

# Frame Order Matters: A Temporal Sequence-Aware Model for Few-Shot Action Recognition

Bozheng Li<sup>1\*</sup>, Mushui Liu<sup>1\*</sup>, Gaoang Wang<sup>2</sup>, Yunlong Yu<sup>1†</sup>

<sup>1</sup>College of Information Science & Electronic Engineering, Zhejiang University

<sup>2</sup>University of Illinois Urbana Champaign Institute, Zhejiang University  
{bozhengli, lms}@zju.edu.cn, gaoangwang@intl.zju.edu.cn, yuyunlong@zju.edu.cn

## Abstract

In this paper, we propose a novel Temporal Sequence-Aware Model (TSAM) for few-shot action recognition (FSAR), which incorporates a sequential perceiver adapter into the pre-training framework, to integrate both the spatial information and the sequential temporal dynamics into the feature embeddings. Different from the existing fine-tuning approaches that capture temporal information by exploring the relationships among all the frames, our perceiver-based adapter recurrently captures the sequential dynamics alongside the timeline, which could perceive the frame order change. To obtain the discriminative representations for each class, we extend a textual corpus for each class derived from the large language models (LLMs) and enrich the visual prototypes by integrating the contextual semantic information. Besides, We introduce an unbalanced optimal transport strategy for feature matching that mitigates the impact of class-unrelated features, thereby facilitating more effective decision-making. Experimental results on five FSAR datasets demonstrate that our method establishes a new benchmark, outperforming the second-best competitors.

## Introduction

Video action recognition (Arnab et al. 2021; Yang et al. 2023), a cornerstone task in video understanding (Tang et al. 2024; Zhao et al. 2024), demands comprehensive understanding and representation of spatial-temporal information. However, the inherent challenges of large data volume and storage complexities of video data often lead to a scarcity of training samples, hindering model development (Dan et al. 2024, 2023; Tang et al. 2022). Few-shot action recognition (FSAR) (Wang et al. 2023a; Cao et al. 2020; Perrett et al. 2021; Xing et al. 2023b) emerges as a promising direction to alleviate this data bottleneck while exploring efficient video encoding architectures. By leveraging a limited number of labeled examples, FSAR aims to enable the rapid recognition of novel action categories.

Learning-based methods have significantly advanced the boundaries of FSAR. Existing approaches have explored diverse avenues, with some focusing on the design of powerful video encoders (Thatipelli et al. 2022a; Zhang, Zhou, and

\*These authors contributed equally.

†Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

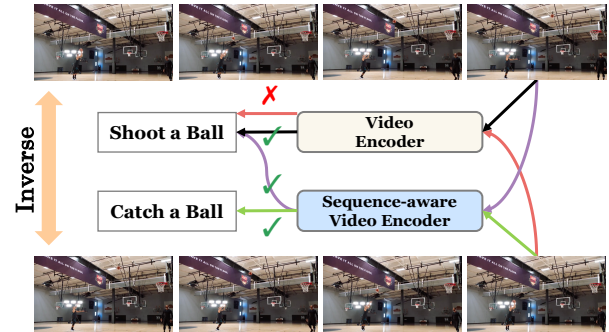


Figure 1: Illustration of the semantic changing when rearranging video frames. The existing model could barely distinguish reversed video due to equally treated frame features while our temporal sequence-aware model could capture the differences between the videos with different frame orders.

He 2021) and others concentrating on the development of effective FSL paradigms (Zheng, Chen, and Jin 2022; Perrett et al. 2021). The advent of Contrastive Language-Image Pre-Training (CLIP) (Radford et al. 2021) has ushered in a more robust visual backbone that supports multi-modal information integration. Consequently, leveraging the pre-trained CLIP image encoder for video representation learning has demonstrated promising potential in various works (Wasim et al. 2023; Ni et al. 2022; Liu, Li, and Yu 2024; Wu et al. 2024b; Lu et al. 2024; Liu et al. 2024; Yang et al. 2024).

Recent endeavors have investigated the adaptation of CLIP for FSAR tasks. Notably, CLIPFSAR (Wang et al. 2023a) pioneered the transfer of the CLIP backbone to FSAR by directly fine-tuning the entire architecture. A series of subsequent works (Pei et al. 2023; Xing et al. 2023a; Cao et al. 2024) employed Parameter-Efficient Fine-Tuning (PEFT) strategies to transfer the pre-trained backbone to FSAR. These approaches prioritize capturing temporal information by delving into the interrelationships among frames, yet they inadequately consider the temporal sequence. In some cases, the category of a video can be significantly altered by merely rearranging the order of its frames. As demonstrated in Fig. 1, reversing the order of input video frames to the models fails in altering the predicted results, while the category semantic has been changed.

Drawing inspiration from the human cognitive process of sequential information processing, we introduce a novel sequence-aware model for FSAR that incorporates a sequential perceiver adapter into the pre-training backbone to integrate both the spatial information and the sequential temporal dynamics into the feature embeddings. Specially, the sequential perceiver adapter captures the inherent order within the video frames, by treating temporal information as a distinct modality alongside spatial features. In each perceiver adapter, a temporal query recurrently compresses temporal information from each frame and injects temporal sequential relation to the pre-trained backbone. To enrich the feature representations of each class, we construct an expanded textual prompt corpus leveraging the Large Language Models (LLMs), which enables discriminative prototype construction with enhanced textual information involved. Besides, we transform the matching of support and query video into an unbalanced optimal transport problem to further improve the matching process. Overall, the highlights are as follows:

- We present a novel temporal sequence-aware model (TSAM) for FSAR that integrates a sequential perceiver adapter into the pre-trained visual backbone to extract both dynamic temporal sequences and spatial information for feature embeddings. The Perceiver Adapter is designed to recurrently integrate and compress spatial information into temporal cues, which could perceive the frame order change.
- To enrich the feature representations of each class, we extend a textual corpus derived from LLMs and integrate the textual semantics into the visual prototypes. Besides, we introduce an unbalanced optimal transport strategy for feature matching, which diminishes the influence of non-discriminative, class-irrelevant features.
- Experiments on five diverse FSAR datasets show that our proposed TSAM sets a new state-of-the-art benchmark, surpassing the second-best competitor by a large margin and demonstrating its effectiveness and superiority.

