

Hierarchical Vector Quantization for Unsupervised Action Segmentation

Federico Spurio^{1,3}, Emad Bahrami^{1,3}, Gianpiero Francesca², Juergen Gall^{1,3}

¹University of Bonn, Germany

²Toyota Motor Europe, Belgium

³Lamarr Institute for Machine Learning and Artificial Intelligence, Germany
{fspurio, bahrami, gall}@iai.uni-bonn.de, gianpiero.francesca@toyota-europe.com

Abstract

In this work, we address unsupervised temporal action segmentation, which segments a set of long, untrimmed videos into semantically meaningful segments that are consistent across videos. While recent approaches combine representation learning and clustering in a single step for this task, they do not cope with large variations within temporal segments of the same class. To address this limitation, we propose a novel method, termed Hierarchical Vector Quantization (HVQ), that consists of two subsequent vector quantization modules. This results in a hierarchical clustering where the additional subclusters cover the variations within a cluster. We demonstrate that our approach captures the distribution of segment lengths much better than the state of the art. To this end, we introduce a new metric based on the Jensen-Shannon Distance (JSD) for unsupervised temporal action segmentation. We evaluate our approach on three public datasets, namely Breakfast, YouTube Instructional and IKEA ASM. Our approach outperforms the state of the art in terms of F1 score, recall and JSD.

Code — <https://github.com/FedeSpu/HVQ>

1 Introduction

Analyzing videos at a temporal level is crucial for applications in domains like healthcare, manufacturing, neuroscience, and robotics. However, fully-supervised methods (Abu Farha and Gall 2019; Huang, Sugano, and Sato 2020; Ishikawa et al. 2021; Yi, Wen, and Jiang 2021; Bahrami, Francesca, and Gall 2023) require expensive frame-wise labels, and weakly-supervised methods (Chang et al. 2019; Ding and Xu 2018; Kuehne, Richard, and Gall 2017; Li, Lei, and Todorovic 2019; Richard, Kuehne, and Gall 2017; Richard et al. 2018; Khan et al. 2022; Li, Abu Farha, and Gall 2021; Rahaman et al. 2022; Fayyaz and Gall 2020; Li and Todorovic 2020; Richard, Kuehne, and Gall 2018; Souri et al. 2022; Behrmann et al. 2022) still rely on prior knowledge of action classes. In order to discover semantically meaningful classes directly from long untrimmed videos without any labeling, unsupervised approaches for temporal action segmentation have received increasing attention




Figure 1: Given a set of long untrimmed video sequences (first row), the goal is to segment each video and cluster the segments across all videos. Our approach for hierarchical vector quantization learns two codebooks \mathbf{Z} (second row) and \mathbf{Q} (third row), which results in a hierarchical representation of actions.

(Kukleva et al. 2019; VidalMata et al. 2021; Li and Todorovic 2021; Swetha et al. 2021; Kumar et al. 2022; Tran et al. 2024; Xu and Gould 2024).

Early approaches for unsupervised action segmentation (Kukleva et al. 2019; VidalMata et al. 2021; Li and Todorovic 2021) learn first a representation and then cluster and segment the frames based on the learned representation in a multi-step fashion. More recent approaches (Swetha et al. 2021; Kumar et al. 2022; Tran et al. 2024; Xu and Gould 2024) combine representation learning and clustering in a single step, since representation learning affects clustering and vice-versa. TOT (Kumar et al. 2022), for example, applies temporal optimal transport for generating pseudo-labels, which are then used to supervise the representation learning. The recent approach ASOT (Xu and Gould 2024) extends TOT by relaxing the constraint of knowing the action order for a video and it improves the performance. However, it introduces implicitly a strong prior on the segment length. It fails to identify very short segments, as shown in Fig. 4, and the distribution over the segment lengths deviates largely from the ground-truth distribution, as shown in Fig. 3. The bias is caused by generating pseudo-labels from a noisy frame-to-cluster assignment matrix, which requires using strong priors like a structural prior.

Since methods for temporal unsupervised action segmentation should not have a strong bias towards segments of certain lengths, we propose a new metric based on the Jensen-Shannon Distance (JSD), which measures such bias and complements existing metrics like the F1 score. To address

this bias, we furthermore propose a novel approach that is based on a novel paradigm compared to previous approaches that either generate pseudo-labels or follow a multi-step approach. Our approach reformulates the problem as a vector quantization problem where a codebook is learned that represents the clusters. It learns an embedding of frames, a codebook and the assignment of frames to the codebook vectors, and thus clusters, jointly. In contrast to previous works, it neither requires the generation of pseudo-labels nor performs the embedding learning, clustering and assignment of frames in three separate steps.

To deal with large intrinsic variations of semantically meaningful segments, we explicitly take in consideration the natural composition of actions, *i.e.*, in order to complete an action it may be also necessary to complete intermediate steps. We thus propose a hierarchical vector quantization approach, termed Hierarchical Vector Quantization (HVQ), that provides a fine-to-coarse clustering as illustrated in Fig. 1. The quantization is performed at two levels. The first quantization deals with subactions, a fine-grained representation of an action, while the second quantization combines the subaction clusters to obtain the action representation. In this way, each coarse cluster can be represented by a varying number of finer clusters to cope with the intrinsic variation within each coarse cluster. HVQ is the first hierarchical approach for unsupervised temporal action segmentation and the first approach that utilizes vector quantization.

We evaluate the proposed approach on the Breakfast (Kuehne, Arslan, and Serre 2014), YouTube Instructional (Alayrac et al. 2016) and IKEA ASM (Ben-Shabat et al. 2021) datasets. On all the datasets, the proposed approach achieves state-of-the-art results in terms of F1 score, recall, and the Jensen-Shannon Distance (JSD), which measures the similarity of the distributions of the segment lengths of the predicted and ground-truth segments. In summary, our contributions are:

- We propose a novel end-to-end architecture for unsupervised temporal action segmentation. The hierarchical approach is able to capture the variability of action clusters and outperforms a non-hierarchical variant by a large margin.
- We introduce a novel metric based on the Jensen-Shannon Distance for evaluating the bias in the predicted segment lengths.
- Our method achieves state-of-the-art results for unsupervised temporal action segmentation in terms of recall and F1-score, and it is less biased.

