

MASTAF: A Model-Agnostic Spatio-Temporal Attention Fusion Network for Few-shot Video Classification

Xin Liu

Huanle Zhang

Hamed Pirsiavash

Xin Liu

University of California, Davis

{rexliu, dtzhang, hpirsiav, xinliu}@ucdavis.edu

Abstract

We propose MASTAF, a Model-Agnostic Spatio-Temporal Attention Fusion network for few-shot video classification. MASTAF takes input from a general video spatial and temporal representation, e.g., using 2D CNN, 3D CNN, and Video Transformer. Then, to make the most of such representations, we use self- and cross-attention models to highlight the critical spatio-temporal region to increase the inter-class variations and decrease the intra-class variations. Last, MASTAF applies a lightweight fusion network and a nearest neighbor classifier to classify each query video. We demonstrate that MASTAF improves the state-of-the-art performance on three few-shot video classification benchmarks (UCF101, HMDB51, and Something-Something-V2), e.g., by up to 91.6%, 69.5%, and 60.7% for five-way one-shot video classification, respectively.

1. Introduction

Few-shot learning has received increasing attention in video classification for its potential to reduce the video annotation cost significantly [3]. In few-shot video classification, the video samples in the training and test sets are from different classes (i.e., unseen classes in the test set). To classify an unlabeled video sample (query), a few-shot video classification model aims to classify the query to the unseen class (support set). Inspired by the development in few-shot image classification [7, 13, 21], recent few-shot video classification approaches using metric-learning-based methods achieve state-of-the-art performance [3, 20]. This paper targets metric-learning-based few-shot video classification.

A metric-learning-based few-shot video learning algorithm classifies a query based on the similarity between the representation of the query video and the representation of each class in the support set. Therefore, the core to metric-learning-based few-shot video classification is to design feature extraction and representation for the support

sets and the query. Many feature embedding networks have been designed for this purpose. Perrett [20] leverages attention mechanism in temporally-ordered frames from support sets to match query frames after extracting representation for each frame with pre-trained 2D Convolutional Neural Network (2D CNN). Zhang [36] introduces permutation-invariant pooling and self-supervised learning tasks to enhance representations after extracting from a 3D Convolutional Neural Network (3D CNN) embedding network.

In few-shot scenarios, prior efforts with a 2D CNN embedding network outperformed those with a 3D CNN embedding network [3, 20, 40]. However, there are two considerable limitations in existing work with a 2D CNN embedding network. The first limitation is that a complex temporal alignment strategy between the video frames for better accuracy increases computational demand and model inference runtime. For example, Perrett [20] achieves SOTA performance on few-shot video classification by exploring all the combinations of two and three ordered sampled frames from a video for temporal information. As the number of sampled frames from a video grows, the computational cost and inference runtime increase significantly.

The second limitation is their inability to maintain high performance when replacing a 2D CNN embedding network with other advanced video representation models such as 3D CNN [5, 11, 28, 33] and Video Transformer [1]. With the release of large-scale video datasets, video classification models' performance based on 3D CNN and Video Transformer surpasses those with 2D CNN [1, 5, 12], which means such models can generate a better representation for discrimination. Therefore, one would expect that if we replace the 2D CNN in the existing few-shot video classification models with an advanced video representation model, performance should improve. However, this did not happen. Instead, Zhu [40] found that 3D CNN models [5, 28, 33] do not perform better than 2D CNN models in PAL [40], a SOTA 2D-CNN based few-shot video classification algorithm. The main reason is that 2D-CNN approaches rely on the frame-level similarity score and temporal alignment, which do not exist in a 3D CNN embedding network.