## Related Works

**Few-shot Action Recognition.** Few-shot action recognition focuses on temporal information modeling and constructing relationships between query and support sets. Previous methods have leveraged metric learning to achieve alignment between query and support sets. OTAM (Cao et al. 2020), TRX (Perrett et al. 2021), CMOT (Lu, Ye, and Zhan 2021a), and HyRSM (Wang et al. 2022) utilize dynamic temporal warping, multi-frame matching, optimal transport, and set matching to calculate the video similarities to build the relationship of videos, respectively. Additionally, other approaches have focused on spatial information in video frames and the fusion of multi-modal information. MTFAN (Wu et al. 2022) learns global motion patterns with a motion modulator. STRM (Thatipelli et al. 2022b) enriches local patch features for spatial-temporal modeling. MoLo (Wang et al. 2023b) explicitly extracts motion dynamics within a unified network. Recently, based on the powerful pre-trained model, *e.g.*, CLIP (Radford et al. 2021), CLITSAM (Wang et al. 2023a) and MA-CLIP (Xing et al. 2023a) have shown

significant progress by transferring pre-trained image encoders to the video domain and constructing prototypes incorporated with textual information.

**Parameter-Efficient Fine-Tuning.** Parameter-Efficient Fine-Tuning (PEFT) (Houlsby et al. 2019; Jia et al. 2022) aims to efficiently fine-tune large pre-trained models by adjusting only a few parameters while freezing the backbone, achieving performance similar to full fine-tuning. This paradigm has shown promise in computer vision. Methods like VPT (Jia et al. 2022), CoOp (Zhou et al. 2022), and CLIP-Adapter (Gao et al. 2023) incorporate learnable visual prompts, textual prompts, and adapters to learn discriminative features for downstream tasks. For video recognition, ST-Adapter (Pan et al. 2022) uses a 3D-convolution-based adapter to capture temporal information, while AIM (Yang et al. 2023) further reuses the attention layer in the pre-trained model to formulate temporal information. Vita-CLIP (Wasim et al. 2023) introduces learnable prompt tokens in both spatial and temporal dimensions. MA-CLIP (Xing et al. 2023a) introduced a textual-guided adapter structure. However, these approaches often overlook the sequential characteristics of video frames and the importance of video frame relationship matching.

**Perceiver.** Perceiver (Jaegle et al. 2021) refers to a transformation of asymmetric attention that can distill multi-modal information into a compressed latent bottleneck. Most former works either utilize perceiver to resample high-volume feature tokens (Alayrac et al. 2022) or bridge the feature space of different modality (Acosta et al. 2022). Recently, Meta utilized perceiver as the video aggregator in LLaMA3.1 (Dubey et al. 2024), proving its ability in temporal modeling and feature compression for the video domain.

In this paper, we utilize an adapter to transfer the pre-trained backbone for FSAR, emphasizing frame order modeling through the proposed sequential perceiver structure.

## Methodology

### Problem Definition

In FSAR tasks, the video datasets are divided into non-overlapping training  $\mathcal{D}_{train}$  and testing  $\mathcal{D}_{test}$  splits. Consistent with the prior work (Snell, Swersky, and Zemel 2017), we employ episodic training by randomly sampling few-shot tasks from the training set. In each task, the video data is divided into a support set  $\mathcal{S}$  and a query set  $\mathcal{Q}$ . The  $N$ -way  $K$ -shot task indicates using a support set comprised of  $N$  classes, with  $K$  samples per class to learn the classifier for the query samples during inference.

### Temporal Sequence-Aware Model

**Overall Framework** Following the sequential nature of video data, we propose a Temporal Sequence-Aware Model (TSAM) framework for FSAR that enriches the frame-level feature representations by incorporating the temporal contexts captured sequentially from the spatial information alongside the temporal dimension. As illustrated in Fig. 2, the framework consists of a perceiver-based video encoder to extract both the visual and temporal information, a textual

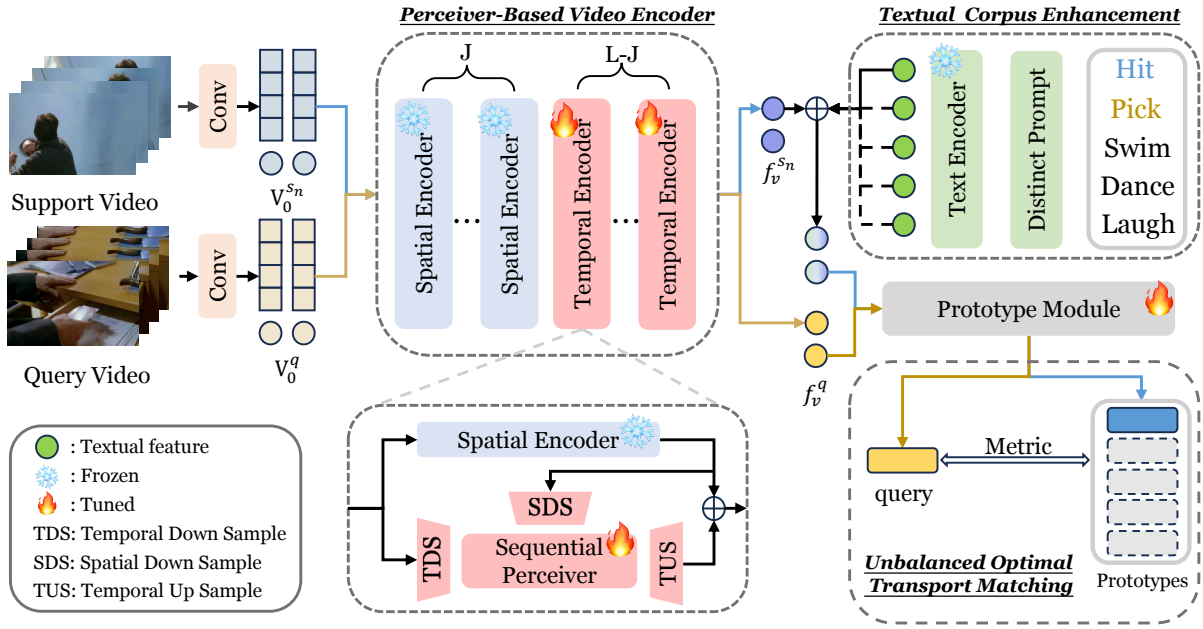


Figure 2: The framework of our proposed temporal sequence-aware model (TSAM), which consists of a perceiver-based video encoder that captures both the temporal sequential dynamics and spatial information, a textual corpus enhancement module that incorporates the class-related semantics into the feature prototype, and an unbalanced optimal transport matching module that enhances the feature matching.

corpus enhancement module to enrich the video representations, and an unbalanced optimal transport matching module to relieve the class-unrelated features for decision-making.

