

Source-free Video Domain Adaptation by Learning Temporal Consistency for Action Recognition^{*}

Yuecong Xu^{1†}, Jianfei Yang^{2†}, Haozhi Cao²,
Keyu Wu¹, Min Wu¹, and Zhenghua Chen¹ (✉)

¹ Institute for Infocomm Research, A*STAR, Singapore
xuyu0014@e.ntu.edu.sg, {wu_keyu, wumin}@i2r.a-star.edu.sg,
chen0832@e.ntu.edu.sg

² School of Electrical and Electronic Engineering,
Nanyang Technological University, Singapore
{yang0478, haozhi001}@e.ntu.edu.sg

Abstract. Video-based Unsupervised Domain Adaptation (VUDA) methods improve the robustness of video models, enabling them to be applied to action recognition tasks across different environments. However, these methods require constant access to source data during the adaptation process. Yet in many real-world applications, subjects and scenes in the source video domain should be irrelevant to those in the target video domain. With the increasing emphasis on data privacy, such methods that require source data access would raise serious privacy issues. Therefore, to cope with such concern, a more practical domain adaptation scenario is formulated as the *Source-Free Video-based Domain Adaptation* (SFVDA). Though there are a few methods for Source-Free Domain Adaptation (SFDA) on image data, these methods yield degenerating performance in SFVDA due to the multi-modality nature of videos, with the existence of additional temporal features. In this paper, we propose a novel Attentive Temporal Consistent Network (ATCoN) to address SFVDA by learning temporal consistency, guaranteed by two novel consistency objectives, namely feature consistency and source prediction consistency, performed across local temporal features. ATCoN further constructs effective overall temporal features by attending to local temporal features based on prediction confidence. Empirical results demonstrate the state-of-the-art performance of ATCoN across various cross-domain action recognition benchmarks. Code is provided at <https://github.com/xuyu0010/ATCoN>.

Keywords: Source-free domain adaptation, video domain adaptation, action recognition, temporal consistency

^{*} This research is jointly supported by A*STAR Singapore under its AME Programmatic Funds (Grant No. A20H6b0151) and Career Development Award (Grant No. C210112046), and by Nanyang Technological University, Singapore, under its NTU Presidential Postdoctoral Fellowship, “Adaptive Multimodal Learning for Robust Sensing and Recognition in Smart Cities” project fund.

[†] Equal Contributions

1 Introduction

Video-based tasks such as action recognition have long been investigated considering their wide applications. Deep neural networks have made remarkable advances with the introduction of large-scale labeled datasets [14,24]. However, due to the expense of laborious video data annotation, sufficient labeled training videos may not be readily available in real-world scenarios. To avoid costly data annotation, various *Video-based Unsupervised Domain Adaptation* (VUDA) methods have been introduced to transfer knowledge from a labeled source video domain to an unlabeled target video domain by reducing discrepancies between source and target video domains [2,4,43]. VUDA methods greatly improve the robustness of video models, enabling them to be applied to action recognition tasks across different environments [41].

Though current VUDA methods [2,4,40,41] enable the transfer of knowledge across video domains, they all require access to source video data during the adaptation process. Yet action information usually contains the private and sensitive information of the actors, including their actions and the relevant scenes. Meanwhile, in real-world applications, such information in the source domain is usually irrelevant to those in the target domain and should be protected from the target domain. Therefore, current VUDA methods would raise serious privacy issues, which is more severe than that raised by image-based domain adaptation. To cope with the video data privacy issue, a more practical domain adaptation scenario is formulated as the *Source-Free Video-based Domain Adaptation* (SFVDA), where only well-trained source video models and unlabeled target domain data would be provided for adaptation.

With the absence of source data, current VUDA methods that mainly align target and source domains statistically [22,34] cannot be applied to the SFVDA problem. Recently, there are a few research efforts [18,21,48] that start exploring Source-Free Domain Adaptation (SFDA) with image data, where SFDA is tackled by adjusting target features to adapt to the source classifier [20]. The key idea is to learn discriminative latent target features while aligning source data distribution embedded within the source classifier. However, aligning videos without source data is even more challenging thanks to the fact that videos are characterized by their multi-modality nature, where temporal features are key components that are excluded in images.

While direct minimization of statistical discrepancy between target and source domains cannot be achieved due to the lack of source data, domain adaptation can also be achieved by aligning the embedded semantic information [39,19] via entropy-based approaches [29,37] such as maximizing mutual information [36] or neighborhood clustering [30]. These methods improve the discriminability of the target features which satisfy the cluster assumption [8], while increasing the source model transferability [45]. However, these methods are insufficient for aligning semantic information in videos. The reason is that overall temporal feature of a video can be constructed with a series of local temporal features, obtained through clips sampled from videos. Each local temporal feature should be discriminative in the first place. However, if each local temporal feature is in-

dividually discriminative yet mutually inconsistent, the local temporal features may not hold similar semantic information. Subsequently, the overall temporal feature may contain indistinct semantic information, and would not be discriminative. Instead, we hypothesize that for source videos, the extracted local temporal features are not only discriminative, but also consistent across each other and possess similar feature distribution patterns, which implies similar semantic information. Such hypothesis is termed as the *cross-temporal hypothesis*. If the target data aligns with the source data distribution, we assume that source-like representations are learned for target data, therefore the *cross-temporal hypothesis* should be satisfied by the target data representation. To this end, our method is designed such that the local temporal features are consistent in their feature representations, which would result in the corresponding overall temporal feature being effective and discriminative.

Meanwhile, since only the source model with the source classifier is available for adaptation, the relevance of the target data to source data distribution is highly correlated to the prediction of target data on the source classifier. Therefore, to better adapt target temporal features to the source classifier, the relevance of the corresponding local temporal features towards source data distribution should also be consistent. Such consistency can be interpreted as the source prediction consistency of local temporal features with respect to the fixed source classifier. Further, to improve the discriminability of the video feature, the overall temporal feature should be built by an attentive combination of local temporal features. The attentive combination builds upon the confidence of each local temporal feature towards its relevance to source data distribution.

To this end, we propose an **Attentive Temporal Consistent Network (ATCoN)** to address SFVDA uniformly. ATCoN leverages temporal features effectively by learning **temporal consistency** via **feature consistency** and **source prediction consistency** for local temporal features in a self-supervised manner. ATCoN further adapts target data to the source data distribution by attending to local temporal features with higher confidence over its relevance towards source data distribution, indicated as higher source prediction confidence.

