

TAEN: Temporal Aware Embedding Network for Few-Shot Action Recognition

Rami Ben-Ari ^{*†1}Mor Shpigel Nacson ^{*§2}Ophir Azulai ^{*3}Udi Barzelay ^{*3}Daniel Rotman ³¹ OriginAI² Technion, Haifa, Israel³ IBM Research AI, Haifa, Israel

ramib@originai.co, mor.shpigel@gmail.com {ophir,udib,danieln}@il.ibm.com

Abstract

Classification of new class entities requires collecting and annotating hundreds or thousands of samples that is often prohibitively costly. Few-shot learning suggests learning to classify new classes using just a few examples. Only a small number of studies address the challenge of few-shot learning on spatio-temporal patterns such as videos. In this paper, we present the Temporal Aware Embedding Network (TAEN) for few-shot action recognition, that learns to represent actions, in a metric space as a trajectory, conveying both short term semantics and longer term connectivity between action parts. We demonstrate the effectiveness of TAEN on two few shot tasks, video classification and temporal action detection and evaluate our method on the Kinetics-400 and on ActivityNet 1.2 few-shot benchmarks. With training of just a few fully connected layers we reach comparable results to prior art on both few shot video classification and temporal detection tasks, while reaching state-of-the-art in certain scenarios.

1. Introduction

Action recognition is one of the fundamental problems in computer vision, with applications such as event detection, clip classification and retrieval in multimedia storage. In this domain, there are two typical tasks, classification and temporal detection. Video classification aims to classify short video clips often *trimmed* from a longer footprint, and usually a few seconds long. In temporal action detection setting, actions appear as short temporal sections within a long untrimmed video that can last even for several minutes. Similar to object detection the goal is to detect the time stamp (temporal location) at which a certain action class

takes place.

Deep learning models have been successfully utilized for action recognition [6, 10, 22, 25]. Although these models obtain remarkable results, they require large amounts of labeled training data that often are prohibitively available. The task of learning new categories from a small number of labeled examples is known as few-shot learning (FSL). Typical FSL methods are based on meta-learning [2, 23], distance metric learning [19, 31] and synthesis methods [5, 7]. While few-shot learning has been extensively studied in the context of visual recognition [8, 17, 27, 28], few studies address the challenge of learning from a few instances to detect spatio-temporal patterns. In particular, only a few works have been proposed for few-shot learning in action recognition and specifically for video classification [1, 4, 11, 32, 37] and temporal action detection [35]. To alleviate the annotation labour, weakly supervised strategies [24, 25] are suggested, where tagging is conducted just over the entire video. However, in cases where new classes may appear on the fly or in rare class types, few-shot learning methods has no alternative. In FSL one can learn *new categories* from just a few video examples, often one to five. This is in contrast to transfer learning where tens or hundreds of new labeled examples are needed to learn a new class. In this paper, we address the task of few-shot learning video classification and show also an extension for temporal action detection.

Action recognition in videos has been greatly advanced thanks to powerful features such as dense trajectories [30], and deep features *e.g.* the Two-Stream networks [26], C3D [29], and I3D [38]. Prior to deep features this was tackled by extracting spatio-temporal local descriptors from space-time interest points [30] (iDTF) presenting trajectories in pixel-time space. The sparse patterns of trajectories were then analyzed to distinguish between different actions. However, these trajectories do not convey the semantic information gained in deep representations of two-stream net-

*Equal contribution

†This work was done while the author was with IBM Research AI

§This work was done while the author was an intern at IBM Research AI

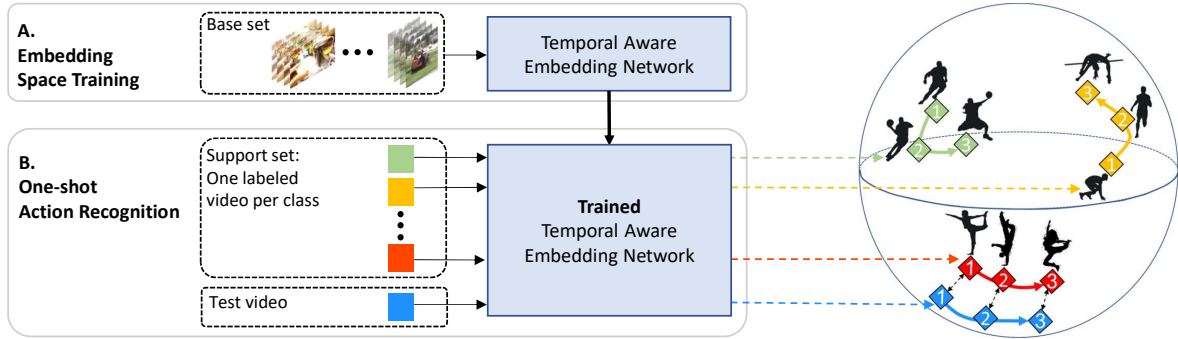


Figure 1. **Action representation as a trajectory in a metric space.** We train an embedding space where actions are represented as trajectories passing through sub-actions. Jointly learning the sub-action prototypes and the embedding function allows our representation to be optimized for both short and long term action granularity. Using the trained embedding network, we embed n videos from different classes (n -way, 1-shot scenario) and the query test video. The query is then classified based on distance between the associated trajectories illustrated by the black arrows between red and blue trajectories.

works, I3D or C3D. Yet, common deep features often miss the long term dependencies involved with complex actions, by processing short video segments (usually 16 consecutive frames).