Using ViT as the visual backbone, we delve into the details of our proposed framework. Given an input video  $v \in \mathbb{R}^{T \times H \times W \times 3}$ , where  $T$  denotes the uniformly sampled frame numbers,  $H$  and  $W$  respectively denote the height and width of each frame, we sequentially input the frames into the pre-trained ViT backbone to obtain the feature representations. As the original ViT blocks inherently struggle to capture the temporal dynamics across frames, we propose to extend the block to make it compress spatial information into the temporal cues, and in turn enhance the frame-level feature representations. Assuming a total of  $L$  blocks in ViT architecture, we propose to freeze the initial  $J$  blocks, preserving their pre-trained weights. For the remaining  $L - J$  blocks, we introduce a perceiver adapter in parallel to each block’s original spatial encoder. These adapters are then fine-tuned during training, allowing the model to learn temporal cues while leveraging the spatial information encoded by the frozen blocks. We incorporate a learnable temporal token as an input to the  $J + 1$ -th block to better capture the temporal information. The spatial information encoded with the spatial encoder and the temporal information captured with the perceiver adapter are fused for the next block.

To obtain discriminative visual representations for each category with a few visual samples, we propose to combine the semantic descriptions related to the categories into visual presentations, powered with the off-the-shelf large language models (LLMs). Once the feature representations of

both support and query samples are obtained, the similarity metric is the key to decision-making. Considering that during the few-shot matching process, irrelevant noise frames will harm the matching of the action class, we introduce an unbalanced optimal matching strategy to filter redundant noise within sampled video frames by punishing many-to-one matching and involving soft condition boundary.

**Perceiver Adapter** The Perceiver Adapter is a pivotal component of video encoder, which is designed to capture the temporal information for feature representations. Perceiver structure was originally introduced in (Jaegle et al. 2021) to enhance feature compression and foster cross-modality communication. Inspired by this approach, we view temporal and spatial features as two distinct modalities and introduce a perceiver-based adapter to compress temporal information within spatial features into a temporal query.

However, when a straightforward perceiver structure (as shown in Fig. 3 (a)) is directly employed as an adapter, two challenges arise. On one hand, utilizing a fully learnable temporal query to condense temporal information within spatial features, while trained on limited data, can inadvertently lead to an overly cohesive prototype space, potentially limiting the model’s ability to capture diverse temporal patterns. On the other hand, the straightforward attention mechanisms employed in the perceiver adapter are prone to an inadvertent issue of information leakage (Wang et al. 2023a; Xing et al. 2023a), where data from subsequent frames inadvertently permeates into preceding ones during processing. This leads to a homogenization effect across temporal dimensions, diminishing the model’s capacity to accu-




Figure 3: Illustration of (a) simple perceiver-based adapter and (b) our sequential perceiver-based adapter.

rately capture and differentiate between temporal dynamics, thereby compromising its overall performance.

To tackle these issues, we propose the sequential perceiver. As shown in Fig. 3(b), instead of learnable query tokens, the sequential perceiver takes the temporal token embedding of the first frame  $\mathbf{f}_t^1$  as the query to interact with the spatial feature embedding  $\mathbf{f}_s^1$  of the first frame to obtain the temporal embedding  $\mathbf{f}_{temp}^1$ . Then, the temporal embedding  $\mathbf{f}_{temp}^1$  is taken as the query to interact with the spatial feature embedding  $\mathbf{f}_s^2$  of the second frame, and so on. Given the feature representations  $\mathbf{V}_i \in \mathbb{R}^{T \times U \times D}$  of  $i$ -th block ( $J \leq i \leq L$ ), the process is first formulated as:

$$\begin{aligned} \mathbf{V}_i^s &= \text{Spatial\_D}(\text{Visual\_B}(\mathbf{V}_i)), \\ \mathbf{V}_i^t &= \text{Temporal\_D}(\mathbf{V}_i), \end{aligned} \quad (1)$$

where Visual\_B denotes the visual block, Spatial\_D and Temporal\_D represent the specific modules tasked with diminishing the dimensionality along the spatial and temporal axes, respectively.

Then, for the  $k$ -th frame, we obtain the temporal feature embedding by:

$$\mathbf{F}_k^t = \text{Cross\_Att}(\mathbf{F}_{k-1}^t, \mathbf{V}_k^s, \mathbf{V}_k^s) + \beta \cdot \mathbf{V}_k^t, \quad (2)$$

where  $\mathbf{F}_{k-1}^t$  performs as the query and  $\beta$  denotes the injection ratio of temporal information. For the first frame, the query is initialized as the temporal embedding, i.e.,  $\mathbf{F}_{-1}^t = \mathbf{V}_0^t$ .

After obtaining the temporal feature embeddings, we fuse the spatial feature embeddings and temporal feature embedding for the next block with:

$$\mathbf{V}_{i+1} = \text{Visual\_B}(\mathbf{V}_i) + \alpha \cdot \mathbf{F}^t, \quad (3)$$

where  $\alpha$  is a scaling factor that controls the incorporation ratio of temporal features.

With this recurrent temporal modeling adapter, temporal information located among visual frames is well captured, which in turn, enrich frame-level feature representations.

**Textual Corpus Enhanced Training** As the limited visual samples hardly provided sufficient information for category learning, prior work (Xing et al. 2023a) has attempted to incorporate class-level semantic information into feature learning. In this work, we propose to integrate the detailed

textual descriptions into the prototype calibration to obtain a more discriminative prototype.

To achieve this, we build a textual corpus based on the class labels with the assistance of the Large Language Models (LLMs) (Achiam et al. 2023; Wu et al. 2024a). Specifically, given a prompt about the specific class that requests to generate some descriptions related to the corresponding class, the off-the-shelf LLMs provide a corpus containing detailed descriptions of the classes. For temporally abstract datasets like SSv2, the extended corpus maintains highly abstract textual representations, where specific objects in the corpus are replaced with generalized terms like ‘‘object’’, e.g. original annotation like ‘‘Bending [something] until it breaks’’ is expanded to ‘‘Bending [something] until it breaks: object bending, snapping into pieces.’’. This enriched corpus expands the original weak textual representation space and enables a more nuanced learning process for video encoder.