In summary, our contributions are threefold. First, we formulated a practical and challenging *Source-Free Video Domain Adaptation* (SFVDA) problem. To the best of our knowledge, this is the first research that studies source-free transfer for video-based tasks, which aims to address data-privacy issues in VUDA. Secondly, we analyze the challenges underlying SFVDA and propose ATCoN to address the challenges uniformly. ATCoN aims to obtain effective and discriminative overall temporal features that satisfies the *cross-temporal hypothesis* by learning temporal consistency which is composed of both feature and source prediction consistency. ATCoN further aligns target data to the source data distribution without source data access by attending to local temporal features with high source prediction confidence. Finally, empirical results demonstrate the efficacy of our proposed ATCoN, achieving state-of-the-art performance across the multiple cross-domain action recognition benchmarks.

2 Related Work

Unsupervised Domain Adaptation (UDA) and Video-based Unsupervised Domain Adaptation (VUDA). Current UDA and VUDA methods aim to distill shared knowledge across labeled source domains and unlabeled target domains. These methods improve the transferability and robustness of models. Generally, they could be divided into three categories: a) reconstruction-based methods [7,46], where domain-invariant features are obtained by encoders trained with data-reconstruction objectives, whose methods are commonly formulated as encoder-decoder networks; b) adversarial-based methods [2,41], where domain-invariant features are extracted by feature generators while leveraging domain discriminators, which are trained jointly in an adversarial manner [11], minimizing adversarial losses [6]; and c) discrepancy-based methods [31,49,44], which mitigate domain shifts across domains by applying metric learning approaches, minimizing metrics such as MMD [22] and CORAL [34]. By comparison, VUDA research lags behind UDA research, mainly due to the challenges brought by aligning temporal features in videos. However, with the introduction of various cross-domain video datasets such as UCF-HMDB_{full} [2] and Sports-DA [43], there has been a significant increase in research interests for VUDA [4,26,3]. Despite the improvements in video model robustness brought by VUDA methods, all such methods require access to source data during the adaptation process. Such requirements could raise serious privacy concerns given the amount of private information of the relevant subject and scene in videos.

Source-Free Domain Adaptation (SFDA). With the increased importance of data privacy, there have been a few recent research efforts that investigate SFDA with images, which enable image models to be adapted to the target domain without access to source data. Among them, 3C-GAN [18] and SDDA [17] seek to produce novel target-style data that are similar to the source domain. Domain invariant features are then obtained by aligning the novel target-style data with the original target data via adversarial-based domain adaptation methods. Similarly, CPGA [28] tackles SFDA by generating avatar feature prototypes for each class, which are trained with the target features in an adversarial manner. Meanwhile, SHOT [21,20] exploits knowledge of source feature distribution by freezing the source classifier and matches target features to the source classifier by leveraging information maximization and pseudo-labeling. More recently, BAIT [47] extends MCD [31] to SFDA. Despite the advances made in the research of SFDA for images, SFVDA has not been tackled. Due to the amount of private data in videos, SFVDA is even more critical, yet is also more challenging given that temporal features must also be aligned. We propose to engage in SFVDA by utilizing temporal features via learning temporal consistency while attending to local temporal features with high confidence.

3 Proposed Method

In the scenario of *Source-Free Video Domain Adaptation* (SFVDA), we are only given a source video model that consists of the spatial feature extractor $G_{S,sp}$, the

temporal feature extractor $G_{S,t}$ and the classifier H_S , and an unlabeled target domain $\mathcal{D}_T = \{V_{iT}\}_{i=1}^{n_T}$ with n_T i.i.d. videos, characterized by a probability distribution of p_T . The source model is generated by training its parameters $\theta_{S,sp}$, $\theta_{S,t}$, and θ_H with the labeled source domain $\mathcal{D}_S = \{(V_{iS}, y_{iS})\}_{i=1}^{n_S}$ containing n_S videos. We assume that both the labeled source domain videos and the unlabeled target domain videos share the same C classes, yet \mathcal{D}_S is inaccessible when adapting the source model to \mathcal{D}_T .

Owing to the absence of the source domain during adaptation, SFVDA is more challenging while current VUDA methods cannot be applied. SFVDA should be tackled by adapting target video features to the source classifier, which contains information regarding source data distribution. The core is to extract source-like representations that satisfy the *cross-temporal hypothesis*, characterized by the consistency across local temporal features. We propose ATCoN, a novel network to transfer source models to the target domain by leveraging temporal features constructed attentively through learning temporal consistency in a self-supervised manner. We start with an introduction to the generation of the source model, followed by a thorough illustration of ATCoN.

3.1 Source Model Generation

A key prior for the transferred model to obtain effective temporal features is that the generated source model could extract precise temporal features. While conventional 3D-CNN-based extractors (e.g., 3D-ResNet [9] or I3D [1]) have been adopted in action recognition due to their performances, they extract spatio-temporal features jointly while temporal features are obtained implicitly by temporal pooling. In contrast, the Temporal Relation Network (TRN) [50] is adopted for SFVDA, thanks to its ability in obtaining more precise temporal features through reasoning over correlations between spatial representations, which corresponds with how humans would recognize actions.

Formally, an input source video with k frames can be expressed as $V_{iS} = \{f_{iS}^{(1)}, f_{iS}^{(2)}, \dots, f_{iS}^{(k)}\}$, where $f_{iS}^{(j)}$ is the spatial representation of the j -th frame in the i -th source video obtained from the source spatial feature extractor $G_{S,sp}$. $G_{S,sp}$ is formulated as a 2D-CNN (e.g., ResNet [10]). The temporal feature of V_{iS} is subsequently obtained from the source temporal feature extractor $G_{S,t}$, constructed by a combination of multiple local temporal features. Each local temporal feature is built upon clips with r temporal-ordered sampled frames where $r \in [2, k]$. Formally, a local temporal feature for V_{iS} , $lt_{iS}^{(r)}$, is defined by:

$$lt_{iS}^{(r)} = \sum_m g_S^{(r)}((V_{iS}^{(r)})_m), \quad (1)$$

where $(V_{iS}^{(r)})_m = \{f_{iS}^{(a)}, f_{iS}^{(b)}, \dots\}_m$ is the m -th clip with r temporal-ordered frames. a and b are the frame indices, which may not be consecutive as the clip with temporal-ordered frames could be extracted with nonconsecutive frames, but should be both in the range of $[1, k]$ with $b > a$. The local temporal feature $lt_{iS}^{(r)}$ is computed by fusing the time ordered frame-level spatial features through the integration function $g_S^{(r)}$, implemented as a Multi-Layer Perceptron (MLP).

