

MoLo: Motion-augmented Long-short Contrastive Learning for Few-shot Action Recognition

Xiang Wang^{1,2*} Shiwei Zhang^{2†} Zhiwu Qing¹ Changxin Gao¹ Yingya Zhang² Deli Zhao² Nong Sang^{1†}

¹Key Laboratory of Image Processing and Intelligent Control,
School of Artificial Intelligence and Automation, Huazhong University of Science and Technology
²Alibaba Group

{wxiang,qzw,cgao,nsang}@hust.edu.cn, {zhangjin.zsw,yingya.zyy}@alibaba-inc.com, zhaodeli@gmail.com

Abstract

Current state-of-the-art approaches for few-shot action recognition achieve promising performance by conducting frame-level matching on learned visual features. However, they generally suffer from two limitations: i) the matching procedure between local frames tends to be inaccurate due to the lack of guidance to force long-range temporal perception; ii) explicit motion learning is usually ignored, leading to partial information loss. To address these issues, we develop a **Motion-augmented Long-short Contrastive Learning (MoLo)** method that contains two crucial components, including a long-short contrastive objective and a motion autodecoder. Specifically, the long-short contrastive objective is to endow local frame features with long-form temporal awareness by maximizing their agreement with the global token of videos belonging to the same class. The motion autodecoder is a lightweight architecture to reconstruct pixel motions from the differential features, which explicitly embeds the network with motion dynamics. By this means, MoLo can simultaneously learn long-range temporal context and motion cues for comprehensive few-shot matching. To demonstrate the effectiveness, we evaluate MoLo on five standard benchmarks, and the results show that MoLo favorably outperforms recent advanced methods. The source code is available at <https://github.com/alibaba-mmai-research/MoLo>.

1. Introduction

Recently, action recognition has achieved remarkable progress and shown broad prospects in many application fields [1, 5, 8, 37, 67]. Despite this, these successes rely heavily on large amounts of manual data annotation, which greatly limits the scalability to unseen categories due to the

* Intern at Alibaba DAMO Academy. † Corresponding authors.

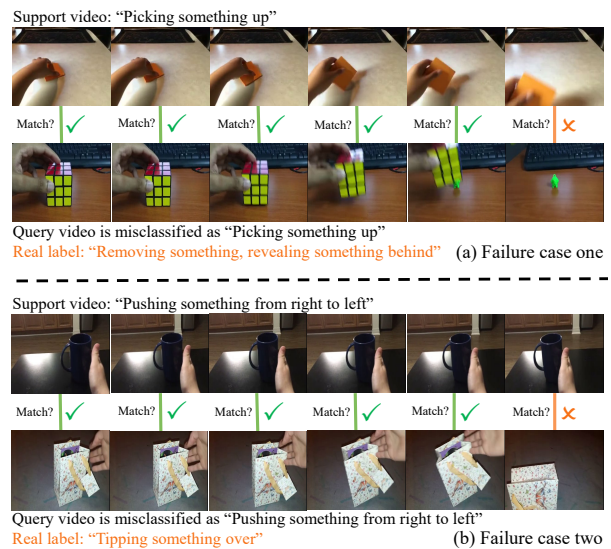


Figure 1. Illustration of our motivation. We show that most existing metric-based local frame matching methods, such as OTAM [4], can be easily perturbed by some similar co-existing video frames due to the lack of forced global context awareness during the support-query temporal alignment process. Example videos come from the commonly used SSv2 dataset [15].

high cost of acquiring large-scale labeled samples. To alleviate the reliance on massive data, few-shot action recognition [88] is a promising direction, aiming to identify novel classes with extremely limited labeled videos.

Most mainstream few-shot action recognition approaches [4, 21, 44, 74] adopt the metric-based meta-learning strategy [61] that learns to map videos into an appropriate feature space and then performs alignment metrics to predict query labels. Typically, OTAM [4] leverages a deep network to extract video features and explicitly estimates an ordered temporal alignment path to match the frames of two videos. HyRSM [74] proposes to explore task-specific semantic correlations across videos and designs a bidirectional Mean Hausdorff Metric (Bi-MHM) to align frames.

Though these works have obtained significant results, there are still two limitations: first, existing standard metric-based techniques mainly focus on local frame-level alignment and are considered limited since the essential global information is not explicitly involved. As shown in Figure 1, local frame-level metrics can be easily affected by co-existing similar video frames. We argue that it would be beneficial to achieve accurate matching if the local frame features can predict the global context in few-shot classification; second, motion dynamics are widely regarded as a vital role in the field of video understanding [6, 22, 25, 42, 50, 70], while the existing few-shot methods do not explicitly explore the rich motion cues between frames for the matching procedure, resulting in a sub-optimal performance. In the literature [5, 67], traditional action recognition works introduce motion information by feeding optical flow or frame difference into an additional deep network, which leads to non-negligible computational overhead. Therefore, an efficient motion compensation method should be introduced to achieve comprehensive few-shot matching.

Inspired by the above observations, we develop a motion-augmented long-short contrastive learning (MoLo) method to jointly model the global contextual information and motion dynamics. More specifically, to explicitly integrate the global context into the local matching process, we apply a long-short contrastive objective to enforce frame features to predict the global context of the videos that belong to the same class. For motion compensation, we design a motion autodecoder to explicitly extract motion features between frame representations by reconstructing pixel motions, *e.g.*, frame differences. In this way, our proposed MoLo enables efficient and comprehensive exploitation of temporal contextual dependencies and motion cues for accurate few-shot action recognition. Experimental results on multiple widely-used benchmarks demonstrate that our MoLo outperforms other advanced few-shot techniques and achieves state-of-the-art performance.

In summary, our contributions can be summarized as follows: (1) We propose a novel MoLo method for few-shot action recognition, aiming to better leverage the global context and motion dynamics. (2) We further design a long-short contrastive objective to reinforce local frame features to perceive comprehensive global information and a motion autodecoder to explicitly extract motion cues. (3) We conduct extensive experiments across five widely-used benchmarks to validate the effectiveness of the proposed MoLo. The results demonstrate that MoLo significantly outperforms baselines and achieves state-of-the-art performance.

2. Related Work

In this section, we will briefly review the works closely related to this paper, including few-shot image classification, motion learning, and few-shot action recognition.