With the class-related descriptions, we obtain the textual embeddings with the existing textual encoder, which are then used to enrich the visual representations as follows:

$$\mathbf{f}_n = \mathbf{f}_n^v + \text{Textenc}(\text{Corpus}(\text{class}_n)), \quad (4)$$

where Textenc denotes the textual encoder and we use CLIP (Radford et al. 2021) in the experiment.  $\mathbf{f}_n^v$  denotes the visual feature representation of class  $n$ . The textual descriptions in the  $n$ -th corpus are concatenated as a sentence to the textual encoder. The augmented feature representations are then fed into a simple prototype merging module that is achieved with a transformer block.

## Model Fine-tuning

**Unbalanced Optimal Transport Matching** Optimal Transport (OT) (Lu, Ye, and Zhan 2021b) uses a Sinkhorn matching algorithm to address the problem of imbalance between video content matching and temporal order matching. Considering that prototype and query videos usually lie under deviational distributions, We introduce unbalanced optimal transport (UOT) to reduce interference of background noise from redundant frames in matching process.

Given the support video features  $\mathbf{S}$  and the query video features  $\mathbf{Q}$ , the distance matrix is defined as  $D = [d(\mathbf{s}_i, \mathbf{q}_j)] \in \mathbb{R}^{N \times N}$ , where the element  $d(\mathbf{s}_i, \mathbf{q}_j)$  represents the Euclidean distance between the latent representation of the  $i$ -th query frame feature and that of the  $j$ -th support frame feature, then the similarity between query and support is denoted as:

$$\text{Dis}(\mathbf{Q}, \mathbf{S}) = \min_{T \in \Pi(\alpha, \beta)} \underbrace{\langle D, T \rangle}_{OT_{problem}} + \lambda \underbrace{\mathcal{H}(T)}_{Reg}. \quad (5)$$

Here  $\lambda \mathcal{H}(T)$  denotes an entropic regularization term that smooths the optimization of matching plan  $T$ . We set the constraints  $(T\mathbf{1}) = \alpha$ ,  $(T^T\mathbf{1}) = \beta$ , which naturally penalize the many-to-one matching in optimal transport. The optimal solution  $T$  is called the optimal transport plan, which is the optimal distribution of latent representation pairs that minimizes the expectation of the distance  $d(\mathbf{S}, \mathbf{Q})$ .

Motivated by the fact that video features can contain background noise, leading to the frames between videos not necessarily requiring one-to-one correspondence, we draw inspiration from unbalanced optimal transport (Zhan et al.

2021). Originally, the score between videos was obtained through matching, as follows:

$$Dis(\mathbf{Q}, \mathbf{S}) = \underbrace{\min_{T \in \Pi(\alpha, \beta)} \langle D, T \rangle}_{\text{OT problem}} + \underbrace{\lambda \mathcal{H}(T)}_{\text{Reg.}} \quad (6)$$

$$+ \underbrace{\tau \text{KL}(T\mathbf{1} \parallel \mathbf{N}) + \tau \text{KL}(T^\top \mathbf{1} \parallel \mathbf{N})}_{\text{SoftCons.}}$$

where KL is the Kullback-Leibler divergence. And  $N$  denotes the correlation distributions generated with the cross-inner product between support and query features. To obtain  $T$ , we followed (Zhan et al. 2021) to approximate the UOT solution by applying the skinhorn algorithm.

**Training Objective** After obtaining the prototype and query features, the metric distance between the query features and each prototype is calculated as:

$$Dis(\mathbf{q}, \mathbf{s}_i) = M(\tilde{f}_q, f_i^s) \quad (7)$$

The probability distribution of query video based on prototypes is given by:

$$p_{\hat{y}=i|q}^{fsl} = \frac{\exp(Dis(\mathbf{q}, \mathbf{s}_i))}{\sum_{j=1}^N \exp(Dis(\mathbf{q}, \mathbf{s}_j))} \quad (8)$$

where  $\hat{y}$  denotes the prediction class of the query video.

Following former works (Wang et al. 2023a; Xing et al. 2023a), we integrate the zero-shot prediction  $p_{\hat{y}=i|q}^{zsl}$  with the few-shot prediction result to get the merged prediction of query videos:

$$p_{\hat{y}=i|q} = (p_{\hat{y}=i|q}^{fsl})^\lambda \cdot (p_{\hat{y}=i|q}^{zsl})^{1-\lambda} \quad (9)$$

The loss function consists of two parts, zero-shot video-text loss that directly predicts the action based on few-shot label and textual feature, and few-shot loss which applies cross-entropy on few-shot probability distribution:

$$\mathcal{L} = CE(p_{\hat{y}=i|q}^{fsl}, gt_{fsl}) + CE(p_{\hat{y}=i|q}^{zsl}, gt_{zsl}). \quad (10)$$

## Experiments

### Experimental Details

**Datasets.** We evaluate our method on four widely used few-shot action recognition datasets: HMDB-51 (Kuehne et al. 2011), UCF101 (Soomro, Zamir, and Shah 2012), Kinetics (Carreira and Zisserman 2017), and SSv2 (Goyal et al. 2017). The first three datasets focus on scene understanding, while SSv2 requires more challenging temporal modeling (Tu et al. 2023). For both HMDB-51 and UCF101, we use the few-shot split from (Zhang et al. 2020), with 31/10/10 and 70/10/21 train/val/test classes, respectively. Following (Zhu and Yang 2018), we select a subset of Kinetics with 64/12/24 train/val/test classes. For SSv2, we report results on two few-shot splits: SSv2-Small (Zhu and Yang 2020) and SSv2-Full (Cao et al. 2020). Both have 64/12/24 train/val/test classes, with SSv2-Full containing  $10\times$  more videos per class in the training set.

**Implementation Details.** We adapt two CLIP pre-trained (Radford et al. 2021) backbones, ResNet50 and ViT-B/16,

for comparison. Each video is uniformly sampled into 8 frames, which are then cropped to a resolution of  $224\times 224$ . During training, we freeze the pre-trained backbone and fine-tune only the proposed adapter. We use SGD as the optimizer. For testing, we sample 10,000 few-shot action tasks to calculate the 5-way 1-shot and 5-way 5-shot average accuracy. Training is conducted on 4 RTX 3090 GPUs.