2 Related Works

Fully-Supervised Action Segmentation. Early fully-supervised action segmentation methods rely on a sliding temporal window and non-maximum suppression (Karaman, Seidenari, and Bimbo 2014; Rohrbach et al. 2012) or model the problem with hidden Markov Models (HMM) (Kuehne, Gall, and Serre 2016; Tang, Fei-Fei, and Koller 2012). Richard and Gall (2016) model the probability of

action sequences with language and length models. Recent methods mainly use temporal convolutional networks (TCNs) with temporal pooling to capture long-range dependencies, though this can result in some loss of fine-grained temporal information. For this reason, Abu Farha and Gall (2019) introduced a multi-stage TCN, which maintains a high temporal resolution and has been further improved by Li et al. (2020). However, frame-wise predictions can lead to over-segmentation. To mitigate this, Ishikawa et al. (2021) introduced a boundary regression branch, and Huang, Sugano, and Sato (2020) proposed a graph-based temporal reasoning module. With the rise of Transformers (Vaswani et al. 2017) in vision tasks (Dosovitskiy et al. 2021; Zheng et al. 2021; Wang et al. 2021; Arnab et al. 2021; Bertasius, Wang, and Torresani 2021), ASFormer (Yi, Wen, and Jiang 2021) combined multi-stage TCNs with transformers, but also faced over-segmentation. Behrmann et al. (2022) addressed this by leveraging segment-level cues, while Bahrami, Francesca, and Gall (2023) used sparse attention to capture a long temporal context. Lastly, Liu et al. (2023) proposed a framework using denoising diffusion models with iterative refinement.

Weakly-Supervised Action Segmentation. Weakly-supervised approaches aim to use a reduced level of supervision. These approaches vary based on the type of weak labels used. Transcript-supervised methods (Chang et al. 2019; Ding and Xu 2018; Kuehne, Richard, and Gall 2017; Li, Lei, and Todorovic 2019; Richard, Kuehne, and Gall 2017; Richard et al. 2018; Lu and Elhamifar 2021) know the action order in training videos, while set-supervised methods (Fayyaz and Gall 2020; Li and Todorovic 2020; Richard, Kuehne, and Gall 2018) only know the set of actions without their order or timing. Timestamp supervision (Khan et al. 2022; Li, Abu Farha, and Gall 2021; Rahaman et al. 2022) requires the labeling of just one frame per action segment, offering improved results and similar annotation cost. Souri et al. (2022) further improved this by handling missing annotations. Recently, Behrmann et al. (2022) achieved competitive results by using a unified fully and timestamp-supervised method. Unlike these methods, our approach does not require any action labels.

Unsupervised Action Segmentation. Two different frameworks have been explored for unsupervised action segmentation. The first framework achieves action segmentation by detecting boundary changes at video level, *i.e.* processing and evaluating the videos individually (Du et al. 2022; Aakur and Sarkar 2019; Sarfraz et al. 2021; Mounir, Vijayaraghavan, and Sarkar 2023). The second framework, instead, processes and evaluates the videos at activity level, with an additional challenge of identifying the same actions across videos. Our work falls in this second framework. Early works in unsupervised action segmentation utilise the narration for the videos for segmentation. This is possible only if the narration is provided. For this reason, recent methods only rely on visual features and exploit the temporal structure of videos. One of the early works, Mallow (Sener and Yao 2018), alternates between

learning a discriminative appearance model and optimizing a generative temporal model of the activity. CTE (Kukleva et al. 2019) introduces a multi-step approach that starts from learning continuous temporal embeddings and performs a clustering on the learned features. JVT (VidalMata et al. 2021) improves CTE by adding a visual embedding. TAEC (Lin et al. 2023) uses an encoder-decoder structure with a TCN for learning temporal embeddings, while ASAL (Li and Todorovic 2021) learns also an action level embedding. These methods use a multi-step approach, representation learning and then offline clustering. UDE (Swetha et al. 2021) and TOT (Tran et al. 2024), on the other hand, allow feedback between these steps, performing a joint representation learning and online clustering. UFSA (Tran et al. 2024) follows these recent approaches, but instead of exploiting frame-level cues only, it uses also segment-level cues and pseudo-labels. The recent work ASOT (Xu and Gould 2024) relaxes the fixed action order constraint of TOT and UFSA, but it suffers from a strong bias on the lengths of estimated segments.

3 Hierarchical Vector Quantization for Unsupervised Action Segmentation

Unsupervised temporal action segmentation is a challenging task since it requires to identify temporal clusters that correspond to semantically meaningful actions not only within a video, but also across videos, as shown in Fig. 4. In particular, the cross-video setting imposes major challenges, since there is a large variation on how an action is executed by different persons and there is a large appearance variation if actions are recorded in different environments. Even though recent works based on generating pseudo-labels using optimal transport (Kumar et al. 2022; Tran et al. 2024; Xu and Gould 2024) have shown good performance, they introduce a segment length bias, which is particularly evident for the most recent approach ASOT (Xu and Gould 2024), as shown in Fig. 3. To quantify such bias, we will introduce a new metric based on the Jensen-Shannon Distance in Section 4.2. To address this bias, we propose a novel approach that is based on a novel paradigm compared to previous approaches that either generate pseudo-labels or follow a multi-step approach as discussed in Section 2. Our approach reformulates the problem as a vector quantization problem where a codebook is learned that represents the clusters. In this way, we learn an embedding of frames, a codebook and the assignment of frames to the codebook vectors, and thus clusters, jointly. We demonstrate that the approach has a lower segment length bias compared to previous works. In order to obtain better clusters, we furthermore introduce a hierarchical approach that learns not one but two codebooks where each codebook represents a different granularity as illustrated in Fig. 1.

