

# Partial Video Domain Adaptation with Partial Adversarial Temporal Attentive Network

Yuecong Xu<sup>1,2,\*</sup>  
 xuyu0014@e.ntu.edu.sg

Jianfei Yang<sup>1,\*</sup>  
 yang0478@ntu.edu.sg

Haozhi Cao<sup>1</sup>  
 haozhi001@e.ntu.edu.sg

Zhenghua Chen<sup>2,†</sup>  
 chen0832@e.ntu.edu.sg

Qi Li<sup>1</sup>  
 liqi0024@e.ntu.edu.sg

Kezhi Mao<sup>1</sup>  
 ekzmao@ntu.edu.sg

<sup>1</sup>School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore

<sup>2</sup>Institute for Infocomm Research, A\*STAR, Singapore

## Abstract

*Partial Domain Adaptation (PDA) is a practical and general domain adaptation scenario, which relaxes the fully shared label space assumption such that the source label space subsumes the target one. The key challenge of PDA is the issue of negative transfer caused by source-only classes. For videos, such negative transfer could be triggered by both spatial and temporal features, which leads to a more challenging Partial Video Domain Adaptation (PVDA) problem. In this paper, we propose a novel Partial Adversarial Temporal Attentive Network (PATAN) to address the PVDA problem by utilizing both spatial and temporal features for filtering source-only classes. Besides, PATAN constructs effective overall temporal features by attending to local temporal features that contribute more toward the class filtration process. We further introduce new benchmarks to facilitate research on PVDA problems, covering a wide range of PVDA scenarios. Empirical results demonstrate the state-of-the-art performance of our proposed PATAN across the multiple PVDA benchmarks. Code will be provided at: <https://github.com/xuyu0010/PATAN>.*

## 1. Introduction

Video-based problems have long been studied thanks to their wide applications in various fields. Neural networks have made notable advances in these problems with the availability of large-scale labeled video data. However, sufficiently large-scale training video data is sometimes unavailable, as annotations of video data are costly. Various

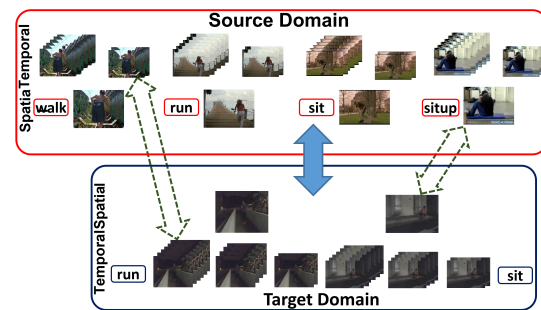


Figure 1. PVDA is a more general setting where the source label space subsumes the target label space. The key challenge of PVDA is the negative transfer caused by outlier source-only classes ('walk' and 'situp'), with extra probability triggered by the incorrect alignment of target temporal features to the source temporal features of the outlier classes, depicted as the left dashed arrow between videos from classes 'run' and 'walk'.

*Video-based Unsupervised Domain Adaptation (VUDA)* methods have been proposed to enable networks to transfer knowledge from a labeled source domain to an unlabeled target domain by learning domain-invariant feature representations in the absence of target labels.

Though existing VUDA methods enable the learning of transferable features across domains, they generally assume that the video source and target domains share an identical label space, which may not hold in real-world applications. With the presence of large-scale labeled public datasets, it is more feasible to transfer representations learned in these datasets to unlabeled small-scale datasets. Such a scenario is defined as *Partial Domain Adaptation (PDA)*, which relaxes the constraint of identical source and target label spaces by assuming that the target label space is a subspace of the source one. This assumption is more practical since large-scale public video datasets can subsume categories of the small-scale target datasets. The PDA

\*Equal Contribution.

†Corresponding Author.

‡This research is supported by A\*STAR under its AME Programmatic Funds (Grant No. A20H6b0151).

problem is more challenging since source-only classes may negatively influence the distribution alignment of target data causing *negative transfer*.

Compared to images that only contain spatial features, videos contain additional temporal features. This leads to a novel *Partial Video Domain Adaptation* (PVDA) problem, with trained networks transferred from video source domain to target domain, with the label space of video target domain being the subspace of the video source domain. Models trained in large-scale video datasets such as Kinetics [17] could therefore be applied to smaller-scale datasets such as ARID [31] without supervision through PVDA. When transferring networks for PVDA, negative transfer would be triggered due to the possible spatial-temporal domain shift as depicted in Figure 1, where the appearances of videos in class ‘walk’ are different from that of videos in class ‘run’, i.e. spatial features are different among videos in the two classes. However, videos from both classes share similar motion patterns where the actor moves further away from the camera in an upright position, indicating similar temporal features among the videos. When performing data distribution alignment, the similarities in temporal features would lead to videos in class ‘run’ of the target domain to incorrectly align with videos in class ‘walk’ of the source domain, triggering negative transfer.

A crucial step for tackling negative transfer in PVDA is the filtration of source-only outlier classes. Different from images, temporal features should be leveraged for PVDA from two perspectives: on one hand, effective temporal features should be constructed such that temporal features in outlier source-only classes discriminate those in target classes, alleviating the possibility of triggering negative transfer by temporal features; on the other hand, the temporal features should also contribute towards the filtration of source-only classes while eliminating possible mistakes caused by mis-classification of spatial features. To this end, we propose a **Partial Adversarial Temporal Attentive Network (PATAN)** to address the two challenges uniformly. **PATAN** first constructs robust overall temporal features by attentive combination of local temporal features which contain different aspects of the whole motion. The attentive combination builds upon the contribution of the local temporal features towards the class filtration process where source-only classes are filtered. The constructed temporal features would therefore have higher discriminability over source-only and target classes. Further, **PATAN** mitigates negative transfer in PDA through a class filtration process by utilizing local and overall temporal features jointly, alleviating possible mistakes during the class filtration process brought by the spatial features.

