

Hybrid Relation Guided Set Matching for Few-shot Action Recognition

Xiang Wang¹ Shiwei Zhang^{2*} Zhiwu Qing¹ Mingqian Tang² Zhengrong Zuo¹
Changxin Gao¹ Rong Jin² 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,zhrzuo,cgao,nsang}@hust.edu.cn, {zhangjin.zsw,mingqian.tmq,jinrong.jr}@alibaba-inc.com

Abstract

Current few-shot action recognition methods reach impressive performance by learning discriminative features for each video via episodic training and designing various temporal alignment strategies. Nevertheless, they are limited in that (a) learning individual features without considering the entire task may lose the most relevant information in the current episode, and (b) these alignment strategies may fail in misaligned instances. To overcome the two limitations, we propose a novel Hybrid Relation guided Set Matching (HyRSM) approach that incorporates two key components: hybrid relation module and set matching metric. The purpose of the hybrid relation module is to learn task-specific embeddings by fully exploiting associated relations within and cross videos in an episode. Built upon the task-specific features, we reformulate distance measure between query and support videos as a set matching problem and further design a bidirectional Mean Hausdorff Metric to improve the resilience to misaligned instances. By this means, the proposed HyRSM can be highly informative and flexible to predict query categories under the few-shot settings. We evaluate HyRSM on six challenging benchmarks, and the experimental results show its superiority over the state-of-the-art methods by a convincing margin. Project page: <https://hyrsm-cvpr2022.github.io/>.

1. Introduction

Action recognition has been witnessing remarkable progress with the evolution of large-scale datasets [5, 9, 18] and video models [14, 32, 52]. However, this success heavily relies on a large amount of manually labeled examples, which are labor-intensive and time-consuming to collect. It actually limits further applications of this task. Few-shot action recognition is promising in reducing manual annota-

* Corresponding authors.

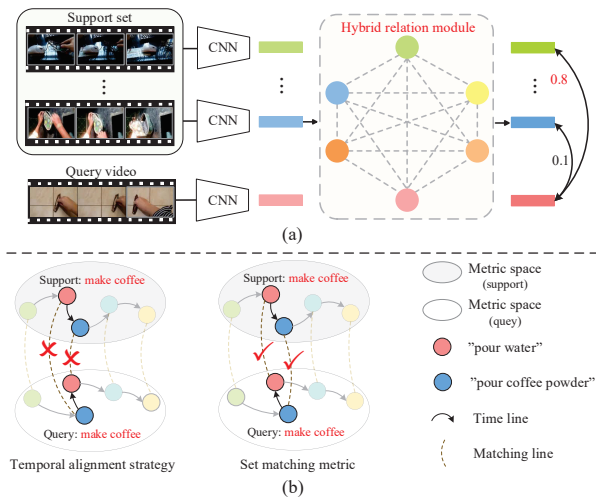


Figure 1. (a) The proposed hybrid relation module. We enhance video representations by extracting relevant discriminative patterns cross videos in an episode, which can adaptively learn task-specific embeddings. (b) Example of *make coffee*, the current temporal alignment metrics tend to be strict, resulting in an incorrect match on misaligned videos. In contrast, the proposed set matching metric is more flexible in finding the best correspondences.

tions and thus has attracted much attention recently [58, 63]. It aims at learning to classify unseen action classes with extremely few annotated examples.

To address the few-shot action recognition problem, current attempts [4, 40, 59, 63] mainly adopt a metric-based meta-learning framework [45] for its simplicity and effectiveness. It first learns a deep embedding space and then designs an explicit or implicit alignment metric to calculate the distances between the query (test) videos and support (reference) videos for classification in an episodic task. For instance, Ordered Temporal Alignment Module (OTAM) [4] extracts features for each video independently and tries to find potential query-support frame pairs only along the ordered temporal alignment path in this feature space. Despite remarkable performance has been reached,

these methods still suffer from two drawbacks. First, discriminative interactive clues cross videos in an episode are ignored when each video is considered independently during representation learning. As a result, these methods actually assume the learned representations are equally effective on different episodic tasks and maintain a fixed set of video features for all test-time tasks, *i.e.*, task-agnostic, which hence might overlook the most discriminative dimensions for the current task. Existing work also shows that the task-agnostic methods tend to suffer inferior generalization in other fields, such as image recognition [29, 56], NLP [35, 38], and information retrieval [33]. Second, actions are usually complicated and involve many subactions with different orders and offsets, which may cause the failure of existing temporal alignment metrics. For example, as shown in Figure 1(b), to *make coffee*, you can *pour water* before *pour coffee powder*, or in a reverse order, hence it is hard for recent temporal alignment strategies to find the right correspondences. Thus a more flexible metric is required to cope with the misalignment.

Inspired by the above observations, we thus propose a novel Hybrid Relation guided Set Matching (HyRSM) algorithm that consists of a hybrid relation module and a set matching metric. In the hybrid relation module, we argue that the considerable relevant relations within and cross videos are beneficial to generate a set of customized features that are discriminative for a given task. To this end, we first apply an intra-relation function to strengthen structural patterns within a video via modeling long-range temporal dependencies. Then an inter-relation function operates on different videos to extract rich semantic information to reinforce the features which are more relevant to query predictions, as shown in Figure 1(a). By this means, we can learn task-specific embeddings for the few-shot task. On top of the hybrid relation module, we design a novel bidirectional Mean Hausdorff Metric to calculate the distances between query and support videos from the set matching perspective. Concretely, we treat each video as a set of frames and alleviate the strictly ordered constraints to acquire better query-support correspondences, as shown in Figure 1(b). In this way, by combining the two components, the proposed HyRSM can sufficiently integrate semantically relational representations within the entire task and provide flexible video matching in an end-to-end manner. We evaluate the proposed HyRSM on six challenging benchmarks and achieve remarkable improvements against current state-of-the-art methods.

Summarily, we make the following three contributions: 1) We propose a novel hybrid relation module to capture the intra- and inter-relations inside the episodic task, yielding task-specific representations for different tasks. 2) We further reformulate the query-support video pair distance metric as a set matching problem and develop a bidirectional

Mean Hausdorff Metric, which can be robust to complex actions. 3) We conduct extensive experiments on six challenging datasets to verify that the proposed HyRSM achieves superior performance over the state-of-the-art methods.

