

SSCAP: Self-supervised Co-occurrence Action Parsing for Unsupervised Temporal Action Segmentation

Zhe Wang¹ * Hao Chen² Xinyu Li² Chunhui Liu²
Yuanjun Xiong² Joseph Tighe² Charless Fowlkes^{1,2}

¹Dept. of CS, UC Irvine

²Amazon Web Services

Abstract

Temporal action segmentation is a task to classify each frame in the video with an action label. However, it is quite expensive to annotate every frame in a large corpus of videos to construct a comprehensive supervised training dataset. Thus in this work we propose an unsupervised method, namely SSCAP, that operates on a corpus of unlabeled videos and predicts a likely set of temporal segments across the videos. SSCAP leverages Self-Supervised learning to extract distinguishable features and then applies a novel Co-occurrence Action Parsing algorithm to not only capture the correlation among sub-actions underlying the structure of activities, but also estimate the temporal path of the sub-actions in an accurate and general way. We evaluate on both classic datasets (Breakfast, 50Salads) and the emerging fine-grained action dataset (FineGym) with more complex activity structures and similar sub-actions. Results show that SSCAP achieves state-of-the-art performance on all datasets and can even outperform some weakly-supervised approaches, demonstrating its effectiveness and generalizability.

1. Introduction

Temporal action segmentation aims to classify each frame in an untrimmed video with an action label. The task is a key step in understanding the structure of complex activities in videos. A recent study [50] also shows that the fine-grained sub-action segmentation can help many of other tasks such as action recognition and detection. Today’s high performing solutions for temporal action segmentation require full frame-level supervision [15, 24, 60]. However, collecting per-frame annotations in untrimmed videos is extremely expensive and often impractical. Some recent work aim to overcome this by leveraging weak supervision to avoid the need of the full frame-level annota-

tions [16, 34, 46]. These solutions assume that at least the number or the order of the sub-actions are given for each video. Even though they do not use full supervision, they still need some annotations that are costly to acquire, especially when expert knowledge on fine-grained atoms of activities (i.e., sub-actions) is required. To further reduce the annotation cost, we ask the question, can we do temporal action segmentation in a fully unsupervised way?

Unsupervised temporal action segmentation is a very challenging problem. To parse long video sequences into sub-actions in a semantically meaningful way without any supervision, it requires a solution to not only extract highly distinguishable visual representations for each individual frame, but also to capture the temporal relations among frames and sub-actions, so that the solution can well estimate the number and the order of the occurrence of each sub-action (i.e., the temporal path), which is typically provided in the weakly-supervised setting. The problem is even more challenging when dealing with videos that contain activities with complex structures and recurrence of sub-actions. Today there are still no effective solutions to handle these challenges and the performance of existing unsupervised temporal action segmentation solutions [1, 30, 48] is still far from desired.

In this work, we propose a Self-Supervised Co-occurrence Action Parsing method, namely SSCAP, for unsupervised temporal action segmentation. Specifically, we introduce self-supervised methods to extract features that are more temporal distinguishable. On top of it, we propose a Co-occurrence Action Parsing (CAP) algorithm to decode the frames into sub-actions by leveraging the estimated prior of the co-occurrence relations of sub-actions. The prior, though estimated over the whole dataset in an unsupervised manner, can statistically depict the natural structure of activities and hence help better estimate the number and the order of the occurrence of each sub-action during inference. Moreover, CAP takes recurrence of sub-actions into account in estimating the temporal path, and hence is able to handle complex structures of activities.

*This work was done during internship at Amazon

Results show that SSCAP achieves state-of-the-art for unsupervised action segmentation, and even outperforms some of the weakly-supervised approaches. Besides evaluating on traditional action segmentation benchmarks such as Breakfast [27], Salad [51], we additionally test SSCAP on the more challenging recent dataset FineGym [49]. Actions in traditional benchmarks are well structured - they are goal-directed (actors have clear goals like “making salad” or “preparing scrambled egg”) and can be decomposed into clear steps, with one step being as the prerequisite for another one (e.g., you have to “get” an egg before “scramble” it). However, actions in FineGym have more complex structures - they are fast moving gymnastics actions, where several actors may perform the same sub-action multiple times in a single video. There are more similar fine-grained sub-actions across different activities as well (for example: both ‘floor exercise’ and ‘vault’ events involve a number of basic sub-action like ‘running’ and ‘jumping’). Results show that comparing to the existing unsupervised action segmentation method, SSCAP can achieve significant performance improvements on this challenging dataset, demonstrating that SSCAP is especially superior in segmenting videos with complex activities and is more generalizable to different scenarios. To summarize, our main contributions are:

(1) We design a new solution called SSCAP to achieve temporal action segmentation in a fully unsupervised manner that can significantly reduce the annotation cost. SSCAP not only achieves state-of-the-art performance under the unsupervised setting, but also outperforms some weakly-supervised solutions;

(2) We introduce self-supervised learning in solving unsupervised action segmentation. While self-supervised learning has been widely exploited on other tasks such as image classification and action recognition [6, 14, 41], the application of it on action segmentation has been largely unexplored. We explore a few of today’s most popular video based self-supervised learning approaches and show superior performance over the current state-of-the-art features in action segmentation. We further show that those self-supervised approaches that learn temporal characteristics outperform other self-supervised methods that focus on spatial transformation, which proves that modeling temporal relations is a key for accurately segmenting the frames into sub-actions;

(3) We propose a co-occurrence action parsing (CAP) algorithm, which takes into account the co-occurrence and recurrence of sub-actions calculated based on global statistics of all the videos, to estimate the temporal path of sub-actions in each video in a more precise and general way, and therefore outperforms today’s solutions. Furthermore, we are the first to study action segmentation on FineGym, the emerging and more challenging dataset that contains more complex activities structures. Results show that the design

of CAP achieves superior performance on this dataset compared to existing solutions.