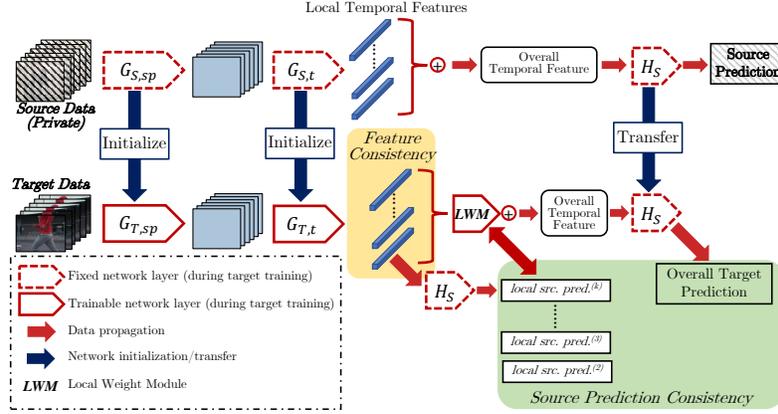


Fig. 1. Structure of the proposed ATCoN. ATCoN adopts the same network architecture for its spatial and temporal feature extractors as the source model, initialized by the source feature extractors. ATCoN extracts overall temporal features by learning *temporal consistency* over its local temporal features which includes both *feature consistency* and *source prediction consistency*. The *local weight module (LWM)* attends to more confident local temporal features. The overall target prediction is obtained by applying the *fixed* source classifier over the overall temporal feature. Dashed shapes indicates fixed network layers during adaptation.

$G_{S,t}$ is therefore a set of all integration functions $g_S^{(r)}$, namely $G_{S,t} = \{\forall_r g_S^{(r)}\}$. The final overall temporal feature \mathbf{t}_{iS} is a simple mean aggregation applied across all local temporal features, defined as: $\mathbf{t}_{iS} = \frac{1}{k-1} \sum_r lt_{iS}^{(r)}$. The source prediction is further computed by applying a source classifier H_S over \mathbf{t}_{iS} . The source model is trained with the standard cross-entropy loss as the objective function, formulated as:

$$\mathcal{L}_{S,ce} = -\frac{1}{n_S} \sum_{i=1}^{n_S} y_{iS} \log \sigma(H_S(\mathbf{t}_{iS})), \quad (2)$$

where σ is the softmax function whose c -th element is defined as $\sigma_c(x) = \exp(x_c) / \sum_{c=1}^C \exp(x_c)$. Inspired by [21], for the source model to be more discriminative and transferrable for better target data alignment, we further adopt the label smoothing technique [35] such that extracted features are encouraged to be distributed in tight clusters evenly separated [25]. By adopting the label smoothing technique, the objective function for training the source model can be further formulated as:

$$\mathcal{L}'_{S,ce} = -\frac{1}{n_S} \sum_{i=1}^{n_S} y'_{iS} \log \sigma(H_S(\mathbf{t}_{iS})), \quad (3)$$

where y'_{iS} is the smoothed label computed as $y'_{iS} = (1 - \epsilon)y_{iS} + \epsilon/C$ with ϵ being the smoothing parameter which is set to 0.1 empirically.

3.2 Attentive Temporal Consistent Network

With the absence of source data, conventional VUDA methods can no longer be applied. Instead, we tackle SFVDA from two perspectives: on the one hand,

extracting effective overall temporal features that are discriminative and comply with the *cross-temporal hypothesis* in a self-supervised manner, without either target label or source data; on the other hand, aligning to the source data distribution via attending to local temporal features with higher confidence in its relevance towards the source data distribution. Following the above strategies, we develop an **Attentive Temporal Consistent Network (ATCoN)**, whose structure is presented in Fig. 1. With the same network architecture adopted for the target spatial and temporal feature extractors $G_{T,sp}$ $G_{T,t}$ as that of $G_{S,sp}$ $G_{S,t}$, $G_{T,sp}$ and $G_{T,t}$ are initialized by $G_{S,sp}$ and $G_{S,t}$ respectively. The overall temporal feature is obtained by learning temporal consistency over the local temporal features as well as the respective local source predictions, resulted by applying the source classifier H_S over the local temporal features directly. Note that the source classifier remains *fixed* throughout the adaptation process. Meanwhile, for attentive aggregation of target local temporal features, a *Local Weight Module (LWM)* is further designed.

Learning Temporal Consistency. As presented in Section 3.1, the different local temporal features are extracted via the multiple temporal-ordered frames, sampled from the input video. For a given input video, these local temporal features should represent the same action even if they may differ in spatial appearances. Therefore, the overall temporal feature is effective and discriminative when the corresponding local temporal features are consistent in feature representations. Given a target input video $V_T \in \mathcal{D}_T$ (with video index i omitted for simplicity), its local temporal features for the set of clips with $r1$ and $r2$ temporal frames ($r1, r2 \in [2, k]$), $lt_T^{(r1)}$ and $lt_T^{(r2)}$, are defined similarly to Eq. 1. If the local temporal features are consistent, then the cross-correlation matrix between $lt_T^{(r1)}$ and $lt_T^{(r2)}$ should be close to the identity matrix. The cross-correlation matrix is formulated by:

$$\mathcal{C}^{r1r2} = \left(\hat{lt}_T^{(r1)} \right)^T \hat{lt}_T^{(r2)}, \quad (4)$$

where \hat{lt} is the normalized local temporal feature computed as:

$$\hat{lt} = \frac{lt - \mathbb{E}(lt)}{\sqrt{\text{Var}(lt) + \varepsilon}}, \quad (5)$$

with ε being a small bias value for numerical stability. The cross-correlation matrix \mathcal{C}^{r1r2} is a square matrix with the size of $d \times d$, where d is the dimension of the local temporal feature. Since \mathcal{C}^{r1r2} should ideally be close to an identity matrix, the feature consistency loss should maximize the similarity of the respective local temporal features while reducing redundancy between the components. Therefore, the feature consistency loss with respect to \mathcal{C}^{r1r2} is expressed as:

$$\mathcal{L}_{fc}^{r1r2} = \sum_i (1 - \mathcal{C}_{ii}^{r1r2})^2 + \lambda \sum_i \sum_{j \neq i} (\mathcal{C}_{ij}^{r1r2})^2, \quad (6)$$

where $i, j \in [0, d-1]$ are indexes of the local temporal feature dimension, while λ is a tradeoff constant. The final feature consistency loss is computed as the mean feature consistency loss over all cross-correlation matrices, with each matrix corresponding to a pair of local temporal features. The final feature consistency

loss can be formulated as:

$$\mathcal{L}_{fc} = \frac{1}{N_{fc}} \left(\sum_{r1} \sum_{r2 \neq r1} \mathcal{L}_{fc}^{r1r2} \right), \quad (7)$$

where $N_{fc} = P_2^{k-1}$ is the total number of local temporal feature pairs.

