

TAMT: Temporal-Aware Model Tuning for Cross-Domain Few-Shot Action Recognition

Yilong Wang^{1,*} Zilin Gao^{2,*} Qilong Wang^{1,†} Zhaofeng Chen³ Peihua Li² Qinghua Hu¹
¹Tianjin University ²Dalian University of Technology ³Yancheng Institute of Technology

Abstract

Going beyond few-shot action recognition (FSAR), cross-domain FSAR (CDF SAR) has attracted recent research interests by solving the domain gap lying in source-to-target transfer learning. Existing CDF SAR methods mainly focus on joint training of source and target data to mitigate the side effect of domain gap. However, such kind of methods suffer from two limitations: First, pair-wise joint training requires retraining deep models in case of one source data and multiple target ones, which incurs heavy computation cost, especially for large source and small target data. Second, pre-trained models after joint training are adopted to target domain in a straightforward manner, hardly taking full potential of pre-trained models and then limiting recognition performance. To overcome above limitations, this paper proposes a simple yet effective baseline, namely Temporal-Aware Model Tuning (TAMT) for CDF SAR. Specifically, our TAMT involves a decoupled paradigm by performing pre-training on source data and fine-tuning target data, which avoids retraining for multiple target data with single source. To effectively and efficiently explore the potential of pre-trained models in transferring to target domain, our TAMT proposes a Hierarchical Temporal Tuning Network (HTTN), whose core involves local temporal-aware adapters (TAA) and a global temporal-aware moment tuning (GTMT). Particularly, TAA learns few parameters to recalibrate the intermediate features of frozen pre-trained models, enabling efficient adaptation to target domains. Furthermore, GTMT helps to generate powerful video representations, improving match performance on the target domain. Experiments on several widely used video benchmarks show our TAMT outperforms the recently proposed counterparts by 13%~31%, achieving new state-of-the-art CDF SAR results.

1. Introduction

Few-shot action recognition (FSAR) aims to develop video recognition models with high generalization ability by using

limited annotated samples, which has achieved remarkable progress with the rapid development of deep models and pre-training techniques [6, 17, 27, 45–48, 50, 52, 54, 55]. Going beyond FSAR, cross-domain FSAR (CDF SAR) has been attracting recent research interests [38, 49], which focuses on transferring knowledge from the well-annotated source domain to target one with few annotated samples. Intuitively, the domain gap between source and target data will clearly impact the performance of transfer learning [1, 15, 29].

As a seminal work, SEEN [49] proposes a joint training paradigm to alleviate side effect of domain gap, where a parameter-shared model is trained on source data and target one with supervision learning and self-supervised contrastive learning objectives, respectively. After the joint training, a simple nearest neighbor classifier is straightforwardly used for inference in target domain. As a parallel work, CDFSL-V [38] proposes a two-stage joint training paradigm, where the model is first pre-trained on source and target data in a self-supervised manner, and then a curriculum learning is developed to further tune the model on source and target data. Subsequently, a few-shot classifier is fine-tuned on the annotated target data for inference.

Although some advanced efforts are made [38, 49], they generally suffer from two limitations. First, both SEEN [49] and CDFSL-V [38] involve joint training paradigms. As illustrated in Fig. 1a (1) and (2), they require to retrain the models B times, given a single source data \mathcal{S}_{CD} and B target data $\{\mathcal{T}_{CD}^1, \dots, \mathcal{T}_{CD}^B\}$ (a commonly used setting [38, 49]). It potentially incurs heavy computation cost due to frequent retraining on source data \mathcal{S}_{CD} , especially for large \mathcal{S}_{CD} and small \mathcal{T}_{CD}^b . Second, during the inference stage, pre-trained models are generally adopted to target domain in a straightforward manner, i.e., simple nearest neighbor classifier [49] or a fine-tuned classifier [38]. They hardly take full advantage of pre-trained models to dynamically fit target data with the frozen backbone, and so potentially limit the final recognition performance.

To address the above limitations, this paper proposes a simple yet effective baseline, namely Temporal-Aware Model Tuning (TAMT). Particularly, as shown in Fig. 1a (3), our TAMT involves a decoupled paradigm by pre-training

*Equal contribution. †Corresponding author: qlwang@tju.edu.cn.

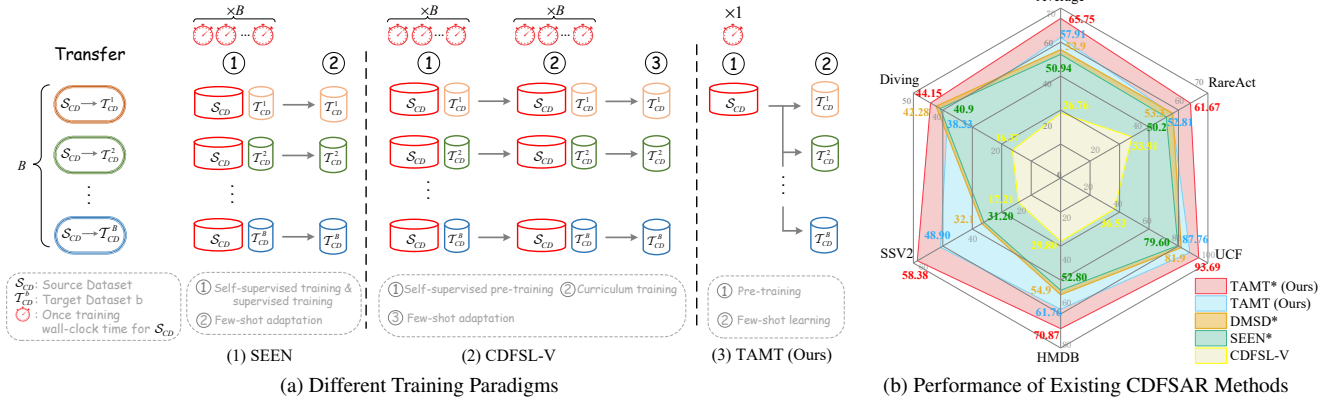


Figure 1. (a) Comparison of existing CDFSAR methods in terms of training paradigm under the case of a single source data S_{CD} and B target data $\{\mathcal{T}_{CD}^1, \dots, \mathcal{T}_{CD}^B\}$. (b) Comparison (%) of existing CDFSAR methods with K-100 as source data. All results are conducted with 112×112 resolution except methods marked by * (224×224 resolution)

the model on source data and subsequently fine-tuning it on target data. For model pre-training, we introduce a self-supervised followed by a supervised learning scheme to consider abilities of both generalization and semantic features extraction. As such, in the case of one source data and multiple target data, our TAMT only requires model pre-training one time, significantly decreasing training cost.