### Performance Comparison

Table 1 provides a comprehensive comparison of our method and 12 recent few-shot action recognition competitors on five datasets. The competitors use either CLIP or ResNet model pre-trained on the ImageNet-21k as the visual backbone. We observe that our TSAM achieves the best performances in both 1-shot and 5-shot tasks. Specifically, for 1-shot scenario, our method surpasses the second-best competitor by significant margins of 1.4%, 2.5%, 1.1%, 1.3%, and 0.5% on SSv2-Small, SSv2-Full, HMDB-51, UCF101, and Kinetics datasets, respectively. In 5-shot setting, our approach respectively outperforms the second-best competitors by 2.2%, 2.3%, 1.0%, 0.2%, and 1.1% on each of these datasets, highlighting its effectiveness.

Besides, we also reach the following observations: **1) Backbone matters:** the visual backbones trained with the Transformer-based architecture (i.e., ViT-B) demonstrate notably superior performance in comparison to those utilizing the ResNet-based architecture (i.e., ResNet50), highlighting the significance of the pre-trained visual backbone as a crucial factor contributing to video recognition capabilities. **2) Fine-tuning manner:** Compared to CLIP-FSAR (Wang et al. 2023a), which fully fine-tunes the backbone, our temporal adapter structure offers a more nuanced representation of temporal information while requiring fewer parameters, effectively balancing the pre-trained model’s capabilities with video feature adaptation. **3) Temporal modeling:** Compared to other adapter-based methods, e.g., MA-CLIP (Xing et al. 2023a), our sequential perceiver structure injects sequential information into the pre-trained backbone, enabling a stronger video encoder that constructs a more aligned video feature space for support and query videos. For example, TSAM outperforms MA-CLIP (Xing et al. 2023a) by 2.2% in the 1-shot on the SSv2-Full dataset, which demands comprehensive temporal understanding.

**Cross-Dataset Comparison.** In this experiment, we conduct cross-dataset evaluations to assess model’s generalization capabilities. Specifically, we train the model on Kinetics dataset and then evaluate the model on UCF-101, HMDB-51, and SSv2 datasets. As depicted in Fig. 4, our method exhibits a notable advantage over CLIP-FSAR (Wang et al. 2023a), achieving impressive improvements of 3.7% on UCF-101, 8.1% on HMDB-51, and a substantial 3.9% gain on challenging SSv2 dataset. This consistent superior performance across diverse video datasets underscores the robust generalization capabilities of our proposed approach, effectively excelling in varied video data distributions.

### Ablation Studies

**Study on Partial Adaption.** In this experiment, we evaluate the impacts of adapter numbers added during the model

Method	Pre-training	SSv2-Small		SSv2-Full		HMDB51		UCF101		Kinetics	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
CMN (Zhu and Yang 2018)	INet-RN50	34.4	43.8	36.2	48.9	-	-	-	-	60.5	78.9
OTAM (Cao et al. 2020)	INet-RN50	-	-	42.8	52.3	-	-	-	-	73.0	85.8
AmeFuNet (Fu et al. 2020)	INet-RN50	-	-	-	-	60.2	75.5	85.1	95.5	74.1	86.8
TRX (Perrett et al. 2021)	INet-RN50	-	59.1	-	64.6	-	75.6	-	96.1	63.6	85.9
SPRN (Wang et al. 2021)	INet-RN50	-	-	-	-	61.6	76.2	86.5	95.8	75.2	87.1
HyRSM (Wang et al. 2022)	INet-RN50	40.6	56.1	54.3	69.0	60.3	76.0	83.9	94.7	73.7	86.1
TA <sup>2</sup> N + Sampler	INet-RN50	-	-	47.1	61.6	59.9	73.5	83.5	96.0	73.6	86.2
MoLo (Wang et al. 2023c)	INet-RN50	41.9	56.2	55.0	69.6	60.8	77.4	86.0	95.5	74.0	85.6
MGCSM (Yu et al. 2023)	INet-RN50	-	-	-	-	61.3	79.3	86.5	97.1	74.2	88.2
SA-CT (Zhang et al. 2023)	INet-RN50	-	-	48.9	69.1	60.4	78.3	85.4	96.4	71.9	87.1
CLIP-FSAR (Wang et al. 2023a)	CLIP-RN50	52.1	55.8	58.7	62.8	69.4	80.7	92.4	97.0	90.1	92.0
<b>TSAM (Ours)</b>	CLIP-RN50	<b>53.1</b>	<b>58.0</b>	<b>60.2</b>	<b>64.3</b>	<b>72.9</b>	<b>82.8</b>	<b>93.5</b>	<b>97.5</b>	<b>93.0</b>	<b>93.5</b>
SA-CT(ViT) (Zhang et al. 2023)	INet-ViT-B	-	-	-	66.3	-	81.6	-	98.0	-	91.2
CLIP-FSAR (Wang et al. 2023a)	CLIP-ViT-B	54.6	61.8	62.1	72.1	77.1	87.7	97.0	99.1	94.8	95.4
MA-CLIP (Xing et al. 2023a)	CLIP-ViT-B	<u>59.1</u>	<u>64.5</u>	<u>63.3</u>	<u>72.3</u>	<u>83.4</u>	<u>87.9</u>	<u>96.5</u>	<u>98.6</u>	<u>95.7</u>	<u>96.0</u>
<b>TSAM (Ours)</b>	CLIP-ViT-B	<b>60.5</b>	<b>66.7</b>	<b>65.8</b>	<b>74.6</b>	<b>84.5</b>	<b>88.9</b>	<b>98.3</b>	<b>99.3</b>	<b>96.2</b>	<b>97.1</b>

Table 1: Comparison results of the proposed method and the existing FSAR competitors on both 5-way 1-shot and 5-way 5-shot tasks on five benchmarks. The best results are highlighted in bold and the second-best results are underlined. The results of the competitors are from the original literature.




