	84.4	99.4	81.0	100.0	92.6
SDM-AG [58]	×	10	95.9	96.4	86.1	100.0	86.5	100.0	94.2
LADA [51]	×	10	<u>97.8</u>	<u>99.1</u>	<u>87.3</u>	99.9	87.6	99.7	<u>95.2</u>
<hr/>									
MHPL [55]	✓	5	97.8	96.7	82.5	99.3	<u>82.6</u>	100.0	93.2
LFTL	✓	5	<u>98.0</u>	<u>98.5</u>	<u>82.6</u>	99.9	82.2	100.0	<u>93.5</u>
LFTL	✓	10	98.9	99.4	87.8	100.0	86.3	100.0	95.4

5% annotation costs, which is comparable with SFUDA methods. We note a continuous enhancement in performance as the budget is raised to 10%, culminating in the highest 85.1% accuracy. Similar trend can also be observed in the small-sized Office-31 dataset. When 5% target samples are labeled, our results surpass MHPL in average performance and in four out of five subtasks, while being on par with the source-accessible LADA.

The results show that, given a minimum amount of annotations, noticeable improvements can be achieved over SFUDA with increased flexibility. Meantime, despite the inavailability of the source domain, learning from what the model already learnt can offset the absence of the source data via our proposed CAS strategy and VPA approach.

4.3 Efficiency Analysis

We compare with open-source SFUDA and ADA methods in terms of actual time consumption, including model training, active sampling and human annotation, as presented in Fig. 3. We take the adaptation on the VisDA-C dataset for example, and provide 1% budget, i.e., 550 images of the target domain for active learning-based methods. Note that here we report the consumption of the whole sample-and-adaptation process, instead of the per-round time. All experiments were conducted on 1 NVIDIA GeForce RTX 3090 GPU, with 80 CPU cores and 256 GB memory.

Considering the **training time for adaptation**, SFUDA SOTAs often require sophisticated schemes (clustering, two-stage training, *etc.*) to adapt to a completely unlabeled domain, which necessitate more computational resource. For example, SHOT and SHOT++ takes 3.7h and 5.8h to converge respectively, and CPGA uses more than five days to attain the reported performance, which we omit in Fig. 3 to prevent it from dominating the figure. On the other hand, when equipped with both source and target

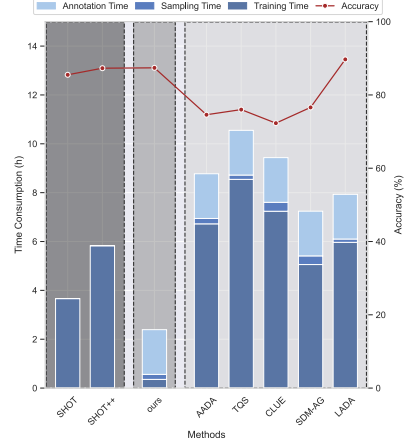


Fig. 3: Comparison on the overall time consumption (training, active sampling and annotation time) and adaptation accuracy with SFUDA (SHOT) and ADA (AADA, TQS, CLUE, SDM-AG, LADA) methods on the VisDA-C dataset.

Table 4: Comparison on complexity and actual query time with ADA methods. Notations are explained in text.

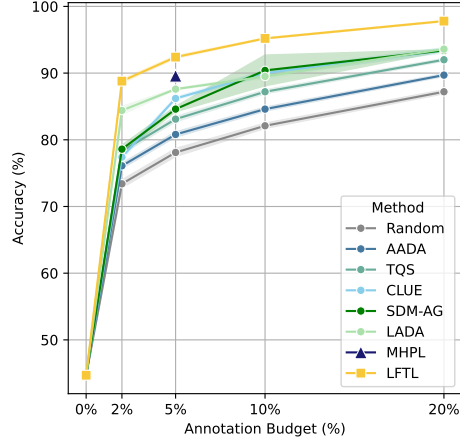
Query Method	Complexity	Time (s/cycle)
AADA [50]	$\mathcal{O}(NC + N \log N)$	79.16
TQS [8]	$\mathcal{O}(NCM + N \log N)$	127.51
CLUE [37]	$\mathcal{O}(tNBD)$	196.92
SDM-AG [58]	$\mathcal{O}(NCD + N \log N)$	127.44
LADA [51]	$\mathcal{O}(ND + N \log N)$	97.6
LFTL	$\mathcal{O}(NC + N \log N)$	73.96

Table 5: Results of LADA minus source loss (SF-LADA) and our LFTL plus a source cross-entropy loss (S-LFTL).