Few-shot image classification. Identifying unseen classes using only a few labels, known as few-shot learning [13], is an important research direction in computer vision. Existing few-shot learning techniques can generally be divided into three groups: augmentation-based, gradient optimization, and metric-based methods. Augmentation-based approaches usually learn to generate samples to alleviate the data scarcity dilemma, mainly including hallucinating additional training examples [7, 17, 52, 77] and adversarial generation [34, 82]. Gradient optimization methods [14, 23, 36, 48, 49, 51] attempt to modify the network optimization process so that the model can be quickly fine-tuned to near the optimal point. For example, MAML [14] explicitly trains the parameters of the network to produce good generalization performance on the current task with a small number of gradient steps. Metric-based paradigm [54] is widely adopted, leveraging a visual encoder to map images into an embedding space and learning a similarity function. This type of methods generally performs few-shot matching to classify query samples through Euclidean distance [43, 54, 76, 80], cosine distance [47, 61], and learnable metrics [16, 56, 78]. Our algorithm also falls into the metric-based line while focusing on taking advantage of temporal context and motion information for accurate few-shot action recognition.

Motion Learning. How to efficiently learn motion cues in videos is a widely researched problem in the community. Previous methods typically exploit optical flow [45, 53, 62, 63, 69, 70, 87] or frame difference [31, 66, 84] to explicitly inject video dynamic cues for action recognition. Due to the high acquisition cost of optical flow [2, 20, 32] and the need for additional networks to extract the motion features [5, 67], recent approaches begin to study how to extract motion features with only raw frame input and achieve promising results [9, 12, 24, 25, 29, 30, 60, 64, 65]. In this paper, we also focus on using raw frames to extract motion information. Different from previous methods, the motion feature in our MoLo is used to facilitate video alignment, and we design a motion autodecoder that learns to recover frame differences to explicitly extract motion dynamics in a unified network without the need for multiple separate networks.

Few-shot action recognition. Compared to images, videos contain richer spatio-temporal context interactions and temporal variations [27, 46, 68, 72]. Existing few-shot action recognition methods [4, 71, 75] mainly belong to the metric-based meta-learning paradigm [54] to temporally match frames, due to its simplicity and effectiveness. Among them, CMN [88, 89] proposes a memory structure and a multi-saliency algorithm to encode variable-length video sequences into fixed-size matrix representations for video matching. OTAM [4] designs a differentiable dynamic time warping [40] algorithm to temporally align two video sequences. ARN [81] builds on a 3D backbone [59] to capture long-range temporal dependencies and learn permutation-




Figure 2. Overview of the proposed MoLo. Given the support and query videos, we first leverage a feature extractor to encode frame features. Subsequently, we feed these frame features into the base head to perform support-query frame-level metric and apply the long-short contrastive objective to improve temporal context awareness during the matching process. Besides, a motion autodecoder is used to recover pixel motions, *i.e.*, frame differences, and the obtained motion features are then entered into a motion head similar to the base head for few-shot matching. Finally, the query video is classified by fusing the matching results of the two heads. For the convenience of illustration, other support videos in a few-shot task are not displayed in the figure.

invariant representations. TRX [44] adopts an attention mechanism to align each query sub-sequence against all sub-sequences in the support set and then fuses their matching scores for classification. ITANet [83] and STRM [58] introduce joint spatio-temporal modeling techniques for robust few-shot matching. HyRSM [74] and MTFAN [79] emphasize the importance of learning task-specific features. Both Nguyen *et al.* [41] and Huang [21] employ multiple similarity functions to compare videos accurately. Despite the remarkable performance, most of these methods are generally limited to local matching without explicitly considering temporal contextual information and motion cues.

3. Method

In this section, we first introduce the problem setup of the few-shot action recognition task. Then, we will elaborate on the proposed MoLo method.

3.1. Problem definition

Given a training set $\mathcal{D}_{train} = \{(v_i, y_i), y_i \in \mathcal{C}_{train}\}$ and a test set $\mathcal{D}_{test} = \{(v_i, y_i), y_i \in \mathcal{C}_{test}\}$, we usually train a deep network on the training set and then verify the performance on the test set. Traditional supervised learning tasks assume that training and testing samples share the same set of classes, *i.e.*, $\mathcal{C}_{train} = \mathcal{C}_{test}$. In contrast, the objective of few-shot action recognition is to identify unseen classes using only a few samples during testing, *i.e.*, $\mathcal{C}_{train} \cap \mathcal{C}_{test} = \emptyset$, and verify the generalization of the trained model. Following previous approaches [4, 44, 74], we adopt the episode-based meta-learning strategy [61] to

optimize the network. In each episode, there is a support set S containing N classes and K videos per class (called the N -way K -shot task) and a query set Q containing query samples to be classified. Few-shot action recognition models are expected to learn to predict labels for videos in the query set with the guidance of the support set. For inference, a large number of episodic tasks will be randomly sampled on the test set, and the average accuracy is used to report the few-shot performance of the learned model.

3.2. MoLo

Overall architecture. An overview of our framework is provided in Figure 2. For clear description, we take the N -way 1-shot task as an example to illustrate our approach. In an N -way 1-shot episodic task, the support set $S = \{s_1, s_2, \dots, s_N\}$ contains N video samples, where $s_i \in \mathbb{R}^{T \times 3 \times H \times W}$ and T is the number of sparsely sampled video frames to obtain the video representation. Given a query video $q \in \mathbb{R}^{T \times 3 \times H \times W}$, the goal is to accurately classify q as a category in the support set. Like previous methods [4, 74], we first adopt a deep feature extraction network [19] to encode the video into feature sequence. Thus we can obtain the support features $F_S = \{f_{s_1}, f_{s_2}, \dots, f_{s_N}\}$ and query feature f_q , where $f_i = \{f_i^1, f_i^2, \dots, f_i^T\}$, $f_i^j \in \mathbb{R}^{C \times H_f \times W_f}$ and C is the channel number. Subsequently, the obtained features are fed into a base head with the long-short contrastive objective for frame matching. Meanwhile, in order to explicitly extract motion information for few-shot matching, we also design a motion autodecoder to reconstruct the frame differences and then input the resulting

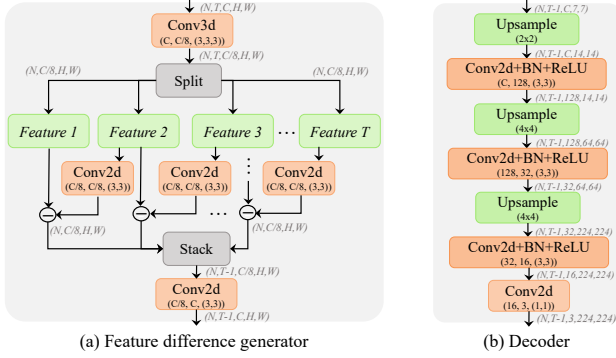


Figure 3. The architecture details of the proposed motion auto-encoder, including a feature difference generator and a decoder. (a) The feature difference generator computes the feature difference between adjacent video frames after a 3D convolution to extract motion information. (b) The decoder involves stacking several up-sampling operations to reconstruct the frame differences.

motion features into a motion head similar to the base head. Based on the output matching scores from the above two heads, we can merge them for query classification.

