

Learning Implicit Temporal Alignment for Few-shot Video Classification

Songyang Zhang^{1,2,4,*}, Jiale Zhou^{1,*}, Xuming He^{1,3}

¹ShanghaiTech University

²University of Chinese Academy of Sciences

³Shanghai Engineering Research Center of Intelligent Vision and Imaging

⁴Shanghai Institute of Microsystem and Information Technology, Chinese Academy of Sciences

{zhangsy1,zhjl,hexm}@shanghaitech.edu.cn

Abstract

Few-shot video classification aims to learn new video categories with only a few labeled examples, alleviating the burden of costly annotation in real-world applications. However, it is particularly challenging to learn a class-invariant spatial-temporal representation in such a setting. To address this, we propose a novel matching-based few-shot learning strategy for video sequences in this work. Our main idea is to introduce an implicit temporal alignment for a video pair, capable of estimating the similarity between them in an accurate and robust manner. Moreover, we design an effective context encoding module to incorporate spatial and feature channel context, resulting in better modeling of intra-class variations. To train our model, we develop a multi-task loss for learning video matching, leading to video features with better generalization. Extensive experimental results on two challenging benchmarks, show that our method outperforms the prior arts with a sizable margin on Something-Something-V2 and competitive results on Kinetics.

1 Introduction

Video classification, which aims to learn spatio-temporal visual concepts in video, is a fundamental task in computer vision. In particular, recognizing action categories plays an important role in a wide range of applications such as video retrieval, behavior analysis and human-computer interaction. Recently, significant progress has been made in video classification thanks to the powerful representation learned by deep neural networks [Wang *et al.*, 2016; Carreira and Zisserman, 2017a; Wang *et al.*, 2018]. While such fully-supervised approaches are capable of modeling class-specific short-term motion patterns and long-range temporal context, they typically require abundant annotated data. This is not only laborious due to extra temporal dimension, but also restrictive for deploying algorithms in real-world settings, which often need to adapt to novel categories.

*Equal contribution. This work was supported by Shanghai NSF Grant (No. 18ZR1425100). Code and model are available: <https://github.com/tonysy/PyAction>

To tackle the issue of label scarcity, a prevailing strategy, inspired by human visual recognition, is to enable the algorithm to learn to recognize a novel video category with only a few annotated clips. Such a learning task, termed as *few-shot video classification*, has attracted much attention recently [Zhu and Yang, 2018; Cao *et al.*, 2020; Zhang *et al.*, 2020]. Most of these attempts adopt the metric-based meta-learning framework [Vinyals *et al.*, 2016], in which they first match learned features from support and query video clips, and then assign the category label for query video based on the matching scores.

Despite their promising results, those methods often suffer from two drawbacks due to the large intra-class variation of video categories and difficulty in matching two mis-aligned video instances. First, most approaches focus on learning a global video representation by exploiting salient memory mechanism [Zhu and Yang, 2018] or introducing the spatial-temporal attention [Zhang *et al.*, 2020], which rarely explores the temporal alignment for matching videos. Second, while few recent method [Cao *et al.*, 2020] attempts to introduce explicit temporal alignment based on dynamic time warping [Müller, 2007], such matching process needs to calculate a dense frame-wise correspondence, which can be computationally expensive (e.g. in the square of video length).

In this work, we aim to tackle the aforementioned limitations by proposing a novel matching-based few-shot video classification strategy, which enables us to leverage the temporal structure of video instances and avoid dense frame correspondence estimation. Our main idea is to learn an implicit alignment of video sequences based on a factorized self-attention mechanism that enhances video representation along temporal, spatial and feature channel dimensions. Such augmented representation allows us to compute the similarity between two videos by directly summing similarity scores of their corresponding frames. More importantly, our attention-based feature augmentation and similarity metric are class-agnostic and can be learned via a meta-learning framework.

Specifically, we represent each video as a fixed-length sequence of frame features, each encoded as a convolutional feature map. Our first step is to introduce a self-attention on each frame’s feature map and a feature channel attention, which encode per-frame spatial context and a global video context, respectively. To achieve implicit temporal alignment, we further adopt a self-attention mechanism in tem-

poral domain, which augments frame features with their temporal context and thus reduces temporal variation in video instances. We sequentially apply those three steps of attention-based feature augmentation, each of which can be computed efficiently. The final augmented video features are then used for computing a frame-based similarity in video matching.

To perform few-shot learning, we develop a meta-learning strategy that employs a multi-task loss to simultaneously exploit the meta-level and semantic-level supervision. This effectively regularizes the feature learning during the meta-training process. We validate our method on two video classification benchmarks, including Kinetics [Carreira and Zisserman, 2017b] and Something-Something-V2 [Goyal *et al.*, 2017]. We also report an empirical study of different meta-learning strategies with varying loss designs.