To explore the potential of pre-trained models in domain adaptation, our TAMT proposes a Hierarchical Temporal Tuning Network (HTTN), whose core involves local Temporal-Aware Adapters (TAA) and a Global Temporal-aware Moment Tuning (GTMT). Particularly, TAA introduces few learnable parameters to recalibrate a part of intermediate features outputted by frozen pre-training models, which helps adapt pre-training models to target data efficiently. By considering the significance of global representations in metric-based few-shot classification, our GTMT proposes to exploit spatio-temporal feature distribution approximated first- and second-order moments to generate powerful video representations. Particularly, GTMT presents an efficient long-short temporal covariance (ELSTC) to effectively compute second-order moments of spatio-temporal features. By equipping with TAA and GTMT, our HTTN dynamically adopts pre-trained models to target data in an effective and efficient way, clearly improving recognition performance. As shown in Fig. 1b, our proposed TAMT can bring significant performance gains over existing methods with lower training cost. To evaluate our TAMT, experiments are conducted on five source datasets (i.e., Kinetics-400 (K-400) [7], Kinetics-100 (K-100) [59], Something-Something V2 (SSV2) [16], Diving48 (Diving) [24] and UCF-101 (UCF) [41]) and five target datasets (i.e., HMDB51 (HMDB) [23], SSV2, Diving, UCF-101 (UCF) [41] and RareAct [30]). The contributions of this work can be summarized as follows:

1) In this paper, we propose a simple yet effective baseline for

the cross-domain few-shot action recognition (CDFSAR) task, namely Temporal-Aware Model Tuning (TAMT). To our best knowledge, TAMT makes the first attempt to introduce a decoupled training paradigm for CDFSAR, effectively avoiding frequent retraining in the case of one source data and multiple target data.

- 2) Unlike previous CDFSAR works, our TAMT pays more attention to effectively and efficiently adopting pre-trained models to target data. Particularly, a lightweight Hierarchical Temporal Tuning Network (HTTN) is proposed to recalibrate intermediate features and generate powerful video representations for the frozen pre-training models via local Temporal-Aware Adapters (TAA) and a Global Temporal-aware Moment Tuning (GTMT), respectively.
- 3) Extensive experiments are conducted on various video benchmarks, and the results show our TAMT significantly outperforms the recently proposed CDFSAR methods.

2. Related Work

2.1. Few-Shot Action Recognition

With the development of large video models and the insurmountable success of action recognition methods [2, 9, 26, 35, 44, 51, 53, 56], few-shot action recognition methods are emerging and thriving. Existing few-shot action recognition methods mainly use pre-trained backbone models on image benchmarks (e.g., ImageNet-1k [10] and CLIP [37]), which focus on the frame-level alignment between query and support videos in few-shot learning (FSL). Some early researches [3, 5, 6] estimate temporal alignment for frame-level features to match the query videos and support set. TRX [33] leverages an attentional mechanism to match each query video with all videos in the support set. HyRSM [45] introduces a hybrid relation module and designs a Bi-MHM for flexible matching. STRM [42] proposes

a spatio-temporal enrichment module to analyze spatial and temporal contexts. MASTAF [27] uses self-attention and cross-attention modules to increase the inter-class variations while decrease the intra-class variations. MoLo [48] learns long-range temporal context and motion cues for comprehensive few-shot matching. CLIP-FSAR [52] devises a video-text contrastive objective and proposes a prototype modulation to fully utilize the rich semantic priors in CLIP. Different from the aforementioned works, our method aims to perform an effective yet efficient temporal-aware model tuning on the pre-trained frozen backbones to realize CDFSAR tasks.

2.2. Cross-Domain Few-Shot Action Recognition

Cross-domain few-shot learning requires base and test data from different domains. BS-CDFSL [19] first introduces an image benchmark for cross-domain few-shot learning, and early studies handle cross-domain action recognition by mainly focusing on deep feature learning and cross-domain learning [12]. Meanwhile, as supplements, previous works [13, 14] also introduce some source-target data pairs to evaluate the performance of CDFSAR. Recently, SEEN [49] proposes to integrate supervised learning with an auxiliary self-supervised contrastive learning to tackle the issue of domain gap lying in CDFSAR task. For above works [13, 14, 49], there exist some shared classes lying in the constructed source-target data pairs, which however, is not expected in CDFSL task. CDFSL-V [38] proposes a new benchmark to solve this problem by removing all overlapping classes between the source and target datasets. DMSD [18] designs two branches called the original-source and the mixed-source branches for meta-training based on the pipeline of CDFSL-V. But different from pair-wise joint training studied in previous CDFSAR methods [18, 38, 49], our proposed TAMT develops a decoupled paradigm to avoid frequent retraining in case of one source data and multiple target data, while proposing an HTTN method to effectively and efficiently adapt pre-training models for the target domain. Experimental comparisons (Sec. 4.2) show our TAMT significantly outperforms existing counterparts.

3. Method

In this section, we first provide a brief definition of CDFSAR task. Then, we show the overview of our decoupled TAMT paradigm, which pre-trains models on source data while fine-tuning the pre-trained models on target one. Finally, we detailedly introduce the proposed hierarchical temporal tuning network (HTTN) for model tuning on target domain, which consists of local Temporal-Aware Adapters (TAA) and Global Temporal-aware Moment Tuning (GTMT).

3.1. Problem Formulation

CDFSAR task aims to develop an FSAR model for mitigating the side effect brought by domain gap between the source

dataset \mathcal{S}_{CD} and the target dataset \mathcal{T}_{CD} . In the context of cross-domain, the model could be trained on well-annotated \mathcal{S}_{CD} and \mathcal{T}_{CD} with few annotated samples. After that, the transferring performance of the proposed method is evaluated on target domain \mathcal{T}_{CD} . In the target-domain \mathcal{T}_{CD} , the pre-trained model is evaluated on its novel (test) set \mathcal{N} under FSL protocol, by providing training samples from its base (training) set \mathcal{B} , w.r.t., $\mathcal{T}_{CD} = \mathcal{B} \cup \mathcal{N}$. For one FSL inference unit (dubbed as task or episode), it consists of unknown query videos $\{Q_1, \dots, Q_U\}$, and an annotated support set Σ . For N -way K -shot setting, each episode involves N categories and each category has K samples in Σ . The final goal of CDFSAR is to accurately classify each query video Q_i by leveraging the limited data available in the support set Σ . Particularly, to assess the transferring performance in a convincing way, the classes are non-overlapping in \mathcal{S}_{CD} and \mathcal{T}_{CD} , i.e., $\mathcal{S}_{CD} \cap \mathcal{T}_{CD} = \emptyset$, and $\mathcal{N} \cap \mathcal{B} = \emptyset$ for FSL.