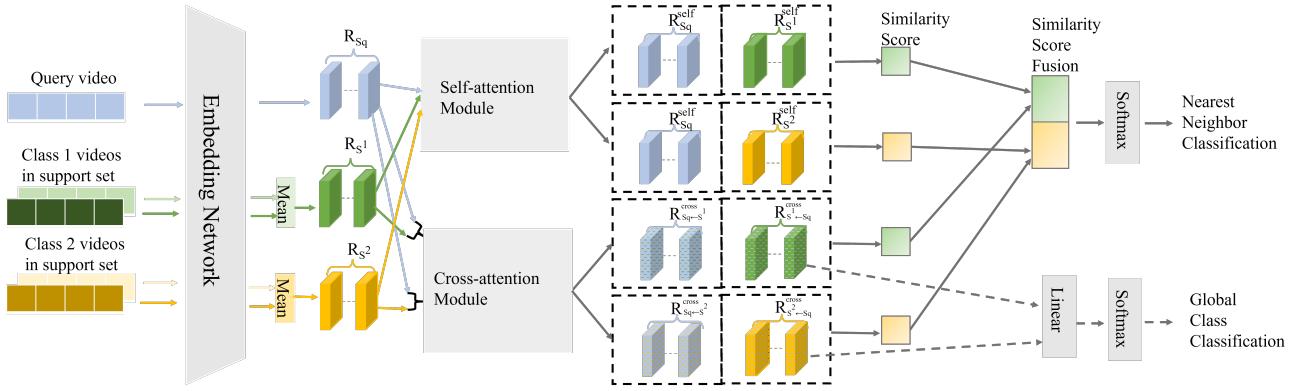


Figure 1. Illustration of the Model-Agnostic Spatio-Temporal Attention Fusion(MASTAF) on a 2-way 2-shot video classification. First, we extract spatio-temporal features with a pre-trained embedding network for each video. Then, we compute a prototypical representation(R_{Sc}) for each class in the support set, which is the mean of all the representations of each class. After that, we use the self-attention module to highlight spatio-temporal features for each query and support class representation and compute the similarity score of each pair of query representation and support class representation using cosine distance. In parallel, we use the cross-attention module to highlight the spatio-temporal correlation features for each pair of query representation and support class representation, and compute the similarity score using the cosine distance. The cross-attention representations of each class in the support set are fed into a global video classifier as a multi-task training set. And the fusion results of similarity scores from the self-attention module and cross-attention module are fed into the nearest neighbor classifier. Details are in Section 3.

In this paper, we propose a model-agnostic few-shot video learning algorithm named Model-Agnostic Spatio-Temporal Attention Fusion network(MASTAF). Our key motivation is to make the most of the rapid advances in video representation learning to build a simple and efficient few-shot video learning framework. To achieve this goal, we have to address the limitations discussed above. Advanced video representation networks, such as 3D CNN and Transformer, extract spatio-temporal representations directly instead of frame-level information. To make good use of such representations, we use self- and cross-attention models to increase the weight of the critical spatio-temporal region to increase the inter-class variations and decrease the intra-class variations as shown in Figure 1. The self-attention network emphasizes the regions of the representation that are essential for representing each class and the query, while the cross-attention network emphasizes the regions of the representation that enhance the discriminability between the query and the unseen classes in the support set. Then, we measure the similarity between the query and each unseen class based on the feature maps from each attention network. Last, we classify the query video by a simple yet efficient fusion network. We also add one multi-task training setting,i.e., global video classification task, to regularize the embedding module and further improve generalization performance. More details are presented in Section 3

Contributions We make the following contributions.

1.We propose MASTAF, a simple and efficient attention-based network compatible with different video classification models for few-shot video classification. MASTAF can

benefit from advanced video classification models such as 3D CNN and Video Transformer that extract good spatial-temporal representations.

2.We design a fusion mechanism to integrate self-attention and cross-attention networks, which greatly enhances the essential spatial and temporal regions of video representation.

3.We extensively evaluate MASTAF using three benchmarks, i.e., UCF101 [25], HMDB51 [15], and Something-Something V2 [10]. Compared to the existing work, MASTAF achieves state-of-the-art performance with a 2D CNN embedding network and improves the state-of-the-art performance with a 3D CNN embedding network without additional computational cost. Our code is available at <https://anonymous.4open.science/r/STAF-30CF1>.

2. Related work

Few shot learning Most existing few-shot learning algorithms can be divided into three categories: model-based methods [18, 23], optimization-based methods [9, 22], and metric-learning-based methods [3, 20, 24]. Metric-learning-based methods are more promising than other two methods in few-shot video classification since the previous work with metric-learning-based achieved better performance [3, 20].

Metric-learning-based method measures the distance between the representation of support samples and query samples and classifies them with the aid of the nearest neighbor

to keep similar classes close and dissimilar classes far away. Particularly, Prototypical Network [24] is based on the idea that each class has a Prototypical representation which is the mean value of support set in embedding space. The few-shot learning problem then becomes the nearest neighbor in the embedding space. Our work is one of the metric-learning-based methods. We can take input from general video spatial and temporal representations extracted from different video representation models. To make the most of representations in the embedding space, we highlight the spatio-temporal features that need attention for each class while increasing the differences from other classes.

Few shot video classification The first module of most metric-learning based methods for a few-shot video classification model is an embedding network that extracts features from each video. The two most commonly used approaches for embedding network are 2D CNN [3, 20, 27, 38, 39, 40] and 3D CNN [2, 8, 36]. After using 2D CNN to extract features from each video frame, Zhu and Yang [38, 39] introduce a memory network structure to learn optimal representations in a larger video representation space. Instead of creating a memory structure to memorize long-term information for video representation, more recent work with 2D CNN embedding networks focus on temporal alignment exploration between the query video and the support set. Cao [3] aligns the frames between the query video and support video by temporal ordering information. Perrett [20] achieves SOTA on 5-way 5-shot video learning by computing the distance of temporal-relational representations between each frame of query video and support video. In comparison, the features extracted from the 3D CNN embedding network already contain temporal information. Therefore, recent work focus on generating general spatio-temporal video representations for unseen classes. Dwivedi [8] leverages GAN to generate the spatio-temporal video representations for the prototype of the unseen classes. Zhang [36] introduces permutation-invariant pooling and self-supervised learning tasks to enhance representations whereas Bishay [2] uses segment-based attention and deep metric learning. Recently, Video Transformers have become a promising option for video representation due to their long-term reasoning ability [1, 16]. Although Video Transformer are not widely used in few-shot video classification, SOTA performance in video classification indicates the promise of being applied in a few-shot scenario.