Long-short contrastive objective. Given the features f_S and f_q , a spatial global-average pooling is applied to collapse the spatial dimension. After that, we append a learnable token to the video feature sequence and send the features to a temporal Transformer for temporal information aggregation. We denote the output features as $\tilde{f}_i = \{\tilde{f}_i^{token}, \tilde{f}_i^1, \dots, \tilde{f}_i^T\}$. The process can be formulated as:

$$\tilde{f}_i = \text{Tformer}([\tilde{f}_i^{token}, \text{GAP}(f_i^1, \dots, f_i^T)] + f_{pos}) \quad (1)$$

where Tformer is the temporal Transformer, $f_i^{token} \in \mathbb{R}^C$ is a learnable token shared cross videos, and $f_{pos} \in \mathbb{R}^{(T+1) \times C}$ means the position embeddings to encode the relative position relationship. Note that the learnable token is used to aggregate global information adaptively, and previous methods [1, 11, 73] have also shown the effectiveness of this manner. In this way, the output feature \tilde{f}_i^{token} corresponding to the learnable token has the global representation ability. We then show how we leverage the global property of the token feature for robust video matching.

Most existing methods [4, 74, 83] usually classify query video sample \tilde{f}_q based on the frame-level alignment metric, which can be expressed as:

$$D(\tilde{f}_i, \tilde{f}_q) = \mathcal{M}([\tilde{f}_i^1, \dots, \tilde{f}_i^T], [\tilde{f}_q^1, \dots, \tilde{f}_q^T]) \quad (2)$$

where \mathcal{M} is a frame-level metric, such as OTAM [4] and Bi-MHM [74]. However, the above metric merely considers the local frame level, ignoring explicitly enhancing global perceptual capabilities of frame features in the matching process. Intuitively, if two videos belong to the same class, the distance between the global feature of one video and the frame feature of the other video is closer than other videos of different categories. We argue that the global context is

much critical for local frame-level matching, especially in the presence of similar co-existing frames. To this end, we leverage a long-short contrastive loss based on [39] to improve the global context awareness of frame features:

$$\mathcal{L}_{LG}^{base} = -\log \frac{\sum_i \text{sim}(\tilde{f}_q^{token}, \tilde{f}_i^i)}{\sum_i \text{sim}(\tilde{f}_q^{token}, \tilde{f}_p^i) + \sum_{j \neq p} \sum_i \text{sim}(\tilde{f}_q^{token}, \tilde{f}_j^i)} - \log \frac{\sum_i \text{sim}(\tilde{f}_q^i, \tilde{f}_p^{token})}{\sum_i \text{sim}(\tilde{f}_q^i, \tilde{f}_p^{token}) + \sum_{j \neq q} \sum_i \text{sim}(\tilde{f}_j^i, \tilde{f}_p^{token})} \quad (3)$$

where \tilde{f}_q is the query feature, \tilde{f}_p is video representations of the same class as the query in an episodic task, and \tilde{f}_j is the negative sample that belongs to the other category. The loss function above is expected to maximize the sum of all positive paired similarity scores and minimize the sum of all negative paired similarity scores. In this way, we explicitly force local features to perceive the global temporal context, enabling more robust few-shot frame alignment.

Motion auto-encoder. To further incorporate motion information into the few-shot matching process, we design a motion auto-encoder consisting of a feature difference generator and a decoder. The motion auto-encoder explicitly extract motion features by operating between frame features. Specifically, we modify the differential operator in [25] to generate features difference, as shown in Figure 3(a). The formula is expressed as follows:

$$f_i^1, \dots, f_i^{T-1} = \mathcal{F}(f_i^1, \dots, f_i^T) \quad (4)$$

where \mathcal{F} means the feature difference generator, and f_i^{j} is the output motion feature. Through the above feature difference generator, we intuitively obtain the motion features between frames. To further explicitly supervise motion feature generation, inspired by MAE [18], we design a decoder (Figure 3(b)) to reconstruct pixel motion dynamics. Subsequently, we input the obtained motion features into a motion head with long-short contrastive objective for video matching. Like the base head, in motion head we first adopt a spatial global-average pooling to motion features and concatenate a learnable token, which are then fed into a temporal Transformer for temporal modeling. We denote the resulting features as $\tilde{f}_i' = \{\tilde{f}_i'^{token}, \tilde{f}_i'^1, \dots, \tilde{f}_i'^{T-1}\}$. For local matching, we apply a frame-level metric to these features:

$$D(\tilde{f}_i', \tilde{f}_q') = \mathcal{M}([\tilde{f}_i'^1, \dots, \tilde{f}_i'^{T-1}], [\tilde{f}_q'^1, \dots, \tilde{f}_q'^{T-1}]) \quad (5)$$

where $D(\tilde{f}_i', \tilde{f}_q')$ is the distance between the support motion feature \tilde{f}_i' and query motion feature \tilde{f}_q' . Similar to \mathcal{L}_{LG}^{base} , we also apply a loss $\mathcal{L}_{LG}^{motion}$ to the motion features.

We express the distance between support and query videos as the weighted sum of the distances obtained by

Method	Reference	SSv2-Full					Kinetics				
		1-shot	2-shot	3-shot	4-shot	5-shot	1-shot	2-shot	3-shot	4-shot	5-shot
MatchingNet [61]	NeurIPS'16	-	-	-	-	-	53.3	64.3	69.2	71.8	74.6
MAML [14]	ICML'17	-	-	-	-	-	54.2	65.5	70.0	72.1	75.3
Plain CMN [88]	ECCV'18	-	-	-	-	-	57.3	67.5	72.5	74.7	76.0
CMN++ [88]	ECCV'18	34.4	-	-	-	43.8	-	-	-	-	-
TRN++ [86]	ECCV'18	38.6	-	-	-	48.9	-	-	-	-	-
TARN [3]	BMVC'19	-	-	-	-	-	64.8	-	-	-	78.5
CMN-J [89]	TPAMI'20	-	-	-	-	-	60.5	70.0	75.6	77.3	78.9
ARN [81]	ECCV'20	-	-	-	-	-	63.7	-	-	-	82.4
OTAM [4]	CVPR'20	42.8	49.1	51.5	52.0	52.3	72.2*	75.9	78.7	81.9	84.2*
ITANet [83]	IJCAI'21	49.2	55.5	59.1	61.0	62.3	73.6	-	-	-	84.3
TRX ($\Omega=\{1\}$) [44]	CVPR'21	38.8	49.7	54.4	58.0	60.6	63.6	75.4	80.1	82.4	85.2
TRX ($\Omega=\{2, 3\}$) [44]	CVPR'21	42.0	53.1	57.6	61.1	64.6	63.6	76.2	81.8	83.4	85.9
TA ² N [35]	AAAI'22	47.6	-	-	-	61.0	72.8	-	-	-	85.8
MTEAN [79]	CVPR'22	45.7	-	-	-	60.4	74.6	-	-	-	87.4
STRM [58]	CVPR'22	43.1	53.3	59.1	61.7	68.1	62.9	76.4	81.1	83.8	86.7
HyRSM [74]	CVPR'22	54.3	<u>62.2</u>	<u>65.1</u>	<u>67.9</u>	69.0	73.7	80.0	<u>83.5</u>	<u>84.6</u>	86.1
Bi-MHM [74]	CVPR'22	44.6*	49.2*	53.1*	54.8*	56.0*	72.3*	77.2*	81.1*	84.1*	84.5*
Nguyen <i>et al.</i> [41]	ECCV'22	43.8	-	-	-	61.1	<u>74.3</u>	-	-	-	87.4
Huang <i>et al.</i> [21]	ECCV'22	49.3	-	-	-	66.7	<u>73.3</u>	-	-	-	86.4
HCL [85]	ECCV'22	47.3	54.5	59.0	62.4	64.9	73.7	79.1	82.4	84.0	85.8
MoLo (OTAM)	-	<u>55.0</u>	61.8	64.8	67.7	69.6	73.8	80.2	83.1	84.2	85.1
MoLo (Bi-MHM)	-	56.6	62.3	67.0	68.5	70.6	74.0	80.4	83.7	84.7	85.6