3.2. Overview of Temporal-aware Model Tuning

Compared to FSAR, CDFSAR is further challenged by domain gap lying in source-to-target transfer learning. Previous works [38, 49] develop some joint training paradigms on source and target data to mitigate side effect of domain gap. As shown in Fig. 1a, joint training paradigm generally suffers from model retraining in case of one source and multiple target data. Besides, they take no full advantage of pre-trained models by using some straightforward few-shot adaptation methods, potentially limiting recognition performance. To solve above issues, we propose a decoupled training paradigm, namely TAMT. As shown in Fig. 2 (a), our TAMT can be summarized as two phases: pre-training on \mathcal{S}_{CD} and fine-tuning on \mathcal{T}_{CD} . Specifically, the model is first pre-trained on \mathcal{S}_{CD} to learn the knowledge from source domain. Subsequently, it is fine-tuned on \mathcal{T}_{CD} to perform transfer learning on target domain. The details are as follows.

Pre-training on Source Data. In this work, our TAMT pre-trains the models only on source data. To consider both generalization and representation abilities of pre-trained model, our TAMT develops a two-stage pre-training strategy. Inspired by success of self-supervised learning (SSL) on video pre-trained models [43], we first introduce the reconstruction-based SSL solution to train our models for capturing general spatio-temporal structures lying videos, helping our pre-trained models can be well generalized to various downstream tasks. However, such SSL solution usually focuses on the fundamental features [20], while neglecting high-level semantic information [31, 36, 44, 58], and limiting representation or discriminative abilities of pre-trained models. Existing works [31, 36, 58] make attempts to combine reconstruction-based SSL with self-supervised contrastive learning to improve discriminative ability of pre-trained models. By considering the samples are well annotated on source data, we simply incorporate a supervised

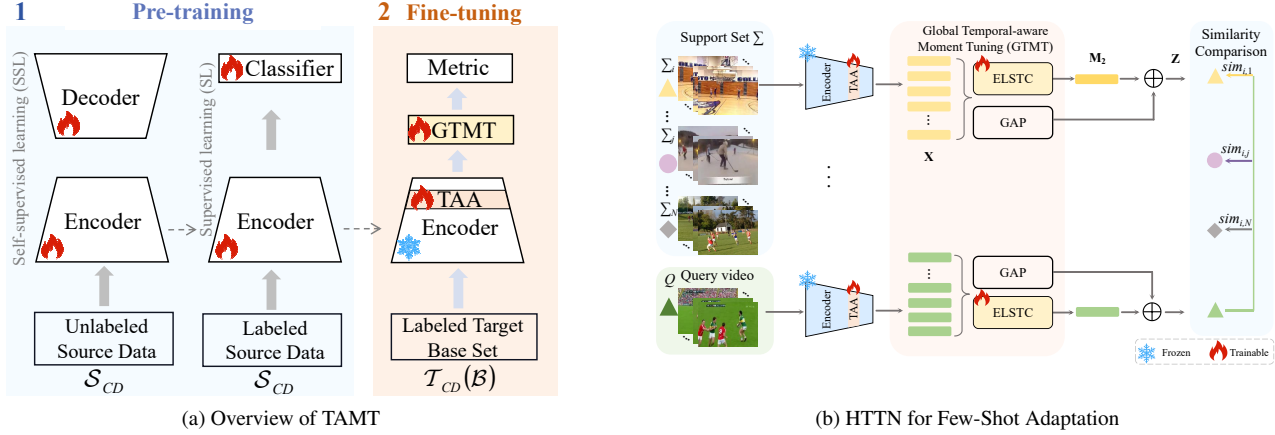


Figure 2. (a) Overview of our TAMT paradigm, which pre-trains the models on \mathcal{S}_{CD} and fine-tunes them on \mathcal{T}_{CD} . Specifically, for pre-training stage, the model is first optimized with a reconstruction-based SSL solution, while the encoder \mathcal{E} is post-trained with the SL objective. Subsequently, the pre-trained \mathcal{E} is fine-tuned for few-shot adaptation on \mathcal{T}_{CD} by using our HTTN. (b) HTTN for few-shot adaptation, where a metric-based is used for few-shot adaptation. Particularly, our HTTN consists of local Temporal-Aware Adapters (TAA) and Global Temporal-aware Moment Tuning (GMT).

learning (SL) after SSL to enhance the representation ability of pre-trained models. To be specific, the encoder \mathcal{E} of the model is first trained with reconstruction-based SSL, and then it is optimized with recognition objectives on annotated \mathcal{S}_{CD} . As such, our two-stage pre-training strategy potentially achieves generalization and representation trade-off, where both SSL and SL play key roles in the final performance. More analysis can refer to Sec. 4.3.

Fine-tuning on Target Data. By considering the issue of domain shift between source data and target one, we propose a hierarchical temporal tuning network (HTTN), aiming to effectively and efficiently adopt the pre-trained model \mathcal{E} to target domain. In particular, we construct our HTTN by using a metric-based FSL pipeline [40]. To fully explore the potential of the frozen pre-trained models in transferring to target domain, we present local temporal-aware adapters (TAA) and a global temporal-aware moment tuning (GMT) to recalibrate the intermediate features and generate powerful video representations according to few annotated samples on target domain, respectively. The details of our HTTN will be described in the following subsection.

3.3. Hierarchical Temporal Tuning Network

To perform few-shot adaptation of pre-trained models on target domain \mathcal{T}_{CD} , we propose a Hierarchical Temporal Tuning Network (HTTN). As depicted in Fig. 2 (b), our HTTN integrates several *local* Temporal-Aware Adapters (TAA) into last- L transformer blocks of pre-trained model \mathcal{E} , and insert a Global Temporal-aware Moment Tuning (GMT) module with efficient long-short temporal covariance (ELSTC) at the end of \mathcal{E} . Given an input video, the features are extracted by the frozen \mathcal{E} , which are recalibrated by our TAA modules and subsequently fed into GMT to generate final

representation. Ultimately, the representations derived from query and support videos are compared using a similarity metric, which serves as logits for training and inference.