2. Related Work

The work related to this paper includes: few-shot image classification, set matching, and few-shot action recognition. In this section, we will briefly review them separately.

Few-shot Image Classification. Recently, the research of few-shot learning [13] has proceeded roughly along with the following directions: data augmentation, optimization-based, and metric-based. Data augmentation is an intuitive method to increase the number of training samples and improve the diversity of data. Mainstream strategies include spatial deformation [39, 41] and semantic feature augmentation [6, 7]. Optimization-based methods learn a meta-learner model that can quickly adopt to a new task given a few training examples. These algorithms include the LSTM-based meta-learner [44], learning efficient model initialization [15], and learning stochastic gradient descent optimizer [31]. Metric-based methods attempt to address the few-shot classification problem by "learning to compare". This family of approaches aims to learn a feature space and compare query and support images through Euclidean distance [45, 56], cosine similarity [51, 55], or learnable non-linear metric [21, 29, 48]. Our work is more closely related to the metric-based methods [29, 56] that share the same spirit of learning task-specific features, whereas we focus on solving the more challenging few-shot action recognition task with diverse spatio-temporal dependencies. In addition, we will further point out the differences and conduct performance comparisons in the supplementary materials.

Set Matching. The objective of set matching is to accurately measure the similarity of two sets, which have received much attention over the years. Set matching techniques can be used to efficiently process complex data structures [1, 2, 42] and has been applied in many computer vision fields, including face recognition [37, 53, 54], object matching [43, 60], *etc.* Among them, Hausdorff distance is an important alternative to handle set matching problems. Hausdorff distance and its variants have been widely used in the field of image matching and achieved remarkable results [12, 22, 23, 47, 49, 60]. Inspired by these great successes, we introduce set matching into the few-shot action recognition field for the first time.

Few-shot Action Recognition. The difference between few-shot action recognition and the previous few-shot learning approaches is that it deals with more complex higher-dimensional video data instead of two-dimensional images. The existing methods mainly focus on metric-based learn-

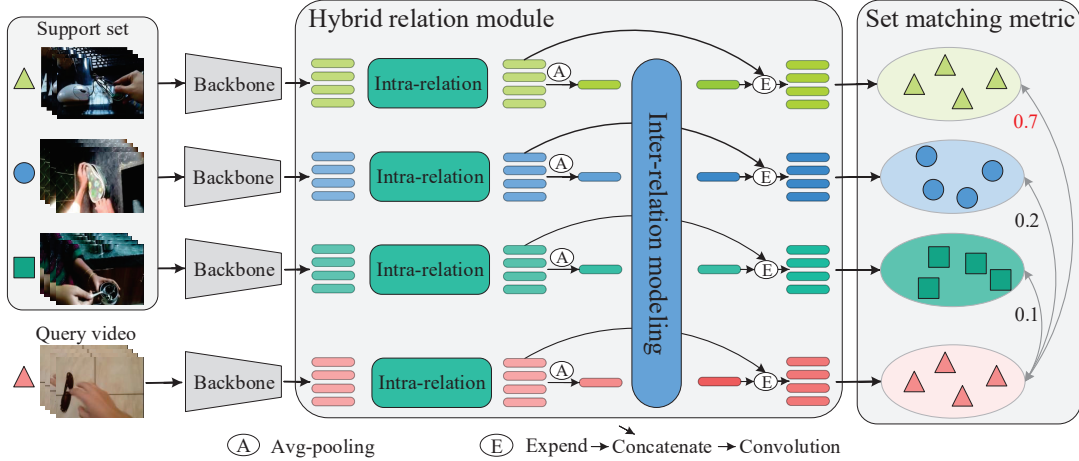


Figure 2. Schematic illustration of the proposed Hybrid Relation guided Set Matching (HyRSM) approach on a 3-way 1-shot problem. Given an episode of video data, a feature embedding network is first employed to extract their feature vectors. A hybrid relation module is then followed to integrate rich information within each video and cross videos with intra-relation and inter-relation functions. Finally, the task-specific features are fed forward into a set matching metric for matching score prediction. Best viewed in color.

ing. OSS-Metric Learning [25] adopts OSS-Metric of video pairs to match videos. TARN [3] learns an attention-based deep-distance measure from an attribute to a class center for zero-shot and few-shot action recognition. CMN [63] utilizes a multi-saliency embedding algorithm to encode video representations. AMeFu-Net [16] uses depth information to assist learning. OTAM [4] preserves the frame ordering in video data and estimates distances with ordered temporal alignment. ARN [58] introduces a self-supervised permutation invariant strategy. ITANet [59] proposes a frame-wise implicit temporal alignment strategy to achieve accurate and robust video matching. TRX [40] matches actions by matching plentiful tuples of different sub-sequences. Note that most above approaches focus on learning video embedding independently. Unlike these previous methods, our HyRSM improves the transferability of embedding by learning intra- and inter-relational patterns that can better generalize to unseen classes.

3. Method

In this section, we first formulate the definition of the few-shot action recognition task. Then we present our Hybrid Relation guided Set Matching (HyRSM) method.

3.1. Problem formulation

The objective of few-shot action recognition is to learn a model that can generalize well to new classes with only a few labeled video samples. To make training more faithful to the test environment, we adopt the episodic training manner [51] for few-shot adaptation as previous work [4, 40, 51, 59]. In each episodic task, there are two sets, *i.e.*, a support set S and a query set Q . The support set S contains $N \times K$ samples from N different action classes, and each class contains K support videos, termed the N -

way K -shot problem. The goal is to classify the query videos in Q into N classes with these support videos.