To further facilitate PVDA research, we propose three sets of benchmarks, built from widely used public datasets and a recent video dataset dedicated to low-

illumination videos. The benchmarks proposed are: (a) *UCF-HMDB<sub>partial</sub>*, (b) *MiniKinetics-UCF*, and (c) *HMDB-ARID<sub>partial</sub>*. The proposed datasets cover a wide range of PVDA scenarios, providing adequate baselines with distinct domain shift.

In summary, our contributions are threefold. First, we formulated a novel and challenging *Partial Video Domain Adaptation* (PVDA) problem. To the best of our knowledge, this is the first research that explores partial transfer in videos. Secondly, we analyze the challenges underlying PVDA and introduce PATAN to address the challenges. PATAN constructs robust temporal features, while utilizing both spatial and temporal features for accurate class filtration. Finally, we introduce several PVDA benchmarks, and demonstrate the effectiveness of our proposed method, achieving state-of-the-art performance across the multiple PVDA benchmarks proposed.

## 2. Related Work

**Unsupervised Domain Adaptation.** Unsupervised Domain Adaptation (UDA) aims to distill shared knowledge across labeled source domain and unlabeled target domain, improving the transferability of models. With the success of Generative Adversarial Network (GAN) [13], researchers have proposed to align the cross-domain data distributions with additional domain discriminators that are trained with the feature generators in an adversarial manner [16], and construct adversarial loss [11] for UDA. Subsequently, various adversarial-based UDA methods [27, 14, 37] have been proposed for a wide range of image-based tasks, such as image recognition [11, 26, 33], object detection [7, 1, 36] and semantic segmentation [38, 29, 8]. More recently, with the wide applications of videos in various fields, there has been increasing research for Video-based Unsupervised Domain Adaptation (VUDA). The success of constructing domain-invariant features with adversarial-based methods extends to VUDA. This gives rise to the introduction of various adversarial-based VUDA approaches for tasks such as action recognition [5, 9, 21] and action segmentation [6].

**Partial Domain Adaptation.** While the approaches for UDA and VUDA advances rapidly, these approaches assume that source and target domains share the same label space. A more general scenario that relaxes such assumption is introduced. Partial Domain Adaptation (PDA) [2] enables models to transfer knowledge from many-class domains to few-class domains. Currently, there are multiple efforts towards the PDA problem. Among these, Selective Adversarial Network (SAN) [2] adopts a multi-discriminator domain adversarial network with a weighting mechanism to select out source-only classes. Partial Adversarial Domain Adaptation (PADA) [3] improves SAN by employing a single discriminator adversarial network and further applies the class weight to the source classi-

fier, while IWAN [32] proposes a two domain classifier strategy to identify the importance score of source samples. More recently, Example Transfer Network (ETN) [4] is introduced to quantify the transferability of source data in a progressive weighting scheme through a discriminative domain discriminator. Generally, the PDA approaches above mitigate negative transfer by filtering out source-only outlier classes during the data alignment process.

Despite the notable advances achieved in PDA, the approaches are all built for image-based PDA problems, whereas Partial Video Domain Adaptation (PVDA) has not been tackled. PVDA is more challenging given that negative transfer could be triggered by temporal features, unique for videos. We propose to tackle PVDA with a novel network that constructs robust temporal features while utilizing spatial-temporal features for accurate class filtration.

### 3. Proposed Method

In the scenario of *Partial Video Domain Adaptation* (PVDA), we are given a source domain  $\mathcal{D}_S = \{(V_{iS}, y_{iS})\}_{i=1}^{n_S}$  with  $n_S$  labeled videos associated with  $|\mathcal{C}_S|$  classes, and a target domain  $\mathcal{D}_T = \{V_{iT}\}_{i=1}^{n_T}$  with  $n_T$  unlabeled videos associated with  $|\mathcal{C}_T|$  classes. The PVDA scenario is more general than VUDA by assuming that the source label space  $\mathcal{C}_S$  is a superset of the target label space  $\mathcal{C}_T$ , i.e.  $\mathcal{C}_T \subset \mathcal{C}_S$ . The source and target domains of PVDA are characterized by two underlying probability distributions  $p$  and  $q$  respectively, where  $p \neq q$ . We also have  $p_{\mathcal{C}_T} \neq q$ , where  $p_{\mathcal{C}_T}$  denotes the distribution of the source domain data in label space  $\mathcal{C}_T$  of target domain.

To tackle the PVDA problem, we aim to construct a network capable of learning transferable features across source and target domains and minimizing the target classification risk. Compared to VUDA, PVDA poses more challenges to the network due to the existence of outlier label space in the source domain  $\mathcal{C}_S \setminus \mathcal{C}_T$ , which causes negative transfer effect to the network’s performance. Meanwhile, during the training of the network, only unlabeled target domain data are accessible. Hence the part of which  $\mathcal{C}_S$  shares with  $\mathcal{C}_T$  is unknown. Therefore the key towards mitigating negative effects lies in the class filtration process which filters out the outlier source-only classes.

Current PDA approaches are built for image-based PDA problems, where the negative transfer could only be triggered by the alignment of spatial features. Whereas for videos, negative transfer could be additionally triggered by the alignment of temporal features (e.g. the alignment of target videos in ‘run’ to source videos in ‘walk’ as depicted in Figure 1). Due to the fact that current feature extractors would pay more attention across the spatial dimension, current PDA approaches may not be sensitive to negative transfer caused by the incorrect alignment of temporal features. Therefore, we propose a novel Partial Adversarial Temporal

Attentive Network (PATAN), to enable partial domain adaptation in an adversarial manner while mitigating negative transfer utilizing attentive temporal features. We begin by reviewing adversarial-based partial domain adaptation approaches, followed by a detailed illustration of PATAN.