Local Temporal-Aware Adapter (TAA). In the decoupled training protocol of TAMT, it is important to utilize target domain for tuning source pre-trained model. However, the conventional full fine-tuning (FFT) strategy will optimize all of the parameters, bringing high training consumption and potentially posing the risk of overfitting, particularly in the few-shot learning scenario. As suggested in previous works for NLP and image classification tasks [21, 25], we introduce a parameter-efficient approach for recalibrating video features in a temporal-aware manner.

Given a certain intermediate-layer feature $\mathbf{F} \in \mathbb{R}^{T \times M \times C}$, where T, M, C denotes temporal length, token number in one frame and channel number, respectively. TAA introduces a few learnable scale and bias parameters for features of each frame, which can be written as follows:

$$\mathbf{F}' = \gamma \odot \mathbf{F} \oplus \beta, \quad (1)$$

where \odot and \oplus represent the multiplication and addition operations, respectively. For the sake of convenience, here we omit the expansion operation along M dimension for scale $\gamma \in \mathbb{R}^{T \times C}$ and bias $\beta \in \mathbb{R}^{T \times C}$. Particularly, γ and β indicate the temporal cues of \mathbf{F} decided by transformation functions \mathcal{W} and \mathcal{G} respectively. Therefore, we have

$$\gamma = \mathcal{W}(\hat{\mathbf{F}}) = \mathbf{W}_{\uparrow}^{\gamma} \otimes g_1(\mathbf{W}_{\downarrow}^{\gamma} \otimes \hat{\mathbf{F}}), \quad (2)$$

$$\beta = \mathcal{G}(\hat{\mathbf{F}}) = \mathbf{W}_{\uparrow}^{\beta} \otimes g_2(\mathbf{W}_{\downarrow}^{\beta} \otimes \hat{\mathbf{F}}), \quad (3)$$

where $\hat{\mathbf{F}} \in \mathbb{R}^{T \times C}$ presents the output of global average pooling on \mathbf{F} . By taking Eqn. (2) as an example, func-

with concurrently considering interaction across all frames (termed as TCov), typically needs $(TC)^2$ -dimensional computation consumption. Thereby, we propose an Efficient Long-Short Temporal Covariance layer (ELSTC),

$$\mathbf{M}_2 = \text{ELSTC}(\mathbf{X}) = \mathcal{A}(\text{LSTC}(\mathbf{X}^{g_1}); \dots; \text{LSTC}(\mathbf{X}^{g_G})). \quad (6)$$

For high efficiency, the sequence feature \mathbf{X} is split into G groups along temporal dimension and can be rewritten as: $\mathbf{X} = [\mathbf{X}^{g_1}; \dots; \mathbf{X}^{g_G}]$, where notation $[\cdot]$ indicates concatenation. For group g_e , $\mathbf{X}^{g_e} \in \mathbb{R}^{T' \times M \times C}$ is with temporal length $T' = \frac{T}{G}$, leading to reduce the computation consumption G times. In each group g_e , aiming for temporal-aware modeling, we devise LSTC to compute long-short temporal covariance as follows:

$$\text{LSTC}(\mathbf{X}^{g_e}) = \left\{ \mathbf{R}_{t,t'}^{g_e} \right\}_{t,t'}^{T'^2}; \quad (7)$$

$$\mathbf{R}_{t,t'}^{g_e} = \frac{1}{M} \sum_{m=1}^M \tilde{\mathbf{X}}_{t,m}^{g_e \top} \tilde{\mathbf{X}}_{t',m}^{g_e}; \tilde{\mathbf{X}} = \mathcal{K}(\mathbf{X}). \quad (8)$$

In group g_e , the covariance matrix $\mathbf{R}_{t,t'}^{g_e} \in \mathbb{R}^{\frac{C}{\tau} \times \frac{C}{\tau}}$ captures temporal correlation between t -th and t' -th frame of $\tilde{\mathbf{X}} \in \mathbb{R}^{T' \times M \times \frac{C}{\tau}}$. The feature $\tilde{\mathbf{X}}$ is a transformation of \mathbf{X} with an MLP layer \mathcal{K} , bringing the dimension reduction by a hyper parameter $\tau: C \rightarrow \frac{C}{\tau}$. \top represents transposition operation.

In particular, for one covariance matrix $\mathbf{R}_{t,t'}^{g_e}$ in group g_e , the timestamps t and t' always have a temporal gap Δ , ranging from 0 to $(T' - G)$ with an interval of G . For $\Delta = 0$, $\mathbf{R}_{t,t}^{g_e}$ indicates the static appearance information of $\tilde{\mathbf{X}}_t^{g_e}$, and for other $\Delta \neq 0$, $\mathbf{R}_{t,t'}^{g_e}$ outputs the temporal cross-covariance of $\tilde{\mathbf{X}}_t^{g_e}$ and $\tilde{\mathbf{X}}_{t+\Delta}^{g_e}$. As a result, the output of $\text{LSTC}(\mathbf{X}^{g_e})$ describes the various temporal correlations from *short-term* (one frame) static appearance to *long-term* crossing $(T' - G)$ frames motion information.

Furthermore, the outputs derived from G groups are ultimately summarized with \mathcal{A} , generating a holistic video representation \mathbf{M}_2 . To the sake of clarity, the output of LSTC for group g_e is rewritten as $\mathbf{Y}^{g_e} \in \mathbb{R}^{T' \times \frac{C}{\tau} \times T' \times \frac{C}{\tau} \times 1}$ by concatenating all $\{\mathbf{R}_{*,*}^{g_e}\}$ in group g_e . And then, \mathbf{M}_2 is:

$$\mathbf{M}_2 = \mathcal{A}(\mathbf{Y}^{g_1}, \dots, \mathbf{Y}^{g_G}), \quad (9)$$

where the indication \mathcal{A} denotes two convolutional layers with $k_c \times k_c$ kernel, with each followed by a BN layer and ReLU activation function. By setting the proper stride and output channel, the dimension of G -group output is changed from $T' \times \frac{C}{\tau} \times T' \times \frac{C}{\tau} \times G$ to $C_M \times C_M \times 1$, and vectorized to $C_M^2 \times 1$ ultimately. By omitting the constant zero-order in Eqn. (4), the final representation of our HTTN is expressed by combining the first and second-order moment as follows:

$$\mathbf{Z} = \mathcal{H}(\mathbf{M}_2) \oplus \mathbf{M}_1, \quad (10)$$

where a linear projection \mathcal{H} is used to align the dimension of \mathbf{M}_2 with that of \mathbf{M}_1 , i.e., $C_M^2 \rightarrow C$. In this work, $k_c = 3$, $\tau = 6$, $C_M = 64$ for all experiments.