In this work, we propose a novel approach based on Distance Metric Learning (DML) for few-shot action recognition, called Temporal Aware Embedding Network - TAEN, as described schematically in Fig. 1. We start with decomposition of the video into *ordered* temporal segments, called *sub-actions*. In contrast to the whole action that can last several seconds, sub-action is only part of the action in the timeline. We consider each action as a consecutive set of sub-actions and train an embedding function that can represent an action as a trajectory in a metric space, considering also the long term sub-action dependencies. This representation preserves the temporal order of the sub-actions while carrying the semantics encoded in deep features. Using a trajectory distance enables our method to distinguish between even fine-grained actions, manifested as similar actions that differ only in some intermediate part (i.e a sub-action).

We evaluate TAEN on Kinetics-400 data set [15]. For extension of TAEN to temporal action detection we build upon the BMN temporal region proposal [20] and evaluate our method on a benchmark previously suggested by [35] on ActivityNet 1.2 [9]. Our method is computationally cheap and scalable, relying on training of only a few fully connected layers for learning the embedding function and the sub-action prototypes. TAEN allows classification of new action categories fairly well, even with a single example.

Our contributions are: **(1)**, we suggest a new metric learning method that encodes the long term sub-action connectivity into an embedding space, eventually representing the actions as trajectories in a *metric space*. **(2)**, Our model

jointly optimizes for sub-action prototypes and the embedding function. **(3)** We suggest a novel loss function that allows an effective learning of class action trajectories. **(4)**, using the same model we present comparable results to prior art in video classification and temporal action detection.

2. Related work

2.1. Sub-actions

Decomposing actions into characteristic sub-actions has been studied before [10, 22]. A common practice in these works consists of aggregating and pooling local features from sub-action segments. While naive approaches use mean or max pooling, recent studies extend the pooling techniques by incorporating them into Deep Neural Network (DNN), namely NetFV [18] and NetVLAD [6]. By looking for correlations between a set of primitive action representations, ActionVLAD [6] has shown state-of-the-art performance in several action recognition benchmarks. However, cluster-and-aggregate based methods such as NetVLAD assign soft-clusters to every frame in the video and therefore also use a too fine granularity, particularly for complex actions. These methods further ignore the temporal ordering that might be crucial for recognizing an action from a single example. Therefore, these methods are not well suited for the few-shot action recognition settings.

Typically a few sub-actions are enough to represent simple actions, while others would need to be represented with more parts. Based on this assumption, Hou *et al.* [10] suggested a temporal action detection method, trying to optimize the number of sub-actions for each action class. However, this method uses the hand crafted dense trajectory features (iDTF) [30] and does not handle the few-shot setting. In [22], the sub-action notion is used in the context of weakly supervised temporal action detection.

2.2. Few-shot video classification

Zhu and Yang [37] suggest a method based on memory networks and meta-learning that require heavy computation and space resources to learn new representations for each episodic task. They further use a single embedding vector to represent the entire video, that overlooks the temporal structure of the action and focuses mainly on the most distinguishable segment of a video. The approach by Bishay *et al.* [1] named TARN, addresses FSL on short video segments by calculating the relation between a query and support video with a similarity measure between *aligned* segments. Their model searches for relations between short-length segments (16 frames) and therefore still relies on an extremely fine action granularity, while in the relation process the temporal order is further lost. Recently [11] showed SOTA results in few shot video classification by action decomposition and temporal alignment using Dynamic Time Warping (DTW). We argue that our method is more efficient by finding the optimal prototypes without the DTW alignment which is $\mathcal{O}(N^2)$. Optimized prototypes yet suggest robustness to slight miss-alignment in the video clips.

2.3. Few-shot temporal action detection

While metric learning approaches have found their way to few-shot learning object detection tasks [13, 12], their use in the few-shot temporal action detection task has been overlooked. The recent attempt taken for this task by Yang *et al.* [35] suggests a meta-learning method based on Matching Networks and uses a Long-Short Term Memory (LSTM) video encoder. Their method is computationally heavy as it requires optimization over many different episodes. Moreover, to tackle action detection, they use a “sliding window” approach for action proposals, which further sets a high computational cost and a high imbalance in the training set, resulting high false-positive rates and inflexible activity boundaries. Recent work shows that a more effective and efficient way is using a temporal region proposal method [3, 33, 36]. In fact, new temporal region proposal methods such as [20, 21] are now suggested as part of temporal action detection pipelines. [34] addresses the problem of effective temporal region proposal for improved detection accuracy.