Table 1. Comparison with recent state-of-the-art few-shot action recognition methods on the SSv2-Full and Kinetics datasets under the 5-way setting. The experimental results are reported as the shot increases from 1 to 5. “-” indicates the result is not available in published works. The best results are bolded and the underline means the second best performance. “*” stands for the results of our implementation.

the base head and motion head:

$$D_{i,q} = D(\tilde{f}_i, \tilde{f}_q) + \alpha D(\tilde{f}'_i, \tilde{f}'_q) \quad (6)$$

where α is a balance coefficient. Then, we can use the output support-query distances as logits for classification. During the training process, the overall framework is trained end-to-end, and the final loss can be denoted as:

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda_1 (\mathcal{L}_{LG}^{base} + \mathcal{L}_{LG}^{motion}) + \lambda_2 \mathcal{L}_{Recons} \quad (7)$$

where \mathcal{L}_{CE} is the cross-entropy loss over the support-query distances based on the ground-truth label, λ_1 and λ_2 are balanced factors, and \mathcal{L}_{Recons} represents the L2 loss of reconstructing frame differences. For few-shot inference, we can leverage the distance values of support-query video pairs in Equation 6 to classify query samples based on the nearest neighbor rule [54], and the decoder can be discarded.

4. Experiments

In this section, we first introduce the experimental settings of our MoLo in Section 4.1 and then compare our approach with previous state-of-the-art methods on multiple commonly used benchmarks in Section 4.2. Finally, comprehensive ablation studies of MoLo are provided in Section 4.3 to demonstrate the effectiveness of each module.

4.1. Datasets and experimental setup

Datasets. Our experiments are conducted on five commonly used few-shot datasets, including SSv2-Full [15], SSv2-Small [15], Kinetics [5], UCF101 [55], and HMDB51 [28]. For SSv2-full and SSv2-Small, we follow

the split settings from [4] and [88], which randomly choose 64 classes from the original dataset [15] as the training set and 24 classes as the test set. The difference between SSv2-Full and SSv2-Small is that the former contains all samples of each category, while the latter selects 100 samples per category. Kinetics used in the few-shot setup [4,74] is also a subset of the original dataset [5]. The UCF101 [55] dataset contains 101 action classes, and we adopt the few-shot split as [74,79,81]. For HMDB51 [28], the common practice is to split 31 classes for training and 10 classes for testing.

Implementation details. For a fair comparison with existing methods [4,44,88], our approach employs a ResNet-50 [19] pre-trained on ImageNet [10] as the basic feature extractor and removes the last global-average pooling layer to retain spatial motion information. We adopt Adam [26] to optimize our MoLo end-to-end and apply an auxiliary semantic loss to stabilize the training process as in [33,38,74,83]. Like previous methods [4,44], 8 frames are uniformly sparsely sampled, i.e. $T = 8$, to obtain the representation of the whole video. We also implement standard data augmentation such as random crop and color jitter in the training stage, and a 224×224 region is cropped at the center of each frame for few-shot testing. For many-shot inference (e.g., 5-shot), we simply embrace the averaged prototype paradigm [54] to classify query video samples. To evaluate few-shot performance on each benchmark, we randomly construct 10,000 episodes from the test set and report the average classification accuracy.

4.2. Comparison with state-of-the-art methods

We compare the performance of MoLo with existing state-of-the-art methods under the 5-way K -shot setting,

Method	Reference	UCF101			SSv2-Small			HMDB51		
		1-shot	3-shot	5-shot	1-shot	3-shot	5-shot	1-shot	3-shot	5-shot
MatchingNet [61]	NeurIPS'16	-	-	-	31.3	39.8	45.5	-	-	-
MAML [14]	ICML'17	-	-	-	30.9	38.6	41.9	-	-	-
Plain CMN [88]	ECCV'18	-	-	-	33.4	42.5	46.5	-	-	-
CMN-J [89]	TPAMI'20	-	-	-	36.2	44.6	48.8	-	-	-
ARN [81]	ECCV'20	66.3	-	83.1	-	-	-	45.5	-	60.6
OTAM [4]	CVPR'20	79.9	87.0	88.9	36.4	45.9	48.0	54.5	65.7	68.0
ITANet [83]	IJCAI'21	-	-	-	39.8	49.4	53.7	-	-	-
TRX [44]	CVPR'21	78.2	92.4	<u>96.1</u>	36.0	51.9	<u>56.7*</u>	53.1	66.8	75.6
TA ² N [35]	AAAI'22	81.9	-	95.1	-	-	-	59.7	-	73.9
MTFAN [79]	CVPR'22	84.8	-	95.1	-	-	-	59.0	-	74.6
STRM [58]	CVPR'22	80.5	92.7	96.9	37.1	49.2	55.3	52.3	67.4	<u>77.3</u>
HyRSM [74]	CVPR'22	83.9	93.0	94.7	40.6	<u>52.3</u>	56.1	<u>60.3</u>	<u>71.7</u>	76.0
Bi-MHM [74]	CVPR'22	81.7*	88.2*	89.3*	38.0*	47.6*	48.9*	58.3*	67.1*	69.0*
Nguyen <i>et al.</i> [41]	ECCV'22	84.9	-	95.9	-	-	-	59.6	-	76.9
Huang <i>et al.</i> [21]	ECCV'22	71.4	-	91.0	38.9	-	61.6	60.1	-	77.0
HCL [85]	ECCV'22	82.5	91.0	93.9	38.7	49.1	55.4	59.1	71.2	76.3
MoLo (OTAM)	-	85.4	93.4	95.1	41.9	50.9	56.2	59.8	71.1	76.1
MoLo (Bi-MHM)	-	86.0	93.5	95.5	42.7	52.9	56.4	60.8	72.0	77.4