#### 3.1. Adversarial-based Partial Domain Adaptation

Domain adaptation (DA) is achieved by matching the feature distributions of the source and target domains. One major line of approaches learn the domain-invariant features  $\mathbf{f}$  in an adversarial manner where additional domain discriminators are trained with the feature generators in a min-max fashion. More specifically, the parameters  $\theta_f$  of the feature extractor  $G_f$  are learned by maximizing the losses of the domain discriminator  $G_d$ , while the parameters  $\theta_d$  of the domain discriminator  $G_d$  are trained by minimizing the losses of the domain discriminators  $G_d$ . Additionally, the loss of the source classifier  $G_y$  is also minimized. The overall objective of adversarial-based DA networks can be formulated as in [12]:

$$\mathcal{L}_0 = \frac{1}{n_S} \sum_{z_i \in \mathcal{D}_S} L_y(G_y(G_f(z_i)), y_i) - \frac{\lambda}{n_A} \sum_{z_i \in \mathcal{D}_A} L_d(G_d(G_f(z_i)), d_i), \quad (1)$$

where  $z_i$  is an input data,  $\mathcal{D}_A = \mathcal{D}_S \cup \mathcal{D}_T$  is the union of source and target domains with  $n_A = |\mathcal{D}_A|$ ,  $d_i$  is the domain label of input  $z_i$ , and  $\lambda$  is the trade-off for the domain loss  $L_d$  with respect to the source classification loss  $L_y$ . Both losses are implemented as cross-entropy losses. The min-max optimization process is achieved by the connecting a Gradient Reverse Layer (GRL) to  $G_d$ .

While the aforementioned adversarial-based networks can be applied to standard DA tasks, yielding reliable results, their performance deteriorates for PDA tasks due to negative transfer caused by outlier source-only classes within the label space of  $\mathcal{C}_S \setminus \mathcal{C}_T$ . Hence a class filtration process is applied to filter out these outlier classes.

The end result of this class filtration process are class weights  $\gamma$  for each source domain label  $l \in \mathcal{C}_S$ , indicating the probability of each class of label space  $\mathcal{C}_S$  overlapping with label space  $\mathcal{C}_T$ . To obtain the class weights  $\gamma$ , it is observed that the output of the source classifier  $G_y$  for data  $z_i$  well represents the probability distribution of  $z_i$  over the source label space  $\mathcal{C}_S$ . The probability of the target data assigned to labels with space overlapped between  $\mathcal{C}_S$  and  $\mathcal{C}_T$  should be significantly larger than the probability of the target data assigned to outlier classes with label space  $\mathcal{C}_S \setminus \mathcal{C}_T$ . Therefore, the class weights  $\gamma$  are generally obtained by the label predictions of the target data through the source classifier  $G_y$ , which indicates the probability of assigning the target data to each source class, and is formulated as:

$$\gamma = \frac{1}{n_T} \sum_{z_i \in \mathcal{D}_T} G_y(G_f(z_i)) = \frac{1}{n_T} \sum_{i=1}^{n_T} \hat{y}_i, \quad (2)$$

where  $\hat{y}_i$  is the label prediction of target input data  $z_i$ .

To down-weight the contributions of the data in outlier classes for PDA tasks, the class weights  $\gamma$  are applied to the adversarial-based networks, yielding the objective of PDA networks formulated as in [3]:

$$\begin{aligned} \mathcal{L}_p = & \frac{1}{n_S} \sum_{z_i \in \mathcal{D}_S} \gamma_{y_i} [L_y(G_y(G_f(z_i)), y_i) - \lambda L_d(G_d(G_f(z_i)), d_i)] \\ & - \frac{\lambda}{n_T} \sum_{z_i \in \mathcal{D}_T} L_d(G_d(G_f(z_i)), d_i), \end{aligned} \quad (3)$$

where  $y_i$  is the ground truth of input  $z_i$  in the source domain, while  $\gamma_{y_i}$  is the corresponding class weight.

### 3.2. PATAN: Partial Adversarial Temporal Attentive Network

Intuitively, when tackling the PVDA problem, the approach in Section 3.1 could be directly integrated into videos, i.e.  $z_i = V_i$ . Given that videos contain both spatial and temporal features, one typical method for obtaining transferable video features is by separating the feature extractor  $G_f$  into a spatial feature  $G_{spf}$  and temporal feature extractor  $G_{tf}$ . The network constructed for PVDA could be formulated by simply substituting  $G_f(z_i)$  in Equation 2 and Equation 3 with  $G_{tf}(G_{spf}(V_i))$ .

One major drawback of direct integration of the above PDA approach into videos lies in the fact that video representations obtained through conventional video feature extractors (e.g. convolution-based networks) are mainly from the spatial features. The overall temporal information is generally encoded implicitly, usually implemented as a temporal pooling mechanism. Without explicit temporal features, the class filtration process in Section 3.1 would depend mainly on the spatial features. Therefore the negative transfer may only be alleviated along the spatial dimension.

In view of such drawback, we propose **Partial Adversarial Temporal Attentive Network (PATAN)** to mitigate negative transfer by utilizing spatial and temporal features jointly, as shown in Figure 2. To utilize the temporal features of videos for negative transfer mitigation, the temporal features should be explicitly extracted first. Given the fact that humans can recognize actions by reasoning over the observations across time, temporal features could be extracted utilizing Temporal Relation Module [35]. We denote an input video with  $k$  frames as  $V_i = \{x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(k)}\}$ , where  $x_i^{(j)}$  is the  $j$ th frame-level feature representation of the  $i$ th video obtained from the spatial feature extractor  $G_{spf}$ . The temporal feature of  $V_i$  is constructed by a combination of multiple local temporal features, each built upon clips with  $r$  temporal-ordered sampled frames where  $r \in [2, k]$ . Formally, a local temporal feature  $\mathbf{f}_i^r$  is defined by:

$$\mathbf{f}_i^r = \sum_m g^r((V_i^r)_m), \quad (4)$$

where  $(V_i^r)_m = \{x_i^{(a)}, x_i^{(b)}, \dots\}_m$  is the  $m$ th clip with  $r$

temporal-ordered frames, with  $a$  and  $b$  denoting the frame indices. The indices  $a$  and  $b$  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 of  $\mathbf{f}_i^r$  is computed by fusing the time ordered frame-level features through function  $g^r$ , implemented as an Multi-Layer Perceptron (MLP).