In this paper we learn to present actions as *ordered* temporal segments, namely *sub-actions*. In contrast to the whole action that can last several seconds, sub-action is only part of the action in the timeline. We consider each action as a consecutive set of sub-actions using sub-actions semantics as well as their long term dependencies for improved action recognition. Our method builds on *jointly* learning the prototypical sub-actions and the embedding function keeping their temporal order, eventually describing actions as a *trajectory* in a metric space. Actions are commonly separated in the embedding space by the semantics of ob-

jects in the frame and the background. We define an action trajectory as a parametrized curve in an embedding space, discretized by sub-actions. Each action signature is then obtained by a *temporally ordered* set of points (prototypes of sub-actions) representing different parts of an action in the embedding space. This type of representation allows to better discriminate between similar actions even when they take place in the same scene. The unique trajectory associated with each action allows recognition of a new action class from only a few examples. Our model requires training of just a few FC layers, making it efficient in both training and inference. Our suggested method called Temporal Aware Embedding Network - TAEN learns to represent videos as trajectories in the feature space.

We evaluate TAEN on Kinetics-400 data set [15]. For extension of TAEN to temporal action detection we build upon the BMN temporal region proposal [20] and evaluate our method on a benchmark previously suggested by [35] on ActivityNet 1.2 [9].

3. TAEN

Fig. 2 presents our model architecture. We define the *trajectory* in the DML space as a collection of sub-action prototypes with temporal order. Actions are then represented in the metric space by a set of ordered prototypes. In the few-shot video classification and temporal action detection, our goal is to train a model that will be able to generalize to new unseen classes. To this end, we propose a novel method that relies on two main steps:

1. Learn an embedding space where actions are represented by well-separated trajectories. Action trajectories are represented by $a \in \mathbb{N}$ temporally ordered centers, one for each sub-action.
2. Learn jointly the prototypes and the embedding function.
3. Classify new videos according to their *trajectory* signature.

3.1. Embedding architecture

Our network architecture is inspired by the RepMet architecture [13] that enables the parallel training of the embedding function and the prototypes. This architecture allows joint learning of multiple prototypes per class, sub-actions in our case, while training them with the embedding to reach optimized embeddings and prototypes.

We design a novel loss function to allow discrimination of trajectories in the embedding space. Our training is done in batches but for simplicity we will describe a single video flow through the architecture. The model architecture is depicted in Fig. 2.

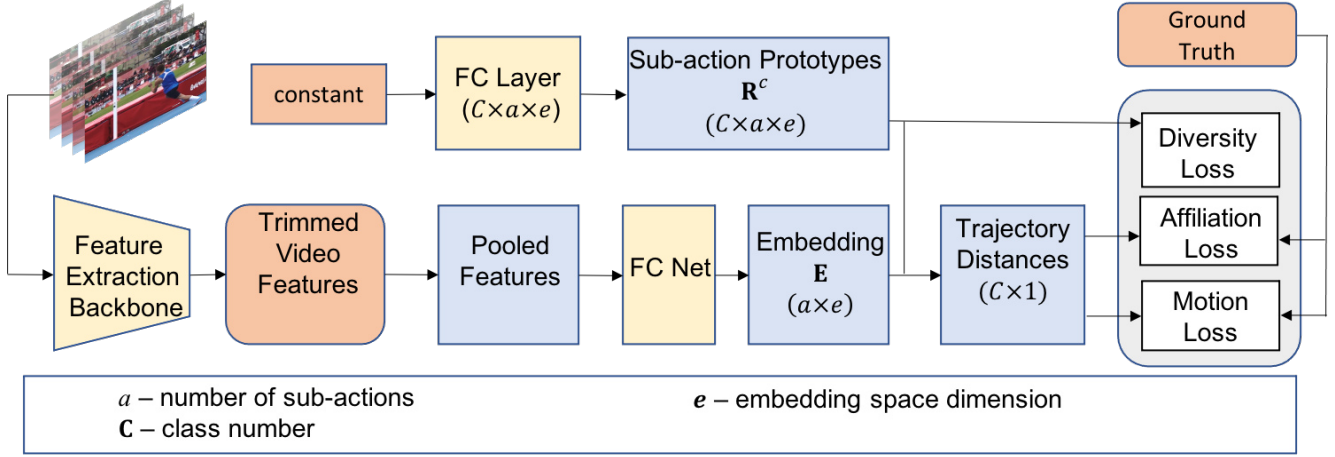


Figure 2. The TAEN model is designed to jointly train the DML-embedding (lower branch) along with the a sub-action prototypes (upper branch).

The model input is a series of video features $\mathbf{X} \in \mathbb{R}^{T \times d_{feat}}$ computed using some pretrained backbone (e.g. C3D or I3D architectures), where T is the number of video segments, namely sub-actions and d_{feat} is the dimension of the backbone output features. We divide the video into a segments and calculate the representation for each segment by average pooling: $\mathbf{X}_p \in \mathbb{R}^{a \times d_{feat}}$. Here, a represents the number of sub-actions. Next, we use \mathbf{X}_p as an input to the DML network, which consists of a few FC layers with RELU activation functions. We denote the embedding network output for certain class c as $\mathbf{E} \in \mathbb{R}^{a \times e}$, where e is the embedding space feature dimension. Typically, $e \leq d_{feat}$.