2. Related Work

Temporal action segmentation is a task to classify each frame in an untrimmed video with a label of action. It is a key for understanding structures of complex activities. While current deep learning techniques [37, 65, 25, 18, 39, 36, 61, 40, 52, 59] have achieved great success on other action understanding tasks such as action recognition [7, 17, 66, 56, 68, 57, 10, 62] and action detection [38, 67, 64, 22], temporal action segmentation is still far from being well-explored. Today’s high performing solutions for temporal action segmentation require full frame-level supervision [15, 24, 28, 60, 32, 31, 54]. However, acquiring annotations on fine-grained sub-actions for each frame is expensive. To reduce the amount of labeled data needed, [9] proposes a domain adaptation approach based on self-supervised auxiliary tasks and achieves good performance. More recent works focus on developing weakly-supervised solutions: instead of having supervision for every frame, only the list (with order information) [4, 33, 46, 44, 12, 23, 8] or the set (without order information) [16, 34, 45] of the sub-actions in each video are required during training. Other work uses script or instructions for weak annotation [2]. Unsupervised temporal action segmentation, which does not have any supervision signal per video, starts to receive growing interests these days, to further reduce the dependence on labeled data [13, 48, 1, 30, 53, 35, 47]. However, these solutions do not fully explore the self-supervised feature, do not take into account semantic correlations among sub-actions and the recurrence of sub-actions in some complex activity structures. The performance of them is still far from desired.

Self-supervised learning of visual representation is a hot topic today, and has shown success in some of action understanding tasks [5, 6, 14, 19]. Many different video properties have been used as supervision to build good representations, such as: cycle consistency between video frames [58]; distinguishing between a video frame sequence and a shuffled version of it [41]; predicting the order of the shuffled clips [63]; solving space-time cubic puzzles [26]; classifying the speed of the video [14, 3]; recognize the rotation of the video [20]. Another common task is predicting the future, either by predicting pixels of future frames [11], or an embedding of a future video segment [21]. While today’s action segmentation solutions [15, 30, 60] still mainly rely on hand-crafted features such as IDT [55], Fisher Vector (FV) [42], or pre-trained features such as I3D on Kinetics dataset [7], we believe that introducing self-supervised methods in feature representation learning for unsupervised action segmentation can significantly boost the performance, as it can not only bring strong discrimi-

native power to the features, but also bring a chance to improve the temporal characteristics of the features.

3. Technical Approach

Temporal action segmentation is a task to classify each frame in the video to a sub-action. Specifically, given M videos, each of them is represented by X_m which contains N_m frames. The task is to predict the action label $l_{mn} \in 1, \dots, K$ for each video frame x_{mn} . In this paper, we study the task in an unsupervised setting, i.e., the label l_{mn} for each frame is unknown during model training.

Fig. 1 is an overview of our SSCAP approach. Briefly: given a set of videos, our feature generation module extracts frame features with a self-supervised representation learning method. Extracted frame features are then sent to the co-occurrence action parsing (CAP) module, which clusters the frames and decodes them into segments of unlabeled sub-actions. Finally, the semantic action label is assigned to each cluster by using Hungarian algorithm [29]. We introduce each module in details in the following sections.

3.1. Feature Generation with Self-supervision

Our co-occurrence action parsing (CAP) module relies on the feature representation to have some semantic meaning so that the clustering step in CAP can effectively separate different semantic concepts in the video. In this work we explore a number of self-supervised feature learning methods that have been developed for image and video classification tasks. We apply these methods for the task of video segmentation and study their relative effectiveness for that task v.s. the classification task they were designed for. Since one essential factor of the feature design for the temporal action segmentation is to capture the temporal characteristics, we want to further understand whether self-supervised learning methods built based on temporal augmentations (e.g, adjusting the frame rate, shuffling the frame or clip order, etc.) are also effective on temporal segmentation or not.

With these questions in mind, we explore three of today’s most popular video self-supervised learning methods and compare them with the feature designs used today in unsupervised action segmentation, including: IDT [55], FV [42], pre-trained I3D features [7], and the continuous temporal features that achieves SOTA performance today [30]. Specifically, we study the following self-supervised methods: (1) *SpeedNet* [14, 3], which augments the frame rate of the video and predict it; (2) *ShuffleLearn* [41], which predict whether the order of frames in the sequence is shuffled or not; (3) *RotationNet* [20], which rotates the whole video sequence by several pre-defined degrees and predict which degree the video is rotated. Note that *RotationNet* is a method based on spatial augmentation, while *ShuffleLearn* and *SpeedNet* are with temporal augmentation that

are expected to have stronger temporal reasoning capability. Briefly we summarize the designs of these three methods in our work:

SpeedNet [14, 3]. Given a clip C with fixed number of N frames: $C = (f_1, f_2, \dots, f_N)$, we pre-define four frame rate settings [2, 4, 8, 15]. Each time when we generate a clip we randomly select one out of these four settings. Note that we keep the number of frames in each clip as the same, hence the time-range of the clip varies according to the frame rate selected. We train the network to predict which frame-rate the clip is sampled from (i.e., a 4-way classification).