Moreover, since the local temporal features of the same input video should be consistent by minimizing \mathcal{L}_{fc} , their relevance towards the source data distribution should also be consistent. With source data inaccessible, such relevance cannot be computed directly through measuring the divergence between source and target data distributions. Since the source classifier contains source data distribution, such relevance could instead be approximated by the prediction of the source classifier over the local temporal features. In other words, the consistency over the relevance of target local temporal features towards source data distribution is equivalent to the consistency over the source prediction of target local temporal features. Meanwhile, the target overall temporal feature is obtained by aggregating the respective local temporal features. It should contain similar motion information as the local temporal features. Therefore, the consistency over source prediction could be extended to the overall temporal feature.

Given local temporal features $lt_T^{(2)}, \dots, lt_T^{(k)}$, the respective local source predictions $p_{lt,T}^{(2)}, \dots, p_{lt,T}^{(k)}$ are obtained via the fixed source classifier H_S , following: $p_{lt,T}^{(r)} = H_S(lt_T^{(r)})$, $\forall r \in [2, k]$. An average local source prediction could be obtained by averaging over the local source predictions $\bar{p}_{lt,T} = \frac{1}{k-1} \sum_{r=2}^k p_{lt,T}^{(r)}$. To achieve source prediction consistency, we aim to minimize the divergence between each local source predictions and the average local source prediction:

$$\mathcal{L}_{pc}^{local} = \frac{1}{k-1} \left(\sum_{r=2}^k KL(\log \sigma(p_{lt,T}^{(r)}) \| \log \sigma(\bar{p}_{lt,T})) \right), \quad (8)$$

where $KL(p||q)$ denotes the Kullback–Leibler (KL) divergence.

Further, the overall target prediction $p_{t,T}$ is computed by applying H_S to the target overall temporal feature \mathbf{t}_T , which is a simple mean aggregation applied across local temporal features $lt_T^{(2)}, \dots, lt_T^{(k)}$. To incorporate $p_{t,T}$ into the source prediction consistency, we aim to minimize the absolute difference between $p_{t,T}$ and $\bar{p}_{lt,T}$, defined as:

$$\mathcal{L}_{pc}^{overall} = \sum_{c=1}^C |\log \sigma_c(p_{t,T}) - \log \sigma_c(\bar{p}_{lt,T})|. \quad (9)$$

The final source prediction consistency is achieved by joint minimization of the prediction divergence between each local source prediction and the average local source prediction, as well as between the overall target prediction and the average local source prediction, formulated as: $\mathcal{L}_{pc} = \alpha_{local} \mathcal{L}_{pc}^{local} + \alpha_{overall} \mathcal{L}_{pc}^{overall}$, where α_{local} and $\alpha_{overall}$ are tradeoff constants. Learning temporal consistency is thus achieved by optimizing both the source prediction consistency loss and feature consistency loss jointly, expressed as: $\mathcal{L}_{tc} = \beta_{fc} \mathcal{L}_{fc} + \beta_{pc} \mathcal{L}_{pc}$, with β_{fc} and β_{pc} being the tradeoff hyperparameters.

Local Weight Module (LWM). While complying with the *cross-temporal hypothesis* via learning temporal consistency with feature and source prediction consistencies enables ATCoN to extract discriminative temporal features, we

observe that the overall temporal feature \mathbf{t}_T is constructed by simply averaging over all local temporal features. This would not be reasonable as the importance of each local temporal feature is commonly uneven. Therefore, we propose the *Local Weight Module (LWM)* to assign *local weights* to the local temporal features for subsequent attentive aggregation.

As mentioned in Section 3.2, ATCoN aims to tackle SFVDA by aligning target videos to the source data distribution. Therefore, LWM is designed such that local temporal features that are more confident towards its relevance to the source data distribution gains more attention, weighted by a *local relevance weight*. More specifically, following Section 3.2, the relevance towards source data distribution for $lt_T^{(r)}$ could be referred to its local source prediction $p_{lt,T}^{(r)} = H_S(lt_T^{(r)})$, from which the confidence score is computed. Subsequently, the confidence of $p_{lt,T}^{(r)}$ is defined as the additive inverse of its entropy computed over probabilities of all classes, formulated as:

$$\mathbb{C}(p_{lt,T}^{(r)}) = \sum_{c=1}^C \sigma_c(p_{lt,T,c}^{(r)}) \log \sigma_c(p_{lt,T,c}^{(r)}). \quad (10)$$

The *local relevance weight* corresponding to the local temporal feature $lt_T^{(r)}$ is finally generated by adding a residual connection for more stable optimization, expressed as: $w_{lt_T^{(r)}} = 1 + \mathbb{C}(p_{lt,T}^{(r)})$. The *local relevance weight* is applied to obtain the weighted overall temporal feature \mathbf{t}'_T , which is the mean aggregation of the corresponding weighted local temporal features, computed as: $\mathbf{t}'_T = \frac{1}{k-1} \sum_r w_{lt_T^{(r)}} lt_T^{(r)}$. Meanwhile, *local relevance weight* is further applied to the local source predictions $p_{lt,T}^{(r)}$, where the source prediction consistency is learnt with relevance-weighted local source predictions $p_{lt,T}^{(r)'} = w_{lt_T^{(r)}} p_{lt,T}^{(r)}$.