3.2. Sub-action prototypes

In order to allow joint training of multiple sub-actions, we build on an architecture that adds a secondary branch as described in Fig. 2. This branch allows computing our sub-action prototypes for each action class. One can initialize the prototypes per-class by K-means clustering. Note that at this point video segments are represented by pre-computed deep features. However we choose to initialize the sub-action prototypes by using a constant scalar as input to a FC layer that yields $C \times a \times e$ parameters, where C is the number of action classes in the training set [13]. This is equivalent to random seed clustering initialization. We reshape the fully connected network output into C segments (one for each class) denoted by $\mathbf{R}^c \in \mathbb{R}^{a \times e}, \forall c = 1, \dots, C$. We also denote $\mathbf{R}_i^c \in \mathbb{R}^e$ as the representation of the i -th sub-action of class c in the DML-embedding space. Note that \mathbf{R}_i^c is the i -th dimension of tensor representation of \mathbf{R}^c .

3.3. Loss function

From the two branches described above, we obtain the video embedding in the DML-space \mathbf{E} and the prototypes for each class \mathbf{R}^c . These representative centers denote the

learned sub-actions for our trajectories in the embedding space. Using the video embedding and the computed prototypes, we can calculate the distance associated with the trajectories in the embedding space as:

$$d(\mathbf{E}, \mathbf{R}^c) = \frac{1}{a} \sum_{i=1}^a d(\mathbf{E}_i, \mathbf{R}_i^c), \forall c = 1, \dots, C \quad (1)$$

where \mathbf{E}_i is the i -th sub-action of \mathbf{E} and $d(\cdot, \cdot)$ is some vector distance metric e.g., euclidean or cosine distance. As the distance metric d we use the cosine distance over a hypersphere, i.e., unit normalized embedding vectors \mathbf{E}, \mathbf{R}^c (see in Fig. 1). Using the prototype representations and the trajectory distance metric, we can calculate the training loss for our model. The loss consists of the following three components:

1. **Affiliation loss:** For a given embedding video \mathbf{E} , this loss minimizes the distance of the embedding sub-actions and their prototypes (jointly learned):

$$\mathcal{L}_{aff} = \sum_{i=1}^a d(\mathbf{E}_i, \mathbf{R}_i^c), \quad (2)$$

where c is the true class index (extracted from labeled data) and \mathbf{R}_i^c denotes the representative center of the i -th sub-action of class c .

2. **Motion loss:** This loss measures the deviation in the trajectory gradient, approximated by the vector of change between consecutive sub-actions:

$$\mathcal{L}_{mot} = \sum_{i=1}^{a-1} -\langle \mathbf{E}_{i+1}^c - \mathbf{E}_i^c, \mathbf{R}_{i+1}^c - \mathbf{R}_i^c \rangle, \quad (3)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product operator. Note that this loss drives the embedding space to higher diversity between sub-actions due to dependency on the

norm of vectors between sub-actions. It further drives the space toward alignment between the model and test motion vectors.

3. **Diversity loss:** This loss term aims to prevent the representation of sub-action prototypes, from collapsing into one point in the embedding space. Since short trajectories determined by close points are less discriminative, we enforce the diversity loss on the corresponding sub-actions. Semantically, this loss ensures that the sub-action prototypes in each class, will be sufficiently apart from each other. This is achieved by penalizing large correlations between different sub-action prototypes:

$$\mathcal{L}_{div} = \sum_{c=1}^C \sum_{i=1}^a \sum_{\substack{j=1 \\ j \neq i}}^a 1 - d(\mathbf{R}_i^c, \mathbf{R}_j^c) \quad (4)$$

Finally, the total loss consists of the weighted sum of the above three terms:

$$\mathcal{L}_{total} = w_{aff} \cdot \mathcal{L}_{aff} + w_{mot} \cdot \mathcal{L}_{mot} + w_{div} \cdot \mathcal{L}_{div}, \quad (5)$$

where w_{aff} , w_{mot} and w_{div} are tuned as part of the hyper-parameters tuning process. In the next section we provide additional details on our method implementation for action classification and for temporal action detection.

At test time we measure similarity between actions as a weighted distance between discrete point-wise similarity and motion:

$$d_{test} = \sum_{i=1}^a w_{aff} d(\mathbf{E}_i, \mathbf{R}_i^c) + w_{mot} \langle \mathbf{E}_{i+1} - \mathbf{E}_i, \mathbf{R}_{i+1}^c - \mathbf{R}_i^c \rangle \quad (6)$$

4. Evaluation

In the few-shot scenario, our goal is to train a network that generalizes well to new action classes. In this setup, we are given a training set that consists of labeled videos from different classes. This set is denoted as the "base set" and is only used during training. For testing, we are given a small number of videos from *new* classes that were not available in the base set. Our task is video classification in trimmed videos or temporal action detection in untrimmed videos.