Method	SF	1% VisDA-C	5% Office-Home	5% Office-31
LADA	\times	<u>89.7</u>	80.0	93.6
S-LFTL	\times	87.6	<u>80.1</u>	<u>94.1</u>
SF-LADA	\checkmark	82.6	72.4	89.0
LFTL	\checkmark	<u>87.4</u>	<u>78.8</u>	<u>93.5</u>

data, ADA methods are expected to directly align different domains, as well as identifying novel samples, which adds to the training burden. For example, AADA, CLUE and TQS consume more than 6h mainly because of the adversarial training. TQS additionally trains five classifiers to construct a voting committee for active learning, culminating in an overall duration of 8.5h. LADA employs local neighbor search for more reliable training guidance, which takes 5.9h. SDM-AG improves on the optimization complexity yet still requires 5.1h. In contrast, drawing upon queried informative data, our LFTL leverages the iterative nature of active learning instead of additional computations. The proposed VPA retrospects on the knowledge gained from the source and intermediate models of previous rounds, which is temporarily ensembled in the VP vault for readily use. As the comparison shows, it significantly outperforms competitors with only 0.3h for adaptation while yielding superior accuracy.

As for the **active sampling efficiency**, we summarize the theoretical complexity and actual training duration for active methods in Tab. 4, where C and N to denote the number of classes and unlabeled instances respectively. AADA and our LFTL operate with $\mathcal{O}(NC + N \log N)$ complexity, which are the most efficient approaches. The clustering-based CLUE works at $\mathcal{O}(tNBD)$, in which D denotes the embedding dimension (512), B is the budget, and t is the number of clustering iterations (300). The ranking-based TQS and SDM-AG uses $\mathcal{O}(NCM + N \log N)$ and $\mathcal{O}(NCD + N \log N)$ respectively, in which M represents the number of members in the classifier committee (5 is adopted). The local neighbor search in LADA leads to a complexity of $\mathcal{O}(ND + N \log N)$. The feature dimension D of SDM-AG and LADA is 256. Contrary to previous active criteria designed for sampling in the target domain, CAS contrasts the hypotheses generated by models from successive rounds to identify target samples that deviate from the source domain and persistently challenging for the recognition task. The comparison shows that the proposed CAS is effective, scalable and efficient.

**Fig. 4:** Results of SFADA (MHPL and LFTL) and ADA (others) methods on VisDA-C (ResNet50) as the annotation budget continuously grows.

Regarding the **human annotation time**, given that image-level labels take 1 second per class [2, 34], 1% target domain subset of VisDA-C would cost $550 \times 12 \approx 1.83\text{h}$. As Fig. 3 shows, compared to non-active methods, although factoring in the annotation time introduces additional burden for ADA methods, we still retain the top place.

Collectively, the comparative analysis of accuracy and efficiency reveals that our LFTL framework, through the synergy of CAS and VPA strategies, can be freed from source-accessibility, specialized architectures or complex optimization techniques while still delivering SOTA performance.

4.4 Promises of Continual Performance Growing

In addition to the experiments presented in main results under varying budget constraints, in this section, we explore the trade-off between the cost and adaptation performance with increased budget and larger strides. The budget is increased to 2%, 5%, 10% and 20% successively, with the corresponding results summarized in Figure 4. We observe that the proposed LFTL framework performs notably better than competitors in terms of both accuracy and robustness. The performance superiority is consistent across different budget scenarios, from tight to generous, indicating its effectiveness in leveraging limited data resources. On the contrary, other methods either gradually deviate from the source domain throughout the process, failing to acquire domain-variant novel knowledge, or demonstrating suboptimal performance in sustaining domain-alignment, or they solely support one-round active sampling as in MHPL. We emphasize that the promise of continual performance growth holds practical value in real-world applications, as it facilitates the effective feedback loop between data reflow and model iteration.