Further, the key for mitigating negative transfer along the temporal dimension lies in the design of an effective class filtration process to down-weight the effects of outlier classes with temporal features. The class filtration process is built upon the observation that the probability of the target data assigned to outlier classes with label space  $\mathcal{C}_S \setminus \mathcal{C}_T$  should be significantly small. To make full use of the local temporal features and to eliminate possible mis-assignment of target classes to spatial features, we apply the above observation to all local temporal features of the target data. The label prediction of local temporal feature  $\mathbf{f}_i^r$  is obtained as  $\hat{y}_{ti}^r = \tilde{C}_{ty}^r(\mathbf{f}_i^r)$ , which gives the probability distribution of  $\mathbf{f}_i^r$  across the source label space  $\mathcal{C}_S$ . Here  $\tilde{C}_{ty}^r$  is the auxiliary source classifier for  $\mathbf{f}_i^r$ , and is trained as cross-entropy loss with source local temporal features, i.e.  $\mathbf{f}_{i'}^r$  where  $V_{i'} \in \mathcal{D}_S$ .

To obtain the overall temporal feature and the class weights of each source class, one straight-forward strategy is to aggregate all local temporal features and their corresponding label predictions. However, not all local temporal features are equally important towards the mitigation of negative transfer. We introduce a *label attention* mechanism to attend to local temporal features that contribute more toward the class filtration process. Specifically, the temporal features would be robust and the class filtration process would be effective only if temporal features in outlier source-only classes discriminate from those in target classes. If the features are of low discriminability and therefore the predictions are uncertain, the class weights of source classes which correlates with the predictions would be similar across all source classes. The network would be unable to filter out outlier source-only classes. Therefore, the proposed network should construct effective overall temporal features which attend to discriminable features that better distinguish if the label of the input data lies within the target label space  $\mathcal{C}_T$  or the outlier label space  $\mathcal{C}_S \setminus \mathcal{C}_T$ . The certainty of the label prediction  $\hat{y}_{ti}^r$  which corresponds to  $\mathbf{f}_i^r$  is quantified by the additive inverse of the entropy of the label prediction as:

$$\mathbb{C}(\hat{y}_{ti}^r) = \sum_{c=1}^{|\mathcal{C}_S|} \hat{y}_{ti,c}^r \log(\hat{y}_{ti,c}^r). \quad (5)$$

For more stable optimization, a residual connection is added towards the formulation of the local temporal feature weight. The weight  $\hat{w}_i^r$  of the local temporal feature  $\mathbf{f}_i^r$

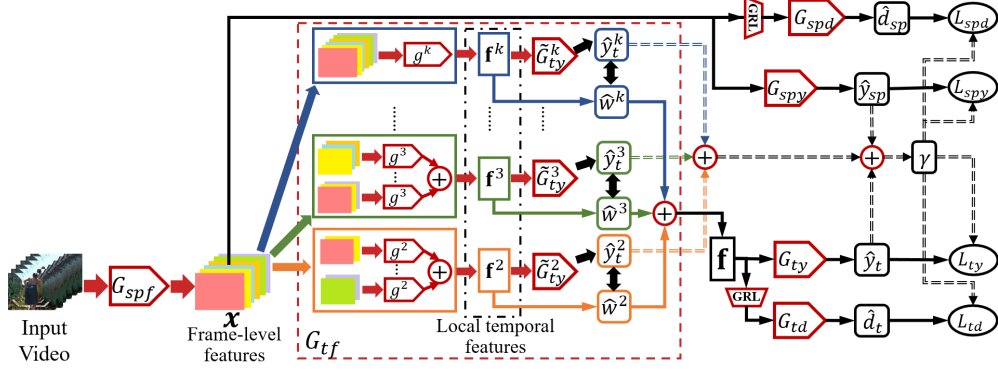


Figure 2. Architecture of the proposed PATAN. To mitigate negative transfer for PVDA effectively, robust overall feature  $\mathbf{f}$  is constructed by weighted combination of local temporal features  $\mathbf{f}^r$ . The local temporal features  $\mathbf{f}^r$  are built by fusing the time ordered frame-level features. The class weights of source domain classes  $\gamma$  averages over the label predictions of the spatial feature, weighted temporal features and the overall temporal feature of target data.  $\gamma$  is applied to both the source domain classifier and spatial/temporal domain discriminators. Dashed arrows indicate how  $\gamma$  is obtained and used in loss functions. *Best viewed in color and zoomed in.*

could therefore be generated as:

$$\hat{w}_i^r = \tanh(1 + \mathbb{C}(\hat{y}_i^r)), \quad (6)$$

where the  $\tanh$  function is applied to ensure that weight  $\hat{w}_i^r$  is constraint within a range of  $[0, 1]$ .