ShuffleLearn [41]. Given a clip C with fixed number of N frames: $C = (f_1, f_2, \dots, f_N)$, we shuffle M frames ($0 < M < N$) in a random order. Each time when we generate the clip, we flip the coin to decide whether to shuffle it or not. We train the network to predict whether a clip is shuffled or not.

RotationNet [20]. Given a clip C with fixed number of N frames: $C = (f_1, f_2, \dots, f_N)$, we pre-define four rotation degree settings: [0, 90, 180, 270]. We also add some randomness when we rotate, i.e., in $[-30, 30]$ degree. Each time when we generate clip we randomly select one out of these four settings and rotate the whole video sequence with the selected degree. We train the network to predict which of the settings is the rotation on (i.e., a 4-way classification).

3.2. Co-occurrence Action Parsing

Our co-occurrence action parsing (CAP) algorithm consists of two steps: clustering and decoding. Extracted frame-level features from Section 3.1 are first clustered. Since we don’t have ground truth annotations on sub-actions, here we aim at generating a number of clusters, with each of them representing a unique but unlabeled sub-action. After clustering, we then generate three outputs: (1) a *video score matrix* that represents the predicting score of each frame to each cluster for the whole video; (2) a *co-occurrence matrix* that captures the chance of any pair of clusters co-existing in videos; and (3) a *cluster temporal location histogram* that captures the distribution of temporal locations of each cluster in videos. In decoding, we first apply the co-occurrence matrix onto the video score matrix to generate a *refined video score matrix*. Then we estimate a *temporal path* with the cluster temporal location histogram and decode the frames to final segments based on the temporal path. The overall algorithm can be seen in Fig. 1. Below we introduce the clustering and decoding in details.

3.2.1 Clustering

Given an activity, and assuming there are K different sub-actions (K is known following the literature [30, 48]), we extract all the frame-level features for all videos labeled with this activity, and cluster them into K clusters using




Figure 1: Overview of our SSCAP approach.

k-means. To get the probability of the frame x_{mn} that belongs to the cluster k , we estimate a Gaussian distribution for each cluster $k \in K$:

$$p(x_{mn}|k) = \mathcal{N}(x_{mn}; \mu_k, \Sigma_k) \quad (1)$$

We generate three outputs from the clustering step: (1) video score matrix; (2) co-occurrence matrix and (3) cluster temporal location histogram.

Video score matrix. Given a video X_m that contains N_m frames, we create a video score matrix $S_m \in \mathbb{R}^{K \times N_m}$ to hold the probabilities calculated from Eq. (1), where K is the number of clusters. Each element $S_m[k, n]$, with $k \in \{1, 2, \dots, K\}$ and $n \in \{1, 2, \dots, N_m\}$ represents the probability of the n -th frame assigned to the cluster k .

Co-occurrence matrix. The purpose of creating the co-occurrence matrix is to capture the correlations among sub-actions underlying the native structure of activities from the global video sets. Given an activity, there are usually strong correlations among its sub-actions across all videos. For example, in the activity ‘‘Fried Egg’’, ‘‘fry egg’’ always occurs with ‘‘take egg’’ and ‘‘crack egg’’. While previous solutions do not build models to capture such global patterns, we leverage a co-occurrence matrix to capture the co-existing relations among sub-actions. Note that, since we are dealing with an unsupervised setting, we don’t have ground truth annotations on sub-actions, hence in the real implementation, what we calculate is the co-occurrence between two clusters i and $j \in \{1, 2, \dots, K\}$, instead of the real sub-actions. The way to generate the co-occurrence matrix is simple: for each pair of cluster i and j , we iterate on all the videos and count the number of times that cluster i and j appear together in each video as $C(i, j)$. We also count the number of videos that contains the cluster i as $C(i)$.

We then calculate a conditional occurrence probability as $P(j|i) = C(i, j)/C(i)$, which represents the probability of cluster j appears in a video given that cluster i appears. We calculate $P(i|j)$ for all pairs of clusters (i, j) in this way and generate the $K \times K$ dimension co-occurrence matrix. Details are shown in Algorithm 1. Since the co-occurrence matrix is learned from all videos, it is capturing the global patterns of the activity structures. The co-occurrence matrix is used for refining the video score matrix, which is introduced in Section 3.2.2.

Cluster temporal location histogram. The purpose of generating the cluster temporal location histogram is to estimate where each cluster generally locates in temporal dimension. Given the frames associated to each cluster from k-means, we compute the histogram of the temporal location for each cluster based on the timestamp of each frame in the video. Specifically, for a frame x_{mn} , i.e., the n -th frame in the video X_m , we calculate its relevant timestamp as $t_{x_{mn}} = n/N_m$, where N_m is the total number of frames in video X_m . Then for each cluster k , we calculate the histogram on $t_{x_{mn}}$ for all the frames x_{mn} from the videos that are clustered to k . The histogram indicates the potential temporal location of the cluster k in the videos, and will be used for temporal path decoding introduced in Section 3.2.2. Note that in real cases a sub-action can appear multiple times in different parts of videos, so each cluster might have more than one significant bin in the cluster temporal location histogram.

3.2.2 Decoding

After clustering all the frames into K clusters, the next step is to decode them into segments of the sub-actions. There

are two main steps, namely, (1) refine the video score matrix with the co-occurrence matrix; (2) estimate the temporal path of the clusters and decode frames to segments.