4.5 Discussion of Source Availability

Considering the performance margin of SFADA over ADA, we take LADA, the ADA SOTA, as an example for comparison to explore the role of source data in adaptation. We first ablate the source domain loss from LADA, denoted as SF-LADA to compare with LFTL. Given that the source data becomes unavailable, each epoch now iterates over the labeled target dataset instead of the labeled source dataset, the same as our LFTL. The averaged results in Tab. 5 show that, without specifically designed hypothesis transfer strategy, removing the source data from ADA methods will cause task models to forget the previously learnt knowledge in the source domain during adaptation. And the newly acquired information from the target domain is too limited to counteract it, resulting in noticeable accuracy gaps. This further validates the necessity of learning from the learnt under the source-free condition.

Then we add a simple cross-entropy loss for the source data to LFTL, namely S-LFTL, while keeping the number of iterations unchanged. As shown in Tab. 5, S-LFTL improves over LFTL, demonstrating that direct source supervision could still provide more useful information than hypothesis transfer and memory recollection. The source data proves especially effective for smaller Office datasets, where the improvements are noticeable while the additional amount of additional time is affordable. For the large dataset, although the performance of S-LFTL does not match that of LADA within just

Table 6: Comparison between the proposed CAS and active learning baselines in terms of classification accuracy (%). Experiments are performed on the large-scale VisDA-C dataset (ResNet101) with 1% budget in total for 10 cycles. The best results are highlighted in bold.

Method	1	2	3	4	5	6	7	8	9	10
Random	67.0	70.6	73.7	76.6	78.7	79.9	81.1	81.8	82.2	82.7
Ent-max	62.2	67.9	73.7	77.6	79.9	80.5	80.2	80.2	80.0	80.5
Ent-min	60.4	66.4	64.5	64.2	63.8	64.0	64.1	63.5	64.6	63.9
Kcenter-greedy	67.5	71.3	74.5	77.4	78.9	80.6	81.4	82.3	82.7	83.5
Kmeans	60.5	67.3	72.9	75.7	77.5	78.7	79.7	80.4	80.9	81.2
Least-confidence	63.2	69.9	74.5	78.5	80.2	80.9	80.7	81.1	81.2	81.9
Bayesian	66.5	70.4	73.9	76.7	78.7	80.2	81.4	82.1	82.8	83.4
BvSB	65.7	72.2	77.0	79.8	81.4	82.5	83.4	84.0	84.6	85.1
LFTL	67.6	74.9	79.1	81.5	83.6	84.3	85.6	86.2	86.9	87.4

6% of its adaptation time, extending the training iterations to access all available data would guarantee further performance improvements.

4.6 Comparison with Active Learning Baselines

To validate the proposed active query strategy, we compare LFTL with active learning baselines. We choose Random, uncertainty-based EntMax [46], Least Confidence [24], BvSB [24], Bayesian [9], distribution-based K-means [62] and K-center greedy [45]. Experiments are performed on the large-scale VisDA-C dataset with 1% annotation budget in total for 10 rounds. The results for each round are listed in Tab. 6. It manifests that the proposed sampling strategy consistently outperforms its substitutes. As the iteration proceeds, information queried by other active learning methods become redundant, failing to fulfill the requirements of the DA task. In comparison, the CAS and class-level transferability estimation we propose efficiently prioritize samples with local uncertainty, global intransferability and temporal stagnancy along the sample-and-adaptation progress, bringing continuous growth to the DA task.