4.1. Training

For the classification task, we follow the experimental set-up in previous works using the same backbone feature extractor-C3D pretrained on Sports-1M [14]. Using the pre-trained backbone on an external dataset is justified to allow fair comparison to previous works [1, 37]. For the temporal action detection experiments we pretrain our I3D backbone only on the *base* set of ActivityNet 1.2 to avoid any chance of feature "contamination". We follow the evaluation process as in [35] detailed in section 5.

4.2. Testing

In the few-shot setting, tests are determined by episodes of n -way, k -shot tasks, where at each test episode we sample k videos from n different classes (total of $n \times k$ videos) to build the support set. In our model, each episode is represented by n trajectories (one representative trajectory per-class). Representation over k -shots is computed by average pooling over corresponding sub-actions.

For the video classification task, each trimmed video is mapped to a trajectory in the embedding space and the class is derived using the nearest trajectory (see Eq. (6)). For the task of action detection, we build upon a *temporal* region proposal network [20] and follow the same process with the candidate proposals. We then filter out background segments by thresholding over the scores. The test pipeline is described in Fig. 3.

TAEN based video classification: Using the trained embedding network, we build the support set trajectories by computing sub-action prototypes for the n classes. Next, given a query video, belonging to one of the n classes, we calculate the video embedding using the trained model. Classification is then based on the nearest trajectory using the trajectory distance metric in Eq. (6).

TAEN based temporal detection: Similarly to the classification task, the base and the support set are trimmed videos that are defined by the ground truth annotations in untrimmed videos. The trajectories for each support class are calculated in the same way as in the classification task. Yet in the detection task, the query is an untrimmed video. We therefore decompose the video into temporal regions using a standard temporal action proposal method. Proposal segments are then classified and scored according to trajectory distance. The probability of an action proposal belonging to certain action class *i.e.* one of the n -way classes is set by:

$$p_m^c = \exp\left(-\frac{d^2(\mathbf{E}, \mathbf{R}^c)}{2\sigma^2}\right), \quad (7)$$

where σ is an hyper-parameter that controls the standard deviation of the probability measure. Then, using the probabilities we associate the m^{th} action proposal to one of the classes \hat{c}_m and calculate the score:

$$\hat{c}_m = \arg \max_c p_m^c, \quad s_i = p_m^{\hat{c}_m} \quad (8)$$

where s is the score of the m^{th} proposal. Eventually, background segments are rejected based on low confidence scores.

5. Experimental results

Implementation details: To train the proposed architecture we use SGD. We set the batch size $B = 30$, learning rate $\eta = 5 \cdot 10^{-4}$ and momentum $m = 0.9$ to minimize

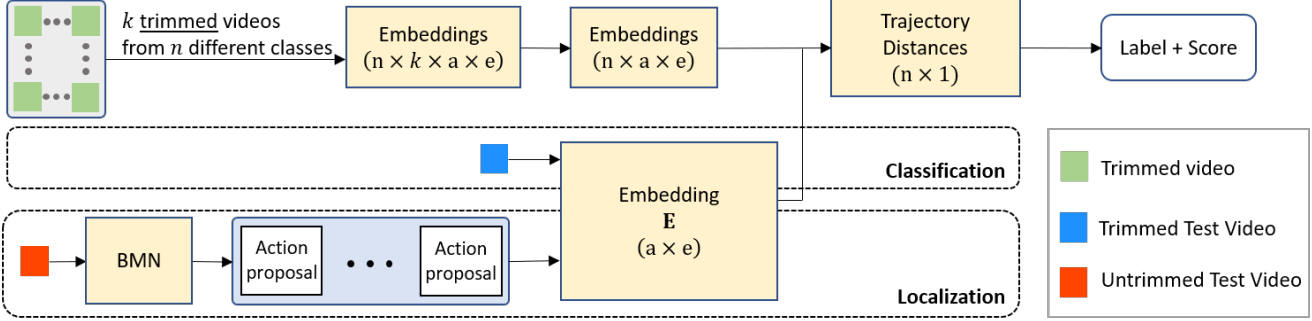


Figure 3. **Test Model Architecture:** Top branch: Using the same flow as in Fig. 2 and the trained DML, we obtain the embedding for all support set videos and the corresponding sub-actions. Middle branch - classification: For the classification task, given a query trimmed test video, the video is decomposed to a sub-actions and mapped into our DML space as a set of ordered points creating a trajectory. Classification is then performed based on the trajectory distance from the support classes using minimal distance. Lower branch - Detection (Localization): For the temporal detection task we are given one untrimmed video. We use the BMN temporal region proposal network to decompose the video into action proposals. Proposals are then treated as trimmed videos and classified in the support set. We use the trajectory distances to compute a confidence score. Rejecting background segments with low confidence score and Non-Maximal Suppression yields our final temporal temporal detection.

the loss defined in Eq. (5), in the video classification task. We used $\alpha = 0.2$ as the margin for the sub-action loss. In the classification experiments we obtained the best results with $w_{aff} = 1$, $w_{mot} = 0.5$ and $w_{div} = 20$. In the detection experiments we used the same parameters except a larger batch size of $B = 50$. In addition, we used $\sigma = 0.5$ as the probability standard deviation hyper-parameter. We obtained best performance for $a = 5$, i.e. decomposing the videos into 5 sub-actions.