3.2. HyRSM

Pipeline. The overall architecture of HyRSM is illustrated in Figure 2. For each input video sequence, we first divide it into T segments and extract a snippet from each segment, as in previous methods [4, 52]. This way, in an episodic task, the support set can be denoted as $S = \{s_1, s_2, \dots, s_{N \times K}\}$, where $s_i = \{s_i^1, s_i^2, \dots, s_i^T\}$. For simplicity and convenience, we discuss the process of the N -way 1-shot problem, *i.e.*, $K = 1$, and consider that the query set Q contains a single video q . Then we apply an embedding model to extract the feature representations for each video sequence and obtain the support features $F_s = \{f_{s_1}, f_{s_2}, \dots, f_{s_N}\}$ and the query feature f_q , where $f_{s_i} = \{f_{s_i}^1, f_{s_i}^2, \dots, f_{s_i}^T\}$ and $f_q = \{f_q^1, f_q^2, \dots, f_q^T\}$. After that, we input F_s and f_q to the hybrid relation module to learn task-specific features, resulting in \tilde{F}_s and \tilde{f}_q . Finally, the enhanced representations \tilde{F}_s and \tilde{f}_q are fed into the set matching metric to generate matching scores. Based on the output scores, we can train or test the total framework.

Hybrid relation module. Given the features F_s and f_q output by the embedding network, current approaches, *e.g.*, OTAM [4], directly apply a classifier \mathcal{C} in this feature space. They can be formulated as:

$$y_i = \mathcal{C}(f_{s_i}, f_q) \quad (1)$$

where y_i is the matching score between f_{s_i} and f_q . During training, $y_i = 1$ if they belong to the same class, otherwise $y_i = 0$. In the testing phase, y_i can be adopted to predict the query label. From the perspective of probability theory, it makes decisions based on the priors f_{s_i} and f_q :

$$y_i = \mathcal{P}((f_{s_i}, f_q) | f_{s_i}, f_q) \quad (2)$$

which is a typical task-agnostic method. However, the task-agnostic embedding is often vulnerable to overfit irrelevant representations [21, 29] and may fail to transfer to unseen classes not yet observed in the training stage.

Unlike the previous methods, we propose to learn task-specific features for each target task. To achieve this goal, we introduce a hybrid relation module to generate task-specific features by capturing rich information from different videos in an episode. Specifically, we elaborately design the hybrid relation module \mathcal{H} in the following form:

$$\tilde{f}_i = \mathcal{H}(f_i, \mathcal{G}); f_i \in [F_s, f_q], \mathcal{G} = [F_s, f_q] \quad (3)$$

That is, we improve the feature f_i by aggregating semantic information cross video representations, *i.e.*, \mathcal{G} , in an episodic task, allowing the obtained task-specific feature \tilde{f}_i to be more discriminative than the isolated feature. For efficiency, we further decompose hybrid relation module into two parts: intra-relation function \mathcal{H}_a and inter-relation function \mathcal{H}_e .

The intra-relation function aims to strengthen structural patterns within a video by capturing long-range temporal dependencies. We express this process as:

$$f_i^a = \mathcal{H}_a(f_i) \quad (4)$$

here $f_i^a \in \mathcal{R}^{T \times C}$ is the output of f_i through the intra-relation function and has the same shape as f_i . Note that the intra-relation function has many alternative implements, including multi-head self-attention (MSA), Transformer [50], Bi-LSTM [19], Bi-GRU [8], *etc.*, which is incredibly flexible and can be any one of them.

Based on the features generated by the intra-relation function, an inter-relation function is deployed to semantically enhance the features cross different videos:

$$f_i^e = \mathcal{H}_e^e(f_i^a, \mathcal{G}^a) = \sum_j^{|G^a|} (\kappa(\psi(f_i^a), \psi(f_j^a)) * \psi(f_j^a)) \quad (5)$$

where $\mathcal{G}^a = [F_s^a, f_q^a]$, $\psi(\cdot)$ is a global average pooling layer, and $\kappa(f_i^a, f_j^a)$ is a learnable function that calculates the semantic correlation between f_i^a and f_j^a . The potential logic is that if the correlation score between f_i^a and f_j^a is high, *i.e.*, $\kappa(f_i^a, f_j^a)$, it means they tend to have the same semantic content, hence we can borrow more information from f_j^a to elevate the representation f_i^a , and vice versa. In the same way, if the score $\kappa(f_i^a, f_i^a)$ is less than 1, it indicates that some irrelevant information in f_i^a should be suppressed.

In this way, we can improve the feature discrimination by taking full advantage of the limited samples in each episodic task. The inter-relation function also has similar implements with the intra-relation function but with a

different target. After the inter-relation function, we employ an Expend-Concatenate-Convolution operation to aggregate information, as shown in Figure 2, where the output feature \tilde{f}_i has the same shape as f_i^e . In the form of prior, our method can be formulated as:

$$y_i = \mathcal{P}((\tilde{f}_{s_i}, \tilde{f}_q) | \mathcal{H}(f_{s_i}, \mathcal{G}), \mathcal{H}(f_q, \mathcal{G})); \mathcal{G} = [F_s, f_q] \quad (6)$$

Intuitively, compared with Equation 2, it can be conducive to making better decisions because more priors are provided. In particular, the hybrid relation module is a plug-and-play unit. In the experiment, we will fully explore different configurations of the hybrid relation module and further investigate its insertability.

Set matching metric. Given the relation-enhanced features \tilde{F}_s and \tilde{f}_q , we present a novel metric to enable efficient and flexible matching. In this metric, we treat each video as a set of T frames and reformulate distance measurement between videos as a set matching problem, which is robust to complicated instances, whether they are aligned or not. Specifically, we achieve this goal by modifying the Hausdorff distance, which is a typical set matching approach. The standard Hausdorff distance \mathcal{D} can be formulated as:

$$\begin{aligned} d(\tilde{f}_i, \tilde{f}_q) &= \max_{\tilde{f}_i^a \in \tilde{f}_i} (\min_{\tilde{f}_q^b \in \tilde{f}_q} \|\tilde{f}_i^a - \tilde{f}_q^b\|) \\ d(\tilde{f}_q, \tilde{f}_i) &= \max_{\tilde{f}_q^b \in \tilde{f}_q} (\min_{\tilde{f}_i^a \in \tilde{f}_i} \|\tilde{f}_q^b - \tilde{f}_i^a\|) \\ \mathcal{D} &= \max(d(\tilde{f}_i, \tilde{f}_q), d(\tilde{f}_q, \tilde{f}_i)) \end{aligned} \quad (7)$$

where $\tilde{f}_i \in \mathcal{R}^{T \times C}$ contains T frame features, and $\|\cdot\|$ is a distance measurement function, which is the cosine distance in our method.