ATCoN learns temporal consistency by learning feature consistency and source prediction consistency of local temporal features jointly. Inspired by prior works in SFDA [21,15,38], we further improve ATCoN from two aspects:

Information Maximization. The ideal overall temporal feature should be both individually certain and globally diverse. Therefore, we apply an Information Maximization (IM) loss over the weighted overall temporal feature as:

$$\begin{aligned} \mathcal{L}_{IM} = & -\mathbb{E}_{V_T \in \mathcal{D}_T} \sum_{c=1}^C \sigma_c(H_S(\mathbf{t}'_T(V_T))) \log \sigma_c(H_S(\mathbf{t}'_T(V_T))) \\ & + \sum_{c=1}^C KL \left(\mathbb{E}_{V_T \in \mathcal{D}_T} [\sigma_c(H_S(\mathbf{t}'_T(V_T)))] \parallel \frac{1}{C} \right), \end{aligned} \quad (11)$$

where $\mathbf{t}'_T(V_T)$ is the weighted overall temporal feature corresponding to target video V_T , while σ_c is the c -th element in the softmax.

Self-supervised Pseudo-label Generation. To further improve class-wise alignment of ATCoN with the lack of target label, we follow [20] and generate pseudo-labels for target videos in a self-supervised manner. Specifically, pseudo-labels are generated through a repeated k-means clustering process over the

overall temporal feature, where the initial centroid for class c is attained by:

$$\mathbf{c}_c^{(0)} = \frac{\sum_{V_T \in \mathcal{D}_T} \sigma_c(H_S(\mathbf{t}'_T(V_T))) \mathbf{t}'_T(V_T)}{\sum_{V_T \in \mathcal{D}_T} \sigma_c(H_S(\mathbf{t}'_T(V_T)))}. \quad (12)$$

Subsequently, the initial pseudo-label of target data V_T is obtained by its nearest centroid, defined by: $\hat{y}_{V_T} = \arg \min_c \cos(\mathbf{t}'_T(V_T), \mathbf{c}_c^{(0)})$, where $\cos(\cdot, \cdot)$ denotes the cosine distance function. The initial centroids are further updated to characterize the category distribution of the target domain more reliably based on the initial pseudo-labels, formulated as:

$$\mathbf{c}_c^{(1)} = \frac{\sum_{V_T \in \mathcal{D}_T} \mathbb{I}(\hat{y}_{V_T}=c) \mathbf{t}'_T(V_T)}{\sum_{V_T \in \mathcal{D}_T} \mathbb{I}(\hat{y}_{V_T}=c)}, \quad (13)$$

with $\mathbb{I}(\cdot)$ being an indicator function. The pseudo-labels are finally renewed following the updated centroids with $\hat{y}_{V_T} = \arg \min_c \cos(\mathbf{t}'_T(V_T), \mathbf{c}_c^{(1)})$. ATCoN is further trained with the cross-entropy loss with respect to the pseudo-labels as:

$$\mathcal{L}_{T,ce} = -\frac{1}{n_T} \sum_{i=1}^{n_T} \hat{y}_{V_T} \log \sigma(H_S(\mathbf{t}'_T(V_T))), \quad (14)$$

where n_T is the total number of target videos.

Overall Objective. In summary, given a trained source model, the overall optimization objective of ATCoN is expressed as: $\mathcal{L} = \beta_{tc}\mathcal{L}_{tc} + \beta_{IM}\mathcal{L}_{IM} + \beta_{ce}\mathcal{L}_{T,ce}$, where β_{tc} , β_{IM} , and β_{ce} are tradeoff hyperparameters.

4 Experiments

In this section, we evaluate our proposed ATCoN across three cross-domain action recognition benchmarks including UCF-HMDB_{full} [2], Daily-DA [43] and Sports-DA [43]. These benchmarks cover a wide range of cross-domain scenarios. We present superior results on all benchmarks. Further, ablation studies and empirical analysis of ATCoN are also presented to validate the architecture of ATCoN. Code is provided at <https://github.com/xuyu0010/ATCoN>.

4.1 Experimental Settings

Among the three benchmarks, **UCF-HMDB_{full}** is one of the most widely used cross-domain video dataset, which contains videos from two public datasets: UCF101 (U101) [33] and HMDB51 (H51) [16], a total of 3,209 videos in 12 action classes, with 2 cross-domain action recognition tasks. Meanwhile, **Daily-DA** is a more challenging dataset that incorporates both normal videos and low-illumination videos. It is constructed from four datasets : ARID (A11) [42], HMDB51 (H51), Moments-in-Time (MIT) [24], and Kinetics (K600) [14]. While HMDB51, Moments-in-Time, and Kinetics are widely used for action recognition benchmarking, ARID is a more recent dark dataset, comprised with videos shot under adverse illumination conditions. In total, **Daily-DA** includes 18,949 videos from 8 classes, with a total of 12 cross-domain action recognition tasks. **Sports-DA** is a large-scale cross-domain video dataset, built from UCF101 (U101), Sports-1M (S1M) [13], and Kinetics (K600), with 23 action classes and

Table 1. Results for SFVDA on UCF-HMDB_{full} and Sports-DA.

Methods	Source-free	UCF-HMDB _{full}			Sports-DA						
		U101→H51	H51→U101	Avg.	K600→U101	K600→SIM	SIM→U101	SIM→K600	U101→K600	U101→SIM	Avg.
TRN	-	72.78	72.15	72.47	86.41	66.95	85.31	71.95	49.29	43.32	67.06
DANN	✗	74.44	75.13	74.79	86.60	66.79	89.32	70.53	61.77	48.73	70.62
MK-MMD	✗	74.72	79.69	77.21	86.49	66.18	87.37	71.43	64.17	49.24	70.81
TA ³ N	✗	78.14	84.83	81.49	88.24	70.56	83.32	75.54	57.51	46.37	70.26
SFDA	✓	69.86	74.98	72.42	86.10	60.02	85.37	68.04	55.75	43.58	66.48
SHOT	✓	74.44	74.43	74.44	91.19	64.95	88.84	72.02	53.93	43.58	69.09
SHOT++	✓	71.11	68.13	69.62	90.01	63.11	88.01	70.34	44.75	40.95	66.20
MA	✓	74.45	67.36	70.91	91.04	65.95	87.84	71.88	60.75	39.41	69.48
BAIT	✓	75.33	76.36	75.85	92.27	66.61	88.33	72.85	57.25	44.67	70.33
CPGA	✓	75.82	68.16	71.99	89.42	66.26	86.49	72.55	55.22	44.53	69.08
ATCoN	✓	79.72	85.29	82.51	93.62	<u>69.70</u>	90.64	75.99	65.24	<u>47.90</u>	73.85

Table 2. Results for SFVDA on Daily-DA.