The main contributions of this work are summarized as the following:

- We propose a simple and yet effective attention-based representation for few-shot video classification, which leverages temporal structure of videos and enables fast estimation of video similarity.
- Our factorized self-attention mechanism captures the spatial, feature channel and temporal context of a video, leading to an implicit alignment effect in matching.
- We adopt a multi-task learning strategy for few-shot video classification that improves model generalization.
- Our model achieves the state-of-the-art or competitive performance on two challenging benchmarks, Something-Something-V2 and Kinetics.

2 Related Work

Few-shot Learning Inspired by data-efficient learning in human cognition, few-shot learning aims to learn a new concept representation from only a few training examples. Most of existing works can be categorized into optimization-learning based [Finn *et al.*, 2017; Nichol and Schulman, 2018; ?; Munkhdalai and Yu, 2017], metric-learning based [Vinyals *et al.*, 2016; Snell *et al.*, 2017; Sung *et al.*, 2018] and graph neural network based methods [Garcia and Bruna, 2017].

Our work is inspired by the metric-learning based methods. In particular, Matching Networks [Vinyals *et al.*, 2016] and Prototypical Networks [Snell *et al.*, 2017; Liu *et al.*, 2020] learn discriminative feature representations of classes under cosine or Euclidean distance metric. Relation Network [Sung *et al.*, 2018] proposes a relation module as a learnable metric. Our work is in line with the Matching Networks, but we adopt this idea in more challenging video recognition task, focusing on a simple and effective design for measuring video similarity.

Video Action Recognition Video classification methods have evolved from using hand-crafted features [Klaser *et al.*, 2008; Scovanner *et al.*, 2007; Wang and Schmid, 2013] to learning representations via deep networks. Particularly, C3D [Tran *et al.*, 2015] utilizes 3D spatio-temporal convolutional filters to extract deep features from sequences of RGB frames. TSN [Wang *et al.*, 2016] and I3D [Carreira and

Zisserman, 2017a] uses two-stream 2D or 3D CNNs on both RGB and optical flow sequences. Prior works [Wang *et al.*, 2016; Tran *et al.*, 2015] on deep learning based video classification focus on modeling short-term action-specific motion information and long-range temporal context [Wang *et al.*, 2018; Wang *et al.*, 2016; Carreira and Zisserman, 2017a], which typically require abundant annotated data to train their models. Hence it is non-trivial to apply them in the few-shot learning setting directly. Moreover, our work is also related to some recent works aiming to reduce the computation cost of the non-local or self-attention [Chen *et al.*, 2019; Zhang *et al.*, 2019].

Few-shot Video Action Classification Few-shot video classification aims to classify new video classes with only a few annotated examples, and attracted much attention in community recently. The previous approaches mainly focus on the metric-learning based methods. OSS-Metric Learning [Kliper-Gross *et al.*, 2011] measures OSS-Metric of video pairs to enable one-shot video classification. [Mishra *et al.*, 2018] learns a mapping function from an attribute to a class center for zero/few-shot video learning. CMN [Zhu and Yang, 2018] improves the feature representation by introducing a multi-saliency embedding algorithm to encode video frames into a fixed-size matrix. [Zhang *et al.*, 2020] propose a model that captures short-range dependencies with a 3D backbone and discards long-range dependencies via a self-trained permutation invariant attention. These works typically use the global averaged feature for the final matching, which leave the temporal alignment rarely explored.

More recent works have benefited from treating the video matching problem as the sequence matching problem. Inspired from text matching in machine translation [Bahdanau *et al.*, 2014; Wang and Jiang, 2016], TARN [Bishay *et al.*, 2019] utilizes attention mechanisms to perform temporal alignment. OTAM [Cao *et al.*, 2020] suggests preserving the temporal order of video data and proposes an ordered temporal alignment module, which utilizes a differentiable variant of the classical DTW [Müller, 2007] algorithm, to perform the matching. However, this kind of explicit alignment design needs to calculate a dense frame-wise similarity, making it inefficient in the video classification tasks. By contrast, we propose a learnable implicit temporal alignment strategy to achieve accurate and robust video matching.

3 Problem Setup and Terminology

We consider the problem of few-shot video classification, which aims to learn to predict video categories from only a few annotated training video clips per category. To this end, we adopt a meta-learning strategy [Vinyals *et al.*, 2016] to build a meta learner \mathcal{M} that solves similar few-shot video recognition tasks $\mathcal{T} = \{T\}$ sampled from an underlying task distribution P_T .