Table 2. Comparison with state-of-the-art few-shot action recognition methods on UCF101, SSv2-Small, and HMDB51 in terms of 1-shot, 3-shot, and 5-shot classification accuracy. "-" stands for the result is not available in published works. The best results are bolded in black, and the underline represents the second best result. "*" indicates the results of our implementation.

with K varying from 1 to 5 on SSV2-Full and Kinetics. As displayed in Table 1, two commonly used alignment metrics are adopted to verify the effectiveness of our framework, *i.e.*, OTAM [4] and Bi-MHM [74]. From the results, we can observe the following: (1) Based on Bi-MHM, our MoLo consistently outperforms all prior works on the SSV2-Full dataset under various shot settings and achieves comparable results on Kinetics, which validates our motivation to introduce global information and motion cues for reliable video matching. (2) Compared with the vanilla OTAM, our method can significantly improve performance. For instance, MoLo improves the performance from 42.8% to 55.0% under the 1-shot SSV2-full setting and obtains 1.6% 1-shot performance gain on Kinetics. The same phenomenon holds for Bi-MHM, which demonstrates the robustness of our framework to the type of local frame metrics. (3) The gains of our MoLo on the SSV2-Full dataset are generally more significant than that on the Kinetics dataset. We attribute this to the fact that the SSV2-Full dataset involves a large amount of temporal reasoning, and explicitly introducing global temporal context and motion compensation yields substantial benefits on this dataset, while the Kinetics dataset is more biased towards appearance scenes.

In order to further verify our MoLo, we also compare with state-of-the-art methods on the UCF101, HMDB51, and SSv2-Small datasets, and the results are shown in Table 2. The conclusions are basically consistent with those in Table 1. Notably, our method outperforms other methods in most cases (*e.g.*, 1-shot and 3-shot). However, the performance on 5-shot lags slightly behind the recent advanced methods on UCF101 and SSv2-Small, since our MoLo is a general framework for different shot settings and is not specially designed for 5-shot like TRX [44] and STRM [58].

Long-short contrastive	Autodecoder	Head		SSv2-Full	
		Base	Motion	1-shot	5-shot
		✓		44.6	56.0
			✓	46.3	60.6
✓		✓		52.2	68.0
✓	✓	✓		53.2	68.1
	✓		✓	47.8	61.8
✓	✓		✓	53.9	69.7
		✓	✓	49.2	63.4
✓		✓	✓	53.3	68.2
	✓	✓	✓	53.2	68.1
✓	✓	✓	✓	56.6	70.6

Table 3. Ablation study on SSV2-Full under 5-way 1-shot and 5-way 5-shot settings. The top line represents the baseline Bi-MHM. To avoid confusion, note that the "motion head without autodecoder" setting contains the feature difference generator by default.

4.3. Ablation study

For the convenience of comparison, we choose Bi-MHM [74] as our baseline and add each component into Bi-MHM to verify the rationality of the module design.

Analysis of network components. We present a detailed ablation analysis of model elements, and the experimental results are summarized in Table 3. Based on the results, we can observe that the long-short contrastive objective and motion autodecoder play a key role in boosting performance. In particular, 1-shot performance can be improved from 44.6% to 52.2% by adding the long-short contrastive objective to baseline BiMHM, reflecting the importance of incorporating global context into the local matching process. The performance gain is also significant by introducing the motion autodecoder. In addition, the experimental results show the complementarity among different modules, indicating the rationality of our model structure. By inte-

Setting	SSv2-Full		Kinetics	
	1-shot	5-shot	1-shot	5-shot
Frame Difference	56.6	70.6	74.0	85.6
RAFT Flow [57]	56.8	71.1	74.4	85.9
TRX [44]	42.0	64.6	63.6	85.9
TRX + Motion autodecoder	45.6	66.1	64.8	86.3

Table 4. Comparison of different motion reconstruction targets and incorporating motion autodecoder into existing TRX.




Figure 4. Performance comparison of varying backbone depth on the SSv2-Full dataset under the 5-way K -shot setting. The experiments are carried out with the shot changing from 1 to 5.




Figure 5. Ablation study on the effect of changing the number of input video frames under the 5-way 1-shot SSv2-Full setting.

grating all modules, our method achieves 56.6% and 70.6% performance under 1-shot and 5-shot settings, respectively. In Table 4, we further change the motion reconstruction target from frame difference to RAFT flow [57] and insert the proposed motion autodecoder into TRX [44]. The results indicate the applicability of the proposed MoLo.

Varying backbone depth. The previous comparison results are all based on ResNet-50. To investigate the impact of backbone depth on few-shot performance, we also conduct detailed comparison experiments on ResNet-18 and ResNet-34, as shown in Figure 4. It can be seen that with the deepening of the backbone, the capacity of the model becomes more extensive, and the performance is also improved. The proposed MoLo consistently outperforms existing state-of-the-art techniques, including OTAM, TRX, and HyRSM, revealing the scalability of our approach.

Effect of the number of input frames. For a fair comparison, our MoLo is compared with the current methods under the condition of 8 frames of input. To analyze the impact of the number of input frames on few-shot perfor-

Setting	SSv2-Full		Kinetics	
	1-shot	5-shot	1-shot	5-shot
Temporal Transformer×1	56.6	70.6	74.0	85.6
Temporal Transformer×2	56.4	71.7	72.5	84.9
Temporal Transformer×3	56.0	71.3	71.6	84.2
Temporal Transformer×4	55.9	69.6	71.1	83.9
Temporal Transformer×5	55.8	69.4	70.5	83.3

Table 5. Ablation study for different number of temporal Transformer layers on the SSv2-Full and Kinetics datasets.

Setting	SSv2-Full		Kinetics	
	1-shot	5-shot	1-shot	5-shot
Temporal Transformer-only	53.2	68.1	72.7	84.6
Temporal Transformer w/ TAP	54.8	69.5	73.3	85.2
Temporal Transformer w/ token (MoLo)	56.6	70.6	74.0	85.6

Table 6. Comparison experiments on the effect of learnable token and other variants on the SSv2-Full and Kinetics datasets.

Setting	SSv2-Full		Kinetics	
	1-shot	5-shot	1-shot	5-shot
Within-video	54.8	70.1	73.9	85.5
Cross-video (MoLo)	56.6	70.6	74.0	85.6

Table 7. Comparison experiments of different long-short contrastive styles on SSv2-Full and Kinetics. “Within video” means maximizing the representation similarity between local frame features and global features from the same video.

mance, we conduct experiments with input video sampling from 2 to 10 frames. As shown in Figure 5, the performance starts to rise and gradually saturate with the increase in the number of input frames. Notably, our MoLo consistently maintains leading performance, even outperforming the results of OTAM with 8 frames of input when only 2 frames are used, indicating the robustness of our method.