Specifically, our model consists of an auto-encoder structure where the hierarchy is composed by two-levels of vector quantization as illustrated in Fig. 2. The encoder maps each frame x_{mt} of a video m with T_m frames to an embedding vector e_{mt} . The first level of the vector quantizer then associates every encoded frame e_{mt} with one embedding vector

$z_j \in \mathbf{Z}$ of the first learned codebook $\mathbf{Z} = \{z_1, \dots, z_{\alpha K}\}$, where α is a hyperparameter. The second stage vector quantizer associates every embedding vector z_j with another embedding vector $q_i \in \mathbf{Q}$ of the second learned codebook $\mathbf{Q} = \{q_1, \dots, q_K\}$. While the second codebook \mathbf{Q} represents the K clusters, the first codebook corresponds to more fine-grained clusters. Since the number of subactions that correspond to a semantically meaningful action may vary between actions, the number of embedding vectors z_j that are assigned to q_i can vary. After the vector quantization, each frame is represented by $q_{mt} \in \mathbf{Q}$. The decoder then aims to reconstruct the input video sequence $\mathbf{X}_m = \{x_{mt}\}_{t=1}^{T_m}$ from $\{q_{mt}\}_{t=1}^{T_m}$. We denote the reconstruction by $\hat{\mathbf{X}}_m = \{\hat{x}_{mt}\}_{t=1}^{T_m}$ and the codebooks \mathbf{Z} and \mathbf{Q} are learned by minimizing the reconstruction error. We now describe the method in detail.

3.1 Model Details

For the encoder and decoder, we use a light version of MS-TCN (Abu Farha and Gall 2019). The encoder and decoder have the same structure. They both consist of two stages, where each stage is composed by 10 layers. Each layer includes a temporal convolution with a dilation factor, which is doubled at each layer. Dilated temporal convolutions are very useful for the decoder since after quantization many neighboring frames are represented by the same vector q_i and the reconstruction requires a large receptive field. We evaluate the impact of the TCN in the GitHub repository.

The prototypes are updated dynamically as a new video is fed to the network, making our clustering online. Given the current codebook \mathbf{Z} , each frame is associated to the closest prototype $z_j \in \mathbf{Z}$:

$$z_{mt} = z_j, \quad \hat{j} = \operatorname{argmax}_j \frac{z_j \cdot e_{mt}}{\|z_j\| \|e_{mt}\|}. \quad (1)$$

The association of a subcluster to a cluster is done in the same way by

$$q_{mt} = q_i, \quad \hat{i} = \operatorname{argmax}_i \frac{q_i \cdot z_{mt}}{\|q_i\| \|z_{mt}\|}, \quad (2)$$

where the association needs only to be computed between $z_j \in \mathbf{Z}$ and $q_i \in \mathbf{Q}$ and not all frames. As in (van den Oord, Vinyals, and Kavukcuoglu 2017), we update the prototypes in the codebooks by the exponential moving average with $\beta = 0.8$:

$$\hat{z}_j = \frac{1}{\hat{N}_{z_j}} \left(\beta z_j + (1 - \beta) \sum_{z_{mt}=z_j} e_{mt} \right) \quad (3)$$

$$\text{where } \hat{N}_{z_j} = \beta N_{z_j} + (1 - \beta) |\{z_{mt} = z_j\}|,$$

$$\hat{q}_i = \frac{1}{\hat{N}_{q_i}} \left(\beta q_i + (1 - \beta) \sum_{q_{mt}=q_i} z_{mt} \right) \quad (4)$$

$$\text{where } \hat{N}_{q_i} = \beta N_{q_i} + (1 - \beta) |\{q_{mt} = q_i\}|.$$

N_{z_j} is the previous estimation of \hat{N}_{z_j} , which represent the number of assigned vectors to prototype z_j , and similarly N_{q_i} .




Figure 2: Overview of our model. The encoder takes as input a video X_m and encodes every frame into an embedding vector e_{mt} . Every e_{mt} is associated with the closest prototype $z_j \in \mathbf{Z}$. The prototypes z_j are then associated with the closest prototype $q_i \in \mathbf{Q}$. The association of a frame to \mathbf{Q} via \mathbf{Z} defines the cluster assignment.

In addition, we also use the strategy proposed by Dhariwal et al. (2020). When a prototype has not been assigned to any input for several batches, we replace the prototype with an input vector randomly sampled within the current batch. Specifically, we apply this when $\hat{N}_{z_j} < 3$ and $\hat{N}_{q_i} < 1$. Furthermore, we apply l_2 normalization on the encoded latent variables e_{mt} and codebook latent variables z_j to improve training stability and reconstruction quality (Yu et al. 2022).

3.2 Training

We train our model using the reconstruction loss. For a video m , it is given by

$$L_{rec} = \sum_{t=1}^{T_m} \|x_{mt} - \hat{x}_{mt}\|_2^2. \quad (5)$$

In addition, we use a commitment loss (van den Oord, Vinyals, and Kavukcuoglu 2017) which encourages the encoding vectors to stay close to the chosen prototype and to prevent frequent fluctuations in the assignment of code vectors (Razavi, van den Oord, and Vinyals 2019):

$$L_{commit_z} = \sum_{t=1}^{T_m} \|e_{mt} - \text{sg}[z_{mt}]\|_2^2, \quad (6)$$

$$L_{commit_q} = \sum_{t=1}^{T_m} \|z_{mt} - \text{sg}[q_{mt}]\|_2^2,$$

where $\text{sg}[\cdot]$ denotes stop-gradient. In this way, the embeddings are pushed towards the closest prototype and not the other way around.