Formally, each few-shot video classification task T (also called an episode) is composed of a set of support data S with ground-truth labels and a set of query video clips Q . For the C -way K -shot setting, the annotated support data consists of K clip-label pairs from each class, denoted as $S = \{(\mathbf{X}_{c,k}^s, y_{c,k}^s)\}_{k=1, \dots, K}^{c \in \mathcal{C}_T}$, where $\{y_{c,k}\}$ are video-wise

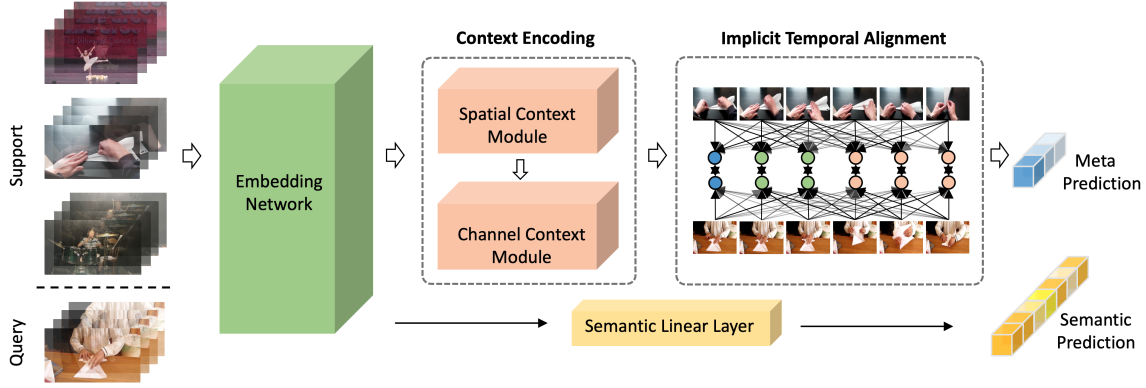


Figure 1: **Illustration of the overall framework.** In training phase, the semantic prediction and meta prediction are used for multi-task learning simultaneously. The semantic prediction is ignored during the inference on the given task.

labels. \mathcal{C}_T is the subset of class sets for the task T and $|\mathcal{C}_T| = C$. Similarly, the query set $Q = \{(\mathbf{X}_j^q, y_j^q)\}_{j=1}^{N_q}$, contains N_q videos from the class set \mathcal{C}_T whose ground-truth labels $\{y_j^q\}$ are provided during training but *unknown* in test.

The meta learner \mathcal{M} aims to learn a functional mapping from the support set \mathcal{S} and a query video \mathbf{X}^q to its category y^q for all the tasks. To achieve this, we construct a training set of video recognition tasks $\mathcal{D}^{tr} = \{(\mathcal{S}_n, \mathcal{Q}_n)\}_{n=1}^{|\mathcal{D}^{tr}|}$ with a class set \mathcal{C}^{tr} , and train the meta learner episodically on the tasks in \mathcal{D}^{tr} . After the meta-training, the model M encodes the knowledge on how to perform video classification on different video categories across tasks. We finally evaluate the learned model on a test set of tasks $\mathcal{D}^{te} = \{(\mathcal{S}_m, \mathcal{Q}_m)\}_{m=1}^{|\mathcal{D}^{te}|}$ whose class set \mathcal{C}^{te} is non-overlapped with \mathcal{C}^{tr} .

4 Our Approach

In this work, we adopt a matching-based few-shot learning framework to build a meta-learner \mathcal{M} for video classification. Our goal is to design a class-agnostic matching strategy capable of measuring the similarity of two video instances in an efficient and robust manner. To achieve this, we develop a video representation based on self-attention mechanism, which provides an implicit temporal alignment effect and yields an efficient matching-based video classifier.

Specifically, we introduce a factorized self-attention mechanism that augment the original frame-based video representation in temporal, spatial and feature channel dimensions. This alleviates the intra-class variations of video instances, particularly along the temporal dimension. Consequently, we are able to compute the similarity of two videos via a direct summation of similarity scores of their corresponding frame-level features and thus bypass explicit video alignment. Such feature augmentation strategy does not depend on specific video class and can be generalized to novel categories.

To this end, we develop a deep conv network, as our meta-learner, to embed each video into its representation and to classify query videos. Our network consists of three main modules: an *embedding network* that computes convolutional feature maps for the video frames within a task (in Sec. 4.1); a *context encoding* network that enrich the feature representa-

tion of support and query videos by encoding the spatial and feature channel context (in Sec. 4.2); and an *implicit temporal alignment network* to align the two videos for accurate and robust matching (in Sec. 4.3).

To train our meta model, we introduce a multi-task learning strategy for the few-shot video recognition that exploits the entire set of known categories for effective learning (in Sec. 4.4). We refer to our model as the **Implicit Temporal Alignment Network (ITANet)**. An overview of our framework is illustrated in Fig. 1 and we will present the model details in the remaining of this section.

4.1 Embedding Network

Given a few-shot task (or episode), the first module of our ITANet is an embedding network that extracts the convolutional feature maps of all frames in each video clip. Following prior work [Cao *et al.*, 2020; Zhu and Yang, 2018], we adopt ResNet [He *et al.*, 2016] as our embedding network. Formally, we denote the embedding network as f_{em} , and compute the feature maps of videos in a task T as $\mathbf{F} = f_{em}(\mathbf{X}), \forall \mathbf{X} \in \mathcal{S} \cup \mathcal{Q}$. Here $\mathbf{F} \in \mathbb{R}^{H_f \times W_f \times n_t \times n_c}$, n_c is the number of feature channels, (H_f, W_f, n_t) is the height, width and temporal length of the feature map.