However, the previous methods [12, 17, 57, 62] pointed out that Hausdorff distance can be easily affected by noisy examples, resulting in inaccurate measurements. Hence they employ a directed modified Hausdorff distance that robust to noise as follows:

$$d_m(\tilde{f}_i, \tilde{f}_q) = \frac{1}{N_i} \sum_{\tilde{f}_i^a \in \tilde{f}_i} (\min_{\tilde{f}_q^b \in \tilde{f}_q} \|\tilde{f}_i^a - \tilde{f}_q^b\|) \quad (8)$$

where N_i is the length of \tilde{f}_i , and equal to T in this paper. Hausdorff distance and its variants achieve great success in image matching [12, 22, 49] and face recognition [17, 47]. We thus propose to introduce the set matching strategy into the few-shot action recognition field and further design a novel bidirectional Mean Hausdorff Metric (Bi-MHM):

$$\mathcal{D}_b = \frac{1}{N_i} \sum_{\tilde{f}_i^a \in \tilde{f}_i} (\min_{\tilde{f}_q^b \in \tilde{f}_q} \|\tilde{f}_i^a - \tilde{f}_q^b\|) + \frac{1}{N_q} \sum_{\tilde{f}_q^b \in \tilde{f}_q} (\min_{\tilde{f}_i^a \in \tilde{f}_i} \|\tilde{f}_q^b - \tilde{f}_i^a\|) \quad (9)$$

where N_i and N_q are the lengths of the support feature \tilde{f}_i and the query feature \tilde{f}_q respectively.

Method	Reference	Dataset	1-shot	2-shot	3-shot	4-shot	5-shot
CMN++ [63]	ECCV'18	SSv2-Full	34.4	-	-	-	43.8
TRN++ [61]	ECCV'18		38.6	-	-	-	48.9
OTAM [4]	CVPR'20		42.8	49.1	51.5	52.0	52.3
TTAN [30]	ArXiv'21		46.3	52.5	57.3	59.3	60.4
ITANet [4]	IJCAI'21		49.2	55.5	59.1	61.0	62.3
TRX ($\Omega=\{1\}$) [40]	CVPR'21		38.8	49.7	54.4	58.0	60.6
TRX ($\Omega=\{2, 3\}$) [40]	CVPR'21		42.0	53.1	57.6	61.1	64.6
HyRSM	-		54.3 (+5.1)	62.2 (+6.7)	65.1 (+6.0)	67.9 (+6.8)	69.0 (+4.4)
MatchingNet [51]	NeurIPS'16	Kinetics	53.3	64.3	69.2	71.8	74.6
MAML [15]	ICML'17		54.2	65.5	70.0	72.1	75.3
Plain CMN [63]	ECCV'18		57.3	67.5	72.5	74.7	76.0
CMN-J [64]	TPAMI'20		60.5	70.0	75.6	77.3	78.9
TARN [3]	BMVC'19		64.8	-	-	-	78.5
ARN [58]	ECCV'20		63.7	-	-	-	82.4
OTAM [4]	CVPR'20		73.0	75.9	78.7	81.9	85.8
ITANet [59]	IJCAI'21		73.6	-	-	-	84.3
TRX ($\Omega=\{1\}$) [40]	CVPR'21		63.6	75.4	80.1	82.4	85.2
TRX ($\Omega=\{2, 3\}$) [40]	CVPR'21		63.6	76.2	81.8	83.4	85.9
HyRSM	-		73.7 (+0.1)	80.0 (+3.8)	83.5 (+1.7)	84.6 (+1.2)	86.1 (+0.2)
OTAM [4]	CVPR'20		Epic-kitchens	46.0	50.3	53.9	54.9
TRX [40]	CVPR'21	43.4		50.6	53.5	56.8	58.9
HyRSM	-	47.4 (+1.4)		52.9 (+2.3)	56.4 (+2.5)	58.8 (+2.0)	59.8 (+0.9)
ARN [58]	ECCV'20	HMDB51	45.5	-	-	-	60.6
OTAM [4]	CVPR'20		54.5	63.5	65.7	67.2	68.0
TTAN [30]	ArXiv'21		57.1	-	-	-	74.0
TRX [40]	CVPR'21		53.1	62.5	66.8	70.2	75.6
HyRSM	-		60.3 (+3.2)	68.2 (+4.7)	71.7 (+4.9)	75.3 (+5.1)	76.0 (+0.4)

Table 1. Comparison to recent few-shot action recognition methods on the meta-testing set of SSv2-Full, Kinetics, Epic-kitchens and HMDB51. The experiments are conducted under the 5-way setting, and results are reported as the shot increases from 1 to 5. "-" means the result is not available in published works, and the underline indicates the second best result.

The proposed Bi-MHM is a symmetric function, and the two items are complementary to each other. From Equation 9, we can find that \mathcal{D}_b can automatically find the best correspondencies between two videos, *e.g.*, \tilde{f}_i and \tilde{f}_q . Note that our Bi-MHM is a non-parametric classifier and does not involve numerous non-parallel calculations, which helps to improve computing efficiency and transfer ability compared to the previous complex alignment classifiers [4, 40]. Moreover, the hybrid relation module and Bi-MHM can mutually reinforce each other, consolidating the correlation between two videos collectively. In the training phase, we take the negative distance for each class as logit. Then we utilize the same cross-entropy loss as in [4, 40] and the regularization loss [28, 34] to train the model. The regularization loss refers to the cross-entropy loss on the real action classes, which is widely used to improve the training stability and generalization. During inference, we select the support class closest to the query for classification.

4. Experiments

The experiments are designed to answer the following key questions: (1) Is HyRSM competitive to other state-of-the-art methods on challenging few-shot benchmarks? (2) What are the essential components and factors that make HyRSM work? (3) Can the hybrid relation module be utilized as a simple plug-and-play component and bring benefits to existing methods? (4) Does the proposed set matching metric have an advantage over other competitors?