Attention-based learning Attention mechanism enhances the learning ability of long-range dependencies in the network to highlight the critical regions of visual representations [29]. These critical regions are useful in discriminating the differences between different classes. Therefore, the recent work with attention mechanisms achieve SOTA accuracy for few-shot learning tasks [20, 27, 32]. Perrett [20] applies a cross transformer with a multi-head

attention mechanism for the representation of each frame to locate the representative frames for similarity computation. Thatipelli [27] proposes a self-attention module for the patch’s representation of each frame in a video to highlight critical regions. These works adopt 2D CNN to extract the features and apply attention mechanisms for temporal alignment and frame-level feature enrichment. However, these works cannot maintain high performance when using a 3D CNN embedding network without the help from complex temporal alignment and frame-level feature enrichment. Our work is compatible with any video classification model and uses attention fusion network to highlight the spatio-temporal features, which help increase the inter-class variations and decrease the intra-class variations. Wang [32] applies multi-head self-attention for video-level feature enrichment, increasing GPU memory and computational consumption. However, our approach adopts a simple and efficient fusion layer for self-attention and cross-attention modules with a lower computational cost.

3. MASTAF: Model-Agnostic Spatio-Temporal Attention Fusion Network

3.1. Problem definition

The few-shot video classification problem aims to classify one unannotated query video into one of several annotated categories set, which we call “support set”. Each category has only a few video instances in this support set, and the model did not see these categories during the training process. Our paper focuses on C-way K-shot video classification, where C denotes the number of categories in the support set and K represents the number of video instances for each category in the support set. We follow the same episodic training as in the previous study [3, 20, 36, 38, 39] that randomly select C classes with K video clips for the support set. Then we select one query video from these C classes, which is different from the K video clips in the support set. For each C-way K-shot episode, the support set contains C classes, and each class has K video clips.

We use $S_k^c = \{f_{k,1}^c, f_{k,2}^c, \dots, f_{k,n}^c\}$ to denote the k^{th} video clip of class c , where c belongs to C and k belongs to K , $f_{k,i}^c$ denotes the i^{th} extracted frame from the video and n denotes the total number of frames extracted from the video. For the query video, we use $S_q = \{f_1, \dots, f_i, \dots, f_n\}$, where f_i denotes the i^{th} frame extracted from the query video and n denotes the total number of frames extracted from the query video. The final goal is to predict S_q to one of the classes.

3.2. The MASTAF Model

The design principle of the MASTAF model is to highlight the critical spatio-temporal region to minimize the intra-class variations while maximizing the inter-class vari-

ations between the query video and support set. To tackle the challenge of only having few samples for the unseen class, we first extract spatio-temporal features using any video classification model. Then, we use the attention fusion module to further highlight the critical spatio-temporal region for metric learning. In parallel, we use a global classification task to regularize the embedding network. Next, we analyze each module in the MASTAF model, which is described in Figure 1.

Embedding module In the MASTAF model, the goal of the embedding module f_φ is to learn the spatio-temporal representations for each video. We evenly extract frames from each video, where n is the total number of frames extracted from each video. We can use any video classification model as the spatio-temporal embedding module.

Given a frame sequence extracted from the video $S_v = \{f_1, f_2, \dots, f_n\}$, let $R_v \in \mathbb{R}^{C' \times T' \times H' \times W'}$ denote the representation learned from the embedding model:

$$R_v = f_\varphi(S_v). \quad (1)$$

For a video clip in the support set S_k^c , we use $R_{S_k^c}$ to denote the representation learned from the embedding module. We use R_{S^c} to denote the representation of the class c , which is the mean of all the representations of video clips for class c in the support set. And since we have only one query video in the few-shot learning task, we use R_{S_q} to denote the representation for the query video clip. After we get the representations for the support set and query video, we go through two separate attention modules in parallel, i.e., the self-attention module and the cross-attention module.

Self-attention module Our goal of the self-attention module is to highlight the critical information in the representation of each class. As shown in Figure 2, we first reshape each representation to $R'_v \in \mathbb{R}^{C' \times L}$, where $L (L = T' \times H' \times W')$ is the number of spatio-temporal positions on each feature cubic map. After that, for each class in the support, R_{S^c} becomes R'_{S^c} , i.e., $[R_1^{S^c}, \dots, R_i^{S^c}, \dots, R_L^{S^c}]$, where $R_i^{S^c}$ denotes the feature vectors at the i^{th} spatio-temporal position in the R'_{S^c} . For each query video, R_{S_q} becomes R'_{S_q} , i.e., $[R_1^{S_q}, \dots, R_i^{S_q}, \dots, R_L^{S_q}]$, where $R_i^{S_q}$ denotes the feature vectors at the i^{th} spatio-temporal position in the R'_{S_q} . Then we compute the self-relation map for each representation as:

$$M^{self} = (R'_v)^\top R'_v, \quad (2)$$

where $M^{self} \in \mathbb{R}^{L \times L}$ that denotes the self-relation map for each video, where M_i^{self} denotes the self-relation at the i^{th} spatio-temporal position in the feature map. Then we apply convolutional operation with a kernel d , i.e., $d \in \mathbb{R}^L$, to fuse each position self-relation vector into an attention scalar, which is in $\mathbb{R}^{T' \times H' \times W'}$. Then we leverage a softmax function to draw self-attention for each i^{th} position:

$$A_i^{self} = \frac{\exp((d^\top M_i^{self})/\tau)}{\sum_{j=1}^L \exp((d^\top M_j^{self})/\tau)}, \quad (3)$$

where τ is the temperature hyperparameter to amplify the variance and A_i^{self} denotes the i^{th} position of self-attention map A^{self} , i.e., $A^{self} \in \mathbb{R}^{T' \times H' \times W'}$.