4.2 Context Encoding Network

To cope with intra-class variation of videos, we first introduce a context encoding network to augment the video features with spatial context in each frame and global context of the entire video. Such a context-aware representation enriches the frame features and focuses more on the discriminative part of the video via attention mechanism. This is particularly important for the few-shot task. Specifically, the context encoding network consists of two components: a *frame-wise spatial context module* and a *video-wise channel context module*, which are described in detail as following:

Frame-wise Spatial Context Module The frame-wise spatial context module augments the feature representation of each frame with its spatial context via a multi-head self-attention. Specifically, we have the feature map of one video clip as $\mathbf{F} = \{\mathbf{f}_1, \dots, \mathbf{f}_{n_t}\}, \mathbf{f}_{n_t} \in \mathbb{R}^{N_f \times n_c}, N_f = H_f \times W_f$. For notation clarity, we focus on a single frame \mathbf{f} here.

We adopt the self-attention mechanism as in [Vaswani *et al.*, 2017], which maps queries and a set of key-value pairs to outputs. Formally, given the query \mathbf{q} , key \mathbf{k} and value \mathbf{v} , we compute the attention as follows:

$$\text{Attention}(\mathbf{q}, \mathbf{k}, \mathbf{v}) = \text{softmax}\left(\frac{\mathbf{q}\mathbf{k}^\top}{\sqrt{d_k}}\right)\mathbf{v} \quad (1)$$

Specifically, we employ the multi-head attention, which linearly projects the frame feature \mathbf{f} by h times with multiple learnable linear transforms to form queries, keys and values of d_q, d_k and d_v dimensions. The output of multi-head attention is generated by concatenating h heads, followed by a linear transform:

$$\text{Head}_i = \text{Attention}(\mathbf{f}\mathbf{W}_i^q, \mathbf{f}\mathbf{W}_i^k, \mathbf{f}\mathbf{W}_i^v), 1 \leq i \leq h \quad (2)$$

$$\text{MultiHead}(\mathbf{f}) = \text{Concat}(\text{Head}_1, \dots, \text{Head}_h)\mathbf{W}^o \quad (3)$$

where $\mathbf{W}_i^q \in \mathbb{R}^{n_c \times d_q}$, $\mathbf{W}_i^k \in \mathbb{R}^{n_c \times d_k}$, $\mathbf{W}_i^v \in \mathbb{R}^{n_c \times d_v}$ are the parameters of linear projections, i is the index of the head, and $\mathbf{W}^o \in \mathbb{R}^{hd_v \times n_c}$, n_c is the feature dimension of the final outputs. We typically set $d_q = d_k = d_v = n_c/h$ and use the residual connection to generate the final spatial context-aware representation as:

$$\mathbf{f}^{sp} = \text{MultiHead}(\mathbf{f}) + \mathbf{f}, \quad \mathbf{F}^{sp} = \{\mathbf{f}_1^{sp}, \dots, \mathbf{f}_{n_t}^{sp}\} \quad (4)$$

Video-wise Channel Context Module Inspired by the SENet [Hu *et al.*, 2018], we develop a video-wise channel context module, which captures the interdependencies between the feature channels of a whole video, as illustrated in Fig. 2. Formally, given feature representation of one video clip $\mathbf{F}^{sp} \in \mathbb{R}^{H_f \times W_f \times n_t \times n_c}$, we first squeeze the global spatial-temporal information into a video-wise channel descriptor as:

$$\mathbf{z} = \frac{1}{H_f \times W_f \times n_t} \sum_{i=1}^{H_f} \sum_{j=1}^{W_f} \sum_{k=1}^{n_t} \mathbf{F}(i, j, k) \in \mathbb{R}^{n_c} \quad (5)$$

Then an adaptive recalibration is applied with excitation operation via a bottleneck structure as following:

$$\mathbf{s}^{ch} = \sigma(\mathbf{W}_2^{ch} \text{ReLU}(\mathbf{W}_1^{ch} \mathbf{z})) \in \mathbb{R}^{n_c} \quad (6)$$

where $\mathbf{W}_2^{ch} \in \mathbb{R}^{n_c \times \frac{n_c}{r}}$ and $\mathbf{W}_1^{ch} \in \mathbb{R}^{\frac{n_c}{r} \times n_c}$ are fully-connected layers, r is the dimensionality reduction ratio. ReLU and σ are the ReLU function and sigmoid function, respectively. With the activation \mathbf{s}^{ch} , the channel context-aware representation are generated by element-wise multiplication along the channel dimension:

$$\tilde{\mathbf{F}} = \{\mathbf{F}_1^{sp} \mathbf{s}_1^{ch}, \dots, \mathbf{F}_{n_t}^{sp} \mathbf{s}_{n_t}^{ch}\} \in \mathbb{R}^{H_f \times W_f \times n_t \times n_c} \quad (7)$$

where $\mathbf{F}_{n_{ch}}^{sp} \in \mathbb{R}^{H_f \times W_f \times n_t}$ is the feature map of a single channel for the entire video.