4.1. Datasets and experimental setups

Datasets. We evaluate our method on six few-shot datasets. For the Kinetics [5], SSv2-Full [18], and SSv2-Small [18] datasets, we adopt the existing splits proposed by [4, 40, 59, 63], and each dataset consists of 64 and 24 classes as the meta-training and meta-testing set, respectively. For UCF101 [46] and HMDB51 [27], we evaluate our method by using splits from [40, 58]. In addition, we also use the Epic-kitchens [9, 10] dataset to evaluate HyRSM. Please see the supplementary materials for more details.

Implementation details. Following previous works [4, 40, 59, 63], we utilize ResNet-50 [20] as the backbone which is initialized with ImageNet [11] pre-trained weights. We sparsely and uniformly sample 8 (*i.e.*, $T = 8$) frames per video, as in previous methods [4, 59]. In the training phase, we also adopt basic data augmentation such as random cropping and color jitter, and we use Adam [24] optimizer to train our model. For inference, we conduct few-shot action recognition evaluation on 10000 randomly sampled episodes from the meta-testing set and report the mean accuracy. For many shot classification, *e.g.*, 5-shot, we follow ProtoNet [45] and calculate the mean features of support videos in each class as the prototypes, and classify the query videos according to their distances against the prototypes.

4.2. Comparison with state-of-the-art

We compare the performance of HyRSM with state-of-the-art methods in this section. As shown in Table 1 and

Method	Reference	UCF101			SSv2-Small		
		1-shot	3-shot	5-shot	1-shot	3-shot	5-shot
MatchingNet [51]	NeurIPS'16	-	-	-	31.3	39.8	45.5
MAML [15]	ICML'17	-	-	-	30.9	38.6	41.9
Plain CMN [63]	ECCV'18	-	-	-	33.4	42.5	46.5
CMN-J [64]	TPAMI'20	-	-	-	36.2	44.6	48.8
ARN [58]	ECCV'20	66.3	-	83.1	-	-	-
OTAM [4]	CVPR'20	79.9	87.0	88.9	36.4	45.9	48.0
TTAN [30]	ArXiv'21	<u>80.9</u>	-	93.2	-	-	-
ITANet [59]	IJCAI'21	-	-	-	<u>39.8</u>	49.4	53.7
TRX [40]	CVPR'21	78.2	<u>92.4</u>	96.1	36.0	51.9	59.1
HyRSM	-	83.9 (+3.0)	93.0 (+0.6)	<u>94.7 (-1.4)</u>	40.6 (+0.8)	52.3 (+0.4)	<u>56.1 (-3.0)</u>

Table 2. Results on 1-shot, 3-shot, and 5-shot few-shot classification on the UCF101 and SSv2-Small datasets. "-" means the result is not available in published works, and the underline indicates the second best result.

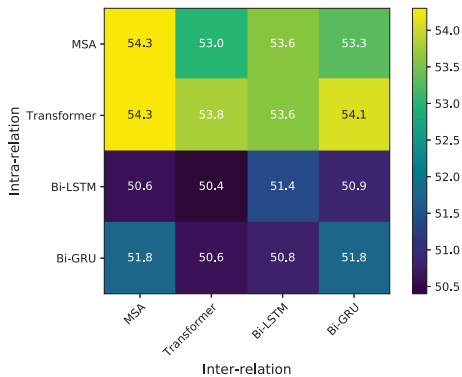


Figure 3. Comparison between different components in hybrid relation module on 5-way 1-shot few-shot action classification. Experiments are conducted on the SSv2-Full dataset.

Table 2, our proposed HyRSM outperforms other methods significantly and achieves new state-of-the-art performance. For instance, HyRSM improves the state-of-the-art performance from 49.2% to 54.3% under the 1-shot setting on SSv2-Full. Specially, compared with the temporal alignment methods [4, 59] and complex fusion methods [30, 40], HyRSM consistently surpasses them under most different shots, which implies that our approach is considerably flexible and efficient. Note that the SSv2-Full and SSv2-Small datasets tend to be motion-based and generally focus on temporal reasoning. While Kinetics and UCF101 are partly appearance-related datasets, and scene understanding is usually important. Besides, Epic-kitchens and HMDB51 are relatively complicated and might involve diverse object interactions. Excellent performance on these datasets reveals that our HyRSM has strong robustness and generalization for different scenes. From Table 2, we observe that HyRSM outperforms current state-of-the-art methods on UCF101 and SSv2-Small under the 1-shot and 3-shot settings, which suggests that our HyRSM can learn rich and effective representations with extremely limited samples. Of note, our HyRSM achieves 94.7% and 56.1% 5-shot performance on UCF101 and SSv2-Small, respectively, which is slightly behind TRX. We attribute this to TRX is an ensemble method specially designed for multiple shots.

Intra-relation	Inter-relation	Bi-MHM	1-shot	5-shot
			35.2	45.3
✓			41.2	55.0
	✓		43.7	55.2
		✓	44.6	56.0
✓	✓		48.1	60.5
	✓	✓	48.3	61.2
✓		✓	51.4	64.6
✓	✓	✓	54.3	69.0

Table 3. Ablation study under 5-way 1-shot and 5-way 5-shot settings on the SSv2-Full dataset.

Method	1-shot	5-shot
OTAM [4]	42.8	52.3
OTAM [4]+ Intra-relation	48.9	60.4
OTAM [4]+ Inter-relation	46.9	57.8
OTAM [4]+ Intra-relation + Inter-relation	51.7	63.9

Table 4. Generalization of hybrid relation module. We conduct experiments on SSv2-Full.

4.3. Ablation study

For ease of comparison, we use a baseline method ProtoNet [45] that applies global-average pooling to backbone representations to obtain a prototype for each class.

Design choices of relation modeling. As shown in Figure 3, we vary the components in the hybrid relation module and systematically evaluate the effect of different variants. The experiments are performed on SSv2-Full under the 5-way 1-shot setting. We can observe that different combinations have quite distinct properties, e.g., multi-head self-attention (MSA) and Transformer are more effective to model intra-class relations than Bi-LSTM and Bi-GRU. Nevertheless, compared with other recent methods [40, 59], the performance of each combination can still be improved, which benefits from the effectiveness of structure design for learning task-specific features. For simplicity, we adopt the same structure to model intra-relation and inter-relation, and we choose multi-head self-attention in the experiments.