The overall loss for training our model is then the combination of the commitment loss for the vector quantization and the reconstruction loss of the auto-encoder:

$$L = L_{commit_z} + L_{commit_q} + \lambda_{rec} \cdot L_{rec} \quad (7)$$

where λ_{rec} weights the two loss terms. We provide additional implementation details in the supp. material, which also includes additional ablation studies on the impact of the hyperparameters.

3.3 Inference

After having trained the model, we obtain for each video m and each frame x_{mt} an assignment to a prototype $q_{mt} \in \{q_1, \dots, q_K\}$, where q_i is a prototype for cluster $i \in \{1, \dots, K\}$. In order to obtain a smooth solution over time, we apply FIFA (Souri et al. 2021). To this end, we convert the hard assignment of a frame to a cluster into a soft assignment. This is achieved by computing for a frame x_{mt} the distances between the encoding e_{mt} and the prototypes in $z_j \in \mathbf{Z}$ and the distances between prototypes $z_j \in \mathbf{Z}$ and $q_i \in \mathbf{Q}$:

$$\text{sim}_{mt,i} = \sum_{j=1}^{|\mathbf{Z}|} (1 - d(e_{mt}, z_j)) + (1 - d(z_j, q_i)) \quad (8)$$

where $d(\cdot, \cdot)$ denotes the cosine distance. Using the softmax over $\text{sim}_{mt,i}$ then gives the soft assignment of frame x_{mt} to cluster i . Besides the assignment likelihood, FIFA takes a cluster order as input. As in (Kukleva et al. 2019), we compute the mean over the timestamps of all frames belonging

to each cluster and order the clusters with respect to this average timestamp. For the length prior, we count how often a cluster has been assigned to a frame in a video and use the average percentage over all videos.

4 Experiments

4.1 Datasets

We evaluate the proposed method on three public datasets: Breakfast (Kuehne, Arslan, and Serre 2014), YouTube Instructional (Alayrac et al. 2016) and IKEA ASM (Ben-Shabat et al. 2021).

- The Breakfast dataset is a large-scale dataset (1712 videos) that includes 70 hours of videos of 10 different common kitchen activities. The duration of videos varies widely from 30 seconds to few minutes. It has a total of 48 subaction classes.
- The YouTube Instructional (YTI) dataset, the smallest used in this work, contains 150 videos from YouTube with an average length of 2 minutes per video. It captures 5 activities with 47 subaction classes. This dataset contains a background class and the fraction of it varies from 46% to 83%.
- The IKEA ASM (IKEA) dataset contains 371 videos recording people assembling IKEA furniture for a total of 35 hours. There are 4 activities with 33 actions. The length of the videos varies from 1 to 8 minutes approximately.

As input, we use Improved Dense Trajectory (IDT) features (Wang and Schmid 2013) for the Breakfast dataset and the features provided by Alayrac et al. (2016) for the YouTube Instructional dataset. For the IKEA dataset, we have extracted the DINO features (Caron et al. 2021).

4.2 Evaluation Protocol

Following the protocol that has been introduced by Kukleva et al. (2019), we apply our approach to all videos of each activity separately. K is set to the max number of subactions that appear for each activity. This is required for a fair comparison of the methods. We establish the mapping between predicted cluster segments and ground-truth segments via Hungarian matching, which is computed over all videos and their frames of one activity.

For all the datasets, we report four metrics: the Mean over Frames (MoF), the F1-score (Kukleva et al. 2019), precision and recall. The MoF is the percentage of frames with correct predictions averaged over all activities. MoF is not a reliable measure for imbalanced datasets because it is dominated by frequent and long action classes, whereas the impact of short segments, which are much more difficult to detect, is very low on MoF. MoF thus prefers methods with a bias towards longer segments. The F1-score (Kukleva et al. 2019), which is the harmonic mean between precision and recall, is a much better metric. In addition to the F1 score, we also report recall, which measures how many segments are missed.

A key aspect of this work is to analyze if methods have a bias in terms of the length of the generated segments.

Method	MOF	F1	Recall*	JSD* \downarrow	Precision*
Mallow (Sener and Yao 2018)	34.6	-	-	-	-
CTE (Kukleva et al. 2019)	41.8	26.4	27.0	87.4	25.8
JVT (VidalMata et al. 2021)	48.1	-	-	-	-
ASAL (Li and Todorovic 2021)	52.5	37.9	-	-	-
UDE \dagger (Swetha et al. 2021)	47.4	31.9	-	-	-
TOT (Kumar et al. 2022)	47.5	31.0	26.3	90.2	37.7
TOT+TCL (Kumar et al. 2022)	39.0	30.3	36.0	<u>85.6</u>	26.2
UFSA (Tran et al. 2024)	52.1	38.0	-	-	-
ASOT (Xu and Gould 2024)	56.1	<u>38.3</u>	<u>40.1</u>	94.9	<u>36.7</u>
Ours (HVQ)	<u>54.4</u>	39.7	44.9	82.5	35.6

Table 1: Results on Breakfast. Best results are in **bold**, while second best ones are underlined. \dagger denotes results based on I3D features instead of IDT. * denotes metric computed by ourselves.