4.3 Implicit Temporal Alignment Network

After integrating the spatial-channel context, we introduce another feature augmentation along temporal dimension to reduce temporal variation in videos, which enables us to achieve an implicit alignment effect on two videos for similarity estimation. To this end, we design an implicit temporal

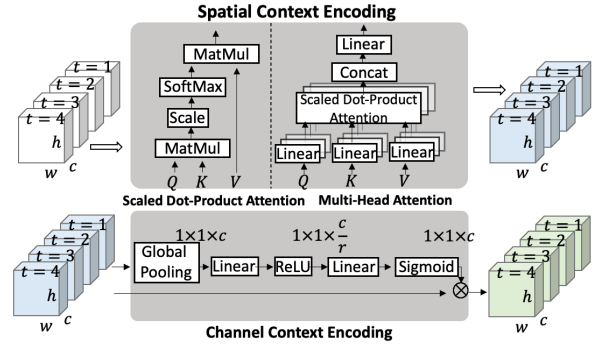


Figure 2: Illustration of the Context Encoding Module.

alignment network for aligning and matching a pair videos, consisting of two modules as follows: a) a *temporal relation encoding module* to enrich each frame with temporal relational context; b) a *similarity metric module* to compute the similarity of a query and a support video.

Temporal Relation Encoding Module We first squeeze the spatial dimension of context encoded feature map $\tilde{\mathbf{F}}$ with average pooling operation to generate frame-level representation as follows:

$$\mathbf{Z} = \text{AvgPool}_{\text{spatial}}(\tilde{\mathbf{F}}) \quad (8)$$

$$\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_{n_t}\} \in \mathbb{R}^{n_t \times n_c} \quad (9)$$

To encode temporal order information, we also add position encoding [Vaswani *et al.*, 2017] to update the frame-level feature \mathbf{Z} as below:

$$\mathbf{z}_{n_t}^p = \mathbf{z}_{n_t} + \mathbf{p}_{n_t}, \quad \mathbf{p} = f_{\text{pos}}(t) \in \mathbb{R}^{n_c} \quad (10)$$

$$\mathbf{Z}^p = \{\mathbf{z}_1^p, \dots, \mathbf{z}_{n_t}^p\} \in \mathbb{R}^{n_t \times n_c} \quad (11)$$

where f_{pos} is position encoding function, which can be cosine/sine function or learnable function.

Given the video representation $\mathbf{Z}^p \in \mathbb{R}^{n_t \times n_c}$, we apply the multi-head self-attention, in a similar form detailed in Eq 2 and Eq. 3, to all the frame of a video to enhance each frame with temporal relational context. We also use the residual structure to produce the final video representation as follows:

$$\tilde{\mathbf{Z}} = \text{MultiHead}(\mathbf{Z}^p) + \mathbf{Z}^p \quad (12)$$

Similarity Metric Module Given our enriched video representation, we measure the similarity of two video clips by direct summing the similarity scores of their corresponding frame features. Such a similarity metric also facilitates learning an implicit temporal alignment for video matching, which is enabled by our multi-dimension feature augmentation.

Concretely, given two videos $\mathbf{X}_i, \mathbf{X}_j$, their augmented representations are denoted as $\tilde{\mathbf{Z}}^i = \{\tilde{\mathbf{z}}_1^i, \dots, \tilde{\mathbf{z}}_{n_t}^i\}$ and $\tilde{\mathbf{Z}}^j = \{\tilde{\mathbf{z}}_1^j, \dots, \tilde{\mathbf{z}}_{n_t}^j\}$. The similarity of these two videos is calculated as following:

$$S(\mathbf{X}_i, \mathbf{X}_j) = \frac{1}{n_t} \sum_{t=1}^{n_t} \frac{(\tilde{\mathbf{z}}_t^i)^\top \tilde{\mathbf{z}}_t^j}{\|\tilde{\mathbf{z}}_t^i\| \|\tilde{\mathbf{z}}_t^j\|} \quad (13)$$

where we use the cosine distance in computing the similarity of each frame pair. We use $S(\mathbf{X}_i, \mathbf{X}_j)$ for the video sequence matching in few-shot classification.