5 Conclusion

In this paper, we investigate the challenging source-free active domain adaptation (SFADA) setting, where source data becomes inaccessible during adaptation, and a minimum amount of annotation budget is available for the target domain. We present a shared framework for both active sampling and domain adaptation to learn from the hypotheses and feature representations of the learnt source model and actively iterated models. The proposed CAS criterion effectively prioritizes samples that are both informative to the current model and persistently challenging throughout the iterative process, and their visual understandings of actively queried samples are temporally preserved and exploited via the VPA strategy during adaptation.

Extensive experiments on three DA benchmarks have shown that, in comparison with ADA and SFUDA settings, SFADA provides a worthy trade-off between annotation costs and model performance in both accuracy and efficiency. In comparison with previous SFADA competitors, our LFTL framework also presents superior performance and exhibits more flexibility than the one-round method.

Acknowledgment

This work was supported by National Science and Technology Major (2021ZD0114703-2), National Natural Science Foundation of China (Nos. 61925107, 62271281, 62021002).

References

1. Baum, E.B., Lang, K.: Query learning can work poorly when a human oracle is used. In: International Joint Conference on Neural Networks. vol. 8, p. 8. Beijing China (1992)
2. Bearman, A., Russakovsky, O., Ferrari, V., Fei-Fei, L.: What's the point: Semantic segmentation with point supervision. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 549–565. Springer (2016)
3. Beluch, W.H., Genewein, T., Nürnberger, A., Köhler, J.M.: The power of ensembles for active learning in image classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9368–9377 (2018)
4. Cao, Z., Ma, L., Long, M., Wang, J.: Partial adversarial domain adaptation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 135–150 (2018)
5. Chang, W.G., You, T., Seo, S., Kwak, S., Han, B.: Domain-specific batch normalization for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7354–7362 (2019)
6. Chuang, Y.S., Xie, Y., Luo, H., Kim, Y., Glass, J., He, P.: Dola: Decoding by contrasting layers improves factuality in large language models. arXiv preprint arXiv:2309.03883 (2023)
7. Ding, N., Xu, Y., Tang, Y., Xu, C., Wang, Y., Tao, D.: Source-free domain adaptation via distribution estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7212–7222 (June 2022)
8. Fu, B., Cao, Z., Wang, J., Long, M.: Transferable query selection for active domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7272–7281 (2021)
9. Gal, Y., Islam, R., Ghahramani, Z.: Deep bayesian active learning with image data. In: International Conference on International Conference on Machine Learning. pp. 1183–1192. PMLR (2017)
10. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* **17**(1), 2096–2030 (2016)
11. Gera, A., Friedman, R., Arviv, O., Gunasekara, C., Sznajder, B., Slonim, N., Shnarch, E.: The benefits of bad advice: Autocontrastive decoding across model layers. In: Rogers, A., Boyd-Graber, J., Okazaki, N. (eds.) Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 10406–10420. Association for Computational Linguistics, Toronto, Canada (Jul 2023)
12. Hao, T., Ding, X., Feng, J., Yang, Y., Chen, H., Ding, G.: Quantized prompt for efficient generalization of vision-language models. In: European Conference on Computer Vision (ECCV). Springer (2024)
13. He, G., Liu, X., Fan, F., You, J.: Classification-aware semi-supervised domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 964–965 (2020)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)