Instead of assigning equal weight to every position, we add a meta-learner to learn the kernel d dynamically to pay attention to the critical positions in the feature cubic map. First, we leverage row-wise global average pooling for M^{self} to get an averaged vector \bar{M}^{self} , which $\bar{M}^{self} \in \mathbb{R}^L$. Then we use a meta-learner to learn the kernel d dynamically:

$$d = f_\gamma(\sigma(f_\delta(\bar{M}^{self}))), \quad (4)$$

where $f_\delta : \mathbb{R}^L \rightarrow \mathbb{R}^l$ and $f_\gamma : \mathbb{R}^l \rightarrow \mathbb{R}^L$, i.e., l denotes the scaled dimension and σ represents the ReLU function[19].

After we get the self-attention cubic map A^{self} , we leverage a residual attention mechanism to weigh each element of the original map R_v with $1 + A^{self}$ to get the self-attention representation R^{self} for each class:

$$R^{self} = R_v(1 + A^{self}), \quad (5)$$

where $R^{self} \in \mathbb{R}^{C' \times T' \times H' \times W'}$.

Cross-attention module While the self-attention module highlights the critical spatio-temporal region in the representation itself, the cross-attention module focuses on the correlation between the query video and the support set. As shown in Figure 3, we follow the same steps as in the self-attention module to reshape each representation to $R'_v \in \mathbb{R}^{C' \times L}$. After that, we compute the correlation map for each pair of the query video and the support class prototype. For example, for the pair of the query video R_{S_q} and support class c , i.e., R_{S^c} , we compute the correlation map for the query video $M_{S_q \leftarrow S^c}^{cross}$ between the query video and support class:

$$M_{S_q \leftarrow S^c}^{cross} = (R'_{S^c})^\top R'_{S_q}. \quad (6)$$

Then for the support class c , the correlation map $M_{S^c \leftarrow S_q}^{cross}$ between the query video and support class is:

$$M_{S^c \leftarrow S_q}^{cross} = (R'_{S_q})^\top R'_{S^c}. \quad (7)$$

After getting the correlation map for query video and support class in each pair, we go through the same steps as in the self-attention module, which are shown in the Figure 3, to get the cross-attention representation for query video and support class in each pair, i.e., $R_{S_q \leftarrow S^c}^{cross}$ and $R_{S^c \leftarrow S_q}^{cross}$.

Attention fusion module After we get the self-attention and cross-attention representation from the two attention

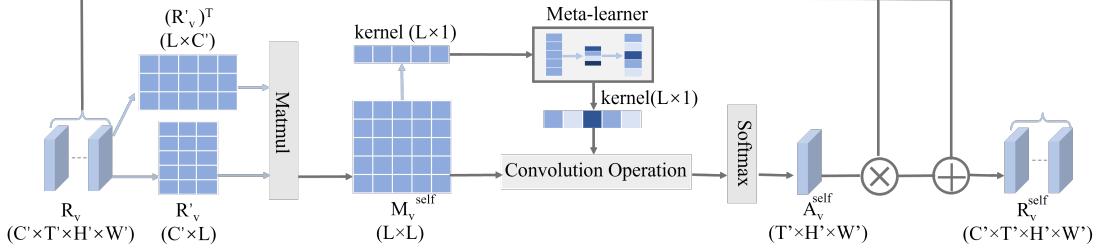


Figure 2. Self-attention module

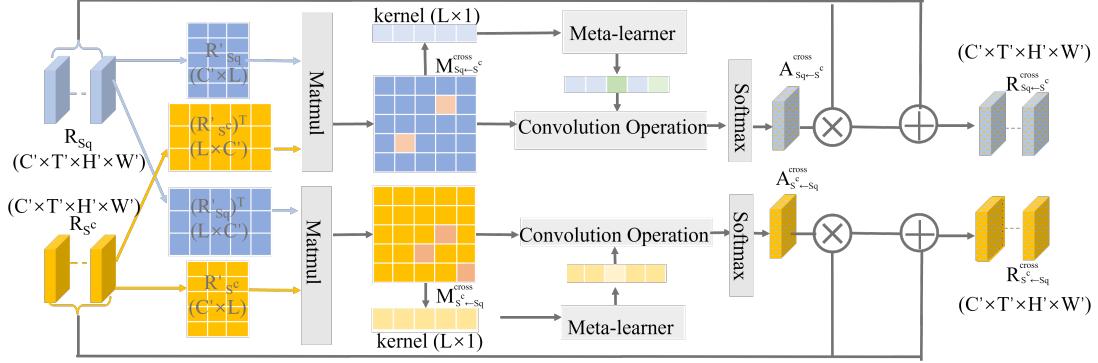


Figure 3. Cross-attention module

modules, we compute the probability of predicting S_q as the class k using self-attention representation:

$$P_{self}(y = k|S_q) = \frac{\exp(-D_{cos}(R_{S_q}^{self}, R_{S^k}^{self}))}{\sum_{j=1}^C \exp(-D_{cos}(R_{S_q}^{self}, R_{S^j}^{self}))}, \quad (8)$$

where D_{cos} denotes the cosine distance and $P_{self}(y = k|S_q)$ denotes the probability of predicting S_q as the class $k \in \{1, 2, \dots, C\}$ using self-attention representations. Then we compute the probability of predicting S_q as the class k using cross-attention representation:

$$P_{cross}(y = k|S_q) = \frac{\exp(-D_{cos}(R_{S_q \leftarrow S^k}^{cross}, R_{S^k \leftarrow S_q}^{cross}))}{\sum_{j=1}^C \exp(-D_{cos}(R_{S_q \leftarrow S^j}^{cross}, R_{S^j \leftarrow S_q}^{cross}))}, \quad (9)$$

where $P_{cross}(y = k|S_q)$ denotes the probability of predicting S_q as the class $k \in \{1, 2, \dots, C\}$ using cross-attention module.

To take advantage of the discriminative information from two attention mechanisms, we leverage the attention fusion module with the nearest neighbor classifier:

$$P(y = k|S_q) = \frac{1}{2}[P_{self}(y = k|S_q) + P_{cross}(y = k|S_q)], \quad (10)$$

where $P(y = k|S_q)$ denotes the final probability of predicting S_q as the class $k \in \{1, 2, \dots, C\}$.

Multi-task training To reduce the risk of overfitting in the training dataset and generate a general representation for

unseen class, we train the MASTAF model in a multi-task setting to regularize the embedding network. We combine the nearest neighbor classifier and the global video classifier.

During the training process, after the attention fusion module computes the probability of predicting query video to one of the classes in the support set, we use a negative log-probability as the loss function of the nearest neighbor classifier based on the actual class label:

$$L_1 = - \sum_{k=1}^C \log P(y = k|S_q). \quad (11)$$

Since the representations after the cross-attention module contain highlighting regions related to the query video, we choose these representations to predict the global class in the whole training dataset. The total class number in the training dataset is Z . We feed these cross-attention representations to a fully connected layer and a softmax layer to get the probability of predicting the global class, i.e., $P(y = z|S^c)$ where $z \in \{1, 2, \dots, Z\}$. Then we define the loss function of the global video classifier as:

$$L_2 = - \sum_{z=1}^Z \log P(y = z|S^c). \quad (12)$$

Finally, the loss function of the MASTAF model is defined as:

$$L = L_1 + \lambda L_2, \quad (13)$$

where we use λ to weigh the impact of different classification tasks. Note that a multi-task training setting is only used during the training process. This setting is discarded at the inference stage.

4. Evaluation

4.1. Experimental Setup

Datasets. We compare MASTAF with existing work on UCF101 [25], HMDB51 [15], and Something-Something V2 (SSv2) [10]. We do not use Kinetics-100 [38] to avoid bias because one of our MASTAF models is pre-trained on Kinetics-700 [4]. In these datasets, SSV2 is more challenging because it focuses on actions related to temporal relationships such as ‘pretending to take something from somewhere’ versus ‘take something from somewhere’ [37]. There are two few-shot splits for SSV2 proposed by CMN [38] and OTAM [3], containing 64, 12, and 24 classes as the training, validation, and test set. We use SSV2-part and SSV2-all denote the split from CMN [38] and the split from OTAM [3]. The difference between these two splits is the number of video samples in each class. For SSV2-part, Zhu and Yang [38] randomly selects 100 samples for each class, whereas for SSV2-all, Cao [3] uses all the samples in the original SSV2. We evaluate our method in these two splits. Additionally, we also follow the split in ARN [36] for HMDB51 and UCF101.

Evaluation and baseline. Following the evaluation process in TRX [20], we evaluate the 5-way 1-shot and 5-way 5-shot video classification task and report the average accuracy over 10,000 randomly selected episodes from the test set. We compare our results with ten SOTA algorithms, i.e., TSN++ [31], CMN-J [39], OTAM [3], FEAT [34], PAL [40], TRX [20], Baseline [6], MatchingNet [30], ProtoGAN [8], ARN [36]. In particular, the idea of Baseline (Notation from [6]) is to train a new classifier with the given labeled examples in novel classes after extracting the representation from the embedding network. For a fair comparison, we use three MASTAF models with three different types of embedding networks, i.e., MASTAF-{TSN}, MASTAF-{R3D} and MASTAF-{ViViT}. For MASTAF-{TSN}, we follow the same embedding network configuration with [3, 20, 34, 40], using an ImageNet pre-trained ResNet-50 as the backbone network. For MASTAF-{R3D}, we use the merged video dataset with Kinetics-700 [4], Moment-in-time [17], and START-action [35] to pre-train 3D ResNet-50 embedding network. We also compare our approach against the previous work based on a 2D CNN embedding network where we replace 2D CNN with 3D CNN. We use Baseline-{R3D}, MatchingNet-{R3D}, TRX-{R3D} as the baselines by replacing the 2D CNN embedding network with a 3D CNN embedding network(same pre-trained R3D model as MASTAF-{R3D}). We extract

one representation using pre-trained R3D from each video and then go through the matching part proposed in Baseline [6], MatchingNet [30] and TRX [20]. For MASTAF-{ViViT} and TRX-{ViViT}, we use ViViT [1] as our embedding network. We initialize ViViT from a ViT [14] image model trained on the JFT [26] dataset. Due to the huge computation demand for ViViT [1], we only perform 5-way 1-shot learning for MASTAF-{ViViT} and TRX-{ViViT}.