4. Experiments

We extensively compare our TAMT with state-of-the-arts on both CDFSAR and FSAR tasks (see supplementary material), and conduct the ablation study on CDFSAR task.

4.1. Experimental Settings

Datasets. For CDFSAR task, we use the K-400 [7], K-100 [59], SSV2 [16], Diving [24] and UCF [41] as the source domains \mathcal{S}_{CD} , which transfer to following five target domains \mathcal{T}_{CD} : HMDB [23], SSV2, Diving, UCF and RareAct [30]. For the source datasets, we follow the non-overlapping setting protocol [38] between \mathcal{S}_{CD} and \mathcal{T}_{CD} in cross-domain scenario. Specifically, source datasets K-400 and K-100 are removed some shared classes with UCF and HMDB, resulting in 364 and 61 categories retained respectively. For the target datasets, we utilize established splits for HMDB, SSV2, Diving and UCF as outlined in previous studies [27, 38, 39, 45, 48, 57]. For RareAct database, we split the base, validation and novel set with 48, 8 and 8 classes. For FSAR task, TAMT is evaluated on SSV2, HMDB and UCF, whose splits follow their configurations in CDFSAR.

Implementation Details. We adopt VideoMAE [43] as the backbone, which is respectively built on ViT-S or ViT-B architectures for CDFSAR and FSAR for fair comparison. If not specified otherwise, the input resolution is 112×112 for ViT-S in CDFSAR and 224×224 in ViT-B for FSAR. The video inputs of the model are set to 16 frames, and then they are reduced to 8 in the patch embedding stage before the first transformer block. For optimization, we use SGD as the optimizer and adopt a cosine decay strategy to schedule the learning rate. The training epochs are set to 400, 140 and 40 for the SSL, SL and fine-tuning, respectively. In the pre-training phase, we employ the mean squared error and cross-entropy (CE) losses for SSL and SL, respectively. For the fine-tuning phase, we use CE loss. Euclidean distance is served as the metric function \mathcal{D} . We report accuracy in 5-way 1-shot and 5-way 5-shot settings on a single view, averaging 10,000 episodes for inference. Source code is available at <https://github.com/TJU-YDragonW/TAMT>.

4.2. Comparison with State-of-the-Arts

To fully evaluate our TAMT in the CDFSAR task, we conduct experiments with 5 source datasets and 5 target datasets, which form a nearly one-vs.-one cross-domain setting. Besides, we compare with state-of-the-art CDFSAR methods, which to our best knowledge cover all published CDFSAR works. The results of different methods in terms of 5-way 5-shot accuracy are reported in Tab. 1, where the best and second best results are highlighted in **red** and **blue** font,

Method	Source	Target					
		HMDB	SSV2	Diving	UCF	RareAct	Average
STARTUP++ [34]	K-400	44.71	39.60	14.92	60.82	45.22	41.05
DD++ [22]		48.04	44.50	16.23	63.26	47.01	43.81
STRM [42]		24.98	35.01	16.69	42.33	39.01	31.60
HYRSM [45]		29.81	40.09	17.57	45.65	44.27	35.49
CDFSL-V [38]		53.23	49.92	17.84	65.42	49.80	47.24
TAMT (Ours)		74.14	59.18	45.18	95.92	67.44	68.37 (+21.13)
TAMT* (Ours)		77.82	64.20	49.16	97.08	73.31	72.31
STARTUP++ [34]	K-100	24.97	15.16	14.55	32.20	31.77	23.73
DD++ [22]		25.99	16.00	16.24	34.10	31.20	24.71
SEEN*† [49]		52.80	31.20	40.90	79.60	50.20	50.94
CDFSL-V [38]		29.80	17.21	16.37	36.53	33.91	26.76
DMSD*† [18]		54.90	32.10	42.28	81.90	53.30	52.90
TAMT (Ours)		61.76	48.90	38.33	87.76	52.81	57.91 (+31.15)
TAMT* (Ours)		70.87	58.38	44.15	93.69	61.67	65.75
CDFSL-V [38]	SSV2	29.86	-	17.60	33.30	35.25	29.00
TAMT (Ours)		63.66	-	38.75	83.45	42.23	57.02 (+28.02)
CDFSL-V [38]	Diving	28.45	17.46	-	31.98	34.11	28.00
TAMT (Ours)		45.18	38.09	-	63.52	36.88	45.92 (+17.92)
STARTUP++ [34]	UCF	23.56	-	14.84	-	31.31	23.24
DD++ [22]		24.06	-	16.15	-	32.00	24.07
CDFSL-V [38]		28.86	-	16.07	-	33.91	26.82
TAMT (Ours)		45.34	-	33.38	-	41.08	39.93 (+13.11)

Table 1. Comparison(%) of state-of-the-arts on CDFSAR setting in terms of 5-way 5-shot accuracy, where five datasets (K-400, K-100, SSV2, Diving and UCF) are used as source data for transferring to five target datasets. All results are conducted with 112×112 resolution by using ViT-S backbone, except methods marked by * (224×224 resolution) and marked by † (backbone of ResNet-18).