Figure 4: Comparison result of our method and the competitor CLIP-FSAR (Wang et al. 2023a) under cross-dataset evaluation.

fine-tuning. Fig. 5 describes the results vary with the number of adapters added on the three datasets. When dealing with datasets such as UCF101 and Kinetics100, which inherently exhibit high accuracy levels, the addition of more adapters results in only marginal improvements in performance. In contrast, for the HMDB-51 dataset, adding a suitable number of adapters leads to significant gains. Specifically, four adapters are shown to be sufficient to achieve a substantial performance boost, effectively reducing the number of parameters that need to be learned and significantly lowering the computational cost during both inference and training. It is worth noting that the adapters are added from the last module backward, as the later modules in the pre-trained backbone extract more abstract semantic information during video frame encoding, making them more efficient for temporal information extraction.

**Impacts of Each Component.** In Table 2, we investigate the role of each designed module in the TSAM framework on the two datasets. The results show the performance improvements as each module is incrementally added to the baseline model. The baseline configuration, without any of the designed modules (S.P, T.C, or S.N), achieves an accuracy of 54.6% on SSv2-Small and 77.1% on HMDB-51, respectively. When the Sequential Perceiver (S.P) module is




Figure 5: Effect of different partial injection numbers

S.P	T.C	UOTM	SSv2-Small	HMDB-51
X	X	X	54.6	77.1
✓	X	X	58.2	82.2
✓	✓	X	59.0	83.9
✓	X	✓	58.7	82.6
✓	✓	✓	<b>60.2</b>	<b>84.5</b>

Table 2: Effectiveness (%) of each component. We compare the performance of different configurations for the 5-way 1-shot task on the SSv2-Small and HMDB-51 datasets.

incorporated, a noticeable improvement is observed, with the accuracy rising to 58.2% on SSv2-Small and 82.2% on HMDB-51, respectively. Adding the Textual Corpus (T.C) module further improves performance, boosting accuracy to 59.0% on SSv2-Small and 83.9% on HMDB-51, respectively. Similarly, incorporating the Unbalanced Optimal Transport Matching (UOTM) module alongside S.P also results in better performance, with an accuracy of 58.7%

	CLIP-FSAR	MA-CLIP	TSAM
<b>Param.</b>	91M	18M	8M
<b>Acc</b>	53.9	56.5	58.3

Table 3: Trainable parameter comparison (Param.) on SSv2-Small 5-way 1-shot task using ViT-B/32 as the backbone.

Dataset	CLIP-FSAR	TSAM
SSv2-Full	62.1	65.5
SSv2-Full-reverse	62.0	62.3 (↓3.2%)
HMDB-51	77.1	84.5
HMDB-51-reverse	77.1	83.4 (↓1.1%)
Kinetics100	94.7	96.2
Kinetics100-reverse	94.8	95.4 (↓0.8%)

Table 4: Temporal reverse comparison between CLIP-FSAR and our TSAM on three datasets.

on SSv2-Small and 82.6% on HMDB-51. The S.N module helps normalize and harmonize the feature representations, ensuring consistency across different video sequences. Finally, when all three modules (S.P, T.C, and S.N) are combined, the model achieves the highest performance, with an accuracy of 60.2% on SSv2-Small and 84.5% on HMDB-51, respectively. This demonstrates the complementary nature of these modules, with each contributing uniquely to the overall effectiveness of the TSAM framework.

### Further Analysis

**Efficiency of TASM.** Take As shown in Table. 3, TSAM require the least trainable parameters compared to former CLIP-based method while achieving the best performance, showing both efficiency and effectiveness of the method.

**Temporal Reverse Comparison.** In this experiment, we assess the temporal modeling capability of the models by subjecting them to both standard and temporally reversed video inputs for a comparative analysis. Intuitively, model performance should decline significantly when the input videos are reversed, as the actions and temporal relationships are disrupted. While many existing methods claim robust temporal modeling abilities, they often fail to exhibit this robustness under such conditions, as shown in Table 4. The results indicate that TSAM experiences notable performance drops, *e.g.*, 3.2%, 1.1%, and 0.8% on the SSv2-Full, HMDB-51, and Kinetics-100 datasets, respectively. Notably, on datasets like SSv2, which heavily rely on temporal information, TSAM shows a significant decline of 3.2%. This suggests that our design effectively captures sequence-dependent temporal information rather than merely relying on recognizing specific objects within frames. For example, on HMDB-51 and UCF-101 datasets, the performance of the competitors remains relatively stable, failing to demonstrate a notable drop when subjected to temporal reversals.

**t-SNE Visualization.** We visualize the feature distribution changes of CLIP-FSAR and our TSAM during the test stage, as illustrated in Fig. 6. Five classes are randomly sampled

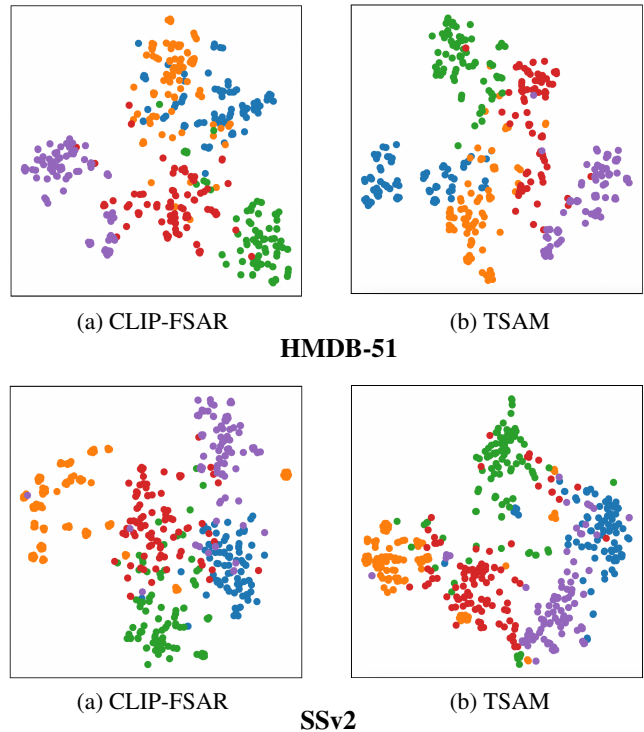


Figure 6: t-SNE visualization of the competitor CLIP-FSAR (Wang et al. 2023a) and our TSAM on both HMDB-51 and SSv2 datasets. For HMDB51 dataset, five classes are ‘fencing’, ‘kick’, ‘kick ball’, ‘pick’ and ‘pour’. For SSv2 dataset, five classes are ‘Approaching [something] with your camera’, ‘Poking a stack of [something] without the stack collapsing’, ‘Pretending to open [something] without actually opening it’, ‘pushing something from right to left’ and ‘showing [something] next to something’.

from the test sets of HMDB-51 and SSv2 for this visualization. By comparing the results, we observe a clear improvement with our method, which yields more compact intra-class distributions and more discriminative inter-class features. Notably, on the challenging SSv2 dataset, our TSAM demonstrates superior performance, resulting in more distinguishable classes and tighter intra-class distributions.