Experimental Configuration. For MASTAF-{TSN}, MASTAF-{ViViT} and TRX-{ViViT}, we evenly sample 8 frames from each video as 8 segments for each video. For 3D CNN-based Baseline, MatchingNet, TRX and MASTAF-{R3D}, we evenly sample 16 frames from each video sample. After that, we resize each frame to 256×256 . Then we randomly flip each frame horizontally and crop the center region of 224×224 to augment the training data. For test data, we only crop the center with the same size without the horizontal flipping. Then for MASTAF-{TSN}, we use an ImageNet pre-trained ResNet-50 as the backbone and average all the frame representations as to the video representation. For 3D CNN-based Baseline, MatchingNet, TRX and MASTAF-{R3D}, we use a 3D ResNet-50[12] with the weights pre-trained on the combined dataset with Kinetics-700[4], Moments in Time[17], and Start Action[35] as the embedding network. After finetuning in the validation dataset, we set 0.025 as the temperature hyperparameter(τ in Eq 3) and set 6 as the meta-learner scaled dimension(l is the scaled dimension of f_γ in Eq 4), and set 2 as the loss weight hyperparameter(λ in Eq 13). We train our model for 128,000 episodes in eight NVIDIA RTX A5000 GPU(except for the larger SSV2-all, we train our model for 256,000 episodes). We optimize the MASTAF model with SGD, in which the learning rate is 0.01. After fine-tuning, we adopt the batch-size of 128, 64, 32, 32 for UCF101, HMDB51, SSV2-part, and SSV2-all, respectively.

4.2. Comparison with State-of-the-art Algorithms

Table 1 tabulates the overall 5-way 1-shot and 5-way 5-shot performance compared with existing methods on two splits of SSV2. We can categorize these comparative methods into three groups based on the embedding network. In 2D CNN embedding group, TSN++ [31], CMN-J [39], FEAT [34] are model agnostic and do not apply any frame-level temporal alignment. Compared with these three methods, the other three methods, i.e., OTAM [3], PAL [40], and TRX [20], adopt frame-level temporal alignment, which further improves the performance of few-shot video classification. MASTAF-{TSN} outperforms existing 5-way 1-shot video classification algorithms in the 2D CNN group. TRX [20] achieves SOTA performance for 5-way 5-shot learning because it leverages the temporal information from different frames in different videos in

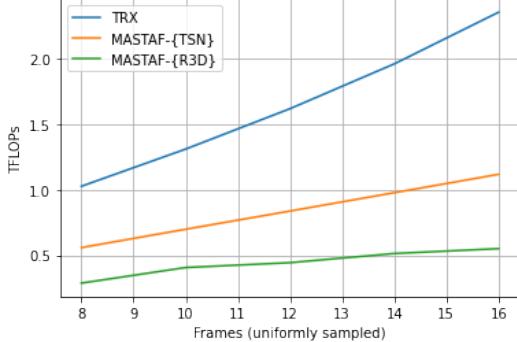


Figure 4. Computational demand analysis for TRX, MASTAF-{TSN} and MASTAF-{R3D} as the number of sampled frames varies from 8 to 16 frames on UCF101

the support set. However, this complex alignment strategy leads to huge computation costs and increases model inference’s runtime. Figure 4 and Figure 5 compare the TFLOPs and model inference’s runtime of TRX [20] and MASTAF-{TSN}. Our approach achieves SOTA accuracy without increased computational cost and is more efficient than TRX [20]. As the number of frames sampled from a video increases, TRX [20] consumes more computational resources and takes longer for the inference process. In the 3D CNN group, the accuracy of the TRX-{R3D} is lower than TRX because it cannot perform the frame-level temporal alignment. For PAL [40], Zhu [40] also mentioned that 3D CNN models [5, 28, 33] do not perform better than 2D CNN models due to the lacking of frame-level similarity scores. In comparison, MASTAF-{R3D} takes advantage of the spatio-temporal representation from R3D and further improves the performance. In the Transformer group, MASTAF-{ViViT} further enhances the performance. These results demonstrate MASTAF works best when spatio-temporal information is well represented in advanced video classification models. In contrast, existing work with a 2D embedding network cannot maintain high performance when replacing a 2D CNN embedding network with other advanced video representation models.