Different number of temporal Transformer layers. In our framework, the temporal Transformer is adopted to model the temporal long-short associations, and we further explore the impact of the number of Transformer layers on performance. The results are displayed in Table 5, and we notice that the best results are achieved on the Kinetics dataset when one Transformer layer is applied, and overfitting starts to appear as the number of layers increases. On the SSv2-Full dataset, the conclusion is slightly different from the above. We attribute this to the temporal dependencies being more complex on this benchmark, and more Transformer layers may be required for discovering comprehensive relations on 5-shot. For the sake of accuracy and model complexity, we utilize one temporal Transformer layer for temporal relationship modeling in our MoLo.

Analysis of long-short contrastive objective. We conduct ablation experiments to show the superiority of our proposed long-short contrastive objective in Table 6. Among the comparison methods, “Temporal Transformer-only” indicates that the frame features after spatial global average-pooling are directly input into the temporal Transformer and

Method	SSv2-Full						Kinetics					
	5-way	6-way	7-way	8-way	9-way	10-way	5-way	6-way	7-way	8-way	9-way	10-way
OTAM [4]	42.8	38.6	35.1	32.3	30.0	28.2	72.2	68.7	66.0	63.0	61.9	59.0
TRX [44]	42.0	41.5	36.1	33.6	32.0	30.3	63.6	59.4	56.7	54.6	53.2	51.1
HyRSM [74]	54.3	50.1	45.8	44.3	42.1	40.0	73.7	69.5	66.6	65.5	63.4	61.0
MoLo	56.6	51.6	48.1	44.8	42.5	40.3	74.0	69.7	67.4	65.8	63.5	61.3

Table 8. N -way 1-shot classification accuracy comparison with recent few-shot action recognition methods on the test sets of SSv2-Full and Kinetics datasets. The experimental results are reported as the way increases from 5 to 10.




Figure 6. Quantitative analysis of 5-way 1-shot class improvements on the SSv2-Full dataset. Top: class performance gain of adding the long-short contrastive objective to the baseline method; Bottom: class performance gain of the proposed MoLo compared to the baseline. ‘‘S’’ is the abbreviation of ‘‘something.’’

then perform local matching of query-support pairs without long-short contrastive loss. ‘‘Temporal Transformer w/ TAP’’ means that when calculating long-short contrastive loss, we replace the global token features with the global temporal average-pooling features. Note that the above variants also contain the motion autodecoder and two few-shot classification heads. Experimental results show that leveraging the temporal Transformer with a learnable token provides superior performance. This suggests the necessity of enforcing the local frame feature to predict the global temporal context during the matching process and the rationality of utilizing the learnable token to adaptively aggregate global temporal features. In addition, we also explore the effect of the long-short contrastive objective within or between videos in Table 7, and the comparison results demonstrate that the cross-video manner gives more benefits than the within-video form [9, 65]. This can be attributed to the fact that the cross-video manner facilitates obtaining a more compact intra-class feature distribution.

N -way few-shot classification. We further conduct ablation experiments to show the N -way 1-shot accuracy, where N ranges from 5 to 10. In Table 8, we compare MoLo with existing methods, including OTAM, TRX, and HyRSM. It can be observed that the larger N leads to greater difficulty in classification, and the performance decreases. For example, the 10-way 1-shot accuracy of MoLo drops by 14.3% compared to the 5-way 1-shot result on SSv2-Full (40.0%

vs. 54.3%). Notably, our method consistently achieves superior results under various settings, demonstrating the generalizability of the proposed framework.

4.4. Visualization analysis

To analyze the impact of the proposed method on the improvements of specific action classes, we perform statistics on the category performance gains compared to the baseline Bi-MHM [74] in Figure 6. From the results, we can observe that by adding our proposed module, each action category has a certain performance improvement. Notably, after adding the long-short contrastive objective, both ‘‘Picking something up’’ and ‘‘Removing something, revealing something behind’’ improve by more than 15.0%. Similar findings can also be observed in ‘‘Pushing something from right to left’’ and ‘‘Tipping something over.’’ This is consistent with our motivation and fully demonstrates that our method can improve the robustness of local frame matching by enforcing long-range temporal perception. In addition, the performance of each class can be further improved by introducing the motion autodecoder.

4.5. Limitations

The 5-shot results of MoLo on some datasets are not satisfactory, and there are marginal gains on several abstract action classes, e.g., ‘‘Showing that something is empty’’ and ‘‘Pretending to sprinkle air onto something.’’ In the future, we will design a more general method to solve these issues.

5. Conclusion

In this paper, we propose a novel MoLo method for few-shot action recognition that consists of a long-short contrastive objective and a motion autodecoder. The long-short contrastive objective encourages local frame features to predict the global context by contrasting representations. The motion autodecoder is leveraged to recover pixel motions for explicit motion information extraction. In this way, our MoLo enables robust and comprehensive few-shot matching. Extensive experiments on five commonly used benchmarks verify the effectiveness of our method and demonstrate that MoLo achieves state-of-the-art performance.

Acknowledgements. This work is supported by the National Natural Science Foundation of China under grant U22B2053 and Alibaba Group through Alibaba Research Intern Program.