Refining the video score matrix. The first step in the decoding is to generate the refined video score matrix $R_m \in \mathbb{R}^{K \times N_m}$ for each video X_m . There are two main issues with directly using the original video score matrix S_m in decoding: (1) the video score matrix is generated purely based on the frame feature representations, it does not capture any global patterns of the activity structures across all videos, and does not contain any correlation information among sub-actions; (2) the clustering is conducted on K clusters, which is the total number of different sub-actions in the given activity. In many cases, only a subset of K will appear in each individual video due to the diversity of the activities being performed by different persons. Directly using video score matrix may lead to the issue of over-segmenting video to the non-existing classes in current video.

To address these issues, we apply the co-occurrence matrix P to refine the score matrix S_m . Specifically, given a video X_m and for each cluster k , we first calculate the number of frames belonging to k in the video, i.e., $c(k)$, and define the ratio of frames for cluster k as $r(k) = c(k)/N_m$, where N_m is the total number of frames in the video.

Then, in each iteration, we pick the cluster k^* with the current largest ratio of frame $r(k)$ and update the video scores $S_m[j, :]$ for the remaining clusters j (i.e., $r(j) \leq r(k^*)$) with the corresponding co-occurrence probability $P(j|k^*)$. Then we pop out k^* and repeat the process for the remaining clusters. Details are in Algorithm 1. By doing so we add global information from the whole dataset gradually into each individual video. Scores of clusters that are not strongly correlated with the dominating clusters in this video will be attenuated. The probabilistic update is inspired by Kalman filter where the new probability $R_m[j, :]$ is updated based on the observation $P(j|k^*)$ and old probability $S_m[j, :]$.

Estimating the temporal path and decoding. Having the refined video score matrix R_m , the final step is to decode the frame sequences into the sub-action segments. Assume that we have a known temporal path of the clusters $k_1 \rightarrow \dots \rightarrow k_T$, then the decoding is to maximize the probability of the sequence following this path based on our refined video score matrix, i.e.,

$$\hat{l}_1^{N_m} = \underset{x_{11}, \dots, x_{N_m}}{\operatorname{argmax}} \prod_{n=1}^{N_m} p(x_{mn}|l_n)p(l_n|l_{n-1}), \quad (2)$$

where, $p(x_{mn}|l_n = k) = R_m$ is the probability of the frame x_{mn} belonging to the cluster k , obtained from the refined video score matrix, and $p(l_n|l_{n-1})$ are the transition probabilities of moving from the label l_{n-1} at frame $n-1$ to the

Algorithm 1: Generating the Co-occurrence Matrix and the Refined Video Score Matrix.

Data: Video score matrix S_m for video X_m , with $S_m[k, n]_{k=1, n=1}^{K, N_m}$, number of clusters: K .

Result: Co-occurrence matrix: $P(i, j)_{i, j=1}^K$, Refined video score matrix R_m , with scores $R_m[k, n]$.

Generate Co-occurrence Matrix

- Iterate all the videos, count the times each cluster appears $C(i)_{i=1}^K$, and the times different clusters co-occur together $C(i, j)_{i, j=1}^K$. Normalize to make it as conditional probability $P(j|i) = C(i, j)/C(i)$

Refined Video Score Matrix

- Initialization: $\mathcal{G} \leftarrow k_0$ (k_0 is the cluster with the largest ratio of frames $r(k_0)$ in current video).

- $k^* = k_0$

while $\text{len}(\mathcal{G}) \leq K$ and $r(k^*) > 0$ **do**

1. For each remaining cluster $j \notin \mathcal{G}$:

- update the video score matrix conditioned on the previous selected cluster k^* :

$$R_m[j, n] = P(j|k^*) \cdot S_m[j, n]$$

2. Select the next cluster: $k^* \leftarrow \arg \max_j r(j)$.

3. Update: $\mathcal{G} \leftarrow \mathcal{G} \cup \{k^*\}$

end

- Return: P, R_m .

next label l_n at frame n . Given a temporal path, Eq. (2) can be solved with a Viterbi algorithm [43].

The temporal path estimation is the key for the decoding. While in weakly-supervised action segmentation the temporal path (i.e., the order of sub-actions) is usually assumed to be known, in unsupervised setting it can only be estimated, which is challenging. Moreover, many of literature work today assume a strict single directional temporal path with each sub-action can appear at most once in the path. Such an assumption may suit for dataset with simple activity structures, e.g., Breakfast [27], however, begins to fail when the structure of activities gets more complex. For instance, in some activities, a single sub-action may occur multiple times, and the order of the sub-actions can be bi-directional (i.e., sub-action i can go to sub-action j and then go back to i). Such sub-action recurrence brings further difficulties to precisely estimate the temporal path.

We design a *multi-occur temporal path* estimation algorithm that is applicable to both simple and complex cases, using the *cluster temporal location histogram* introduced in Section 3.2.1. For each cluster, we select the bins from temporal location histogram with counts (normalized to the total counts across all the bins) larger than a certain value. Then we concatenate the selected bins from all the clusters and order them into a time sequence based on the temporal location of each bin, which is calculated as $T_i = \sum t_{x_{mn}^i} / N_i$, where $t_{x_{mn}^i}$ is the relevant timestamp of frame x_{mn}^i that fall into bin i , and N_i is the total number of frames

in bin i . The order of bins with their corresponding cluster labels forms our temporal path. Our multi-occur temporal path is able to capture the multi-time occurrence for each cluster as well as the bi-direction transition among different clusters, because more than one bins from each cluster can be selected.