Method	Data Split	Something-Something-V2				
		1-shot	2-shot	3-shot	4-shot	5-shot
RGB*	Small	20.0	25.3	29.2	30.9	33.6
Flow*		21.2	26.0	30.1	31.8	33.8
LSTM*		19.8	24.9	28.6	30.6	32.5
Nearest-FT*		27.5	32.0	35.9	37.8	41.0
Nearest*		28.1	33.3	37.2	39.2	43.8
MatchingNet[Vinyals <i>et al.</i> , 2016]*		31.3	35.9	39.8	40.5	45.5
MAML[Finn <i>et al.</i> , 2017]		30.9	35.1	38.6	40.0	41.9
Plain CMN[Zhu and Yang, 2018]		33.4	38.9	42.5	44.0	46.5
LSTM-EMB*		33.0	38.5	41.8	43.8	46.2
CMN[Zhu and Yang, 2020]		36.2	42.1	44.6	47.0	48.8
ITANet			39.8±0.2	45.8±0.2	49.4±0.4	51.8±0.3
TSN++ [‡]	Full	33.6	-	-	-	43.0
CMN++ [‡]		34.4	-	-	-	43.8
TRN++ [‡]		38.6	-	-	-	48.9
OTAM[Cao <i>et al.</i> , 2020] [‡]		42.8	-	-	-	52.3
ITANet			49.2±0.2	55.5±0.3	59.1±0.2	61.0±0.3

Table 1: **Accuracy on Something-Something-V2 dataset.** We use ResNet-50 as the backbone for fair comparison with previous work. * means the results are copied from [Zhu and Yang, 2020] directly. ‡ means the results are copied from [Cao *et al.*, 2020].

4.4 Meta Learning with Multi-task Loss

To estimate the model parameters, we train our network in the meta-learning framework. Inspired by previous works [Liu *et al.*, 2020; Yan *et al.*, 2019], we propose a multi-task learning paradigm to train the meta model in order to obtain better video representation.

Specifically, our learning objective consists of two losses: a) the per-task classification loss defined on the episode label space, b) the semantic video classification loss on the global video class space \mathcal{C}^{tr} . For the per-task classification, the final prediction for a query video \mathbf{X}^q in C -way- K -shot setting is generated by:

$$\tilde{y}^q = \sum_{i=1}^{CK} S(\mathbf{X}_i^s, \mathbf{X}^q) y_i^s \quad (14)$$

where y_i^s is the label of the i -th support video. For the semantic video classification, we add a linear classification head after the context encoding network to predict the global label \tilde{y}_{sem} for all videos in each task.

We adopt the standard cross-entropy loss in both loss terms for training our model on the meta-training set \mathcal{D}^{tr} , and the loss for each task can be written as:

$$\mathcal{L}_{all} = \mathcal{L}_{meta} + \beta \mathcal{L}_{sem} \quad (15)$$

where $\mathcal{L}_{sem} = \mathcal{L}_{CE}(\tilde{y}_{sem}^s, y_{sem}^s) + \mathcal{L}_{CE}(\tilde{y}_{sem}^q, y_{sem}^q)$ and $\mathcal{L}_{meta} = \mathcal{L}_{CE}(\tilde{y}^q, y^q)$. Here β is the hyper-parameters to balance the ratio of the meta loss and semantic loss. To compare different learning paradigms, we perform extensive experiments and detailed analysis in Sec 5.2.

5 Experiments

In this section, we conduct a series of experiments to validate the effectiveness of our method. Below we first give a brief introduction of experimental configurations and report the quantitative results on two benchmarks in Sec. 5.1. Then we conduct ablation experiments to show the efficacy of our model design in Sec. 5.2.

5.1 Results on SthSth-V2 and Kinetics

Datasets Following previous works, we use the Kinetics [Carreira and Zisserman, 2017b] and Something-Something V2 [Goyal *et al.*, 2017] as the benchmarks. For the Kinetics dataset, we follow the same split as CMN [Zhu and Yang, 2018], which samples 64 classes for meta training, 12 classes for validation, and 24 classes for meta testing. Each class consists of round 100 samples.

For the Something-Something V2 dataset, there exist two splits proposed by [Cao *et al.*, 2020] and [Zhu and Yang, 2020]. The two splits differ in the number of samples for each class. There are about 100 samples per class in the [Zhu and Yang, 2020]’s split, denoted as ”SthV2-*small*” in the following paragraphs. And [Cao *et al.*, 2020]’s split uses all annotated videos of the original Something-Something V2, which is denoted as ”SthV2-*full*” below.

Experimental Configuration We follow the same video preprocessing procedure as OTAM [Cao *et al.*, 2020]. During training, we first resize each frame in the video to 256×256 and then randomly crop a 224×224 region from the video clip. For the Something-Something V2 dataset, as pointed out in [Cao *et al.*, 2020], the dataset is sensitive to concepts of left and right, hence we do not use horizontal flip for this dataset.

Following the experiment settings and learning schedule from [Zhu and Yang, 2018] [Cao *et al.*, 2020], we perform different C -way K -shot experiments on the two datasets, with 95% confidence interval in the meta-test phase. Specifically, the final results are reported over 5 runs and we randomly sample 20,000 episodes for each run.

Quantitative Comparison We compare our algorithm with a wide range of recent methods, as shown in Tab 1 and Tab 2, and report the 5-way performance on 2 benchmarks, including 3 splits: SthV2-*small*, SthV2-*full* and Kinetics-100.