Pre-tr		Tuning											
SSL	SL	SSV2			Diving			UCF			Average		
		Frozen	FFT	TAMT	Frozen	FFT	TAMT	Frozen	FFT	TAMT	Frozen	FFT	TAMT
✓		29.27	48.99	47.21	22.10	35.13	33.59	55.30	80.92	77.59	35.56	55.01	52.80
	✓	34.91	41.39	45.15	27.15	33.27	37.96	88.41	89.36	89.73	50.16	54.67	56.48
✓	✓	40.45	55.99	59.18	28.09	42.85	45.18	94.69	94.95	95.92	54.41	64.30	66.76

Table 2. Results (%) of various pre-training (Pre-tr) schemes and tuning strategies in terms of 5-way 5-shot accuracy (ViT-S as backbone). Memory: Training memory cost of FFT and TAMT. Parameter: Training parameters of FFT and TAMT.

respectively. From Tab. 1 we can conclude that (1) our TAMT outperforms existing methods by 13%~31% across all settings, leading to large performance gains. (2) On two widely used source datasets (K-400 and K-100), TAMT outperforms CDFSL-V [38] by an average of 21.13% and 31.15% across five target datasets, respectively. For the source datasets SSV2, Diving and UCF, our TAMT achieves improvements of 28.02%, 17.92% and 13.11% over CDFSL-V on average across different target datasets. (3) TAMT achieves performance improvements of 18.07%, 27.18%, 3.25%, 14.09%, and 11.47% over SEEN [49] across five target domains, while outperforming DMSD [18] with 15.97%, 26.28%, 1.87%, 11.79% and 8.37%. Furthermore, compared to CDFSL-V, our TAMT with decoupled training paradigm has nearly $5\times$ less training computational cost¹. These results above clearly demonstrate that our TAMT provides a promising baseline for the CDFSAR task in terms of both efficiency and effectiveness.

¹TAMT consumes 19 GPU days compared to CDFSL-V’s 88 GPU days when training on the K-400 source dataset across five target datasets.

Method	Memory	Parameters	Training Time
FFT	17.5G	29.9M	10.6h
TAMT	1.9G	2.8M	7.3h

Table 3. Efficiency Comparison for FFT and TAMT.




Figure 4. Convergence curves of FFT and TAMT on SSV2 dataset.

4.3. Ablation Studies

In this subsection, we conduct ablation studies to evaluate the effect of various components on three target datasets, i.e.,

Adapter	Moment	SSV2	Diving	UCF	Average
None	GAP	40.45	28.09	94.69	54.41
	GTMT	53.41	42.87	94.97	63.58
TAA	GAP	54.66	43.26	95.23	64.38
	TCov	56.72	43.68	95.16	65.19
	ELSTC	58.56	43.90	95.37	65.94
	GTMT	59.18	45.18	95.92	66.76

(a) Separated local or global modules.

Method	SSV2	Diving	UCF	Average
Adapter [21]	52.01	40.72	95.05	62.59
SSF [25]	50.46	31.51	95.25	59.07
TAM [28]	53.74	42.67	94.96	63.79
ST-Adapter [32]	54.26	43.12	95.16	64.18
TAMT (Ours)	59.18	45.18	95.92	66.76

(b) Adapter counterparts.

Table 4. Effect of core components (i.e., TAA & GTMT) on HTTN, where the results (%) in terms of 5-way 5-shot accuracy are reported.

SSV2, Diving and UCF, with K-400 serving as the source dataset. If not specified otherwise, the group number G is 4, and TAA block number L is 2 for the default option, whose ablation studies are reported in the supplemental material. The results highlighted with the same color indicate the identical deployment and experimental setting. To compare the training efficiency, we compare the GPU memory allocation required for training one episode (with 85 instances) on a server equipped with single NVIDIA TiTAN RTX A6000 GPU and Intel Xeon 8358 @ 2.6GHz CPU.

Pre-training and Fine-Tuning Strategies. We first assess the effect of various pre-training protocols (i.e., SSL, SL, and their combinations) and different fine-tuning approaches. For fine-tuning methods, Frozen, FFT, and our TAMT approaches learn none, full, or partial parameters of the pre-trained model, respectively. Particularly, Frozen method employs a learnable logistic regression classifier for few-shot inference [38]. As shown in Tab. 2, combining SSL with SL as a pre-training protocol consistently yields superior performance² over SSL or SL alone, across different fine-tuning methods and various target benchmarks, achieving gains of about 10%~20%. By integrating SSL with SL for pre-training, our TAMT consistently outperforms FFT by about 1%~3%. These results verify the effectiveness of our TAMT in mitigating the potential risk of overfitting under the optimal SSL & SL pre-training protocol. Furthermore, as shown in Tab. 3, our TAMT requires only $\sim 0.1\times$ GPU memory, $0.1\times$ learnable parameters, and $0.7\times$ training time compared to FFT, indicating the high efficiency of our TAMT decoupled training paradigm. Additionally, the convergence curves of models trained by FFT and our TAMT on SSV2 dataset [16] are illustrated in Fig. 4, where we can observe that FFT reaches earlier performance saturation than TAMT at the training stage, and validation performance degrades in the later training period (after ~ 30 epochs). This may suggest that FFT, with fine-tuning all parameters, suffers from the issue of overfitting again.

Effect of TAA and GTMT. We explore the individual contributions of TAA and GTMT within our HTTN in Tab. 4.

²When a single SSL is used for pre-training, FFT shows a 2% improvement over TAMT, likely due to SSL focusing on structural information, while FFT better assimilates semantic information.

When evaluated separately, both TAA and GTMT exhibit notable performance enhancements, as detailed in Tab. 4 (a). Specifically, in the absence of any adapter, GTMT outperforms the conventional GAP with an overall accuracy improvement of 9.17%, showcasing remarkable superiority on the challenging SSV2 and Diving with improvements of 12.96% and 14.78%, respectively. In addition, by integrating TAA with a variety of global statistical methods, our GTMT surpasses its counterparts, achieving the highest average performance of 66.76% across three datasets. It is evident that the approaches utilizing second-order statistics (TCov, ELSTC and GTMT) generally outperform first-order methodology (GAP). And among various second-order fashions, our ELSTC notably exceeds TCov with 0.75% average gains in a more efficient manner (with 262K vs. 4K dimension gap, see supplementary material for detailed analysis). Moreover, by adopting first-order for ELSTC, GTMT achieves a further improvement of 0.82%. These findings prove that TAA and GTMT can serve as competitive solutions for local and global tuning strategies. Furthermore, our TAMT also consistently surpasses its adapter counterparts—Adapter [21], SSF [25], TAM [28] and ST-Adapter [32]—with average performance improvements of 4.17%, 7.69%, 2.97% and 2.58% as detailed in Tab. 4 (b), respectively.