15. Hu, W., Miyato, T., Tokui, S., Matsumoto, E., Sugiyama, M.: Learning discrete representations via information maximizing self-augmented training. In: International Conference on International Conference on Machine Learning. pp. 1558–1567. PMLR (2017)
16. Huang, J., Guan, D., Xiao, A., Lu, S.: Model adaptation: Historical contrastive learning for unsupervised domain adaptation without source data. *Advances in Neural Information Processing Systems* **34** (2021)
17. Hwang, U., Lee, J., Shin, J., Yoon, S.: SF(DA)²: Source-free domain adaptation through the lens of data augmentation. In: The Twelfth International Conference on Learning Representations (2024)
18. Jin, Y., Wang, X., Long, M., Wang, J.: Minimum class confusion for versatile domain adaptation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 464–480. Springer (2020)
19. Joshi, A.J., Porikli, F., Papanikolopoulos, N.: Multi-class active learning for image classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2372–2379. IEEE (2009)
20. Kim, Y., Cho, D., Han, K., Panda, P., Hong, S.: Domain adaptation without source data. *IEEE Transactions on Artificial Intelligence* **2**(6), 508–518 (2021)
21. Kothandaraman, D., Shekhar, S., Sancheti, A., Ghuman, M., Shukla, T., Manocha, D.: Salad: Source-free active label-agnostic domain adaptation for classification, segmentation and detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 382–391 (2023)
22. Kundu, J.N., Venkat, N., Babu, R.V., et al.: Universal source-free domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4544–4553 (2020)
23. Leng, S., Zhang, H., Chen, G., Li, X., Lu, S., Miao, C., Bing, L.: Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
24. Lewis, D.D., Gale, W.A.: A sequential algorithm for training text classifiers. In: Proceedings of the Seventeenth Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval. pp. 3–12. Springer (1994)
25. Li, R., Jiao, Q., Cao, W., Wong, H.S., Wu, S.: Model adaptation: Unsupervised domain adaptation without source data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9641–9650 (2020)
26. Li, S., Liu, C.H., Lin, Q., Wen, Q., Su, L., Huang, G., Ding, Z.: Deep residual correction network for partial domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**(7), 2329–2344 (2020)
27. Li, X.L., Holtzman, A., Fried, D., Liang, P., Eisner, J., Hashimoto, T., Zettlemoyer, L., Lewis, M.: Contrastive decoding: Open-ended text generation as optimization. In: Rogers, A., Boyd-Graber, J., Okazaki, N. (eds.) Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 12286–12312. Association for Computational Linguistics, Toronto, Canada (Jul 2023)
28. Liang, J., Hu, D., Wang, Y., He, R., Feng, J.: Source data-absent unsupervised domain adaptation through hypothesis transfer and labeling transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021)
29. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 740–755. Springer (2014)
30. Luo, Y., Liu, P., Guan, T., Yu, J., Yang, Y.: Adversarial style mining for one-shot unsupervised domain adaptation. *Advances in Neural Information Processing Systems* **33**, 20612–20623 (2020)

31. Melville, P., Mooney, R.J.: Diverse ensembles for active learning. In: Proceedings of the Twenty-First International Conference on Machine Learning. p. 74. ICML '04, Association for Computing Machinery, New York, NY, USA (2004)
32. Motiian, S., Jones, Q., Iranmanesh, S., Doretto, G.: Few-shot adversarial domain adaptation. *Advances in Neural Information Processing Systems* **30** (2017)
33. Panareda Busto, P., Gall, J.: Open set domain adaptation. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 754–763 (2017)
34. Papadopoulos, D.P., Clarke, A.D., Keller, F., Ferrari, V.: Training object class detectors from eye tracking data. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 361–376. Springer (2014)
35. Peng, K.C., Wu, Z., Ernst, J.: Zero-shot deep domain adaptation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 764–781 (2018)
36. Peng, X., Usman, B., Kaushik, N., Wang, D., Hoffman, J., Saenko, K.: Visda: A synthetic-to-real benchmark for visual domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 2021–2026 (2018)
37. Prabhu, V., Chandrasekaran, A., Saenko, K., Hoffman, J.: Active domain adaptation via clustering uncertainty-weighted embeddings. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8505–8514 (2021)
38. Qiu, Z., Zhang, Y., Lin, H., Niu, S., Liu, Y., Du, Q., Tan, M.: Source-free domain adaptation via avatar prototype generation and adaptation. In: International Joint Conference on Artificial Intelligence (2021)
39. Rai, P., Saha, A., Daumé III, H., Venkatasubramanian, S.: Domain adaptation meets active learning. In: Proceedings of the NAACL HLT 2010 Workshop on Active Learning for Natural Language Processing. pp. 27–32 (2010)
40. Saenko, K., Kulis, B., Fritz, M., Darrell, T.: Adapting visual category models to new domains. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 213–226. Springer (2010)
41. Saha, A., Rai, P., Daumé, H., Venkatasubramanian, S., DuVall, S.L.: Active supervised domain adaptation. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases. pp. 97–112. Springer (2011)
42. Saito, K., Kim, D., Sclaroff, S., Darrell, T., Saenko, K.: Semi-supervised domain adaptation via minimax entropy. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8050–8058 (2019)
43. Saito, K., Kim, D., Sclaroff, S., Saenko, K.: Universal domain adaptation through self supervision. *Advances in Neural Information Processing Systems* **33**, 16282–16292 (2020)
44. Saito, K., Yamamoto, S., Ushiku, Y., Harada, T.: Open set domain adaptation by backpropagation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 153–168 (2018)
45. Sener, O., Savarese, S.: Active learning for convolutional neural networks: A core-set approach. In: International Conference on Learning Representations (2018)
46. Settles, B.: Active learning literature survey (2009)
47. Settles, B.: Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison (2009)
48. Shi, Y., Sha, F.: Information-theoretical learning of discriminative clusters for unsupervised domain adaptation. In: Proceedings of the 29th International Conference on International Conference on Machine Learning. pp. 1275–1282 (2012)
49. Shu, R., Bui, H.H., Narui, H., Ermon, S.: A dirt-t approach to unsupervised domain adaptation. In: Proc. 6th International Conference on Learning Representations (2018)
50. Su, J.C., Tsai, Y.H., Sohn, K., Liu, B., Maji, S., Chandraker, M.: Active adversarial domain adaptation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 739–748 (2020)