The actions in SthSth-v2 are typically order-sensitive (*e.g.* *pushing something from left to right* and *pushing something from right to left*), and more challenging to recognize, which

Method	5-Way 1-Shot	5-Way 5-Shot
MatchingNet[Vinyals <i>et al.</i> , 2016]	53.3	74.6
MAML[Finn <i>et al.</i> , 2017]	54.2	75.3
LSTM-EMB[Zhu and Yang, 2020]	57.6	76.2
CMN[Zhu and Yang, 2020]	60.5	78.9
ARN[Zhang <i>et al.</i> , 2020]	63.7	82.4
Mean[Cao <i>et al.</i> , 2020]†	67.8	78.9
Plain DTW[Cao <i>et al.</i> , 2020]†	69.2	80.6
OTAM[Cao <i>et al.</i> , 2020]†	73.0	85.8
ITANet	73.6±0.2	84.3±0.3

Table 2: Results on Kinetics-100 dataset. † means the results are copied from [Cao *et al.*, 2020].

requires more accurate and robust temporal alignment for classification. We report the performance of small/full split in Tab 1. We achieve **39.8%/53.7%** for the 5-way 1/5-shot setting on SthV2-*small*, which outperform the state-of-the-art methods by a sizable margin of **3.6** and **4.9**, respectively. On the full split, our method benefits from more labeled samples in meta-train phase, and achieve **49.2%/62.3%** for 5-way 1/5-shot in meta-test phase. The significant improvements demonstrate the effectiveness of our proposed method.

Further, we evaluate our method on the Kinetics dataset, in which the videos are not order-sensitive, and report the performance in Tab 2. The results show our method is also competitive compared with recent works.

5.2 Ablation Study

We conduct several ablative experiments to evaluate the effectiveness of our model on the SthV2-*small*.

Analysis of Model Components We first show the importance of each component used in our model by a series of ablation studies. Specifically, we show the effectiveness of each component under the 5-way 1-shot setting in Tab 3. We can see that the spatial and channel context can improve the baseline method with an improvement of **0.5%** and **1.2%**, respectively, which indicate the frame-level context and video-level context make the representation more discriminative. The proposed implicit temporal alignment contributes to a significant improvement of **5.4%** over the baseline by exploring the temporal order information effectively. The multi-task learning paradigms introduced for few-shot video learning is able to improve the final performance with a sizable margin at **1.9%**, which demonstrates the effectiveness of exploiting semantic information in our method. All components collaborate and complement each other, reaching **39.8%** finally.

The Effect of Learning Paradigms To study the impact of different learning strategies, we conduct several experiments on a MatchingNet baseline and report the results in Tab 4. First, we only use the semantic loss to train the embedding network or the standard meta loss and learning protocol to train the baseline model. Then, we explore learning with semantic loss in the first stage, followed by learning with meta loss in the second stage. Finally, multi-task learning strategy are used to verify the effective of the semantics information used in a joint learning manner. From the experimental results, we find the multi-task learning achieve the best results among the different learning strategy, which indicates it can

Spatial	Channel	ImAlign	MultiTask	5-Way 1-Shot
				33.2±0.3
✓				33.7±0.2
	✓			34.4±0.2
		✓		38.6±0.3
			✓	35.1±0.3
✓	✓	✓	✓	39.8±0.2

Table 3: Ablation of the Model Components on SthV2-*small*. We use MatchingNet as the baseline(our re-implemented matching network without FCE), which is exactly the first row.

Learning Method		5-Way 1-Shot
Stage-1	Stage-2	
Meta Loss	✗	33.2±0.3
Semantic Loss	✗	32.2±0.2
Semantic Loss	Meta Loss	34.4±0.2
Multi Task	✗	35.1±0.3

Table 4: Different Learning paradigms on SthV2-*small*.

TempRelation	PosEncoding	Frame-wise	5-Way 1-Shot
			33.2±0.3
✓			35.0±0.2
✓	✓		35.9±0.3
✓	✓	✓	38.6±0.3

Table 5: Ablation of the Implicit Temporal Alignment on SthV2-*small*. We use MatchingNet as the baseline.

exploit the global semantic label effectively and benefit from the joint optimization.

Design of Implicit Temporal Alignment Here we show the effectiveness of each component of implicit temporal alignment by adding them progressively onto our baseline model. From the Tab 5, we find the temporal relation improves the baseline from 33.2 to 35.0, which indicates the temporal contextual relation helps the similarity measurement. Moreover, adding position encoding enables the implicit alignment process temporal order-aware, and has a gain of 0.9. Finally, we use the frame-wise similarity to reduce the noise in the similarity of global averaged feature, which achieves the further improvement and reach **38.6%**.