5.1. Classification

We start by evaluating our architecture performance on the classification task. In [37], the authors introduced a dataset for few-shot classification which is a modification of the original Kinetics-400 dataset [16], consisting of 400 categories and 306,245 videos, covering videos from a wide range of actions and events, e.g., “dribbling basketball”, “robot dancing”, “shaking hands” and “playing violin”. The modified dataset contains videos from 100 categories out of 400 available ones randomly selected from the original Kinetics-400 dataset. Each class category contains 100 videos. In addition, the authors divided the dataset into 64, 12, and 24 non over-lapping classes for training, validation and testing respectively. In our classification evaluation, we use the same dataset and evaluation protocol as defined in [37] and followed by [1].

Feature extraction We follow the feature extraction protocol of C3D trained on Sports-1M [14], as in [1, 37] to perform a fair comparison to the previous works. Note that the feature extraction network is trained on a different dataset than Kinetics. Using the pre-trained C3D architecture, the 4096D features are extracted from the the last FC layer (i.e. FC7) of the network, corresponding to 16 consecutive frames, and then used as input for our architecture

as illustrated in Fig. 2. We divide the video into “ a ” non-overlapping segments and use average pooling to obtain the representation for each segment. Note that each video segment represents a different sub-action in the original video. \mathbf{X}_p is used as an input to our embedding network. In these experiments we used a network with two hidden layers and RELU activation function to obtain the embedding network output $\mathbf{E} \in \mathbb{R}^{a \times e}$ with $e = 2048$.

Evaluation protocol We compare our model with several baselines and two previous methods. For each class, we randomly choose k videos where $k = 1, \dots, 5$ is the number of shots. We use these $n \times k$ videos to calculate class trajectories as described in Section 3. Then, we randomly draw one additional test video belonging to one of the n classes, and associate the video embedding to one of the n classes based on the lowest trajectory distance (Eq. (6)). We repeat this episode for 20,000 iterations and evaluate our performance based on the average accuracy.

Results: Table 1 shows our results compared to previous work on the Kinetics few-shot dataset and under the same protocol, as well as a baseline and an ablation study. As our baseline we use the C3D feature space without our embedding, which represents the classification accuracy when each video is represented by a single C3D feature using max-pooling. Without our metric learning C3D shows inferior results (see Table 1). Our model benefit is manifested when actions are represented by their trajectories in the embedding space. The gaps with respect to C3D of absolute 6.85 points for 1-shot and 4.22 points for 5-shots further show the benefit of our method. Our model also outperforms previous methods of [1] and [37] while being inferior with respect to temporal alignment based method of [11]. The results from our method present a growing gap for higher number of shots, with respect to previous methods of




Figure 4. Kinetics 5-way Accuracy vs. number of sub actions. The dashed cyan line indicates the results from [1] for 5-shots.




Figure 5. Accuracy variation with number of shots and categories (n-ways) with comparison to baseline (without embedding). Shaped points show the results from [1].

[1, 37] as adding more samples (shots) allows smoother and more reliable trajectories. The ablation study shows the impact of each term in our loss function, showing that the contribution of diversity and motion terms is more significant for higher shots. The break down with vanishing affiliation weight is expected since it directly associates the actions with their relevant class.

Next, we depict in Fig. 4 the variation of accuracy with respect to sub-action granularity and number of shots. We observe that 3-5 sub-actions are sufficient to capture the action structure and enable improved recognition capability. We further show for comparison the performance of [1] for 5-shots (dashed cyan) as reference. An interesting comparison is to the raw C3D features, *i.e.* without using our embedding network. Fig. 4 shows a large margin in 1-shot compared to raw C3D features (see bold and dashed purple lines in Fig. 4). The gap between TAEN and raw C3D 1-shot results is maintained for any number of sub-actions, emphasizing the contribution of our embedding model in the 1-shot scenario.

For detailed analysis, we show in Fig. 5 the variation in accuracy of our model, with respect to number of cate-

gories n (n-ways). As the number of categories grows, the chance for error is raised, causing the decrease in accuracy. In fact, this is a typical scenario in few-shot benchmarks. The figure also shows the sharp decrease in accuracy with n for low number of shots (1,2), while for higher shots (4,5) there is a moderate degradation. This indicates the need for over 4-5 samples to learn a more reliable trajectory for an action class. The dashed purple line shows the lower bound obtained from raw C3D features. Note again the higher impact of our model especially in low shots.

5.2. Temporal Detection

In this section we illustrate an extension of our model to the temporal few-shot action detection task. To this end, we evaluate our model on the ActivityNet1.2 dataset [9]. This dataset contains roughly 10k untrimmed videos with 100 activity classes. The 100 activity classes are randomly split into 80 classes (ActivityNet1.2-train-80) for training and 20 classes (ActivityNet1.2-test-20) for testing. Our few-shot action detection network is trained on videos containing only the 80 classes in the training set, denoted by ActivityNet-train-80, and is tested on the other 20 classes in the validation set, denoted by ActivityNet-val-20, following the protocol in [35].