To quantify this, we propose a new metric based on the Jensen-Shannon Distance (JSD). For each video within the same activity, we compute the histogram of the predicted segment lengths, using a bin width of 20 frames. We then compare this histogram with the corresponding ground-truth histogram using the Jensen-Shannon Distance. These JSD scores are averaged across all videos for each activity. Finally, we calculate a weighted average across all activity, where the weights are the number of frames in each activity. In particular:

$$\text{JSD} = \frac{\sum_{a \in A} F_a \cdot \frac{1}{|M_a|} \sum_{m \in M_a} \text{JSDist}(H_m^{\text{pred}}, H_m^{\text{gt}})}{\sum_{a \in A} F_a}, \quad (9)$$

where $F_a = \sum_{m=1}^{M_a} T_m$ is the total number of frames and M_a is the set of all the videos for activity $a \in A$. JSDist is the Jensen-Shannon Distance:

$$\text{JSDist}(P, Q) = \sqrt{\frac{D_{\text{KL}}(P \parallel M) + D_{\text{KL}}(Q \parallel M)}{2}}, \quad (10)$$

where $M = \frac{P+Q}{2}$ and D_{KL} is the Kullback-Leibler divergence. The input of JSDist is normalized such that the sum of each histogram H_m is one. The metric measures how dissimilar the distributions of the lengths of the predicted and ground-truth segments are. A high value indicates a strong bias in segment lengths. Note that we do not report the metric for YTI since the video sequences are masked based on the ground-truth background for evaluation, but we report it for the other datasets that are evaluated without masking.

4.3 Comparison to State of the Art

Result on Breakfast. Table 1 presents the results of our method compared to state-of-the-art approaches for unsupervised temporal action segmentation. Our approach outperforms the state of the art in terms of F1 score, recall and JSD, and it achieves the second best MoF. Only the recent work ASOT (Xu and Gould 2024) achieves a higher MoF. Our approach, however, outperforms ASOT for all the other metrics except of precision. For instance, it achieves a higher recall, *i.e.*, it misses less action segments than ASOT. Furthermore, ASOT has also the highest JSD indicating the strongest bias of segment lengths compared to other approaches. The relatively high MoF of ASOT compared to F1




Figure 3: Histograms of the segment lengths across the whole milk activity for ground-truth, HVQ, TOT and ASOT. Segments longer than 580 frames are in the last bin.

score, recall and JSD can be explained by a bias of ASOT to longer segments, while missing short ground-truth segments that have a very low impact on MoF.

To analyze this further, Fig. 3 plots the histograms of the segment lengths across the whole milk activity, using a bin width of 20 frames. The plots show the strong length bias of ASOT. While there are many short segments (first two bins) in the ground-truth, ASOT generates only very few short segments. These segments have a very low impact on MoF, in contrast to the other metrics, and explain the high MoF of ASOT. TOT (Kumar et al. 2022) is less biased as shown by the histograms and JSD, but short segments are underrepresented as well. Our approach follows the ground-truth distribution of segment lengths best, and thus achieves the lowest JSD. While TOT+TCL achieves the second best JSD, the segments are often misclassified, which results in a low F1-score and recall in Table 1. Our approach achieves +9.4 and +8.9 higher F1-score and recall than TOT+TCL, respectively.

Fig. 4 shows qualitative segmentation results for a video from the Breakfast dataset. The activity is *friedegg* and it contains short and long actions. We compare to CTE (Kukleva et al. 2019), TOT (Kumar et al. 2022) and ASOT (Xu and Gould 2024). V1, V2, V3 and V4 represent four different videos of the same sequence (“P22_friedegg”) recorded with different cameras, thus the ground-truth segmentation is the same for all the four videos. For CTE and TOT, the predictions are very different across videos, therefore the predictions are not consistent even though the sequence is the same. In contrast, the predictions of ASOT do not change much across videos, but the network assigns the first four actions to a unique cluster, and the same for the last three actions. As discussed, ASOT is strongly biased towards longer segments, often overlooking shorter ones. Our method is not

Method	MOF	F1	Recall*	Precision*
Frank-Wolfe (Alayrac et al. 2016)	-	24.4	-	-
Mallow (Sener and Yao 2018)	27.8	27.0	-	-
CTE (Kukleva et al. 2019)	39.0	28.3	22.1	39.3
VTE (VidalMata et al. 2021)	-	29.9	-	-
ASAL (Li and Todorovic 2021)	44.9	32.1	-	-
UDE† (Swetha et al. 2021)	43.8	29.6	-	-
TOT (Kumar et al. 2022)	40.6	30.0	<u>31.4</u>	28.7
TOT+TCL (Kumar et al. 2022)	45.3	<u>32.9</u>	<u>27.9</u>	<u>40.1</u>
UFSA (Tran et al. 2024)	<u>49.6</u>	32.4	-	-
ASOT (Xu and Gould 2024)	52.9	35.1	27.8	47.6
Ours (HVQ)	50.3	35.1	38.7	32.1

Table 2: Results on YouTube Instructions. Best results are in **bold**, while second best ones are underlined. †denotes results based on I3D features instead of IDT. * denotes metric computed by ourselves.

only consistent across V1, V2, V3 and V4, but it is also capable of correctly clustering the long dark green action and the shorter actions. However, our approach fails as the other methods to identify the small red segment of *pour oil*.

Results on YouTube Instructional. In Tab. 2, we report the comparison between our method and the state of the art on the YouTube Instructional (YTI) dataset. We follow the protocol of prior works and evaluate the results excluding background frames. As for Breakfast, we outperform UDE (Swetha et al. 2021) and TOT (Kumar et al. 2022) by a large margin. While the temporal coherence loss (TCL) decreases MoF and F1-score of TOT on Breakfast, it improves the results on YTI. Nevertheless, our approach outperforms TOT+TCL on both datasets. Our approach also outperforms UFSA (Tran et al. 2024). As on Breakfast, ASOT achieves higher MOF and precision, but lower recall. The F1-scores

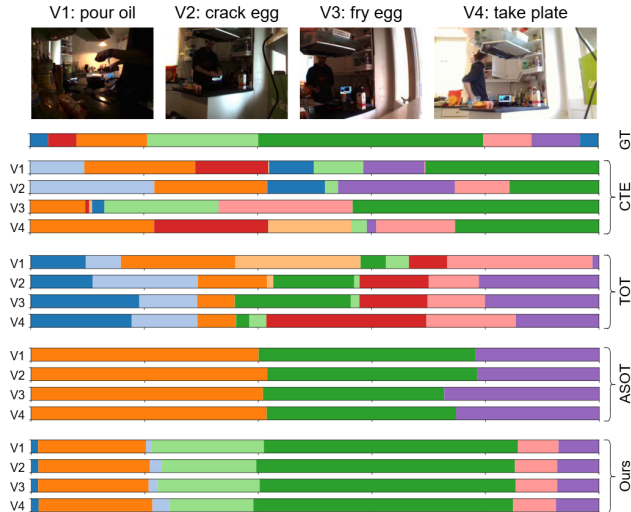


Figure 4: Segmentation results for a sample from the Breakfast dataset (*P22.friedegg*). Our approach delivers highly consistent results across multiple videos (V1, V2, V3, V4) recorded from different cameras, but with the same ground truth.