4. Experiments

Datasets. We evaluate our approach on two most widely used datasets in action segmentation: Breakfast [27] and 50Salads [51], and an emerging fine-grained action dataset FineGym [49]. **Breakfast** [27] is a large-scale dataset that consists of ten different complex activities of performing common kitchen activities with approximately eight sub-actions per activity class. The duration of the videos varies significantly, e.g. ‘coffee’ has an average duration of 30 seconds while ‘cooking pancake’ takes roughly 5 minutes. **50Salads** [51] dataset contains 4.5 hours of different people performing a single complex activity, making mixed salad. Compared to the other datasets, the videos are much longer with an average video length of 10k frames. We also observe that a single sub-action can occur multiple times along the video in this dataset, which seldom happens in Breakfast. **FineGym** [49] is an emerging dataset provided with fine-grained sports action annotations. Though it is originally proposed for fine-grained action recognition, it contains rich hierarchical information including annotations of start and end for the fine-grained sub-actions, which is suitable for action segmentation evaluation. These sports videos contain several athletes performing the same action and their replays. The structure of the videos are more complex and it includes many recurrences of a single sub-action, and includes more similar fine-grained sub-actions (up to 35), which makes the dataset more challenging for action segmentation than typical ones used today. We use Gym99 setting in original paper.

Evaluation Metrics. We use the Hungarian algorithm [29] to get a one-to-one matching that maximizes the similarity between all the clusters and sub-action classes for all the videos in evaluation. The number of total sub-actions within each dataset is known and used to define the number of clusters K . We report frame-wise accuracy as the *mean over frames (MoF)*, indicating the percentage of frames correctly labelled. We also report *F1 score* the same as [30] to measure the quality of the temporal segmentation.

4.1. Comparison with SOTA

We compare our SSCAP with SOTA solutions on Breakfast in Table 1, 50Salads in Table 2 and FineGym in Table 3. For all the datasets, We use SpeedNet [14] trained on the Kinetics-400 [7] as the self-supervised feature in our

<i>Breakfast</i>	MoF	F1 score
Unsupervised setting		
GMM [48]	0.346	-
LSTM + AL [1]	0.429*	-
CTE [30]	0.418	0.264
VTE-UNET [53]	0.481	-
ASAL [35]	0.525	0.379
Our SSCAP	0.511	0.392
Weakly-supervised setting		
Action Sets [45]	0.284	-
NNviterbi [46]	0.430	-
SCT [16]	0.304	-
SetViterbi [34]	0.408	-
EnergySeg [33]	0.630	-
Fully-supervised setting		
HTK [28]	0.259	-
GTRM [24]	0.650	-
MS-TCN [15]	0.663	-
BCN [60]	0.704	-

Table 1: Results on Breakfast.

<i>50Salads</i>	MoF	F1 score
Unsupervised setting		
LSTM + AL [1]	0.606*	-
CTE [30]	0.355	-
VTE-UNET [53]	0.306	-
ASAL [35]	0.392	-
Our SSCAP	0.414	0.303
Weakly-supervised setting		
NNviterbi [46]	0.494	-
EnergySeg [33]	0.547	-
Fully-supervised setting		
HTK [28]	0.247	-
GTRM [24]	0.826	-
MS-TCN [15]	0.734	-
BCN [60]	0.844	-

Table 2: Results on 50Salads. Note that ‘*’ uses a per-video instance level ‘cluster to ground-truth’ mapping, which uses additional video-level information, while we follow other work to use ‘cluster to ground-truth’ mapping in a global manner in evaluation and do not use any additional information [53].

SSCAP. Details on feature comparison can be found in Section 4.3. We can see SSCAP achieves SOTA performance on both Breakfast and 50Salads datasets in the unsupervised setting. SSCAP on Breakfast even outperforms most of the weakly-supervised solutions.

Since FineGym is new and we are the first to conduct action segmentation on it, there is no existing related results

<i>FineGym</i>	MoF	F1 score
Baseline [30]	0.294	0.167
Our SSCAP	0.666	0.297

Table 3: Results on FineGym.

yet. We generate the baseline by running the SOTA approach [30] on FineGym. We can see that on this challenging dataset, SSCAP achieves significant improvement on both MoF and F1 compared to the baseline solution, demonstrating the effectiveness of our designed solution in handling videos with more complex structures. Overall the consistent improvement across all the benchmarks shows that our approach is able to generalize to different cases.

4.2. Ablation Study on SSCAP

Next, we conduct ablation studies on SSCAP. Three main units are investigated: (1) Self-supervised Feature Learning; (2) Co-occurrence Matrix; (3) Multi-occur Temporal Path. Results are summarized in Table 4.

Effectiveness of Self-supervised Learning. From Table 4 we observe that applying self-supervised feature learning consistently improves performance over the baseline using continuous temporal features [30], on all of the three datasets, demonstrating the value of self-supervised feature learning on the action segmentation task.

Effectiveness of Co-occurrence Matrix. We evaluate the influence of using the co-occurrence matrix to refine the video score matrix in Table 4 on all datasets. From the table, our co-occurrence matrix brings performance improvement over using original video score on all of the three datasets and on both *MoF* and *F1*. The improvement on FineGym dataset is more notable than the other two, indicating the importance of using the co-occurrence matrix while handling more complex scenarios.