Evaluation: In order to fit this task to our model we use the Boundary Matching Network (BMN), a temporal region proposal method [20]. We train the BMN on ActivityNet1.2-train-80 for a class-agnostic temporal region proposal. Note that our temporal region proposal is also trained only on the base set to avoid any train-test contamination. For each test video, BMN outputs 100 proposals represented by the start and end time of the actions and associated with a confidence score. We use this score for eliminating some proposals with a score threshold below 0.2. The remaining regions are then processed by TAEN in the same way as trimmed videos. We then reject background segments based on low TAEN confidence scores obtained from (8), and use standard Non-maximum Suppression (NMS) for final temporal detection. We show results on I3D trained *only on the base set*.

For evaluation we follow the typical protocol in few shot object detection task [13]. Our benchmark contains multiple random episodes (instances of the few-shot detection tasks). We randomly choose $n = 5$ test classes in ActivityNet1.2-test-20. For each test class we randomly sample k videos (number of shots). As standard procedure, we extract the action segments from the ground truth annotations to build the support set. The rest of the videos from these classes are used as query examples with this episodic benchmark (namely, using the same support set). To get reliable test results, we randomly sample 100 different episodes, creating over 1500 query samples per class. Then we compute multi-class average precision for each

Method	Benchmark					Ablation		
	C3D ¹	CMN [37]	TARN [1]	TAM [11]	TAEN	$w_{aff} = 0$	$w_{div=0}$	$w_{mot} = 0$
1-shot	60.42	60.50	66.50	73.0	67.27	24.65	67.16	67.61
2-shot	70.10	70.00	74.56	-	74.87	25.30	74.15	74.19
3-shot	75.09	75.60	77.33	-	79.06	24.70	77.97	78.32
4-shot	77.80	77.30	78.89	-	81.78	24.52	80.01	80.55
5-shot	79.76	78.90	80.66	85.8	83.12	24.79	81.04	81.55

Table 1. Benchmark and ablation study: Video classification accuracy for different shots, compared to state of the art. Our method outperforms the C3D baseline and the prior art of [1, 37]. Yet, it is inferior to [11], that performs alignment between action segments and associated with $\mathcal{O}(N^2)$ computational cost. C3D means directly using the backbone features in test, without using DML. Ablation: ($w_{\cdot} = 0$) shows the impact for different parts of the loss function, with the affiliation loss as the main part. All of our tests were conducted with the same random seed.

Method	mAP@0.5	Avg. mAP
SMN@1 [35]	22.3	9.8
SMN@5 [35]	23.1	10
TAEN@1(I3D)	21.99	9.78
TAEN@5(I3D)	33.64	17.39

Table 2. Few-shot temporal action detection results on ActivityNet1.2 (in percentage). mAP @ tIoU threshold=0.5 and average mAP for the range [0.5, 0.95] are reported. @1 and @5 denote “1-shot” and “5-shot” respectively.

episodic task, and average over all episodes. The detection performance is measured via the standard mean average precision (mAP). The evaluation of our approach and SOTA are reported in Table 2. The results show that TAEN can provide high results also for the temporal detection task. While we achieve comparable results to the previous work in [35] in 1-shot, we outperform [35] in 5-shots by a large margin. This shows again the higher performance expected when larger samples yield more reliable trajectories.

6. Summary and future directions

In this work, we suggest a novel idea of representing actions as trajectories in a learned feature space for few shot action recognition. In this model, actions are encoded as trajectories in a metric space (in opposed to pixel-space) by a collection of temporally ordered sub-actions. The proposed architecture and loss function learn the coarse sub-action connectivity, by jointly learning the representation of the sub-actions per-class and the embedding function. The associated loss function optimizes for sub-action affiliation and motion between consecutive sub-actions, in the deep feature space. For effective *few-shot* class discrimination we suggest a trajectory distance that combines the affiliation and motion in the feature space. The proposed network requires no additional resources or fine-tuning on the target. We further extend our model to few-shot action detection in untrimmed videos. A recent study of [11] shows that using sub-action alignment improves the results. We argue that our method is more efficient by finding the optimal proto-

types without the Dynamic Time Warping alignment with $\mathcal{O}(N^2)$ cost. Our method, requiring learning of only few FC layers, runs in average at rate of 6.1 [*msec/video*] classification (excluding feature computation), on Tesla K-40, 12G.

Combining our proposed approach with fine-tuning on the few support examples of the novel categories is a good orthogonal direction that could be interesting to explore in a follow up work.

Acknowledgement

We want to thank Eli Schwartz and Leonind Karlinsky for providing us the RepMet repo.