The weight  $\hat{w}_i^r$  computed represents the contribution of the corresponding local temporal feature towards the class filtration process, which ultimately computes the class weights  $\gamma$  for each source domain label  $l \in \mathcal{C}_S$ . The above *label attention* weight is applied to both the generation of temporal attentive class weights utilizing the local temporal features and also the construction of the overall temporal feature. Formally, the overall temporal feature of input video  $V_i$  with  $k$  frames are constructed by:

$$\mathbf{f}_i = G_{tf}(V_i) = \sum_{r=2}^k \hat{w}_i^r \mathbf{f}_i^r, \quad (7)$$

where  $G_{tf}$  denotes the overall temporal feature extractor as shown in Figure 2. Meanwhile, the temporal attentive class weights generated for filtering out outlier source-only classes is formulated as:

$$\gamma = \frac{1}{n_T(k+1)} \sum_{i=1}^{n_T} (\hat{y}_{ti} + \hat{y}_{spi} + \sum_{r=2}^k \hat{w}_i^r \hat{y}_i^r), \quad (8)$$

where the  $\hat{y}_{ti}$  and  $\hat{y}_{spi}$  are the label predictions of the  $i$ th input target video with temporal feature  $\mathbf{f}_i$  and spatial feature  $\mathbf{x}_i$ , computed as  $\hat{y}_{ti} = G_{ty}(\mathbf{f}_i)$  and  $\hat{y}_{spi} = G_{spy}(\mathbf{x}_i)$ .

Finally, PATAN enables partial domain adaptation for videos by down-weighting the contributions of all source data belonging to the outlier label space  $\mathcal{C}_S \setminus \mathcal{C}_T$ . This is achieved by applying the temporal attentive class weight  $\gamma$  to the source label classifier as well as the spatial and temporal domain discriminators over the source domain data. The overall optimization objective of the proposed PATAN is formulated as:

$$\mathcal{L} = \frac{1}{n_S} \sum_{V_i \in \mathcal{D}_S} \gamma_{y_i} L_{ty}(G_{ty}(G_{tf}(G_{spf}(V_i))), y_i)$$

$$\begin{aligned} & + \frac{1}{n_S} \sum_{V_i \in \mathcal{D}_S} \gamma_{y_i} L_{spy}(G_{spy}(G_{spf}(V_i)), y_i) \\ & - \frac{\lambda_{sp}}{n_S} \sum_{V_i \in \mathcal{D}_S} \gamma_{y_i} L_{spd}(G_{spd}(G_{spf}(V_i)), d_i) \\ & - \frac{\lambda_t}{n_S} \sum_{V_i \in \mathcal{D}_S} \gamma_{y_i} L_{td}(G_{td}(G_{tf}(G_{spf}(V_i))), d_i) \\ & - \frac{\lambda_{sp}}{n_T} \sum_{V_i \in \mathcal{D}_T} L_{spd}(G_{spd}(G_{spf}(V_i)), d_i) \\ & - \frac{\lambda_t}{n_T} \sum_{V_i \in \mathcal{D}_T} L_{td}(G_{td}(G_{tf}(G_{spf}(V_i))), d_i), \quad (9) \end{aligned}$$

where  $y_i$  is the ground truth of input  $V_i$  in the source domain, while  $\gamma_{y_i}$  is the corresponding class weight, and  $\lambda_{sp}$  and  $\lambda_t$  are the trade-offs for the domain loss  $L_{spd}$  and  $L_{td}$  with respect to the source classification losses  $L_{spy}$  and  $L_{ty}$ .

## 4. PVDA Benchmarks

There are very limited cross-domain benchmark datasets for VUDA. Current cross-domain VUDA datasets are designed for the standard VUDA tasks, with the source label space constraint to be the same as target label space. To further facilitate PVDA research, we propose three sets of benchmarks, UCF-HMDB<sub>partial</sub>, MiniKinetics-UCF, and HMDB-ARID<sub>partial</sub>, which cover a wide range of PVDA scenarios and provide adequate baselines with distinct domain shift to facilitate PVDA research.

**UCF-HMDB<sub>partial</sub>.** UCF-HMDB<sub>partial</sub> is constructed from two widely used video datasets: UCF101 (**U**) [25] and HMDB51 (**H**) [18]. The overlapping classes between the two datasets are collected, resulting in 14 classes with 2,780 videos. The first 7 categories in alphabetic order of the target domain are chosen as target categories, and we construct two PVDA tasks: **U-14**→**H-7** and **H-14**→**U-7**. We follow the official split for the training and validation sets.

**MiniKinetics-UCF.** MiniKinetics-UCF is built from two large-scale video datasets: MiniKinetics-200 (**M**) [30] and UCF101 (**U**) [25], which contains 45 overlapping classes. Similar to the construction of UCF-HMDB<sub>partial</sub>, the first 18 categories in alphabetic order of the target domain are chosen as target categories, resulting in two PVDA tasks: **M-45**→**U-18** and **U-45**→**M-18**. In this dataset, there are a total of 22,102 videos, nearly 8 times larger than that of UCF-HMDB<sub>partial</sub>. Thus this dataset could validate the effectiveness of PVDA approaches on large-scale dataset.

**HMDB-ARID<sub>partial</sub>.** HMDB-ARID<sub>partial</sub> is constructed with the goal of leveraging current video datasets to boost performance on videos shot in adverse environments. It incorporates both HMDB51 (**H**) [18] and a more recent dark dataset, ARID (**A**) [31], with videos shot under adverse illumination conditions. Statistically, videos in ARID possess much lower RGB mean value and standard deviation (std), which leads larger domain shift between ARID and HMDB51 compared to other cross-domain datasets. The overlapping classes between the two datasets are collected, resulting in 10 classes with 3,252 videos. The first 5 categories in alphabetic order of the target domain are chosen as target categories, resulting in two PVDA tasks: **H-10**→**A-5** and **A-10**→**H-5**. For all the aforementioned benchmarks, the training and validation sets are separated following the official split methods.

## 5. Experiments

In this section, we evaluate our proposed PATAN by performing cross-domain action recognition on PVDA benchmarks introduced in Section 4. We present state-of-the-art results on all proposed benchmarks. We also present ablation studies and empirical analysis of our proposed network to verify our design.