Effectiveness of Multi-occur Temporal Path. We then assess the impact of using the proposed multi-occur temporal path in decoding. Results on all three datasets are in Table 4. We can see the designed multi-occur temporal path decoding improves both *MoF* and *F1* on 50Salads and FineGym notably, but does not on Breakfast. We found that the recurrence of sub-actions in a video happens more frequently in 50Salads and FineGym, while in Breakfast, most of the sub-actions only occur once in each video, therefore, the multi-occur temporal path is not as helpful on Breakfast as on the other datasets. The improvement on FineGym is significant, representing that multi-occur temporal path is superior in segmenting activities with complex structures. It is worth to mention that on the Breakfast dataset, our multi-occur temporal path algorithm generates exactly the same temporal path (with each sub-action occurring only once) automatically as the path generated from turning off the multi-occur algorithm, representing that our algorithm

Dataset	SS	C-Matrix	M-T-Path	MoF	F1
<i>Breakfast</i>				0.418	0.264
	✓			0.508	0.391
	✓	✓	✓	0.511	0.392
				0.511	0.392
<i>50Salads</i>				0.355	-
	✓			0.372	0.281
	✓	✓	✓	0.378	0.290
				0.414	0.303
<i>FineGym</i>				0.294	0.167
	✓			0.425	0.246
	✓	✓	✓	0.442	0.248
				0.666	0.297

Table 4: Ablation study on all datasets. ‘SS’ denotes whether to use the *self-supervised feature learning* (SpeedNet feature trained on K400); ‘C-Matrix’ denotes whether to use the *co-occurrence matrix* to refine the video score matrix; and ‘M-T-Path’ denotes whether to use our designed *multi-occur temporal path* in decoding. If none of them is selected then it’s a baseline setting with the continuous temporal feature [30].

is self-applicable to different scenarios.

4.3. Comparison of Feature Learning

In the section, we explore what type of self-supervised design best fits for the unsupervised action segmentation task. We study three different self-supervised feature learning methods: (1) SpeedNet [14], (2) RotationNet [20], and (3) ShuffleLearn [41], and compare the results on Breakfast dataset in Table 5. All the numbers in the table are generated without using the co-occurrence matrix and the multi-occur temporal path, so that we can purely evaluate the effectiveness of features without bringing any additional factors.

We conduct the experiments on self-supervised training (1) only on Kinetics dataset (i.e., row (d) - (f)), (2) only on Breakfast dataset (i.e., row (g) - (i)) and (3) first on Kinetics dataset and then on Breakfast dataset (i.e., row (j) - (l)), to further explore that whether we need to train on target dataset (i.e., Breakfast in our case) to get a good feature representation for the task. From the table we can find that:

(1) Based on experiment (a) - (l), we conclude that SpeedNet trained on Kinetics shown in row (d) performs the best;

(2) Based on experiment (d) - (l), RotationNet consistently performs worse than SpeedNet and ShuffleLearn. This may due to the fact that RotationNet is using self-supervised signal from spatial dimension, which may lack of ability to model the temporal characteristics between frames. SpeedNet and ShuffleLearn, however, use temporal augmentation as the self-supervised signal, so that they

	Feature Description	MoF	F1
Baseline			
(a)	IDT [30]	0.316	-
(b)	K400 I3D [30]	0.251	-
(c)	CTE [30]	0.418	0.264
Self-supervised on K400			
(d)	K400 SpeedNet	0.508	0.391
(e)	K400 RotationNet	0.328	0.317
(f)	K400 shuffleLearn	0.339	0.328
Self-supervised on Breakfast			
(g)	Breakfast SpeedNet	0.344	0.327
(h)	Breakfast RotationNet	0.307	0.319
(i)	Breakfast shuffleLearn	0.315	0.309
Self-supervised first on K400, then on Breakfast			
(j)	K400, Breakfast, SpeedNet	0.501	0.337
(k)	K400, Breakfast, RotationNet	0.279	0.290
(l)	K400, Breakfast, shuffleLearn	0.292	0.318

Table 5: Comparison of features for temporal action segmentation on the Breakfast [27] dataset. ‘K400’ denotes features trained on Kinetics [7] dataset while ‘Breakfast’ represents features trained on Breakfast [27] dataset. ‘SpeedNet’, ‘RotationNet’ and ‘ShuffleLearn’ denote features trained using the self-supervised methods from [14], [20] and [41].

may have stronger capability in modeling the temporal relations among frames and thus are more suitable for temporal action segmentation. The results also indicate that not all the self-supervised learning methods can help improve the performance, it is important to capture temporal relations among frames in feature learning for the segmentation task. In addition, SpeedNet, as one of the most emerging video self-supervised learning approaches, consistently outperforms the ShuffleLearn, showing the effectiveness of it;

(3) Based on experiment (b), (d), (e), (f), self-supervised features always perform better than classical I3D feature pre-trained on Kinetics. We argue that while the I3D feature trained for action recognition task optimizes feature representation on video clip level, self-supervised learning representation implicitly learns better frame-level representation and can better capture the temporal structure inside the clip;

(4) Based on the experiments (d) - (l), we can see larger dataset like Kinetics can help build better self-supervised representation, while smaller ones contain less variety. It’s not needed to use target dataset to get a good feature representation for the temporal action segmentation task.

4.4. Qualitative Results

Fig. 2 visualizes two examples from the Breakfast dataset. From both Fig. 2a and Fig. 2b, we can see that SSCAP is able to notably improve the segmentation quality.



(a) P24 friedegg sequence from Breakfast dataset. The semantics order are ‘background’, ‘crack egg’, ‘fry egg’, ‘take plate’, ‘put egg2plate’, ‘background’.

(b) P38 juice sequence from Breakfast dataset. The semantics order are ‘background’, ‘cut orange’, ‘squeeze orange’, ‘take glass’, ‘pour juice’, ‘background’.

Figure 2: Two qualitative examples (a) and (b). The first row is the frame-wise ground-truth (x-axis stands for the time in the video while different sub-action classes are colored differently); the second row is the prediction from [30]; the third row is our prediction; and the fourth row is the uniform sampled 8 frames from the original video. Best viewed zoom in and in color.