5. Conclusion

This paper proposed a novel Temporal-Aware Model Tuning (TAMT) method for cross-domain few-shot action recognition (CDFSAR) task. Particularly, to our best knowledge, our TAMT makes the first attempt to introduce a decoupled training paradigm for CDFSAR, effectively avoiding model retraining in the case of single source data and multiple target data. Moreover, from the perspectives of local feature recalibration and global (powerful) representation generation, a Hierarchical Temporal Tuning Network (HTTN) is proposed to effectively transfer the pre-trained models to target domain in a memory- and parameter-efficient manner. Extensive comparisons on CDFSAR tasks verify the effectiveness of our TAMT. We believe our TAMT provides a strong baseline for CDFSAR, and potentially contributes to push CDFSAR forward.

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grant 62276186, 62471083, 61971086, in part by the Haihe Lab of ITAI under Grant 22HHXCJC00002.

References

- [1] Aibek Alanov, Vadim Titov, Maksim Nakhodnov, and Dmitry Vetrov. StyleDomain: Efficient and lightweight parameterizations of StyleGAN for one-shot and few-shot domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2184–2194, 2023. 1
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. ViViT: A video vision transformer. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6816–6826, 2021. 2
- [3] Mina Bishay, Georgios Zoumpourlis, and Ioannis Patras. TARN: Temporal attentive relation network for few-shot and zero-shot action recognition. *British Machine Vision Conference (BMVC)*, pages 130.1–130.14, 2019. 2
- [4] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*. Springer, 2006. 5
- [5] Congqi Cao, Yajuan Li, Qinyi Lv, Peng Wang, and Yanning Zhang. Few-shot action recognition with implicit temporal alignment and pair similarity optimization. *Computer Vision and Image Understanding (CVIU)*, 210:103250, 2021. 2
- [6] Kaidi Cao, Jingwei Ji, Zhangjie Cao, Chien Yi Chang, and Juan Carlos Niebles. Few-shot video classification via temporal alignment. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10615–10624, 2020. 1, 2
- [7] Joao Carreira and Andrew Zisserman. Quo vadis, Action Recognition? A new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6299–6308, 2017. 2, 6
- [8] Guangyi Chen, Xiao Liu, Guangrun Wang, Kun Zhang, Philip HS Torr, Xiao-Ping Zhang, and Yansong Tang. Tempadapter: Adapting image-text pretraining for video question answer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13945–13955, 2023. 5
- [9] Zixuan Chen, Junhong Zou, and Xiaotao Wang. Semantic segmentation on VSPW dataset through aggregation of transformer models. *arXiv preprint arXiv:2109.01316*, 2021. 2
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009. 2
- [11] Mingze Gao, Qilong Wang, Zhenyi Lin, Pengfei Zhu, Qinghua Hu, and Jingbo Zhou. Tuning pre-trained model via moment probing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11803–11813, 2023. 5
- [12] Zan Gao, Tao-tao Han, Lei Zhu, Hua Zhang, and Yinglong Wang. Exploring the cross-domain action recognition problem by deep feature learning and cross-domain learning. *IEEE Access*, 6:68989–69008, 2018. 3
- [13] Zan Gao, Leming Guo, Weili Guan, Anan Liu, and Shengyong Chen. A pairwise attentive adversarial spatiotemporal network for cross-domain few-shot action recognition-R2. *IEEE Transactions on Image Processing*, 30:767–782, 2020. 3
- [14] Zan Gao, Yibo Zhao, Hua Zhang, Da Chen, An-An Liu, and Shengyong Chen. A novel multiple-view adversarial learning network for unsupervised domain adaptation action recognition. *IEEE Transactions on Cybernetics*, 52(12):13197–13211, 2022. 3
- [15] Zhiqiang Gao, Kaizhu Huang, Rui Zhang, Dawei Liu, and Jieming Ma. Towards better robustness against common corruptions for unsupervised domain adaptation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 18882–18893, 2023. 1
- [16] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The “Something Something” video database for learning and evaluating visual common sense. In *IEEE International Conference on Computer Vision (ICCV)*, pages 5843–5851, 2017. 2, 6, 8
- [17] Fei Guo, YiKang Wang, Han Qi, Li Zhu, and Jing Sun. Consistency prototype module and motion compensation for few-shot action recognition (CLIP-CPM²C). *Neurocomputing*, 611:128649, 2025. 1
- [18] Fei Guo, Yi Kang Wang, Han Qi, Li Zhu, and Jing Sun. DMSD-CDFSAR: Distillation from mixed-source domain for cross-domain few-shot action recognition. *Expert Systems With Applications*, 270, 2025. 3, 7
- [19] Yunhui Guo, Noel C. Codella, Leonid Karlinsky, James V. Codella, and Rogerio Feris. A broader study of cross-domain few-shot learning. In *European Conference on Computer Vision (ECCV)*, pages 124–141, 2020. 3
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16000–16009, 2022. 3
- [21] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning (ICML)*, pages 2790–2799, 2019. 4, 5, 8
- [22] Ashraf Islam, Chun-Fu Richard Chen, Rameswar Panda, Leonid Karlinsky, Rogerio Feris, and Richard J Radke. Dynamic distillation network for cross-domain few-shot recognition with unlabeled data. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:3584–3595, 2021. 7
- [23] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: A large video database for human motion recognition. In *International Conference on Computer Vision (ICCV)*, pages 2556–2563, 2011. 2, 6
- [24] Yingwei Li, Yi Li, and Nuno Vasconcelos. RESOUND: Towards action recognition without representation bias. In *Pro-*