6 Conclusion

In this paper, we have presented a novel matching-based method for few-shot video classification. To cope with large intra-class variation in videos, we developed a novel video representation that augments conv features of videos with spatial, feature channel and temporal context based on a factorized self-attention mechanism. This enables us to achieve an implicit temporal alignment in video feature space and adopt a simple frame-wise similarity metric for video matching. We also study a multi-task learning strategy and demonstrate its superiority in few-shot video learning. Our method achieves competitive results on the Kinetics and outperforms the previous works with a sizable margin on the more challenging Something-Something-V2 benchmark.

References

- [Bahdanau *et al.*, 2014] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint*, 2014.
- [Bishay *et al.*, 2019] Mina Bishay, Georgios Zoumpourlis, and Ioannis Patras. Tarn: Temporal attentive relation network for few-shot and zero-shot action recognition. *arXiv preprint*, 2019.
- [Cao *et al.*, 2020] Kaidi Cao, Jingwei Ji, Zhangjie Cao, Chien-Yi Chang, and Juan Carlos Niebles. Few-shot video classification via temporal alignment. In *CVPR*, 2020.
- [Carreira and Zisserman, 2017a] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [Carreira and Zisserman, 2017b] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [Chen *et al.*, 2019] Yunpeng Chen, Marcus Rohrbach, Zhicheng Yan, Yan Shuicheng, Jiashi Feng, and Yannis Kalantidis. Graph-based global reasoning networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [Finn *et al.*, 2017] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint*, 2017.
- [Garcia and Bruna, 2017] Victor Garcia and Joan Bruna. Few-shot learning with graph neural networks. *arXiv preprint*, 2017.
- [Goyal *et al.*, 2017] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The” something something” video database for learning and evaluating visual common sense. In *ICCV*, 2017.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [Hu *et al.*, 2018] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.
- [Klaser *et al.*, 2008] Alexander Klaser, Marcin Marszałek, and Cordelia Schmid. A spatio-temporal descriptor based on 3d-gradients. In *BMVC 2008-19th British Machine Vision Conference*, pages 275–1. British Machine Vision Association, 2008.
- [Kliper-Gross *et al.*, 2011] Orit Kliper-Gross, Tal Hassner, and Lior Wolf. One shot similarity metric learning for action recognition. In *International Workshop on Similarity-Based Pattern Recognition*, 2011.
- [Liu *et al.*, 2020] Yongfei Liu, Xiangyi Zhang, Songyang Zhang, and Xuming He. Part-aware prototype network for few-shot semantic segmentation. In *European Conference on Computer Vision*, 2020.
- [Mishra *et al.*, 2018] Ashish Mishra, Vinay Kumar Verma, M Shiva Krishna Reddy, S Arulkumar, Piyush Rai, and Anurag Mittal. A generative approach to zero-shot and few-shot action recognition. In *WACV*, 2018.
- [Müller, 2007] Meinard Müller. Dynamic time warping. *Information retrieval for music and motion*, 2007.
- [Munkhdalai and Yu, 2017] Tsendsuren Munkhdalai and Hong Yu. Meta networks. *Proceedings of machine learning research*, 2017.
- [Nichol and Schulman, 2018] Alex Nichol and John Schulman. Reptile: a scalable metalearning algorithm. *arXiv preprint*, 2018.
- [Scovanner *et al.*, 2007] Paul Scovanner, Saad Ali, and Mubarak Shah. A 3-dimensional sift descriptor and its application to action recognition. In *ACMMM*, 2007.
- [Snell *et al.*, 2017] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017.
- [Sung *et al.*, 2018] 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*, 2018.
- [Tran *et al.*, 2015] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [Vaswani *et al.*, 2017] 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*, 2017.
- [Vinyals *et al.*, 2016] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *NeurIPS*, 2016.
- [Wang and Jiang, 2016] Shuohang Wang and Jing Jiang. A compare-aggregate model for matching text sequences. *arXiv preprint*, 2016.
- [Wang and Schmid, 2013] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *ICCV*, 2013.
- [Wang *et al.*, 2016] 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*, 2016.
- [Wang *et al.*, 2018] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [Yan *et al.*, 2019] Shipeng Yan, Songyang Zhang, and Xuming He. A dual attention network with semantic embedding for few-shot learning. In *AAAI*, volume 33, 2019.
- [Zhang *et al.*, 2019] Songyang Zhang, Xuming He, and Shipeng Yan. Latentgcn: Learning efficient non-local relations for visual recognition. In *International Conference on Machine Learning*, 2019.
- [Zhang *et al.*, 2020] Hongguang Zhang, Li Zhang, X Qui, Hongdong Li, Philip HS Torr, and Piotr Koniusz. Few-shot action recognition with permutation-invariant attention. In *ECCV*. Springer, 2020.
- [Zhu and Yang, 2018] Linchao Zhu and Yi Yang. Compound memory networks for few-shot video classification. In *ECCV*, 2018.
- [Zhu and Yang, 2020] Linchao Zhu and Yi Yang. Label independent memory for semi-supervised few-shot video classification. *TPAMI*, 2020.