Method	MOF	F1	Recall*	JSD* \downarrow	Precision*
CTE* (Kukleva et al. 2019)	23.1	22.6	18.9	<u>73.7</u>	<u>28.1</u>
TOT* (Kumar et al. 2022)	21.0	20.1	17.1	80.0	24.4
TOT+TCL* (Kumar et al. 2022)	23.8	20.9	17.7	79.5	25.5
ASOT* (Xu and Gould 2024)	<u>34.0</u>	<u>27.9</u>	<u>24.0</u>	88.7	21.1
Ours (HVQ)	51.2	30.7	25.9	64.8	37.7

Table 3: Results on IKEA ASM. Best results are in **bold**, while second best ones are underlined. * denotes metric computed by ourselves.

of our method and ASOT are the same.

Results on IKEA ASM. In Tab. 3, we report our results compared to the state-of-the-art approaches on the IKEA ASM dataset. We compare our approach to CTE, TOT, considering also the version with temporal consistency loss TOT+TCL, and ASOT. Our approach outperforms all methods in all metrics and the predicted segment lengths are closest to the ground-truth as measured by JSD.

4.4 Ablation Studies

Impact of Loss Terms. We report the impact of the loss terms for the Breakfast dataset in Tab. 4. Using only the reconstruction loss results in a relatively low performance, which is expected since the reconstruction loss does not enforce compact clusters. This is achieved by the commitment loss. Using the commitment loss for both levels of the hierarchy $L_{commit_z} + L_{commit_q}$ already performs well. From the table, we can also see that L_{commit_z} has a larger impact than L_{commit_q} , which is expected since L_{commit_z} maps the frame-wise embeddings e_{mt} to the fine-grained clusters whereas L_{commit_q} associates the fine-grained clusters to the coarse clusters. Adding the reconstruction loss to the com-

	MOF	F1	JSD \downarrow
L_{rec}	41.6	29.7	87.2
L_{commit_q}	46.0	35.1	84.4
L_{commit_z}	51.8	37.4	83.6
$L_{commit_z} + L_{commit_q}$	54.4	38.0	82.5
$L_{rec} + L_{commit_q}$	47.8	36.4	84.1
$L_{rec} + L_{commit_z}$	51.3	38.4	83.7
$L_{rec} + L_{commit_z} + L_{commit_q}$	54.4	39.7	82.5

Table 4: Impact of loss terms on the Breakfast dataset. Best results are in **bold**.

λ_{rec}	0.0005	0.001	0.002	0.005	0.01
Breakfast	36.4	38.8	39.7	38.9	37.8
YTI	33.2	33.0	35.1	34.1	33.2
IKEA	30.0	29.6	30.7	29.1	30.3

Table 5: F1 on Breakfast and YTI for different values of λ_{rec}

mitment loss improves the F1 score further.

Impact of λ_{rec} . We also study the impact of the weight parameter λ_{rec} for the loss term (7) and we report the results in Tab. 5. Using $\lambda_{rec} = 0.002$ performs well and we use it for all other experiments. The parameter is not sensitive to the dataset.

Impact of α . In Tab. 6, we show the results on the Breakfast dataset considering different numbers of prototypes in \mathbf{Z} , which is steered by different values of α . While there is an improvement for $\alpha > 1$ compared to $\alpha = 1$, the performance saturates at $\alpha = 2$. For larger values of α , the performance decreases except for JSD since the subclusters become too fine-grained.

Impact of the Hierarchy. To show the benefits of having coarse and fine clustering, we report in Tab. 7 a comparison between single vector quantization and hierarchical vector quantization, with two and three levels. The results show that the hierarchical vector quantization with two levels outperforms the vector quantization with one level by a large margin in terms of F1 score. JSD is also better in all cases where the biggest improvement can be observed for IKEA ASM. Adding a third layer to the hierarchy does not improve the results, except for the F1 score on IKEA ASM. Compared to Breakfast and YTI, IKEA ASM has the highest average action length. To analyze the relation of hierarchy level and average action length more in detail, we group the

	$\alpha = 1$	$\alpha = 2$	$\alpha = 3$	$\alpha = 4$
MOF	53.6	54.4	52.7	51.8
F1	38.2	39.7	38.3	38.2
JSD \downarrow	83.7	82.5	83.0	82.2

Table 6: Impact of α on the Breakfast dataset. Best results are in **bold**.

Dataset	Metric	Single	Double	Triple
YTI	F1	33.0	35.1	31.9
IKEA ASM	F1	25.8	27.6	30.7
	JSD	81.9	62.2	64.8
Breakfast	F1	37.1	39.7	38.2
	JSD	83.1	82.5	84.1
BF short	F1	38.3	39.9	39.0
BF medium	F1	41.4	42.5	40.0
BF long	F1	35.7	38.4	38.0

Table 7: Impact of the number of hierarchy levels.

Breakfast	
Model	Time (minutes)
TOT	43
CTE	23
ASOT	18
Ours (Single)	18:48
Ours (Double, $\alpha = 2$)	19:04
Ours (Double, $\alpha = 5$)	19:15

Table 8: Training and inference time for processing the entire Breakfast dataset.