Methods	Source-free	Daily-DA												
		K600→A11	K600→H51	K600→MIT	MIT→A11	MIT→H51	MIT→K600	H51→A11	H51→MIT	H51→K600	A11→H51	A11→MIT	A11→K600	Avg.
TRN	-	20.87	36.66	29.00	22.11	43.75	53.10	13.81	22.00	37.10	17.20	14.75	24.38	27.89
DANN	✗	21.18	37.50	21.75	22.81	43.33	58.76	14.20	29.50	38.24	20.11	19.75	27.03	29.51
MK-MMD	✗	21.66	36.25	24.00	21.02	50.42	58.48	20.35	25.75	33.79	18.75	18.00	26.07	29.55
TA ³ N	✗	19.87	37.67	31.53	21.57	43.01	55.47	14.38	25.71	38.59	14.92	15.56	23.42	28.49
SFDA	✓	12.57	44.95	27.50	15.96	35.19	49.23	13.08	24.25	24.86	16.29	13.25	25.22	25.19
SHOT	✓	12.03	44.58	29.50	15.28	36.67	51.04	13.58	24.25	21.24	17.08	14.00	24.35	25.30
SHOT++	✓	12.57	40.83	28.75	14.90	41.67	46.34	15.98	22.25	33.10	15.42	12.50	21.76	24.42
MA	✓	12.76	45.82	30.00	17.75	37.36	53.54	12.90	25.00	22.19	16.67	15.25	24.29	26.13
BAIT	✓	12.69	45.73	30.00	16.93	39.64	53.00	13.65	25.50	21.17	15.70	14.50	25.52	26.17
CPGA	✓	13.06	46.02	30.75	18.08	39.21	55.09	13.14	26.25	25.54	19.19	16.50	26.72	26.46
ATCoN	✓	<u>12.21</u>	48.25	32.50	27.23	<u>42.35</u>	<u>57.66</u>	<u>17.92</u>	30.75	48.55	26.67	<u>17.25</u>	31.05	33.53

a total of 40,718 videos. With three different domains, **Sports-DA** contains 6 cross-domain action recognition tasks. For fair comparison, all methods adopt the TRN [50] as the backbone for video feature extraction, with the source model pre-trained on ImageNet [5]. Following [21], a Batch Normalization [12] and an additional fully connected layer are inserted while weight normalization [32] is applied to the last fully connected layer. All experiments are implemented with PyTorch [27] library. *More specifications on benchmark details and network implementation are provided in the Appendix.*

4.2 Overall Results and Comparisons

We compare ATCoN with state-of-the-art SFDA approaches, as well as several competitive UDA/VUDA approaches. These include: SFDA [15], SHOT [20], SHOT++ [21], MA [18], BAIT [47] and CPGA [28] which are designed for source-free adaptation; as well as DANN [6], MK-MMD [22] and TA³N that are designed for UDA/VUDA scenario. We also report the results of the source-only model (TRN), which is obtained by applying the model trained with source data directly to the target data. We report the top-1 accuracy on the target domains, averaged on 5 runs with identical settings for each approach. Table 1 and Table 2 show the performance of our proposed ATCoN compared with the above methods in the three cross-domain action recognition benchmarks.

The results in Table 1 and Table 2 show that the novel ATCoN achieves the best results among source-free methods on all 20 cross-domain tasks across the three cross-domain benchmarks, and outperforms previous source-free approaches considerably by noticeable margins. Notably, ATCoN exceeds all prior SFDA approaches designed for the image-based SFDA task (e.g., SHOT, MA, and CPGA) consistently by an average of more than 10% relative improvements on mean accuracy over the second-best performances across all 18 cross-domain tasks. The consistent improvements empirically justify the effectiveness of learning temporal consistency for obtaining discriminative overall temporal features

Table 3. Ablations studies of ATCoN on UCF-HMDB_{full}

Methods	U101→H51	H51→U101
Source-only (TRN)	72.78	72.15
ATCoN	79.72	85.29
ATCoN- <i>FC</i>	77.78	83.36
ATCoN- <i>PC</i> [†]	76.67	82.83
ATCoN- <i>PC</i>	77.50	83.01
ATCoN- <i>TC</i>	78.89	84.59

(a) Components of temporal consistency

(b) Application of *local relevance weight*

while attending to local temporal features with high source prediction confidence. Our proposed ATCoN even exceeds the performance of VUDA methods which are trained with accessible source data under 13 cross-domain tasks, while the mean accuracies of our method are consistently higher than all VUDA methods evaluated across the three benchmarks. This further validates the capability of ATCoN in constructing effective temporal features.

Further, it could be observed that prior SFDA approaches could not tackle SFVDA well. Specifically, in 11 out of the 20 cross-domain tasks, more than half of the evaluated SFDA approaches result in performances inferior to that of the source-only model trained without any adaptation approaches. Prior SFDA approaches could only handle spatial features, while unable to obtain discriminative and transferrable temporal features, resulting in little or negative improvements compared to the source-only baseline. This further demonstrates the challenges faced when adapting video models under the source-free scenario. In particular, all tasks that involve ARID as the source or target domain in **Daily-DA** would lead to inferior results by prior SFDA approaches. This scenario could be further owed to the fact that videos in ARID are collected in adverse illumination with distinct statistical characteristics, leading to larger cross-domain gaps.

4.3 Ablation Studies and Feature Visualization

To dive deeper into the effectiveness of ATCoN and validate its architecture, we perform detailed ablation studies and feature visualization. The ablation studies investigate ATCoN from two perspectives: firstly, the components of temporal consistency; and secondly the application of *local relevance weight* generated by *LWM*. All ablation studies are conducted utilizing the **UCF-HMDB_{full}** dataset with 2 cross-domain action recognition tasks, while TRN is adopted as the feature extractor backbone.