- ceedings of the European Conference on Computer Vision (ECCV)*, pages 513–528, 2018. 2, 6
- [25] Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & shifting your features: A new baseline for efficient model tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 109–123, 2022. 4, 5, 8
- [26] Liqun Lin, Zheng Wang, Jiachen He, Weiling Chen, Yiwen Xu, and Tiesong Zhao. Deep quality assessment of compressed videos: A subjective and objective study. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(6): 2616–2626, 2023. 2
- [27] Xin Liu, Huanle Zhang, Hamed Pirsiavash, and Xin Liu. MASTAF: A model-agnostic spatio-temporal attention fusion network for few-shot video classification. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2508–2517, 2023. 1, 3, 6
- [28] Zhaoyang Liu, Limin Wang, Wayne Wu, Chen Qian, and Tong Lu. TAM: Temporal adaptive module for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13708–13718, 2021. 8
- [29] Kang Ma, Ying Fu, Dezhi Zheng, Yunjie Peng, Chunshui Cao, and Yongzhen Huang. Fine-grained unsupervised domain adaptation for gait recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11313–11322, 2023. 1
- [30] Antoine Miech, Jean Baptiste Alayrac, Ivan Laptev, Josef Sivic, and Andrew Zisserman. RareAct: A video dataset of unusual interactions. *arXiv preprint arXiv:2008.01018*, 2020. 2, 6
- [31] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. DINOv2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 3
- [32] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. ST-Adapter: Parameter-efficient image-to-video transfer learning. *Advances in Neural Information Processing Systems (ICCV)*, 35:26462–26477, 2022. 5, 8
- [33] Toby Perrett, Alessandro Masullo, Tilo Burghardt, Majid Mirmehdi, and Dima Damen. Temporal-relational cross Transformers for few-shot action recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 475–484, 2021. 2
- [34] Cheng Perng Phoo and Bharath Hariharan. Self-training for few-shot transfer across extreme task differences. In *International Conference on Learning Representations (ICLR)*, 2021. 7
- [35] A. J. Piergiovanni, Weicheng Kuo, and Anelia Angelova. Rethinking Video ViTs: Sparse video tubes for joint image and video learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2214–2224, 2023. 2
- [36] Zekun Qi, Runpei Dong, Guofan Fan, Zheng Ge, Xiangyu Zhang, Kaisheng Ma, and Li Yi. Contrast with reconstruct: Contrastive 3D representation learning guided by generative pretraining. In *International Conference on Machine Learning (ICML)*, 2023. 3
- [37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, and Jack Clark. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, pages 8748–8763. PMLR, 2021. 2
- [38] Sarinda Samarasinghe, Mamshad Nayeem Rizve, Navid Kardan, and Mubarak Shah. CDFSL-V: Cross-domain few-shot learning for videos. In *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)*, pages 11643–11652, 2023. 1, 3, 5, 6, 7, 8
- [39] Yuheng Shi, Xinxiao Wu, Hanxi Lin, and Jiebo Luo. Commonsense knowledge prompting for few-shot action recognition in videos. *IEEE Transactions on Multimedia*, 26:8395–8405, 2024. 6
- [40] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS)*, pages 4080–4090, 2017. 4
- [41] 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. 2, 6
- [42] Anirudh Thatipelli, Sanath Narayan, Salman Khan, Rao Muhammad Anwer, Fahad Shahbaz Khan, and Bernard Ghanem. Spatio-temporal relation modeling for few-shot action recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19926–19935, 2022. 2, 7
- [43] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems (NeurIPS)*, 35:10078–10093, 2022. 3, 6
- [44] Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. VideoMAE V2: Scaling video masked autoencoders with dual masking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14549–14560, 2023. 2, 3
- [45] Xiang Wang, Shiwei Zhang, Zhiwu Qing, Mingqian Tang, Zhengrong Zuo, Changxin Gao, Rong Jin, and Nong Sang. Hybrid relation guided set matching for few-shot action recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19916–19925, 2022. 1, 2, 6, 7
- [46] Xixi Wang, Xiao Wang, Bo Jiang, and Bin Luo. Few-shot learning meets transformer: Unified query-support transformers for few-shot classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(12):7789–7802, 2023.
- [47] Xiao Wang, Weirong Ye, Zhongang Qi, Guangge Wang, Jianping Wu, Ying Shan, Xiaohu Qie, and Hanzi Wang. Task-aware dual-representation network for few-shot action recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(10):5932–5946, 2023.
- [48] Xiang Wang, Shiwei Zhang, Zhiwu Qing, Changxin Gao, Yingya Zhang, Deli Zhao, and Nong Sang. MoLo: Motion-augmented long-short contrastive learning for few-shot action

- recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18011–18021, 2023. [1](#), [3](#), [6](#)
- [49] Xiang Wang, Shiwei Zhang, Zhiwu Qing, Yiliang Lv, Changxin Gao, and Nong Sang. Cross-domain few-shot action recognition with unlabeled videos. *Computer Vision and Image Understanding (CVIU)*, 233:103737, 2023. [1](#), [3](#), [5](#), [7](#)
- [50] Xiang Wang, Shiwei Zhang, Hangjie Yuan, Yingya Zhang, Changxin Gao, Deli Zhao, and Nong Sang. Few-shot action recognition with captioning foundation models. *arXiv preprint arXiv:2310.10125*, 2023. [1](#)
- [51] Xiao Wang, Yang Lu, Wanchuan Yu, Yanwei Pang, and Hanzhi Wang. Few-shot action recognition via multi-view representation learning. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1, 2024. [2](#)
- [52] Xiang Wang, Shiwei Zhang, Jun Cen, Changxin Gao, Yingya Zhang, Deli Zhao, and Nong Sang. CLIP-guided prototype modulating for few-shot action recognition. *International Journal of Computer Vision*, 132(6):1899–1912, 2024. [1](#), [3](#)
- [53] Yi Wang, Kunchang Li, Yizhuo Li, Yanan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, Zun Wang, Sen Xing, Guo Chen, Junting Pan, Jiashuo Yu, Yali Wang, Limin Wang, and Yu Qiao. InternVideo: General video foundation models via generative and discriminative learning. *arXiv preprint arXiv:2212.03191*, 2022. [2](#)
- [54] Yuyang Wanyan, Xiaoshan Yang, Chaofan Chen, and Changsheng Xu. Active exploration of multimodal complementarity for few-shot action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6492–6502, 2023. [1](#)
- [55] Jiazheng Xing, Mengmeng Wang, Yudi Ruan, Bofan Chen, Yaowei Guo, Boyu Mu, Guang Dai, Jingdong Wang, and Yong Liu. Boosting few-shot action recognition with graph-guided hybrid matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1740–1750, 2023. [1](#)
- [56] Shen Yan, Xuehan Xiong, Anurag Arnab, Zhichao Lu, Mi Zhang, Chen Sun, and Cordelia Schmid. Multiview Transformers for video recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3323–3333, 2022. [2](#)
- [57] Songyang Zhang, Jiale Zhou, and Xuming He. Learning implicit temporal alignment for few-shot video classification. *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1309–1315, 2021. [6](#)
- [58] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. Image BERT pre-training with online tokenizer. In *International Conference on Learning Representations (ICLR)*, 2021. [3](#)
- [59] Linchao Zhu and Yi Yang. Compound memory networks for few-shot video classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 751–766, 2018. [2](#), [6](#)