10 activities of Breakfast into three sets, based on their average action segment lengths. Activities with short actions (*coffee, tea, cereals*) have an average action length below 150 frames. Activities with medium-range actions (*sandwich, juice, milk*) have average action length between 150 and 250 frames. Finally, activities with long action segments (*fried-egg, salat, pancake, scrambled-egg*) have average action length above 250 frames. The results in Tab. 7 show that the difference between having two or three layers is smallest for the activities with long average length. While using more than two layers depends on the expected average segment lengths, using a hierarchical approach with 2 layers always outperforms a non-hierarchical approach.

In Tab. 8, we report the runtime of our method without hierarchy (single), with hierarchy (double), and with a large value of α , and compare it to TOT, CTE and ASOT. We measured the runtime on the same workstation with Intel i9-13900k CPU with 24 cores and one NVIDIA RTX 3090 GPU. The overhead added by the hierarchy or by a larger codebook size ($\alpha = 5$) is negligible. The runtime is the same as ASOT, but lower than CTE and TOT. We expect that the runtime of UFSA (Tran et al. 2024) is higher than the runtime of TOT, since UFSA extends TOT by additional steps. The reported runtime of ASAL (Li and Todorovic 2021) is roughly 10 hours on Breakfast.

Additional ablation studies are provided in the supp. material.

5 Conclusion

In this paper, we introduced Hierarchical Vector Quantization (HVQ), a novel paradigm for the problem of un-

supervised temporal action segmentation. While the first level of the hierarchy captures fine-grained subactions, the second one refines these predictions and produces the desired action-level granularity. Furthermore, we introduced a new metric that measures the bias in the generated segment lengths. We showed that our approach is less biased than previous works and that it is capable of dealing with the large variations of temporal segments. Additionally, the resulting vector quantization clusters are consistently assigned to the same actions across different videos. We provided an extensive set of ablations studies to show the efficiency of the proposed approach, and showed that our method achieves state-of-the-art results over previous methods on the Breakfast, YouTube Instructional and IKEA ASM datasets.

Acknowledgments

The work has been supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) GA 1927/4-2 (FOR 2535 Anticipating Human Behavior), the project iBehave (receiving funding from the programme “Netzwerke 2021”, an initiative of the Ministry of Culture and Science of the State of Northrhine Westphalia), and the ERC Consolidator Grant FORHUE (101044724). The sole responsibility for the content of this publication lies with the authors.

References

- Aakur, S. N.; and Sarkar, S. 2019. A perceptual prediction framework for self supervised event segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1197–1206.
- Abu Farha, Y.; and Gall, J. 2019. MS-TCN: Multi-Stage Temporal Convolutional Network for Action Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Alayrac, J.-B.; Bojanowski, P.; Agrawal, N.; Sivic, J.; Laptev, I.; and Lacoste-Julien, S. 2016. Unsupervised Learning From Narrated Instruction Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Arnab, A.; Dehghani, M.; Heigold, G.; Sun, C.; Lučić, M.; and Schmid, C. 2021. ViViT: A Video Vision Transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Bahrami, E.; Francesca, G.; and Gall, J. 2023. How Much Temporal Long-Term Context is Needed for Action Segmentation? In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Behrmann, N.; Golestaneh, S. A.; Kolter, Z.; Gall, J.; and Noroozi, M. 2022. Unified fully and timestamp supervised temporal action segmentation via sequence to sequence translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 52–68. Springer.
- Ben-Shabat, Y.; Yu, X.; Saleh, F.; Campbell, D.; Rodriguez-Opazo, C.; Li, H.; and Gould, S. 2021. The ikea asm dataset: Understanding people assembling furniture through actions,

- objects and pose. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 847–859.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is Space-Time Attention All You Need for Video Understanding? In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 9650–9660.
- Chang, C.-Y.; Huang, D.-A.; Sui, Y.; Fei-Fei, L.; and Niebles, J. C. 2019. D3tw: Discriminative differentiable dynamic time warping for weakly supervised action alignment and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3546–3555.
- Dhariwal, P.; Jun, H.; Payne, C.; Kim, J. W.; Radford, A.; and Sutskever, I. 2020. Jukebox: A Generative Model for Music. arXiv:2005.00341.
- Ding, L.; and Xu, C. 2018. Weakly-supervised action segmentation with iterative soft boundary assignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6508–6516.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations (ICLR)*.
- Du, Z.; Wang, X.; Zhou, G.; and Wang, Q. 2022. Fast and Unsupervised Action Boundary Detection for Action Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3323–3332.
- Fayyaz, M.; and Gall, J. 2020. Sct: Set constrained temporal transformer for set supervised action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 501–510.
- Huang, Y.; Sugano, Y.; and Sato, Y. 2020. Improving action segmentation via graph-based temporal reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ishikawa, Y.; Kasai, S.; Aoki, Y.; and Kataoka, H. 2021. Alleviating Over-Segmentation Errors by Detecting Action Boundaries. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.
- Karaman, S.; Seidenari, L.; and Bimbo, A. D. 2014. Fast saliency based pooling of fisher encoded dense trajectories. In *ECCV Workshops*.
- Khan, H.; Haresh, S.; Ahmed, A.; Siddiqui, S.; Konin, A.; Zia, M. Z.; and Tran, Q.-H. 2022. Timestamp-supervised action segmentation with graph convolutional networks. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 10619–10626. IEEE.
- Kuehne, H.; Arslan, A.; and Serre, T. 2014. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kuehne, H.; Gall, J.; and Serre, T. 2016. An end-to-end generative framework for video segmentation and recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.
- Kuehne, H.; Richard, A.; and Gall, J. 2017. Weakly supervised learning of actions from transcripts. *Computer Vision and Image Understanding*, 163: 78–89.
- Kukleva, A.; Kuehne, H.; Sener, F.; and Gall, J. 2019. Unsupervised learning of action classes with continuous temporal embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12066–12074.
- Kumar, S.; Haresh, S.; Ahmed, A.; Konin, A.; Zia, M. Z.; and Tran, Q.-H. 2022. Unsupervised Action Segmentation by Joint Representation Learning and Online Clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20174–20185.
- Li, J.; Lei, P.; and Todorovic, S. 2019. Weakly supervised energy-based learning for action segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 6243–6251.
- Li, J.; and Todorovic, S. 2020. Set-constrained viterbi for set-supervised action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10820–10829.
- Li, J.; and Todorovic, S. 2021. Action shuffle alternating learning for unsupervised action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12628–12636.
- Li, S.-J.; Abu Farha, Y.; Liu, Y.; Cheng, M.-M.; and Gall, J. 2020. MS-TCN++: Multi-Stage Temporal Convolutional Network for Action Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Li, Z.; Abu Farha, Y.; and Gall, J. 2021. Temporal Action Segmentation from Timestamp Supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8365–8374.
- Lin, W.; Kukleva, A.; Possegger, H.; Kuehne, H.; and Bischof, H. 2023. TAEC: Unsupervised Action Segmentation with Temporal-Aware Embedding and Clustering. In *Proceedings of the Computer Vision Winter Workshop (CVWW)*.
- Liu, D.; Li, Q.; Dinh, A.-D.; Jiang, T.; Shah, M.; and Xu, C. 2023. Diffusion action segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 10139–10149.
- Lu, Z.; and Elhamifar, E. 2021. Weakly-Supervised Action Segmentation and Alignment via Transcript-Aware Union-of-Subspaces Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 8085–8095.