## Conclusion

In this paper, we have proposed a novel FSAR framework named Temporal Sequential-Aware Model (TSAM), utilizing sequential perceiver as an adapter to transfer the pre-trained image backbone to the video encoder. To further improve the few-shot matching process, we expanded the textual corpus to construct multimodal prototypes and unbalanced optimal transport matching metrics to filter extra background noise within video frames. Our framework emphasizes temporal order during the video encoding process and metric matching, greatly outperforms the former method across all FSAR benchmarks, leading to 2.5% improvement on the temporal-challenging SSv2 datasets.

## Acknowledgments

This research was supported in part by the Key R&D Program of Zhejiang Province, China 2024C01G1752215, 2023C01043, Zhejiang Provincial Natural Science Foundation of China under Grant LD24F020016, Science and Technology Innovation 2025 Major Project of Ningbo (2023Z236).

## References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Alteschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Acosta, J. N.; Falcone, G. J.; Rajpurkar, P.; and Topol, E. J. 2022. Multimodal biomedical AI. *Nature Medicine*, 28(9): 1773–1784.
- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 23716–23736.
- Arnab, A.; Dehghani, M.; Heigold, G.; Sun, C.; Lučić, M.; and Schmid, C. 2021. Vivit: A video vision transformer. In *ICCV*, 6836–6846.
- Cao, C.; Zhang, Y.; Yu, Y.; Lv, Q.; Min, L.; and Zhang, Y. 2024. Task-Adapter: Task-specific Adaptation of Image Models for Few-shot Action Recognition. *arXiv preprint arXiv:2408.00249*.
- Cao, K.; Ji, J.; Cao, Z.; Chang, C.-Y.; and Niebles, J. C. 2020. Few-shot video classification via temporal alignment. In *CVPR*, 10618–10627.
- Carreira, J.; and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 6299–6308.
- Dan, J.; Liu, W.; Liu, M.; Xie, C.; Dong, S.; Ma, G.; Tan, Y.; and Xing, J. 2024. HOGDA: Boosting Semi-supervised Graph Domain Adaptation via High-Order Structure-Guided Adaptive Feature Alignment. In *ACM MM*, 11109–11118.
- Dan, J.; Liu, Y.; Xie, H.; Deng, J.; Xie, H.; Xie, X.; and Sun, B. 2023. Transface: Calibrating transformer training for face recognition from a data-centric perspective. In *ICCV*, 20642–20653.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Fu, Y.; Zhang, L.; Wang, J.; Fu, Y.; and Jiang, Y.-G. 2020. Depth guided adaptive meta-fusion network for few-shot video recognition. In *ACM MM*, 1142–1151.
- Gao, P.; Geng, S.; Zhang, R.; Ma, T.; Fang, R.; Zhang, Y.; Li, H.; and Qiao, Y. 2023. Clip-adapter: Better vision-language models with feature adapters. *IJCV*, 1–15.
- Goyal, R.; Ebrahimi Kahou, S.; Michalski, V.; Materzynska, J.; Westphal, S.; Kim, H.; Haenel, V.; Freund, I.; Yianilos, P.; Mueller-Freitag, M.; et al. 2017. The” something something” video database for learning and evaluating visual common sense. In *ICCV*, 5842–5850.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *ICML*, 2790–2799.
- Jaegle, A.; Gimeno, F.; Brock, A.; Vinyals, O.; Zisserman, A.; and Carreira, J. 2021. Perceiver: General perception with iterative attention. In *ICML*, 4651–4664.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning. In *ECCV*, 709–727.
- Kuehne, H.; Jhuang, H.; Garrote, E.; Poggio, T.; and Serre, T. 2011. HMDB: a large video database for human motion recognition. In *ICCV*, 2556–2563.
- Liu, M.; Li, B.; and Yu, Y. 2024. OmniCLIP: Adapting CLIP for Video Recognition with Spatial-Temporal Omni-Scale Feature Learning. *arXiv preprint arXiv:2408.06158*.
- Liu, M.; Wu, F.; Li, B.; Lu, Z.; Yu, Y.; and Li, X. 2024. Envisioning Class Entity Reasoning by Large Language Models for Few-shot Learning. *arXiv preprint arXiv:2408.12469*.
- Lu, S.; Ye, H.-J.; and Zhan, D.-C. 2021a. Few-shot action recognition with compromised metric via optimal transport. *arXiv preprint arXiv:2104.03737*.
- Lu, S.; Ye, H.-J.; and Zhan, D.-C. 2021b. Few-shot action recognition with compromised metric via optimal transport. In *CVPR*.
- Lu, Z.; Shen, F.; Liu, M.; Yu, Y.; and Li, X. 2024. Improving Zero-Shot Generalization for CLIP with Variational Adapter. In *ECCV*.
- Ni, B.; Peng, H.; Chen, M.; Zhang, S.; Meng, G.; Fu, J.; Xiang, S.; and Ling, H. 2022. Expanding language-image pretrained models for general video recognition. In *ECCV*, 1–18. Springer.
- Pan, J.; Lin, Z.; Zhu, X.; Shao, J.; and Li, H. 2022. St-adapter: Parameter-efficient image-to-video transfer learning. In *NeurIPS*.
- Pei, W.; Tan, Q.; Lu, G.; and Tian, J. 2023. D<sup>2</sup> ST-Adapter: Disentangled-and-Deformable Spatio-Temporal Adapter for Few-shot Action Recognition. *arXiv preprint arXiv:2312.01431*.
- Perrett, T.; Masullo, A.; Burghardt, T.; Mirmehdi, M.; and Damen, D. 2021. Temporal-relational crosstransformers for few-shot action recognition. In *CVPR*, 475–484.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. In *NeurIPS*.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Tang, J.; Liu, Q.; Ye, Y.; Lu, J.; Wei, S.; Lin, C.; Li, W.; Mahmood, M. F. F. B.; Feng, H.; Zhao, Z.; Wang, Y.; Liu,