### 5.1. Experimental Settings

We perform action recognition tasks on all three benchmarks: UCF-HMDB<sub>partial</sub>, MiniKinetics-UCF and HMDB-ARID<sub>partial</sub>, with a total of six cross-domain settings as suggested in Section 4. For all six settings, we use all labeled source videos and all unlabeled target videos for PVDA following standard evaluation protocols [23, 20]. We report the top-1 accuracy on the target datasets. Our experiments are implemented using PyTorch [22] library. All methods utilize the Temporal Relation Network (TRN) [35] as the backbone for video feature extraction, with the model pre-trained on ImageNet [10].

### 5.2. Overall Results and Comparisons

We compare the performance of PATAN with competitive and state-of-the-art UDA/VUDA approaches and state-of-the-art PDA approaches. These include: (a) adversarial-based methods: DANN [11], TA<sup>3</sup>N [5], PADA [3], ETN

[4], SAVA [9], BA<sup>3</sup>US [19] and DPDAN [15]; and (b) discrepancy-based methods: MK-MMD [20], MCD [24] and MDD [34]. We also report the results of the backbone feature extractor TRN, which is trained with supervised source data only and tested on the target data. Table 1 shows the performance of our proposed PATAN compared with the above methods in all six PVDA settings.

The results in Table 1 show that our proposed PATAN achieves the best results on all six settings, and substantially outperforms previous approaches by noticeable margins. It can be observed that for all UDA/VUDA approaches, i.e. DANN, TA<sup>3</sup>N, MK-MMD, MCD, and MDD, there exists at least three settings where their performance are inferior to that of TRN trained without any domain adaptation methods. This suggests that these methods suffer from the negative transfer issue of PVDA.

Compare to previous PDA approaches PADA and ETN, our proposed PATAN exceeds both approaches consistently, with an average 11.27% relative improvement towards PADA, and an average 11.57% relative improvement towards ETN. These large improvements imply the effectiveness of building temporal attentive features and incorporating local and overall temporal features for class filtration. In particular, the improvement of PATAN with respect to PADA and ETN is most significant for HMDB-ARID<sub>partial</sub>, with a relatively average improvement of 17.12% and 21.85% towards PADA and ETN respectively. HMDB-ARID<sub>partial</sub> possesses the largest domain shift across the source and target domains, with the lowest source-only accuracies. This suggests that the class filtration process may not be accurate by utilizing spatial or temporal features alone, which explains the relatively small improvement in performance of PADA and ETN compared to UDA/VUDA approaches. On the contrary, by constructing temporal features by *label attention*, the effectiveness of the class filtration process is improved with features of higher certainty attended, thus large improvements are brought by PATAN.

### 5.3. Ablation Studies

To go deeper with the efficacy of the proposed PATAN network, we perform ablation studies by evaluating PATAN against its variants: (a) **PATAN w/o attentive** is the variant where the *label attention* weight is not computed, therefore the overall temporal feature and class weights of source classes are the result of the aggregation of all local temporal features and the label predictions of all local and overall temporal features with that of the spatial feature; (b) **PATAN w/o local weights** is the variant where the class weights  $\gamma$  do not incorporate all the local weights; (c) **PATAN w/o classifier** is the variant without the class weights  $\gamma$  applied on the source spatial and temporal classifiers, and (d) **PATAN w/o adversarial** is the variant without the class weights applied on the spatial and temporal

Methods		UCF-HMDB <sub>partial</sub>		MiniKinetics-UCF		HMDB-ARID <sub>partial</sub>	
		U-14→H-7	H-14→U-7	M-45→U-18	U-45→M-18	H-10→A-5	A-10→H-5
Source-only	TRN [35]	62.85%	78.95%	78.77%	54.14%	14.10%	26.00%
Adversarial-based	DANN [11]	60.95%	74.44%	79.21%	52.25%	20.77%	12.00%
	TA <sup>3</sup> N [5]	50.49%	70.68%	75.70%	48.23%	18.30%	24.00%
	PADA [3]	65.71%	82.33%	82.43%	61.23%	21.79%	30.67%
	ETN [4]	67.88%	82.89%	83.33%	62.51%	21.40%	28.82%
	SAVA [9]	56.67%	75.19%	78.55%	54.61%	20.77%	26.67%
	BA <sup>3</sup> US [19]	71.90%	85.34%	84.48%	64.07%	23.59%	31.33%
	DPDAN [15]	68.10%	81.58%	82.72%	61.93%	22.31%	28.67%
Discrepancy-based	MK-MMD [20]	58.57%	82.71%	79.79%	55.79%	21.28%	14.00%
	MCD [24]	55.71%	73.31%	75.13%	52.48%	12.56%	14.67%
	MDD [34]	62.58%	80.45%	80.12%	50.35%	15.13%	9.33%
Ours	<b>PATAN</b>	<b>73.81%</b>	<b>89.85%</b>	<b>86.82%</b>	<b>65.25%</b>	<b>26.41%</b>	<b>34.67%</b>

Table 1. Results for Partial Video Domain Adaptation on UCF-HMDB<sub>partial</sub>, MiniKinetics-UCF and HMDB-ARID<sub>partial</sub>.

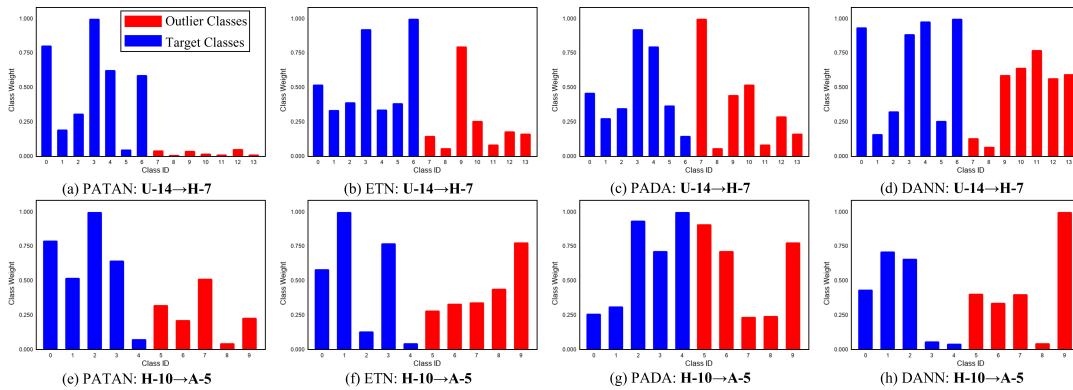


Figure 3. Histograms of class weights learned by PATAN, ETN, PADA and DANN on settings **U-14→H-7** and **H-10→A-5**.