Table 2 tabulates the overall 5-way 1-shot and 5-way 5-shot performance compared with existing methods on UCF101 and HMDB51. Our MASTAF with a 2D embedding network achieves decent performance while TRX and PAL achieve SOTA accuracy on these two datasets. The reason is that TSN does not provide enough spatio-temporal information for MASTAF to distinguish the query video from the videos in the support set. So to benefit the most from MASTAF, we explore our MASTAF with a 3D CNN embedding network and Video Transformer. As shown in Table 2, our MASTAF-{R3D} outperforms other methods based on 3D models and MASTAF-{ViViT} outperforms TRX-{ViViT} and achieves new SOTA performance. Com-

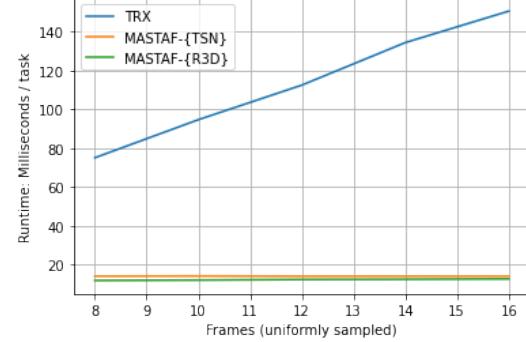


Figure 5. Model inference’s runtime analysis for TRX, MASTAF-{TSN} and MASTAF-{R3D} in one NVIDIA RTX A5000 GPU as the number of sampled frames varies from 8 to 16 frames on UCF101

pared with MASTAF-{TSN}, MASTAF-{R3D} has significantly lower resource consumption and running time, as shown in Figure 4 and Figure 5.

4.3. Ablation study

We have shown in Section 4.2 that our MASTAF can make the most of the advanced video classification model to improve the accuracy without more computational cost. We now perform detailed ablation studies on two dataset UCF101 and SSV2-all to show each module’s influence. The ablation studies about multi-task learning setting, meta-learner and residual structure are in the Appendix. In these ablation studies, all MASTAF models use the 3D ResNet-50 model pre-trained on the merged video dataset with Kinetics-700 [4], Moment-in-time [17], and START-action [35] as the embedding network.

4.3.1 Attention fusion mechanism

To explore the effectiveness of the attention fusion mechanism, we introduce three comparison models, i.e., MASTAF-Neighbor, MASTAF-Self, MASTAF-Cross. In MASTAF-Neighbor, representations learned from the embedding network are fed into the nearest neighbor classifier and a global video classifier directly without our attention mechanisms. For MASTAF-Self and MASTAF-Cross, before being fed into two classifiers, representations go through the self-attention and cross-attention mechanism, respectively. Table 3 shows the comparison results. Compared with MASTAF-Neighbor, after adding the attention mechanism, all three other models have a significant performance improvement, demonstrating that representations after the embedding network have some spatio-temporal features related to the non-target action region. The cross-attention mechanism in MASTAF-Cross aid in highlighting the spatio-temporal features associated with the target action region among the query video and support

Table 1. Comparison on 5-way 1-shot and 5-shot benchmarks of SSv2-part, and SSv2-all. The best performance in each group is highlighted. †: Results from [3]. *: Results from [40]

Method	Embedding Groups	SSv2-part		SSv2-all	
		1-shot	5-shot	1-shot	5-shot
TSN++ [†] [31]	2D CNN	-	-	34.4	43.8
CMN-J [39]		36.2	48.8	-	-
FEAT* [34]		-	-	45.3	61.2
OTAM [3]		-	-	42.8	52.3
PAL [40]		-	-	46.4	62.6
TRX [20]		36.0	59.1	42.0	64.6
MASTAF-{TSN}	3D CNN	37.5	50.2	46.9	62.4
Baseline-{R3D}		24.9	36.1	25.6	39.8
MatchingNet-{R3D}		34.1	45.2	43.2	54.4
TRX-{R3D}		26.1	47.0	34.9	58.9
MASTAF-{R3D}		39.9	52.2	50.3	66.7
TRX-{ViViT}	Transformer	34.7	-	42.7	-
MASTAF-{ViViT}		45.6	-	60.7	-

Table 2. Comparison on 5-way 1-shot and 5-shot benchmarks of UCF101, and HMDB51. The best performance in each group is highlighted. *: Results from [40]

Method	Embedding Groups	UCF101		HMDB51	
		1-shot	5-shot	1-shot	5-shot
FEAT* [34]	2D CNN	83.9	94.5	60.4	75.2
PAL [40]		85.3	95.2	60.9	75.8
TRX [20]		-	96.1	-	75.6
MASTAF-{TSN}		79.3	90.3	54.8	67.7
ProtoGAN [8]		57.8	80.2	34.7	54
ARN [36]		66.3	83.1	45.5	60.6
Baseline-{R3D}	3D CNN	53.4	88.7	40.1	68.1
MatchingNet-{R3D}		82.7	93.5	61.8	75.6
TRX-{R3D}		82.5	94.1	57.0	74.3
MASTAF-{R3D}		90.6	97.6	67.9	81.2
TRX-{ViViT}		84.8	-	58.1	-
MASTAF-{ViViT}		91.6	-	69.5	-

Table 3. Comparison results with three variants of MASTAF for 5-way 1-shot video classification

Method	UCF101	SSv2-all
MASTAF-Neighbor	82.7	43.2
MASTAF-Self	90.3	49.4
MASTAF-Cross	90.5	49.2
MASTAF	90.6	50.3

set. MASTAF-Self’s self-attention module helps highlight spatio-temporal features related to the action in each video itself. Therefore, combining two different attention modules can take advantage of each module to further extract more discriminative spatio-temporal representations. The results in Table 3 demonstrate our argument.