Temporal Consistency. We assess ATCoN against 4 variants to validate the design of the proposed temporal consistency loss \mathcal{L}_{tc} : **ATCoN-FC**, where only the feature consistency is learnt; **ATCoN-PC**[†] and **ATCoN-PC**, where only the source prediction consistency is learnt, with the overall target prediction not included for **ATCoN-PC**[†]; and finally **ATCoN-TC**, where only the temporal consistency loss is learnt with both feature consistency and source prediction consistency. The above 4 variants would not apply both the IM loss and pseudo-label generation as proposed in Eq. 11 and 14 during training, while the *local*

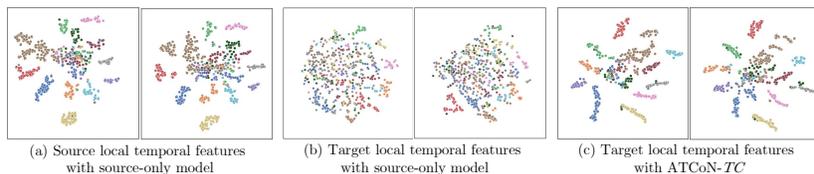


Fig. 2. t-SNE visualizations of local temporal features with class information. Different colors represent different classes.

relevance weight in Sec. 3.2 is applied. Results in Table 3(a) demonstrate the efficacy of learning temporal consistency for constructing discriminative overall temporal features for tackling SFVDA. By learning either feature consistency or source prediction consistency, the network is able to outperform all prior SFDA approaches on both cross-domain tasks. Meanwhile, extending the source prediction consistency to the overall temporal feature further improves its efficacy. The superior performance of ATCoN-TC justifies that learning feature consistency and source prediction consistency complements each other.

Further, it could be observed that ATCoN performs slightly better than ATCoN-TC, thanks to the inclusion of both IM loss and pseudo-labeling in training the full ATCoN. However, compared to the improvements towards the baseline model performance brought by learning temporal consistency, the performance gain by applying both IM loss and pseudo-labeling is marginal. The comparison empirically proves that the key towards ATCoN’s success lies more in the learning of temporal consistency.

Applying Local Relevance Weight. We propose the *local relevance weight* w_{lt} obtained from the LWM which attends to the local temporal features with high confidence over their relevance to the source data distribution. To justify the necessity of the w_{lt} , we compare ATCoN against 3 variants: **ATCoN-NA**, where the LWM is not inserted thus w_{lt} is not obtained at all; **ATCoN-A@F**, where w_{lt} is only applied for obtaining the overall temporal feature \mathbf{t}'_T ; and **ATCoN-A@P**, where w_{lt} is only applied to obtain the weighted local source prediction $p_{lt,T}^{(r) \prime}$. Both the IM loss and pseudo-label generation are adopted during the training of the three aforementioned variants. As illustrated in Table 3(b), applying *local relevance weight* brings consistent improvements wherever it has been applied, which justifies the necessity for such a weight. By employing the *local relevance weight* w_{lt} , ATCoN is able to obtain more discriminative temporal features. While w_{lt} bring further improvements on network performance, it should be noted that the improvement is relatively marginal compared to that brought by learning temporal consistency, which indicates that the proposed temporal consistency plays a more vital role in tackling SFVDA effectively.

Feature Visualization. To further understand the characteristics of ATCoN, we plot the t-SNE embeddings [23] of the features extracted. Specifically, we first prove our *cross-temporal hypothesis* by visualizing local temporal features learned by the source-only model on the source data and the target data, and local temporal features learned by ATCoN-TC for the H51→U101 task, as pre-

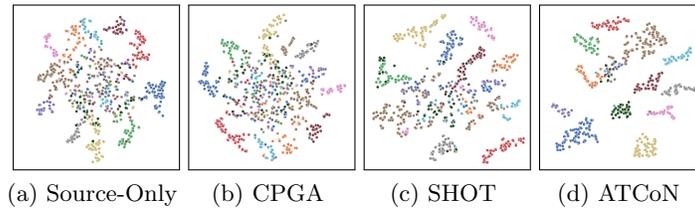


Fig. 3. Visualization of features extracted by the (a) source-only model, (b) CPGA, (c) SHOT, and (d) ATCoN with class information. Different classes are marked by different colors.

sented in Fig. 2. The local temporal features of the source data share similar distribution patterns, which confirms that they are both discriminative and consistent, with similar semantic information embedded. Meanwhile, the data distribution patterns of target data with the source model are inconsistent. In comparison, by learning temporal consistency, ATCoN-TC is able to extract discriminative and relatively consistent local temporal features, satisfying the *cross-temporal hypothesis*. This implies that learning temporal consistency enables the learning of source-like representations for target data, and therefore is effective in aligning target data to source data distribution.

We further plot the t-SNE embeddings of the overall temporal features learnt by ATCoN, CPGA, and SHOT for the H51→U101 task with class information in the target domain. The results are presented in Fig. 3, where we can clearly observe that the features learned by ATCoN are much more clustered than those learned by other networks. This verifies that features learned by ATCoN are of higher discriminability, resulting in better SFVDA performance. In contrast, features learned by CPGA are even less clustered and discriminative than those learned by the source-only backbone, which corresponds to its inferior performance over the backbone in this task. The above observation implies the superiority of our ATCoN in tackling SFVDA while reflecting the challenges faced by prior SFDA approaches in tackling SFVDA.

5 Conclusion

In this work, we pioneer in formulating the challenging yet realistic Source-Free Video Domain Adaptation (SFVDA) problem, which addresses data-privacy issues in videos. We proposed a novel ATCoN to tackle SFVDA effectively. With source video data inaccessible, ATCoN tackles SFVDA via obtaining effective and discriminative overall temporal features satisfying the *cross-temporal hypothesis*, achieved by learning temporal consistency, guaranteed by both feature consistency and source prediction consistency. ATCoN further aims to align target data to the source distribution through attending to local temporal features with higher source prediction confidence. Extensive experiments and detailed ablation studies across multiple cross-domain action recognition benchmarks validate the superiority of our proposed ATCoN in tackling SFVDA.