Analysis of the proposed components. Table 3 summarizes the effects of each module in HyRSM. We take ProtoNet [45] as our baseline method. From the results, we observe that each component is highly effective. In particular,




Figure 4. N-way 1-shot performance trends of our HyRSM and other state-of-the-art methods with different N on SSv2-Full. The comparison results prove the superiority of our HyRSM




Figure 5. (a) Performance on SSv2-Full using a different number of frames under the 5-way 1-shot setting. (b) The effect of the number of heads on SSv2-Full.

compared to baseline, intra-relation modeling can respectively bring 6% and 9.7% performance gain on 1-shot and 5-shot, and inter-relation function boosts the performance by 8.5% and 9.9% on 1-shot and 5-shot. In addition, the proposed set matching metric improves on 1-shot and 5-shot by 9.4% and 10.7%, respectively, which indicates the ability to find better corresponding frames in the video pair. Moreover, stacking modules can further improve performance, indicating the complementarity between components.

Pluggability of hybrid relation module. In Table 4, we experimentally show that the hybrid relation module generalizes well to other methods by inserting it into the recent OTAM [4]. In this study, OTAM with our hybrid relation module benefits from relational information and finally achieves 8.9% and 11.6% gains on 1-shot and 5-shot. This fully evidences that mining the rich information among videos to learn task-specific features is especially valuable.

N-way few-shot classification. In the previous experiments, all of our comparative evaluation experiments were carried out under the 5-way setting. In order to further explore the influence of different N, in Figure 4, we compare N-way ($N \geq 5$) 1-shot results on SSv2-Full and Kinetics. Results show that as N increases, the difficulty becomes higher, and the performance decreases. Nevertheless, the performance of our HyRSM is still consistently ahead of the recent state-of-the-art OTAM [4] and TRX [40], which shows the feasibility of our method to boost performance by introducing rich relations among videos and the power of the set matching metric.

Varying the number of frames. To demonstrate the scalability of HyRSM, we also explore the impact of different video frame numbers on performance. Of note, previous

Metric	Bi-direction	1-shot	5-shot
Diagonal	-	38.3	48.7
Plain DTW [36]	-	39.6	49.0
OTAM [4]	✗	39.3	47.7
OTAM [4]	✓	<u>42.8</u>	<u>52.3</u>
Bi-MHM (ours)	✓	44.6	56.0

Table 5. Comparison with recent temporal alignment methods on the SSv2-Full dataset under the 5-way 1-shot and 5-way 5-shot settings. Diagonal means matching frame by frame.

Metric	Bi-direction	1-shot	5-shot
Hausdorff distance	✗	32.4	38.2
Hausdorff distance	✓	34.5	39.1
Modified Hausdorff distance	✗	<u>44.2</u>	<u>50.0</u>
Bi-MHM (ours)	✓	44.6	56.0

Table 6. Comparison of different set matching strategies on the SSv2-Full dataset.

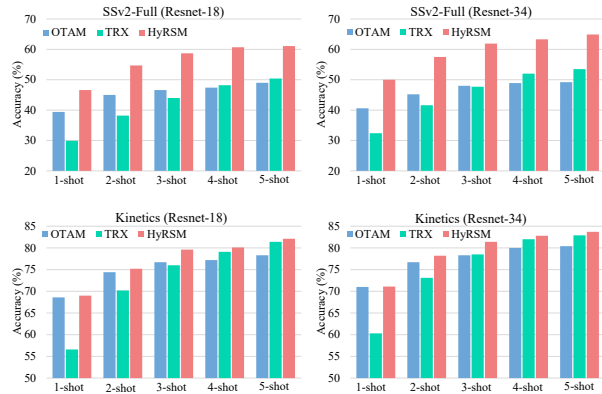


Figure 6. Comparison of the backbone with different depths.

comparisons are performed under 8 frames of input. Results in Figure 5(a) show that as the number of frames increases, the performance improves. HyRSM gradually tends to be saturated when more than 8 frames.

Influence of head number. Previous analyses have shown that multi-head self-attention can focus on different patterns and is critical to capturing diverse features [26]. We investigate the effect of varying the number of heads in multi-head self-attention on performance in Figure 5(b). Results indicate that the effect of multi-head is significant, and the performance starts to saturate beyond a particular point.

Varying depth of the backbone. The previous methods all utilize ResNet-50 as backbone by default for a fair comparison, and the impact of backbone’s depth on performance is still under-explored. As presented in Figure 6, we attempt to answer this question by adopting ResNet-18 and ResNet-34 pre-trained on ImageNet as alternative backbones. Results demonstrate that the deeper network clearly benefits from greater learning capacity and results in better performance. In addition, we notice that our proposed HyRSM consistently outperforms the competitors (*i.e.*, OTAM and TRX), which indicates that our HyRSM is a general framework.

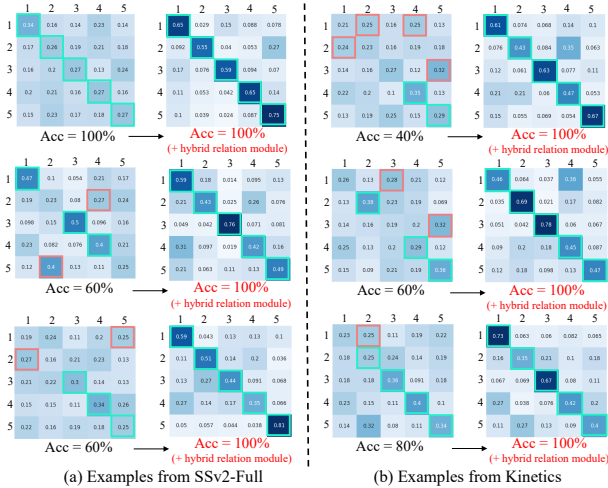


Figure 7. Similarity visualization of how query videos (rows) match to support videos (columns). The boxes of different colors correspond to: **correct match** and **incorrect match**.