We can also see that the CAP algorithm is able to effectively suppressing the over-segmentation issue by introducing the co-occurrence relations among sub-actions in decoding.

5. Conclusion

In this work, we have proposed SSCAP, an unsupervised temporal action segmentation solution that uses self-supervised methods in feature learning and a co-occurrence action parsing algorithm that helps model the correlation among sub-actions and better handle complex activity structures in videos. SSCAP has achieved SOTA performance on three public benchmarks in unsupervised setting, and has even outperformed several recently proposed weakly-supervised methods. The future direction can be leveraging the deep and hierarchical clustering approaches in further improving the accuracy of the method.

Acknowledgement

We’d like to give our special thanks to Fanyi Xiao, Bing Shuai and Uta Büchler for their helpful discussions.

References

- [1] Sathyanarayanan N. Aakur and Sudeep Sarkar. A perceptual prediction framework for self supervised event segmentation. In *CVPR*, 2019.
- [2] Jean-Baptiste Alayrac, Piotr Bojanowski, Nishant Agrawal, Josef Sivic, Ivan Laptev, and Simon Lacoste-Julien. Unsupervised learning from narrated instruction videos. In *CVPR*, 2016.
- [3] Sagie Benaim, Ariel Ephrat, Oran Lang, Inbar Mosseri, William T. Freeman, Michael Rubinstein, Michal Irani, and Tali Dekel. Speednet: Learning the speediness in videos. In *CVPR*, 2020.
- [4] Piotr Bojanowski, Rémi Lajugie, Francis Bach, Ivan Laptev, Jean Ponce, Cordelia Schmid, and Josef Sivic. Weakly supervised action labeling in videos under ordering constraints. In *ECCV*, 2014.
- [5] Biagio Brattoli, Uta Büchler, Anna-Sophia Wahl, Martin E. Schwab, and Björn Ommer. Lstm self-supervision for detailed behavior analysis. In *CVPR*, 2017.
- [6] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018.
- [7] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [8] Chien-Yi Chang, De-An Huang, Yanan Sui, Li Fei-Fei, and Juan Carlos Niebles. D3tw: Discriminative differentiable dynamic time warping for weakly supervised action alignment and segmentation. In *CVPR*, 2019.
- [9] Min-Hung Chen, Baopu Li, Yingze Bao, Ghassan AlRegib, and Zsolt Kira. Action segmentation with joint self-supervised temporal domain adaptation. In *CVPR*, 2020.
- [10] George E. Dahl, Dong Yu, Li Deng, and Alex Acero. Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition. In *IEEE Transactions on Audio, Speech, and Language Processing*, 2012.
- [11] Ali Diba, Vivek Sharma, Luc Van Gool, and Rainer Stiefelhagen. Dynamonet: Dynamic action and motion network. In *Arxiv*, 2019.
- [12] Li Ding and Chenliang Xu. Weakly-supervised action segmentation with iterative soft boundary assignment. In *CVPR*, 2018.
- [13] E. Elhamifar and Z. Naing. Unsupervised procedure learning via joint dynamic summarization. In *ICCV*, 2019.
- [14] Dave Epstein, Boyuan Chen, and Carl Vondrick. Oops! predicting unintentional action in video. In *CVPR*, 2020.
- [15] Yazan Abu Farha and Juergen Gall. Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In *CVPR*, 2019.
- [16] Mohsen Fayyaz and Juergen Gall. Sct: Set constrained temporal transformer for set supervised action segmentation. In *CVPR*, 2020.
- [17] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019.
- [18] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, 2016.
- [19] Basura Fernando, Hakan Bilen, Efstratios Gavves, and Stephen Gould. Self-supervised video representation learning with odd-one-out networks. In *CVPR*, 2017.
- [20] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018.
- [21] Tengda Han, Weidi Xie, and Andrew Zisserman. Video representation learning by dense predictive coding. In *ICCV*, 2019.
- [22] Rui Hou, Chen Chen, and Mubarak Shah. Tube convolutional neural network (t-cnn) for action detection in videos. In *ICCV*, 2017.
- [23] De-An Huang, Li Fei-Fei, and Juan Carlos Niebles. Connectionist temporal modeling for weakly supervised action labeling. In *ECCV*, 2016.
- [24] Yifei Huang, Yusuke Sugano, and Yoichi Sato. Improving action segmentation via graph based temporal reasoning. In *CVPR*, 2020.
- [25] Pan Ji, Runze Li, Bir Bhanu, and Yi Xu. Monoindoor: Towards good practice of self-supervised monocular depth estimation for indoor environments. In *ICCV*, 2021.
- [26] Dahun Kim, Donghyeon Cho, and In So Kweon. Self-supervised video representation learning with space-time cubic puzzles. In *AAAI*, 2019.
- [27] Hilde Kuehne, Ali Arslan, and Thomas Serre. Recovering the syntax and semantics of goal-directed human activities. In *CVPR*, 2014.
- [28] Hilde Kuehne, Juergen Gall, and Thomas Serre. An end-to-end generative framework for video segmentation and recognition. In *WACV*, 2016.
- [29] H. W. Kuhn. The hungarian method for the assignment problem. In *Wiley Online Library*, 1955.
- [30] Anna Kukleva, Hilde Kuehne, Fadime Sener, and Juergen Gall. Unsupervised learning of action classes with continuous temporal embedding. In *CVPR*, 2019.
- [31] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In *CVPR*, 2017.
- [32] Colin Lea, Austin Reiter, René Vidal, and Gregory D Hager. Segmental spatiotemporal cnns for fine-grained action segmentation. In *ECCV*, 2016.
- [33] Jun Li and Sinisa Todorovic. Weakly supervised energy-based learning for action segmentation. In *ICCV*, 2019.
- [34] Jun Li and Sinisa Todorovic. Set-constrained viterbi for set-supervised action segmentation. In *CVPR*, 2020.
- [35] Jun Li and Sinisa Todorovic. Action shuffle alternating learning for unsupervised action segmentation. In *CVPR*, 2021.
- [36] Runze Li, Srikrishna Karanam, Ren Li, Terrence Chen, Bir Bhanu, and Ziyang Wu. Learning local recurrent models for human mesh recovery, 2021.
- [37] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, 2019.
- [38] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In *ECCV*, 2018.