References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. ViVit: A video vision Transformer. In *ICCV*, pages 6836–6846, 2021. 1, 4
- [2] Steven S. Beauchemin and John L. Barron. The computation of optical flow. *ACM computing surveys*, 27(3):433–466, 1995. 2
- [3] Mina Bishay, Georgios Zoumpourlis, and Ioannis Patras. TARN: temporal attentive relation network for few-shot and zero-shot action recognition. In *BMVC*, page 154, 2019. 5
- [4] Kaidi Cao, Jingwei Ji, Zhangjie Cao, Chien-Yi Chang, and Juan Carlos Niebles. Few-shot video classification via temporal alignment. In *CVPR*, pages 10618–10627, 2020. 1, 2, 3, 4, 5, 6, 8
- [5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017. 1, 2, 5
- [6] Jun Cen, Di Luan, Shiwei Zhang, Yixuan Pei, Yingya Zhang, Deli Zhao, Shaojie Shen, and Qifeng Chen. The devil is in the wrongly-classified samples: Towards unified open-set recognition. In *ICLR*, 2023. 2
- [7] Zitian Chen, Yanwei Fu, Yu-Xiong Wang, Lin Ma, Wei Liu, and Martial Hebert. Image deformation meta-networks for one-shot learning. In *CVPR*, pages 8680–8689, 2019. 2
- [8] Dima Damen, Hazel Doughty, Giovanni Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):4125–4141, 2020. 1
- [9] Ishan Dave, Rohit Gupta, Mamshad Nayeem Rizve, and Mubarak Shah. TCLR: Temporal contrastive learning for video representation. *Computer Vision and Image Understanding*, 219:103406, 2022. 2, 8
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 5
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 4
- [12] Linxi Fan, Shyamal Buch, Guanzhi Wang, Ryan Cao, Yuke Zhu, Juan Carlos Niebles, and Li Fei-Fei. RubiksNet: Learnable 3d-shift for efficient video action recognition. In *ECCV*, pages 505–521. Springer, 2020. 2
- [13] Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4):594–611, 2006. 2
- [14] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pages 1126–1135. PMLR, 2017. 2, 5, 6
- [15] Raghav Goyal, Vincent Michalski, Joanna Materzy, Susanne Westphal, Heuna Kim, Valentin Haenel, Peter Yianilos, Moritz Mueller-freitag, Florian Hoppe, Christian Thureau, Ingo Bax, and Roland Memisevic. The “Something” Video Database for Learning and Evaluating Visual Common Sense. In *ICCV*, 2017. 1, 5
- [16] Fusheng Hao, Fengxiang He, Jun Cheng, Lei Wang, Jianzhong Cao, and Dacheng Tao. Collect and select: Semantic alignment metric learning for few-shot learning. In *ICCV*, pages 8460–8469, 2019. 2
- [17] Bharath Hariharan and Ross Girshick. Low-shot visual recognition by shrinking and hallucinating features. In *ICCV*, pages 3018–3027, 2017. 2
- [18] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 4
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 3, 5
- [20] Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. 2
- [21] Yifei Huang, Lijin Yang, and Yoichi Sato. Compound prototype matching for few-shot action recognition. In *ECCV*, pages 351–368, 2022. 1, 3, 5, 6
- [22] Mihir Jain, Hervé Jégou, and Patrick Bouthemy. Better exploiting motion for better action recognition. In *CVPR*, pages 2555–2562, 2013. 2
- [23] Muhammad Abdullah Jamal and Guo-Jun Qi. Task agnostic meta-learning for few-shot learning. In *CVPR*, pages 11719–11727, 2019. 2
- [24] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):221–231, 2012. 2
- [25] Boyuan Jiang, MengMeng Wang, Weihao Gan, Wei Wu, and Junjie Yan. STM: Spatiotemporal and motion encoding for action recognition. In *ICCV*, pages 2000–2009, 2019. 2, 4
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [27] Orit Kliper-Gross, Tal Hassner, and Lior Wolf. One shot similarity metric learning for action recognition. In *SIMBAD*, pages 31–45. Springer, 2011. 2
- [28] H Kuehne, T Serre, H Jhuang, E Garrote, T Poggio, and T Serre. HMDB: A large video database for human motion recognition. In *ICCV*, pages 2556–2563, nov 2011. 5
- [29] Heeseung Kwon, Manjin Kim, Suha Kwak, and Minsu Cho. MotionSqueeze: Neural motion feature learning for video understanding. In *ECCV*, pages 345–362. Springer, 2020. 2
- [30] Heeseung Kwon, Manjin Kim, Suha Kwak, and Minsu Cho. Learning self-similarity in space and time as generalized motion for video action recognition. In *ICCV*, pages 13065–13075, 2021. 2
- [31] Daeho Lee and SeungGwan Lee. Vision-based finger action recognition by angle detection and contour analysis. *ETRI journal*, 33(3):415–422, 2011. 2
- [32] Myunggi Lee, Seungeui Lee, Sungjoon Son, Gyutae Park, and Nojun Kwak. Motion feature network: Fixed motion filter for action recognition. In *ECCV*, pages 387–403, 2018. 2

- [33] Aoxue Li, Tiange Luo, Tao Xiang, Weiran Huang, and Liwei Wang. Few-shot learning with global class representations. In *ICCV*, pages 9715–9724, 2019. [5](#)
- [34] Kai Li, Yulun Zhang, Kunpeng Li, and Yun Fu. Adversarial feature hallucination networks for few-shot learning. In *CVPR*, pages 13470–13479, 2020. [2](#)
- [35] Shuyuan Li, Huabin Liu, Rui Qian, Yuxi Li, John See, Mengjuan Fei, Xiaoyuan Yu, and Weiyao Lin. TA2N: Two-stage action alignment network for few-shot action recognition. In *AAAI*, pages 1404–1411, 2022. [5](#), [6](#)
- [36] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-SGD: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017. [2](#)
- [37] Ji Lin, Chuang Gan, and Song Han. TSM: Temporal shift module for efficient video understanding. In *ICCV*, pages 7083–7093, 2019. [1](#)
- [38] Yongfei Liu, Xiangyi Zhang, Songyang Zhang, and Xuming He. Part-aware prototype network for few-shot semantic segmentation. In *ECCV*, pages 142–158, 2020. [5](#)
- [39] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*, pages 9879–9889, 2020. [4](#)
- [40] Meinard Müller. Dynamic time warping. *Information Retrieval for Music and Motion*, pages 69–84, 2007. [2](#)
- [41] Khoi D Nguyen, Quoc-Huy Tran, Khoi Nguyen, Binh-Son Hua, and Rang Nguyen. Inductive and transductive few-shot video classification via appearance and temporal alignments. In *ECCV*, pages 471–487, 2022. [3](#), [5](#), [6](#)
- [42] Juan Carlos Niebles, Chih-Wei Chen, and Li Fei-Fei. Modeling temporal structure of decomposable motion segments for activity classification. In *ECCV*, pages 392–405. Springer, 2010. [2](#)
- [43] Boris Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. TADAM: Task dependent adaptive metric for improved few-shot learning. *NeurIPS*, 31, 2018. [2](#)
- [44] Toby Perrett, Alessandro Masullo, Tilo Burghardt, Majid Mirmehdi, and Dima Damen. Temporal-relational CrossTransformers for few-shot action recognition. In *CVPR*, pages 475–484, 2021. [1](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [45] AJ Piergiovanni and Michael S Ryoo. Representation flow for action recognition. In *CVPR*, pages 9945–9953, 2019. [2](#)
- [46] Ronald Poppe. A survey on vision-based human action recognition. *Image and Vision Computing*, 28(6):976–990, 2010. [2](#)
- [47] Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan L Yuille. Few-shot image recognition by predicting parameters from activations. In *CVPR*, pages 7229–7238, 2018. [2](#)
- [48] Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *NeurIPS*, volume 32, 2019. [2](#)
- [49] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017. [2](#)
- [50] Yong Rui and P Anandan. Segmenting visual actions based on spatio-temporal motion patterns. In *CVPR*, pages 111–118, 2000. [2](#)
- [51] Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. In *ICLR*, 2019. [2](#)
- [52] Eli Schwartz, Leonid Karlinsky, Joseph Shtok, Sivan Harary, Mattias Marder, Abhishek Kumar, Rogerio Feris, Raja Giryes, and Alex Bronstein. Delta-Encoder: an effective sample synthesis method for few-shot object recognition. *NeurIPS*, 31, 2018. [2](#)
- [53] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *NeurIPS*, 27, 2014. [2](#)
- [54] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *NeurIPS*, 30, 2017. [2](#), [5](#)
- [55] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. *arXiv*, 2012. [5](#)
- [56] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, pages 1199–1208, 2018. [2](#)
- [57] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *ECCV*, pages 402–419. Springer, 2020. [7](#)
- [58] 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 *CVPR*, pages 19958–19967, 2022. [3](#), [5](#), [6](#)
- [59] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015. [2](#)
- [60] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, pages 6450–6459, 2018. [2](#)
- [61] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. *NeurIPS*, 9, 2016. [1](#), [2](#), [3](#), [5](#), [6](#)
- [62] Heng Wang, Dan Oneata, Jakob Verbeek, and Cordelia Schmid. A robust and efficient video representation for action recognition. *International Journal of Computer Vision*, 119(3):219–238, 2016. [2](#)
- [63] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *ICCV*, pages 3551–3558, 2013. [2](#)
- [64] Heng Wang, Du Tran, Lorenzo Torresani, and Matt Feiszli. Video modeling with correlation networks. In *CVPR*, pages 352–361, 2020. [2](#)
- [65] Jue Wang, Gedas Bertasius, Du Tran, and Lorenzo Torresani. Long-short temporal contrastive learning of video Transformers. In *CVPR*, pages 14010–14020, 2022. [2](#), [8](#)
- [66] Limin Wang, Zhan Tong, Bin Ji, and Gangshan Wu. TDN: Temporal difference networks for efficient action recognition. In *CVPR*, pages 1895–1904, 2021. [2](#)
- [67] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment

- networks: Towards good practices for deep action recognition. In *ECCV*, pages 20–36. Springer, 2016. 1, 2
- [68] Xiang Wang, Zhiwu Qing, Ziyuan Huang, Yutong Feng, Shiwei Zhang, Jianwen Jiang, Mingqian Tang, Changxin Gao, and Nong Sang. Proposal relation network for temporal action detection. *arXiv preprint arXiv:2106.11812*, 2021. 2
- [69] Xiang Wang, Zhiwu Qing, Ziyuan Huang, Yutong Feng, Shiwei Zhang, Jianwen Jiang, Mingqian Tang, Yuanjie Shao, and Nong Sang. Weakly-supervised temporal action localization through local-global background modeling. *arXiv preprint arXiv:2106.11811*, 2021. 2
- [70] Xiang Wang, Huaxin Zhang, Shiwei Zhang, Changxin Gao, Yuanjie Shao, and Nong Sang. Context-aware proposal network for temporal action detection. *arXiv preprint arXiv:2206.09082*, 2022. 2
- [71] Xiang Wang, Shiwei Zhang, Jun Cen, Changxin Gao, Yingya Zhang, Deli Zhao, and Nong Sang. CLIP-guided prototype modulating for few-shot action recognition. *arXiv preprint arXiv:2303.02982*, 2023. 2
- [72] Xiang Wang, Shiwei Zhang, Zhiwu Qing, Yuanjie Shao, Changxin Gao, and Nong Sang. Self-supervised learning for semi-supervised temporal action proposal. In *CVPR*, pages 1905–1914, 2021. 2
- [73] Xiang Wang, Shiwei Zhang, Zhiwu Qing, Yuanjie Shao, Zhengrong Zuo, Changxin Gao, and Nong Sang. OadTR: Online action detection with transformers. In *ICCV*, pages 7565–7575, 2021. 4
- [74] 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 *CVPR*, pages 19948–19957, 2022. 1, 3, 4, 5, 6, 8
- [75] Xiang Wang, Shiwei Zhang, Zhiwu Qing, Zhengrong Zuo, Changxin Gao, Rong Jin, and Nong Sang. HyRSM++: Hybrid relation guided temporal set matching for few-shot action recognition. *arXiv preprint arXiv:2301.03330*, 2023. 2
- [76] Yan Wang, Wei-Lun Chao, Kilian Q Weinberger, and Laurens van der Maaten. SimpleShot: Revisiting nearest-neighbor classification for few-shot learning. *arXiv preprint arXiv:1911.04623*, 2019. 2
- [77] Yu-Xiong Wang, Ross Girshick, Martial Hebert, and Bharath Hariharan. Low-shot learning from imaginary data. In *CVPR*, pages 7278–7286, 2018. 2
- [78] Zeyuan Wang, Yifan Zhao, Jia Li, and Yonghong Tian. Cooperative bi-path metric for few-shot learning. In *ACMMM*, pages 1524–1532, 2020. 2
- [79] Jiamin Wu, Tianzhu Zhang, Zhe Zhang, Feng Wu, and Yongdong Zhang. Motion-modulated temporal fragment alignment network for few-shot action recognition. In *CVPR*, pages 9151–9160, 2022. 3, 5, 6
- [80] Sung Whan Yoon, Jun Seo, and Jaekyun Moon. TAPNet: Neural network augmented with task-adaptive projection for few-shot learning. In *ICML*, pages 7115–7123. PMLR, 2019. 2
- [81] Hongguang Zhang, Li Zhang, Xiaojuan Qi, Hongdong Li, Philip HS Torr, and Piotr Koniusz. Few-shot action recognition with permutation-invariant attention. In *ECCV*, pages 525–542. Springer, 2020. 2, 5, 6
- [82] Ruixiang Zhang, Tong Che, Zoubin Ghahramani, Yoshua Bengio, and Yangqiu Song. MetaGAN: An adversarial approach to few-shot learning. *NeurIPS*, 31, 2018. 2
- [83] Songyang Zhang, Jiale Zhou, and Xuming He. Learning implicit temporal alignment for few-shot video classification. In *IJCAI*, 2021. 3, 4, 5, 6
- [84] Zhang Zhang and Dacheng Tao. Slow feature analysis for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(3):436–450, 2012. 2
- [85] Sipeng Zheng, Shizhe Chen, and Qin Jin. Few-shot action recognition with hierarchical matching and contrastive learning. In *ECCV*. Springer, 2022. 5, 6
- [86] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, pages 803–818, 2018. 5
- [87] Guangyu Zhu, Changsheng Xu, Wen Gao, and Qingming Huang. Action recognition in broadcast tennis video using optical flow and support vector machine. In *ECCV*, pages 89–98. Springer, 2006. 2
- [88] Linchao Zhu and Yi Yang. Compound memory networks for few-shot video classification. In *ECCV*, pages 751–766, 2018. 1, 2, 5, 6
- [89] Linchao Zhu and Yi Yang. Label independent memory for semi-supervised few-shot video classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):273–285, 2020. 2, 5, 6