References

- [1] Mina Bishay, Georgios Zoumpourlis, and Ioannis Patras. Tarn: Temporal attentive relation network for few-shot and zero-shot action recognition. *BMVC*, 2019.
- [2] P. Abbeel C. Finn and S. Levine. Model-Agnostic. Meta-learning for fast adaptation of deep networks. In *arXiv:1703.03400*, 2017.
- [3] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A. Ross, Jia Deng, and Rahul Sukthankar. Re-thinking the faster r-cnn architecture for temporal action localization. *CVPR*, 2018.
- [4] S. K. Dwivedi, V. Gupta, R. Mitra, S. Ahmed, and A. Jain. ProtoGAN: towards few shot learning for action recognition. *ICCVW*, 2019.
- [5] J. Shtok S. Harary M. Marder A. Kumar R. Feris R. Giryes E. Schwartz, L. Karlinsky and A. M. Bronstein. Delta-Encoder: an effective sample synthesis method for few-shot object recognition. In *NIPS*, 2018.
- [6] Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, and Bryan Russell. Actionvlad: Learning spatio-temporal aggregation for action classification. *CVPR*, 2017.
- [7] B. Hariharan and R. Girshick. Low-shot visual recognition by shrinking and hallucinating features. In *ICCV*, 2017.
- [8] B. Hariharan and R. Girshick. Low-shot visual recognition by shrinking and hallucinating features. *ICCV*, 2017.

- [9] B. G. Fabian Caba Heilbron, Victor Escorcia, and J. C. Niebles. A large-scale video benchmark for human activity understanding. In *CVPR*, 2015.
- [10] Rui Hou, Rahul Sukthankar, and Mubarak Shah. Real-time temporal action localization in untrimmed videos by sub-action discovery. *BMVC*, 2017.
- [11] K.C. Jingwei, Z. Cao, C. Chang, and J. Niebles. Few-shot video classification via temporal alignment. *CVPR*, 2020.
- [12] Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. Few-shot object detection via feature reweighting. In *ICCV*, 2019.
- [13] Leonid Karlinsky, Joseph Shtok, Sivan Harary, Eli Schwartz, Amit Aides, Rogerio Feris, Raja Giryes, and Alex M Bronstein. Repmet: Representative-based metric learning for classification and few-shot object detection. *CVPR*, 2019.
- [14] Andrej Karpathy, George Toderici, Sanketh Shetty, Rahul Sukthankar Thomas Leung, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.
- [15] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, and Paul Natsev. The Kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [16] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv:1705.06950*, 2017.
- [17] Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. *CVPR*, 2019.
- [18] Guy Lev, Gil Sadeh, Benjamin Klein, and Lior Wolf. Rnn fisher vectors for action recognition and image annotation. *ECCV*, 2016.
- [19] Wenbin Li, Lei Wang, Jinglin Xu, Jing Huo, Yang Gao, and Jiebo Luo. Revisiting local descriptor based image-to-class measure for few-shot learning. In *CVPR*, 2019.
- [20] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. BMN: Boundary-Matching Network for Temporal Action Proposal Generation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3889–3898, 2019.
- [21] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. BSN: Boundary Sensitive Network for Temporal Action Proposal Generation. In *ECCV*, 2018.
- [22] Daochang Liu, Tingting Jiang, and Yizhou Wang. Completeness modeling and context separation for weakly supervised temporal action localization. *CVPR*, 2019.
- [23] X. Chen N. Mishra, M. Rohaninejad and P. Abbeel. Meta-learning with temporal convolutions. In *arXiv:1707.03141*, 2017.
- [24] Phuc Nguyen, Ting Liu, Gautam Prasad, and Bohyung Han. Weakly supervised action localization by sparse temporal pooling network. *ECCV*, 2018.
- [25] Sujoy Paul, Sourya Roy, and Amit K Roy-Chowdhury. W-TALC: weakly-supervised temporal activity localization and classification. *ECCV*, 2018.
- [26] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014.
- [27] J. Snell, K. Swersky, and R. S. Zemel. Prototypical networks for few-shot learning. *NeurIPS*, 2017.
- [28] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H.S. Torr, and Timothy M. Hospedales. Learning to compare: Relation network for few-shot learning. *CVPR*, 2018.
- [29] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [30] Heng Wang, Alexander Klaser, Cordelia Schmid, and Cheng-Lin Liu. Dense trajectories and motion boundary descriptors for action recognition. In *International Journal of Computer Vision*, 2013.
- [31] Davis Wertheimer and Bharath Hariharan. Few-shot learning with localization in realistic settings. In *CVPR*, 2019.
- [32] Y. Xian, B. Korbar, M. Douze, B. Schiele, Z. Akata, and L. Torresani. Generalized many-way few-shot video classification. *ECCV Workshops*, 2020.
- [33] Huijuan Xu, Abir Das, and Kate Saenko. R-C3D: Region Convolutional 3D Network for Temporal Activity Detection. *ICCV*, 2017.
- [34] H. Xu, B. Kang, X. Sun, J. Feng, K. Saenko, and T. Darrel. Similarity r-c3d for few-shot temporal activity detection. *arxiv.org/pdf/1812.10000*, 2018.
- [35] Hongtao Yang, Xuming He, and Fatih Porikli. One-shot action localization by learning sequence matching network. *CVPR*, 2018.
- [36] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. *ICCV*, 2017.
- [37] Linchao Zhu and Yi Yang. Compound memory networks for few-shot video classification. *ECCV*, 2018.
- [38] Joao Carreira ; Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.