References

1. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017) [5](#)
2. Chen, M.H., Kira, Z., AlRegib, G., Yoo, J., Chen, R., Zheng, J.: Temporal attentive alignment for large-scale video domain adaptation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 6321–6330 (2019) [2](#), [4](#), [10](#)
3. Chen, M.H., Li, B., Bao, Y., AlRegib, G., Kira, Z.: Action segmentation with joint self-supervised temporal domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9454–9463 (2020) [4](#)
4. Choi, J., Sharma, G., Schuler, S., Huang, J.B.: Shuffle and attend: Video domain adaptation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 678–695. Springer (2020) [2](#), [4](#)
5. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009) [11](#)
6. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: International conference on machine learning. pp. 1180–1189. PMLR (2015) [4](#), [11](#)
7. Ghifary, M., Kleijn, W.B., Zhang, M., Balduzzi, D., Li, W.: Deep reconstruction-classification networks for unsupervised domain adaptation. In: European conference on computer vision. pp. 597–613. Springer (2016) [4](#)
8. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. *Advances in neural information processing systems* **17** (2004) [2](#)
9. Hara, K., Kataoka, H., Satoh, Y.: Learning spatio-temporal features with 3d residual networks for action recognition. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 3154–3160 (2017) [5](#)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) [5](#)
11. Huang, L., Joseph, A.D., Nelson, B., Rubinstein, B.I., Tygar, J.D.: Adversarial machine learning. In: Proceedings of the 4th ACM workshop on Security and artificial intelligence. pp. 43–58 (2011) [4](#)
12. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International conference on machine learning. pp. 448–456. PMLR (2015) [11](#)
13. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 1725–1732 (2014) [10](#)
14. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., Suleyman, M., Zisserman, A.: The kinetics human action video dataset (2017) [2](#), [10](#)
15. 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). <https://doi.org/10.1109/TAI.2021.3110179> [9](#), [11](#)
16. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: Hmdb: a large video database for human motion recognition. In: 2011 International Conference on Computer Vision. pp. 2556–2563. IEEE (2011) [10](#)

17. Kurmi, V.K., Subramanian, V.K., Namboodiri, V.P.: Domain impression: A source data free domain adaptation method. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 615–625 (2021) [4](#)
18. 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) [2](#), [4](#), [11](#)
19. Li, S., Xie, M., Lv, F., Liu, C.H., Liang, J., Qin, C., Li, W.: Semantic concentration for domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9102–9111 (2021) [2](#)
20. Liang, J., Hu, D., Feng, J.: Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In: International Conference on Machine Learning. pp. 6028–6039. PMLR (2020) [2](#), [4](#), [9](#), [11](#)
21. 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) [2](#), [4](#), [6](#), [9](#), [11](#)
22. Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: International conference on machine learning. pp. 97–105. PMLR (2015) [2](#), [4](#), [11](#)
23. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(11) (2008) [13](#)
24. Monfort, M., Andonian, A., Zhou, B., Ramakrishnan, K., Bargal, S.A., Yan, T., Brown, L., Fan, Q., Gutfreund, D., Vondrick, C., et al.: Moments in time dataset: one million videos for event understanding. *IEEE transactions on pattern analysis and machine intelligence* **42**(2), 502–508 (2019) [2](#), [10](#)
25. Müller, R., Kornblith, S., Hinton, G.: When does label smoothing help? arXiv preprint arXiv:1906.02629 (2019) [6](#)
26. Pan, B., Cao, Z., Adeli, E., Niebles, J.C.: Adversarial cross-domain action recognition with co-attention. In: AAAI. pp. 11815–11822 (2020) [4](#)
27. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. In: Advances in neural information processing systems. pp. 8026–8037 (2019) [11](#)
28. 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) [4](#), [11](#)
29. 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) [2](#)
30. Saito, K., Kim, D., Sclaroff, S., Saenko, K.: Universal domain adaptation through self supervision. *Advances in neural information processing systems* **33**, 16282–16292 (2020) [2](#)
31. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3723–3732 (2018) [4](#)
32. Salimans, T., Kingma, D.P.: Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems* **29** (2016) [11](#)
33. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 (2012) [10](#)

34. Sun, B., Feng, J., Saenko, K.: Return of frustratingly easy domain adaptation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 30 (2016) [2](#), [4](#)
35. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2818–2826 (2016) [6](#)
36. Viola, P., Wells III, W.M.: Alignment by maximization of mutual information. *International journal of computer vision* **24**(2), 137–154 (1997) [2](#)
37. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2517–2526 (2019) [2](#)
38. 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) [9](#)
39. Xie, S., Zheng, Z., Chen, L., Chen, C.: Learning semantic representations for unsupervised domain adaptation. In: International conference on machine learning. pp. 5423–5432. PMLR (2018) [2](#)
40. Xu, Y., Yang, J., Cao, H., Chen, Z., Li, Q., Mao, K.: Partial video domain adaptation with partial adversarial temporal attentive network. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9332–9341 (2021) [2](#)
41. Xu, Y., Yang, J., Cao, H., Mao, K., Yin, J., See, S.: Aligning correlation information for domain adaptation in action recognition (2021) [2](#), [4](#)
42. Xu, Y., Yang, J., Cao, H., Mao, K., Yin, J., See, S.: Arid: A new dataset for recognizing action in the dark. In: International Workshop on Deep Learning for Human Activity Recognition. pp. 70–84. Springer (2021) [10](#)
43. Xu, Y., Yang, J., Cao, H., Wu, K., Wu, M., Zhao, R., Chen, Z.: Multi-source video domain adaptation with temporal attentive moment alignment. arXiv preprint arXiv:2109.09964 (2021) [2](#), [4](#), [10](#)
44. Yang, J., Yang, J., Wang, S., Cao, S., Zou, H., Xie, L.: Advancing imbalanced domain adaptation: Cluster-level discrepancy minimization with a comprehensive benchmark. *IEEE Transactions on Cybernetics* (2021) [4](#)
45. Yang, J., Zou, H., Zhou, Y., Zeng, Z., Xie, L.: Mind the discriminability: Asymmetric adversarial domain adaptation. In: European Conference on Computer Vision. pp. 589–606. Springer (2020) [2](#)
46. Yang, J., An, W., Wang, S., Zhu, X., Yan, C., Huang, J.: Label-driven reconstruction for domain adaptation in semantic segmentation. In: European Conference on Computer Vision. pp. 480–498. Springer (2020) [4](#)
47. Yang, S., Wang, Y., van de Weijer, J., Herranz, L., Jui, S.: Unsupervised domain adaptation without source data by casting a bait. arXiv preprint arXiv:2010.12427 (2020) [4](#), [11](#)
48. Yeh, H.W., Yang, B., Yuen, P.C., Harada, T.: Sofa: Source-data-free feature alignment for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 474–483 (2021) [2](#)
49. Zhang, Y., Liu, T., Long, M., Jordan, M.: Bridging theory and algorithm for domain adaptation. In: International Conference on Machine Learning. pp. 7404–7413. PMLR (2019) [4](#)
50. Zhou, B., Andonian, A., Oliva, A., Torralba, A.: Temporal relational reasoning in videos. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 803–818 (2018) [5](#), [11](#)