- Mounir, R.; Vijayaraghavan, S.; and Sarkar, S. 2023. STREAMER: Streaming Representation Learning and Event Segmentation in a Hierarchical Manner. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 45694–45715.
- Rahaman, R.; Singhania, D.; Thiery, A.; and Yao, A. 2022. A Generalized and Robust Framework for Timestamp Supervision in Temporal Action Segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 279–296. Springer.
- Razavi, A.; van den Oord, A.; and Vinyals, O. 2019. Generating Diverse High-Fidelity Images with VQ-VAE-2. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Richard, A.; and Gall, J. 2016. Temporal action detection using a statistical language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Richard, A.; Kuehne, H.; and Gall, J. 2017. Weakly supervised action learning with rnn based fine-to-coarse modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 754–763.
- Richard, A.; Kuehne, H.; and Gall, J. 2018. Action sets: Weakly supervised action segmentation without ordering constraints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5987–5996.
- Richard, A.; Kuehne, H.; Iqbal, A.; and Gall, J. 2018. Neuralnetwork-viterbi: A framework for weakly supervised video learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7386–7395.
- Rohrbach, M.; Amin, S.; Andriluka, M.; and Schiele, B. 2012. A database for fine grained activity detection of cooking activities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sarfraz, S.; Murray, N.; Sharma, V.; Diba, A.; Van Gool, L.; and Stiefelhagen, R. 2021. Temporally-weighted hierarchical clustering for unsupervised action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11225–11234.
- Sener, F.; and Yao, A. 2018. Unsupervised learning and segmentation of complex activities from video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8368–8376.
- Souri, Y.; Abu Farha, Y.; Despinoy, F.; Francesca, G.; and Gall, J. 2021. FIFA: Fast Inference Approximation for Action Segmentation. In *Proceedings of the German Conference on Pattern Recognition (GCPR)*.
- Souri, Y.; Farha, Y. A.; Bahrami, E.; Francesca, G.; and Gall, J. 2022. Robust Action Segmentation from Timestamp Supervision. In *Proceedings of the British Machine Vision Conference (BMVC)*.
- Swetha, S.; Kuehne, H.; Rawat, Y. S.; and Shah, M. 2021. Unsupervised discriminative embedding for sub-action learning in complex activities. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, 2588–2592. IEEE.
- Tang, K.; Fei-Fei, L.; and Koller, D. 2012. Learning latent temporal structure for complex event detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tran, Q.-H.; Mehmood, A.; Ahmed, M.; Naufil, M.; Zafar, A.; Konin, A.; and Zia, Z. 2024. Permutation-Aware Activity Segmentation via Unsupervised Frame-To-Segment Alignment. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 6426–6436.
- van den Oord, A.; Vinyals, O.; and Kavukcuoglu, K. 2017. Neural Discrete Representation Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- VidalMata, R. G.; Scheirer, W. J.; Kukleva, A.; Cox, D.; and Kuehne, H. 2021. Joint visual-temporal embedding for unsupervised learning of actions in untrimmed sequences. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 1238–1247.
- Wang, H.; and Schmid, C. 2013. Action recognition with improved trajectories. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 3551–3558.
- Wang, H.; Zhu, Y.; Adam, H.; Yuille, A.; and Chen, L.-C. 2021. MaX-DeepLab: End-to-End Panoptic Segmentation With Mask Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xu, M.; and Gould, S. 2024. Temporally Consistent Unbalanced Optimal Transport for Unsupervised Action Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14618–14627.
- Yi, F.; Wen, H.; and Jiang, T. 2021. ASFormer: Transformer for Action Segmentation. In *Proceedings of the British Machine Vision Conference (BMVC)*.
- Yu, J.; Li, X.; Koh, J. Y.; Zhang, H.; Pang, R.; Qin, J.; Ku, A.; Xu, Y.; Baldrige, J.; and Wu, Y. 2022. Vector-quantized Image Modeling with Improved VQGAN. In *International Conference on Learning Representations (ICLR)*.
- Zheng, S.; Lu, J.; Zhao, H.; Zhu, X.; Luo, Z.; Wang, Y.; Fu, Y.; Feng, J.; Xiang, T.; Torr, P. H.; and Zhang, L. 2021. Rethinking Semantic Segmentation From a Sequence-to-Sequence Perspective With Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.