4.4. Comparison with other matching approaches

Our proposed set matching metric Bi-MHM aims to accurately find the corresponding video frames between video pairs by relaxing the strict temporal ordering constraints. The following comparative experiments in Table 5 are carried out under the identical experimental setups, *i.e.*, replace the OTAM directly with our Bi-MHM while keeping other settings unchanged. Results show that our Bi-MHM performs well and outperforms other temporal alignment methods (*e.g.*, OTAM). We further analyze different set matching approaches in Table 6, and the results indicate Hausdorff distance is susceptible to noise interference, resulting in the mismatch and relatively poor performance. However, our Bi-MHM shows stability to noise and obtains better performance. Furthermore, compared with the single directional metric, our proposed bidirectional metric is more comprehensive to reflect the actual distances between videos and achieves better performance on few-shot tasks.

4.5. Visualization results

To qualitatively show the discriminative capability of the learned task-specific features in our proposed method, we visualize the similarities between query and support videos with and without the hybrid relation module. As depicted in Figure 7, by adding the hybrid relation module, the discrimination of features is significantly improved, contributing to predicting more accurately. Additionally, the matching results of the set matching metric are visualized in Figure 8, and we can observe that our Bi-MHM is considerably flexible in dealing with alignment and misalignment.

4.6. Limitations

In order to further understand HyRSM, Table 7 illustrates its differences with OTAM and TRX in terms of pa-




Figure 8. Visualization of matching results with the proposed set matching metric on SSV2-Full and Kinetics.

Method	Backbone	Param	FLOPs	Latency	Acc
HyRSM	ResNet-18	13.8M	3.64G	36.5ms	46.6
HyRSM	ResNet-34	23.9M	7.34G	67.5ms	50.0
OTAM [4]	ResNet-50	23.5M	8.17G	116.6ms	42.8
TRX [40]	ResNet-50	47.1M	8.22G	94.6ms	42.0
HyRSM	ResNet-50	65.6M	8.36G	83.5ms	54.3

Table 7. Complexity analysis for 5-way 1-shot SSV2-Full evaluation. The experiments are carried out on one Nvidia V100 GPU.

rameters, computation, and runtime. Notably, HyRSM introduces extra parameters (*i.e.*, hybrid relation module), resulting in increased GPU memory and computational consumption. Nevertheless, without complex non-parallel classifier heads, the whole inference speed of HyRSM is faster than OTAM and TRX. We will further investigate how to reduce complexity with no loss of performance in the future.

5. Conclusion

In this work, we have proposed a hybrid relation guided set matching (HyRSM) approach for few-shot action recognition. Firstly, we design a hybrid relation module to model the rich semantic relevance within one video and cross different videos in an episodic task to generate task-specific features. Secondly, built upon the representative task-specific features, an efficient set matching metric is proposed to be resilient to misalignment and match videos accurately. Experimental results demonstrate that our HyRSM achieves the state-of-the-art performance on the six standard benchmarks, including Kinetics, SSV2-Full, SSV2-Small, HMDB51, UCF101, and Epic-kitchens.

Acknowledgment

This work is supported by the National Natural Science Foundation of China under grant 61871435, Fundamental Research Funds for the Central Universities no.2019kfyXKJC024, 111 Project on Computational Intelligence and Intelligent Control under Grant B18024, and Alibaba Group through Alibaba Research Intern Program.

References

- [1] Yunsheng Bai, Hao Ding, Ken Gu, Yizhou Sun, and Wei Wang. Learning-based efficient graph similarity computation via multi-scale convolutional set matching. In *AAAI*, volume 34, pages 3219–3226, 2020. 2
- [2] Yunsheng Bai, Hao Ding, Yizhou Sun, and Wei Wang. Convolutional set matching for graph similarity. *arXiv preprint arXiv:1810.10866*, 2018. 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. BMVA Press, 2019. 3, 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, 3, 5, 6, 7, 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, 5
- [6] Zitian Chen, Yanwei Fu, Yinda Zhang, Yu-Gang Jiang, Xiangyang Xue, and Leonid Sigal. Semantic feature augmentation in few-shot learning. *arXiv preprint arXiv:1804.05298*, 86:89, 2018. 2
- [7] Zitian Chen, Yanwei Fu, Yinda Zhang, Yu-Gang Jiang, Xiangyang Xue, and Leonid Sigal. Multi-level semantic feature augmentation for one-shot learning. *IEEE Transactions on Image Processing*, 28(9):4594–4605, 2019. 2
- [8] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 4
- [9] 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*, (01):1–1, 2020. 1, 5
- [10] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos, Jian Ma, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Rescaling egocentric vision. *arXiv preprint arXiv:2006.13256*, 2020. 5
- [11] 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
- [12] M-P Dubuisson and Anil K Jain. A modified hausdorff distance for object matching. In *ICPR*, volume 1, pages 566–568. IEEE, 1994. 2, 4
- [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] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, pages 6202–6211, 2019. 1
- [15] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Mearning for Fast Adaptation of Deep Networks. In *ICML*, 2017. 2, 5, 6
- [16] Yuqian Fu, Li Zhang, Junke Wang, Yanwei Fu, and Yu-Gang Jiang. Depth guided adaptive meta-fusion network for few-shot video recognition. In *ACMMM*, pages 1142–1151, 2020. 3
- [17] Yongsheng Gao. Efficiently comparing face images using a modified hausdorff distance. *IEEE Proceedings-Vision, Image and Signal Processing*, 150(6):346–350, 2003. 4
- [18] 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 Something” Video Database for Learning and Evaluating Visual Common Sense. In *ICCV*, 2017. 1, 5
- [19] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *ICASSP*, pages 6645–6649, 2013. 4
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5
- [21] Ruibing Hou, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Cross attention network for few-shot classification. In *NeurIPS*, pages 4003–4014, 2019. 2, 4
- [22] Daniel P Huttenlocher, Gregory A. Klanderman, and William J Rucklidge. Comparing images using the hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):850–863, 1993. 2, 4
- [23] Oliver Jesorsky, Klaus J Kirchberg, and Robert W Frischholz. Robust face detection using the hausdorff distance. In *AVBPA*, pages 90–95. Springer, 2001. 2
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [25] Orit Kliper-Gross, Tal Hassner, and Lior Wolf. One shot similarity metric learning for action recognition. In *SIMBAD*, pages 31–45. Springer, 2011. 3
- [26] Yuma Koizumi, Kohei Yatabe, Marc Delcroix, Yoshiki Masuyama, and Daiki Takeuchi. Speech enhancement using self-adaptation and multi-head self-attention. In *ICASSP*, pages 181–185, 2020. 7
- [27] H Kuehne, T Serre, H Jhuang, E Garrote, T Poggio, and T Serre. HMDB: A large video database for human motion recognition. In *ICCV*, nov 2011. 5
- [28] 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
- [29] Hongyang Li, David Eigen, Samuel Dodge, Matthew Zeiler, and Xiaogang Wang. Finding task-relevant features for few-shot learning by category traversal. In *CVPR*, pages 1–10, 2019. 2, 4
- [30] Shuyuan Li, Huabin Liu, Rui Qian, Yuxi Li, John See, Mengjuan Fei, Xiaoyuan Yu, and Weiyao Lin. Ttan: Two-stage temporal alignment network for few-shot action recognition. *arXiv preprint arXiv:2107.04782*, 2021. 5, 6