- [39] Wenqian Liu, Runze Li, Meng Zheng, Srikrishna Karanam, Ziyang Wu, Bir Bhanu, Richard J. Radke, and Octavia Camps. Towards visually explaining variational autoencoders. In *CVPR*, 2020.
- [40] Haoyu Ma, Liangjian Chen, Deying Kong, Zhe Wang, Xingwei Liu, Hao Tang, Xiangyi Yan, Yusheng Xie, Shih-Yao Lin, and Xiaohui Xie. Transfusion: Cross-view fusion with transformer for 3d human pose estimation. In *BMVC*, 2021.
- [41] Ishan Misra, C. Lawrence Zitnick, and Martial Hebert. Shuffle and learn: Unsupervised learning using temporal order verification. In *ECCV*, 2016.
- [42] Florent Perronnin, Jorge Sánchez, and Thomas Mensink. Improving the fisher kernel for large-scale image classification. In *ECCV*, 2010.
- [43] T. Quach and M Farooq. Maximum likelihood track formation with the viterbi algorithm. In *IEEE Conference on Decision and Control*, 1994.
- [44] Alexander Richard, Hilde Kuehne, and Juergen Gall. Weakly supervised action learning with rnn based fine-to-coarse modeling. In *CVPR*, 2017.
- [45] Alexander Richard, Hilde Kuehne, and Juergen Gall. Action sets: Weakly supervised action segmentation without ordering constraints. In *CVPR*, 2018.
- [46] Alexander Richard, Hilde Kuehne, Ahsan Iqbal, and Juergen Gall. Neuralnetwork-viterbi: A framework for weakly supervised video learning. In *CVPR*, 2018.
- [47] M. Saquib Sarfraz, Naila Murray, Vivek Sharma, Ali Diba, Luc Van Gool, and Rainer Stiefelhagen. Temporally-weighted hierarchical clustering for unsupervised action segmentation. In *CVPR*, 2021.
- [48] Fadime Sener and Angela Yao. Unsupervised learning and segmentation of complex activities from video. In *CVPR*, 2018.
- [49] Dian Shao, Yue Zhao, Bo Dai, and Dahua Lin. Finegym: A hierarchical video dataset for fine-grained action understanding. In *CVPR*, 2020.
- [50] Dian Shao, Yue Zhao, Bo Dai, and Dahua Lin. Intra- and inter-action understanding via temporal action parsing. In *CVPR*, 2020.
- [51] Sebastian Stein and Stephen J. McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *UBICOMP*, 2013.
- [52] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [53] Rosaura G. VidalMata, Walter J. Scheirer, Anna Kukleva, David Cox, and Hilde Kuehne. Joint visual-temporal embedding for unsupervised learning of actions in untrimmed sequences. In *WACV*, 2021.
- [54] Dong Wang, Di Hu, Xingjian Li, and Dejing Dou. Temporal relational modeling with self-supervision for action segmentation. In *AAAI*, 2021.
- [55] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *ICCV*, 2013.
- [56] 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.
- [57] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [58] Xiaolong Wang, Allan Jabri, and Alexei A Efros. Learning correspondence from the cycle-consistency of time. In *CVPR*, 2019.
- [59] Zhe Wang, Liyan Chen, Shauray Rathore, Daeyun Shin, and Charless Fowlkes. Geometric pose affordance: 3d human pose with scene constraints. In *Arxiv*, 2019.
- [60] Zhenzhi Wang, Ziteng Gao, Limin Wang, Zhifeng Li, and Gangshan Wu. Boundary-aware cascade networks for temporal action segmentation. In *ECCV*, 2020.
- [61] Zhe Wang, Daeyun Shin, and Charless Fowlkes. Predicting camera viewpoint improves cross-dataset generalization for 3d human pose estimation. In *ECCVW*, 2020.
- [62] Zhe Wang, Yali Wang, Limin Wang, and Yu Qiao. Codebook enhancement of vlad representation for visual recognition. In *ICASSP*, 2016.
- [63] Dejing Xu, Jun Xiao, Zhou Zhao, Jian Shao, Di Xie, and Yueting Zhuang. Self-supervised spatiotemporal learning via video clip order prediction. In *CVPR*, 2019.
- [64] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *ICCV*, 2017.
- [65] Xiangyi Yan, Hao Tang, Shanlin Sun, Haoyu Ma, Deying Kong, and Xiaohui Xie. After-unet: Axial fusion transformer unet for medical image segmentation. In *WACV*, 2022.
- [66] Ceyuan Yang, Yinghao Xu, Jianping Shi, Bo Dai, and Bolei Zhou. Temporal pyramid network for action recognition. In *CVPR*, 2020.
- [67] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *ICCV*, 2017.
- [68] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, 2018.