Methods	U-14→H-7	H-14→U-7
<b>PATAN</b>	<b>73.81%</b>	<b>89.85%</b>
PATAN w/o attentive	71.43%	85.34%
PATAN w/o local weights	70.47%	84.21%
PATAN w/o classifier	69.52%	82.71%
PATAN w/o adversarial	67.14%	81.58%

Table 2. Ablation studies of PATAN on UCF-HMDB<sub>partial</sub>.

domain discriminators. The results of the variants are presented in Table 2.

Specifically, PATAN outperforms PATAN w/o attentive by a noticeable margin proves the necessity of combining local temporal features and class weights with *label attention*, which constructs more discriminable overall temporal features. Similarly, PATAN’s superior performance over PATAN w/o local weights proves that the class filtration process could be improved by utilizing label prediction of local temporal features. We note that the results of both PATAN w/o attentive and PATAN w/o local weights outperform that of PADA and ETN. This further justifies the effectiveness of both utilizing local temporal features for class filtration and constructing attentive overall temporal features with *label attention*. Further, PATAN outperforms both PATAN w/o classifier and PATAN w/o adversarial by huge margins. This strongly suggests that the class weights applied can assign small weights on outlier classes

and down-weight the source data of the outlier classes effectively, the class weights applied thus mitigates negative transfer and boost the performance for PVDA.

## 5.4. Empirical Analysis

To further understand our proposed PATAN, we perform empirical analysis focusing on four areas of interest: class weights visualization, effect of number of target classes, and feature visualization.

**Class weights visualization.** We first illustrate and compare the learned class weights  $\gamma$  generated by methods PATAN, ETN, PADA and DANN for settings **U-14→H-7** and **H-10→A-5** in Figure 3. It could be observed that our proposed PATAN assigns much smaller weights to the outlier source only classes than to the shared target classes, which shows that PATAN could effectively filter out the outlier classes. It is noted that the difference of class weights between target and outlier classes is less significant for **H-10→A-5**, given the much larger cross-domain shift for dataset HMDB-ARID<sub>partial</sub>. Despite the difficulty brought by the large domain shift, our proposed PATAN still assigns significantly larger weights for target classes compared to other methods. The much larger weights assigned to target classes show that our network could effectively filter out

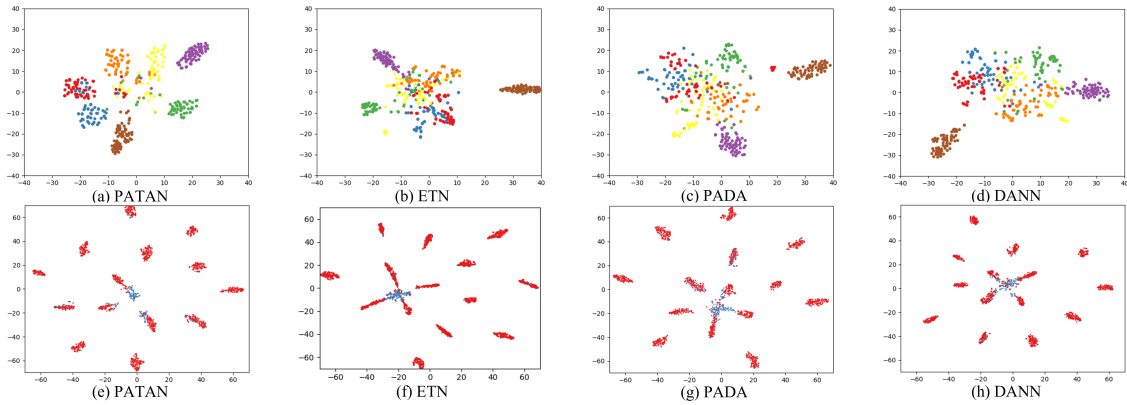


Figure 4. Visualization of features learned by PATAN, ETN, PADA, and DANN, with class ((a)-(d)) and domain information ((e)-(h)). Different classes are denoted by different colors. The red and blue dots represent data from the source and target domains respectively.




Figure 5. Accuracy with different number of target classes.

outlier classes, therefore explains the strong performance of PATAN on both datasets. Though both PADA and ETN incorporate class filtration processes, the effectiveness of such a process is hindered by failing to incorporate temporal features. This results in poorer performance. However, it is noted that there are target classes not weighted correctly by PATAN (class 5 in Figure 3a and class 4 in Figure 3e). We observe that the predicted probability of the ground-truth classes are low, while the probabilities are spread evenly (i.e. of low certainties) for target data in both classes.

**Effect of number of target classes.** We investigate a wider spectrum of PVDA by varying the number of target classes, conducted with the UCF-HMDB<sub>partial</sub> dataset. The result of the accuracy of the target dataset against the different numbers of target classes is shown in Figure 5. It is observed that the performance of DANN, PADA, and ETN degrades noticeably with fewer target classes. This is a clear indication of negative transfer brought by the increasing outlier classes. Comparatively, the performance of PATAN is more stable and is consistently better than all compared methods. The stability of performance suggests that PATAN effectively alleviates the influence of outlier classes. It could also be observed that when the number of target classes is equivalent to that of source classes (in this

case 14), the PVDA task is turned into a standard VUDA task. Under this condition, our PATAN also performs better than DANN. This shows that the class filtration process will not degrade performance when there are no outlier classes.