51. Sun, T., Lu, C., Ling, H.: Local context-aware active domain adaptation. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 18588–18597 (2023)
52. Tan, S., Jiao, J., Zheng, W.S.: Weakly supervised open-set domain adaptation by dual-domain collaboration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5394–5403 (2019)
53. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7167–7176 (2017)
54. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5018–5027 (2017)
55. Wang, F., Han, Z., Zhang, Z., He, R., Yin, Y.: Mhpl: Minimum happy points learning for active source free domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20008–20018 (2023)
56. Wang, J., Jiang, J.: Conditional coupled generative adversarial networks for zero-shot domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3375–3384 (2019)
57. Xia, H., Zhao, H., Ding, Z.: Adaptive adversarial network for source-free domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9010–9019 (2021)
58. Xie, M., Li, Y., Wang, Y., Luo, Z., Gan, Z., Sun, Z., Chi, M., Wang, C., Wang, P.: Learning distinctive margin toward active domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7993–8002 (2022)
59. Xiong, Y., Chen, H., Lin, Z., Zhao, S., Ding, G.: Confidence-based visual dispersal for few-shot unsupervised domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11621–11631 (2023)
60. Xu, X., Zhou, X., Venkatesan, R., Swaminathan, G., Majumder, O.: d-sne: Domain adaptation using stochastic neighborhood embedding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2497–2506 (2019)
61. Xu, X., Chen, H., Lin, Z., Han, J., Gong, L., Wang, G., Bao, Y., Ding, G.: Tad: A plug-and-play task-aware decoding method to better adapt llms on downstream tasks. In: Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, 3-9 August 2024. ijcai.org (2024)
62. Yan, D., Huang, L., Jordan, M.I.: Fast approximate spectral clustering. In: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 907–916 (2009)
63. Yang, B., Yeh, H.W., Harada, T., Yuen, P.C.: Model-induced generalization error bound for information-theoretic representation learning in source-data-free unsupervised domain adaptation. *IEEE Transactions on Image Processing* **31**, 419–432 (2021)
64. You, K., Long, M., Cao, Z., Wang, J., Jordan, M.I.: Universal domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2720–2729 (2019)
65. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: International Conference on Learning Representations (2018)
66. Zhang, Z., Chen, W., Cheng, H., Li, Z., Li, S., Lin, L., Li, G.: Divide and contrast: Source-free domain adaptation via adaptive contrastive learning. In: Oh, A.H., Agarwal, A., Belgrave, D., Cho, K. (eds.) *Advances in Neural Information Processing Systems* (2022)