- Y.; Liu, H.; Bai, X.; and Huang, C. 2024. MTVQA: Benchmarking Multilingual Text-Centric Visual Question Answering. *arXiv:2405.11985*.
- Tang, J.; Zhang, W.; Liu, H.; Yang, M.; Jiang, B.; Hu, G.; and Bai, X. 2022. Few could be better than all: Feature sampling and grouping for scene text detection. In *CVPR*, 4563–4572.
- Thatipelli, A.; Narayan, S.; Khan, S.; Anwer, R. M.; Khan, F. S.; and Ghanem, B. 2022a. Spatio-temporal relation modeling for few-shot action recognition. In *CVPR*, 19958–19967.
- Thatipelli, A.; Narayan, S.; Khan, S.; Anwer, R. M.; Khan, F. S.; and Ghanem, B. 2022b. Spatio-temporal relation modeling for few-shot action recognition. In *CVPR*, 19958–19967.
- Tu, S.; Dai, Q.; Wu, Z.; Cheng, Z.-Q.; Hu, H.; and Jiang, Y.-G. 2023. Implicit Temporal Modeling with Learnable Alignment for Video Recognition. In *ICCV*, 19936–19947.
- Wang, X.; Ye, W.; Qi, Z.; Zhao, X.; Wang, G.; Shan, Y.; and Wang, H. 2021. Semantic-guided relation propagation network for few-shot action recognition. In *ACM MM*, 816–825.
- Wang, X.; Zhang, S.; Cen, J.; Gao, C.; Zhang, Y.; Zhao, D.; and Sang, N. 2023a. CLIP-guided Prototype Modulating for Few-shot Action Recognition. *IJCV*.
- Wang, X.; Zhang, S.; Qing, Z.; Gao, C.; Zhang, Y.; Zhao, D.; and Sang, N. 2023b. Molo: Motion-augmented long-short contrastive learning for few-shot action recognition. In *CVPR*, 18011–18021.
- Wang, X.; Zhang, S.; Qing, Z.; Gao, C.; Zhang, Y.; Zhao, D.; and Sang, N. 2023c. Molo: Motion-augmented long-short contrastive learning for few-shot action recognition. In *CVPR*, 18011–18021.
- Wang, X.; Zhang, S.; Qing, Z.; Tang, M.; Zuo, Z.; Gao, C.; Jin, R.; and Sang, N. 2022. Hybrid relation guided set matching for few-shot action recognition. In *CVPR*, 19948–19957.
- Wasim, S. T.; Naseer, M.; Khan, S.; Khan, F. S.; and Shah, M. 2023. Vita-CLIP: Video and text adaptive CLIP via Multimodal Prompting. In *CVPR*, 23034–23044.
- Wu, J.; Zhang, T.; Zhang, Z.; Wu, F.; and Zhang, Y. 2022. Motion-modulated temporal fragment alignment network for few-shot action recognition. In *CVPR*, 9151–9160.
- Wu, Y.; Hu, X.; Sun, Y.; Zhou, Y.; Zhu, W.; Rao, F.; Schiele, B.; and Yang, X. 2024a. Number it: Temporal Grounding Videos like Flipping Manga. *arXiv preprint arXiv:2411.10332*.
- Wu, Y.; Zhu, W.; Cao, J.; Lu, Y.; Li, B.; Chi, W.; Qiu, Z.; Su, L.; Zheng, H.; Wu, J.; et al. 2024b. Video Repurposing from User Generated Content: A Large-scale Dataset and Benchmark. *arXiv preprint arXiv:2412.08879*.
- Xing, J.; Wang, M.; Hou, X.; Dai, G.; Wang, J.; and Liu, Y. 2023a. Multimodal Adaptation of CLIP for Few-Shot Action Recognition. *arXiv preprint arXiv:2308.01532*.
- Xing, J.; Wang, M.; Liu, Y.; and Mu, B. 2023b. Revisiting the spatial and temporal modeling for few-shot action recognition. In *AAAI*, 3001–3009.
- Yang, T.; Zhu, Y.; Xie, Y.; Zhang, A.; Chen, C.; and Li, M. 2023. Aim: Adapting image models for efficient video action recognition. *arXiv preprint arXiv:2302.03024*.
- Yang, X.; Wu, Y.; Yang, M.; Chen, H.; and Geng, X. 2024. Exploring diverse in-context configurations for image captioning. In *NeurIPS*.
- Yu, T.; Chen, P.; Dang, Y.; Huan, R.; and Liang, R. 2023. Multi-Speed Global Contextual Subspace Matching for Few-Shot Action Recognition. In *ACM MM*, 2344–2352.
- Zhan, F.; Yu, Y.; Cui, K.; Zhang, G.; Lu, S.; Pan, J.; Zhang, C.; Ma, F.; Xie, X.; and Miao, C. 2021. Unbalanced feature transport for exemplar-based image translation. In *CVPR*, 15028–15038.
- Zhang, H.; Zhang, L.; Qi, X.; Li, H.; Torr, P. H.; and Koniusz, P. 2020. Few-shot action recognition with permutation-invariant attention. In *ECCV*.
- Zhang, S.; Zhou, J.; and He, X. 2021. Learning implicit temporal alignment for few-shot video classification. *arXiv preprint arXiv:2105.04823*.
- Zhang, Y.; Fu, Y.; Ma, X.; Qi, L.; Chen, J.; Wu, Z.; and Jiang, Y.-G. 2023. On the Importance of Spatial Relations for Few-shot Action Recognition. In *ACM MM*, 2243–2251.
- Zhao, W.; Feng, H.; Liu, Q.; Tang, J.; Wei, S.; Wu, B.; Liao, L.; Ye, Y.; Liu, H.; Zhou, W.; Li, H.; and Huang, C. 2024. TabPedia: Towards Comprehensive Visual Table Understanding with Concept Synergy. In *NeurIPS*.
- Zheng, S.; Chen, S.; and Jin, Q. 2022. Few-shot action recognition with hierarchical matching and contrastive learning. In *ECCV*, 297–313.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Learning to prompt for vision-language models. *IJCV*, 2337–2348.
- Zhu, L.; and Yang, Y. 2018. Compound memory networks for few-shot video classification. In *ECCV*, 751–766.
- Zhu, L.; and Yang, Y. 2020. Label independent memory for semi-supervised few-shot video classification. *T-PAMI*, 44(1): 273–285.