**Feature visualization.** We further plot the t-SNE embeddings [28] of the features learned by PATAN, ETN, PADA, and DANN for the U-14→H-7 with class information in the target domain as shown in Figure 4 (a)-(d), and with domain information as shown in Figure 4 (e)-(h). From Figure 4 (a), it is observed that the features learned by PATAN are more clustered. This proves that features extracted by PATAN with *label attention* have higher discriminability. Meanwhile, Figure 4 (f)-(h) shows that other methods align target data to all source classes, which includes outlier ones, triggering negative transfer. We note that though ETN and PADA include class filtration processes, the negative transfer is still triggered due to misalignment of temporal features not utilized in their class filtration processes. Comparatively, PATAN only aligns target data to the shared classes (7 classes), alleviating the effects of the outlier classes.

## 6. Conclusion

In this work, we propose a novel approach for partial video domain adaptation (PVDA). Unlike previous approaches where only spatial features are utilized for mitigating negative transfer in partial domain adaptation, the new PATAN tackles PVDA with full utilization of both spatial and temporal features, filtering out outlier source-only classes effectively. The proposed PATAN also attends to local temporal features that contribute more towards the class filtration process. We further introduce novel PVDA benchmarks to facilitate PVDA research, which are the first PVDA benchmarks introduced. Our proposed PATAN addresses the PVDA problem well, justified by extensive experiments across the proposed PVDA benchmarks.



## References

- [1] Qi Cai, Yingwei Pan, Chong-Wah Ngo, Xinmei Tian, Lingyu Duan, and Ting Yao. Exploring object relation in mean teacher for cross-domain detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11457–11466, 2019. 2
- [2] Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Michael I Jordan. Partial transfer learning with selective adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2724–2732, 2018. 2
- [3] Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 135–150, 2018. 2, 4, 6, 7
- [4] Zhangjie Cao, Kaichao You, Mingsheng Long, Jianmin Wang, and Qiang Yang. Learning to transfer examples for partial domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2985–2994, 2019. 3, 6, 7
- [5] Min-Hung Chen, Zsolt Kira, Ghassan AlRegib, Jaekwon Yoo, Ruxin Chen, and Jian Zheng. Temporal attentive alignment for large-scale video domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6321–6330, 2019. 2, 6, 7
- [6] Min-Hung Chen, Baopu Li, Yingze Bao, Ghassan Al-Regib, and Zsolt Kira. Action segmentation with joint self-supervised temporal domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9454–9463, 2020. 2
- [7] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3339–3348, 2018. 2
- [8] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. Crdoco: Pixel-level domain transfer with cross-domain consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1791–1800, 2019. 2
- [9] Jinwoo Choi, Gaurav Sharma, Samuel Schulter, and Jia-Bin Huang. Shuffle and attend: Video domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 678–695. Springer, 2020. 2, 6, 7
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [11] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 2, 6, 7
- [12] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. 3
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 2
- [14] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. PMLR, 2018. 2
- [15] Jian Hu, Hongya Tuo, Chao Wang, Lingfeng Qiao, Haowen Zhong, Junchi Yan, Zhongliang Jing, and Henry Leung. Discriminative partial domain adversarial network. In *Proceedings of the European conference on computer vision (ECCV)*, pages 632–648, 2020. 6, 7
- [16] Ling Huang, Anthony D Joseph, Blaine Nelson, Benjamin IP Rubinstein, and J Doug Tygar. Adversarial machine learning. In *Proceedings of the 4th ACM workshop on Security and artificial intelligence*, pages 43–58, 2011. 2
- [17] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 2
- [18] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *2011 International Conference on Computer Vision*, pages 2556–2563. IEEE, 2011. 5, 6
- [19] Jian Liang, Yunbo Wang, Dapeng Hu, Ran He, and Jiashi Feng. A balanced and uncertainty-aware approach for partial domain adaptation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 123–140, 2020. 6, 7
- [20] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International conference on machine learning*, pages 97–105. PMLR, 2015. 6, 7
- [21] Boxiao Pan, Zhangjie Cao, Ehsan Adeli, and Juan Carlos Niebles. Adversarial cross-domain action recognition with co-attention. In *AAAI*, pages 11815–11822, 2020. 2
- [22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019. 6
- [23] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 213–226. Springer, 2010. 6
- [24] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3723–3732, 2018. 6, 7

- [25] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 5, 6
- [26] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4068–4076, 2015. 2
- [27] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7167–7176, 2017. 2
- [28] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 8
- [29] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Dada: Depth-aware domain adaptation in semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7364–7373, 2019. 2
- [30] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning for video understanding. *arXiv preprint arXiv:1712.04851*, 1(2):5, 2017. 6
- [31] Yuecong Xu, Jianfei Yang, Haozhi Cao, Kezhi Mao, Jianxiong Yin, and Simon See. Arid: A new dataset for recognizing action in the dark. *arXiv preprint arXiv:2006.03876*, 2020. 2, 6
- [32] Jing Zhang, Zewei Ding, Wanqing Li, and Philip Ogunbona. Importance weighted adversarial nets for partial domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8156–8164, 2018. 3
- [33] Jing Zhang, Wanqing Li, and Philip Ogunbona. Joint geometrical and statistical alignment for visual domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1859–1867, 2017. 2
- [34] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *International Conference on Machine Learning*, pages 7404–7413. PMLR, 2019. 6, 7
- [35] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 803–818, 2018. 4, 6, 7
- [36] Xinge Zhu, Jiangmiao Pang, Ceyuan Yang, Jianping Shi, and Dahua Lin. Adapting object detectors via selective cross-domain alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 687–696, 2019. 2
- [37] Han Zou, Yuxun Zhou, Jianfei Yang, Huihan Liu, Hari Prasanna Das, and Costas J Spanos. Consensus adversarial domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5997–6004, 2019. 2
- [38] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289–305, 2018. 2