5. Conclusion

This paper proposes a Model-Agnostic Spatio-Temporal Attention Fusion network(MASTAF) for few-shot video classification. MASTAF is a simple and efficient few-shot video classification framework compatible with different video classification models. MASTAF make the most of the knowledge learned from the advanced video classification model and uses self- and cross-attention to highlight the spatio-temporal features. MASTAF works best when spatio-temporal information is well represented in advanced video classification models and improves the state-of-the-art performance of 5-way 1-shot, and 5-shot video classification on UCF101, HMDB51, and SSv2, e.g., MASTAF improves the accuracy of 5-way 1-shot video classification to 91.6%, 69.5%, and 60.7% for UCF101, HMDB51, and SSv2, respectively.

References

[1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. 2021.

[2] Mina Bishay, Georgios Zoumpourlis, and Ioannis Patras. Tarn: Temporal attentive relation network for few-shot and zero-shot action recognition. 2019.

[3] K. Cao, J. Ji, Z. Cao, C. Chang, and J. Niebles. Few-shot video classification via temporal alignment. pages 10615–10624, 2020.

[4] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *ArXiv*, 2019.

[5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. 2017.

[6] W.Y. Chen, Y.C. Liu, Z. Kira, Y.C. Wang, and J.B. Huang. A closer look at few-shot classification. In *ICLR*, 2019.

[7] Carl Doersch, Ankush Gupta, and Andrew Zisserman. Crosstransformers: spatially-aware few-shot transfer. 2021.

[8] Sai Dwivedi, Vikram Gupta, Rahul Mitra, Shuaib Ahmed, and Arjun Jain. Protagon: Towards few shot learning for action recognition. 2019.

[9] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1126–1135, 2017.

[10] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, Florian Hoppe, Christian Thurau, Ingo Bax, and Roland Memisevic. The “something something” video database for learning and evaluating visual common sense. pages 5843–5851, 10 2017.

[11] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Learning spatio-temporal features with 3d residual networks for action recognition. In *ICCV Workshops*, pages 3154–3160, 10 2017.

[12] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? pages 6546–6555, 2018.

[13] Ruibing Hou, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Cross attention network for few-shot classification. 2019.

[14] Alexander Kolesnikov, Alexey Dosovitskiy, Dirk Weissenborn, Georg Heigold, Jakob Uszkoreit, Lucas Beyer, Matthias Minderer, Mostafa Dehghani, Neil Houlsby, Sylvain Gelly, Thomas Unterthiner, and Xiaohua Zhai. An image is worth 16x16 words: Transformers for image recognition at scale. 2021.

[15] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: a large video database for human motion recognition. 2011.

[16] Xinyu Li, Yanyi Zhang, Chunhui Liu, Bing Shuai, Yi Zhu, Biagio Brattoli, Hao Chen, Ivan Marsic, and Joseph Tighe. Vidtr: Video transformer without convolutions. 2021.

[17] Mathew Monfort, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, Tom Yan, Lisa Brown, Quanfu Fan, Dan Gutfreund, Carl Vondrick, et al. Moments in time dataset: one million videos for event understanding. 2019.

[18] Tsendsuren Munkhdalai and Hong Yu. Meta networks. *Proceedings of machine learning research*, 70, 03 2017.

[19] Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, page 807–814, 2010.

[20] Toby Perrett, Alessandro Masullo, Tilo Burghardt, Majid Mirmehdi, and Dima Damen. Temporal-relational crosstransformers for few-shot action recognition. 2021.

[21] Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of maml. 2020.

[22] Sachin Ravi and H. Larochelle. Optimization as a model for few-shot learning. 2017.

[23] Adam Santoro, Sergey Bartunov, Matthew M. Botvinick, Daan Wierstra, and Timothy P. Lillicrap. One-shot learning with memory-augmented neural networks. In *Proceedings of the International Conference on Machine Learning*, 2016.

[24] Jake Snell, Kevin Swersky, and Richard S. Zemel. Prototypical networks for few-shot learning. *CoRR*, 2017.

[25] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. 2012.

[26] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. pages 843–852, 10 2017.

[27] 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, June 2022.

[28] Du Tran, L. Bourdev, Rob Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. pages 4489–4497, 01 2015.

[29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. volume 30, 2017.

[30] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.

[31] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Val Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016.

[32] 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, June 2022.

[33] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. 2018.

- [34] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. 2020.
- [35] Yuya Yoshikawa, Jiaqing Lin, and Akikazu Takeuchi. Stair actions: A video dataset of everyday home actions. *ArXiv*, 2018.
- [36] Hongguang Zhang, Li Zhang, Xiaojuan Qi, Hongdong li, Philip Torr, and Piotr Koniusz. Few-shot action recognition with permutation-invariant attention. 2020.
- [37] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. 2018.
- [38] Linchao Zhu and Yi Yang. Compound memory networks for few-shot video classification. September 2018.
- [39] Linchao Zhu and Yi Yang. Label independent memory for semi-supervised few-shot video classification. 2020.
- [40] Xiatian Zhu, Antoine Toisoul, Juan-Manuel Perez-Rua, Li Zhang, Brais Martinez, and Tao Xiang. Few-shot action recognition with prototype-centered attentive learning. 2021.