- [31] 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](#)
- [32] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, pages 7083–7093, 2019. [1](#)
- [33] Xiaodong Liu, Jianfeng Gao, Xiaodong He, Li Deng, Kevin Duh, and Ye-yi Wang. Representation learning using multi-task deep neural networks for semantic classification and information retrieval. In *NAACL*, pages 912–921, 2015. [2](#)
- [34] 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](#)
- [35] Peng Lu, Ting Bai, and Philippe Langlais. Sc- lstm: Learning task-specific representations in multi-task learning for sequence labeling. In *NAACL*, pages 2396–2406, 2019. [2](#)
- [36] Meinard Müller. Dynamic time warping. *Information Retrieval for Music and Motion*, pages 69–84, 2007. [7](#)
- [37] Masashi Nishiyama, Mayumi Yuasa, Tomoyuki Shibata, Tomokazu Wakasugi, Tomokazu Kawahara, and Osamu Yamaguchi. Recognizing faces of moving people by hierarchical image-set matching. In *CVPR*, pages 1–8, 2007. [2](#)
- [38] Minlong Peng, Qi Zhang, Xiaoyu Xing, Tao Gui, Jinlan Fu, and Xuanjing Huang. Learning task-specific representation for novel words in sequence labeling. 2019. [2](#)
- [39] Luis Perez and Jason Wang. The effectiveness of data augmentation in image classification using deep learning. *arXiv preprint arXiv:1712.04621*, 2017. [2](#)
- [40] 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](#)
- [41] Alexander J Ratner, Henry R Ehrenberg, Zeshan Hussain, Jared Dunnmon, and Christopher Ré. Learning to compose domain-specific transformations for data augmentation. *NeurIPS*, 30:3239, 2017. [2](#)
- [42] Mohammad Rezaei and Pasi Fränti. Set matching measures for external cluster validity. *IEEE Transactions on Knowledge and Data Engineering*, 28(8):2173–2186, 2016. [2](#)
- [43] Yuki Saito, Takuma Nakamura, Hirotaka Hachiya, and Kenji Fukumizu. Exchangeable deep neural networks for set-to-set matching and learning. In *ECCV*, pages 626–646. Springer, 2020. [2](#)
- [44] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *ICML*, pages 1842–1850. PMLR, 2016. [2](#)
- [45] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *NeurIPS*, 30:4077–4087, 2017. [1](#), [2](#), [5](#), [6](#)
- [46] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. *arXiv*, 2012. [5](#)
- [47] N Sudha et al. Robust hausdorff distance measure for face recognition. *Pattern Recognition*, 40(2):431–442, 2007. [2](#), [4](#)
- [48] 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](#)
- [49] Barnabas Takacs. Comparing face images using the modified hausdorff distance. *Pattern recognition*, 31(12):1873–1881, 1998. [2](#), [4](#)
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. [4](#)
- [51] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching Networks for One Shot Learning. In *NeurIPS*, 2016. [2](#), [3](#), [5](#), [6](#)
- [52] 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](#), [3](#)
- [53] Renliang Weng, Jiwen Lu, Junlin Hu, Gao Yang, and Yap-Peng Tan. Robust feature set matching for partial face recognition. In *ICCV*, pages 601–608, 2013. [2](#)
- [54] Renliang Weng, Jiwen Lu, and Yap-Peng Tan. Robust point set matching for partial face recognition. *IEEE Transactions on Image Processing*, 25(3):1163–1176, 2016. [2](#)
- [55] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *CVPR*, pages 8808–8817, 2020. [2](#)
- [56] 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, 2019. [2](#)
- [57] Cheng-Bo Yu, Hua-Feng Qin, Yan-Zhe Cui, and Xiao-Qian Hu. Finger-vein image recognition combining modified hausdorff distance with minutiae feature matching. *Interdisciplinary Sciences: Computational Life Sciences*, 1(4):280–289, 2009. [4](#)
- [58] 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. [1](#), [3](#), [5](#), [6](#)
- [59] Songyang Zhang, Jiale Zhou, and Xuming He. Learning implicit temporal alignment for few-shot video classification. In *IJCAI*, 2021. [1](#), [3](#), [5](#), [6](#)
- [60] Chunjiang Zhao, Wenkang Shi, and Yong Deng. A new hausdorff distance for image matching. *Pattern Recognition Letters*, 26(5):581–586, 2005. [2](#)
- [61] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, pages 803–818, 2018. [5](#)
- [62] Zhi-Qiang Zhou and Bo Wang. A modified hausdorff distance using edge gradient for robust object matching. In *IASP*, pages 250–254. IEEE, 2009. [4](#)
- [63] Linchao Zhu and Yi Yang. Compound memory networks for few-shot video classification. In *ECCV*, pages 751–766, 2018. [1](#), [3](#), [5](#), [6](#)
- [64] Linchao Zhu and Yi Yang. Label independent memory for semi-supervised few-shot video classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (01):1–1, 2020. [5](#), [6](#)