

CLOT: Closed Loop Optimal Transport for Unsupervised Action Segmentation

Elena Bueno-Benito, Mariella Dimiccoli

Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain

{ebueno, mdimiccoli}@iri.upc.edu

Abstract

Unsupervised action segmentation has recently pushed its limits with ASOT, an optimal transport (OT)-based method that simultaneously learns action representations and performs clustering using pseudo-labels. Unlike other OT-based approaches, ASOT makes no assumptions about action ordering and can decode a temporally consistent segmentation from a noisy cost matrix between video frames and action labels. However, the resulting segmentation lacks segment-level supervision, limiting the effectiveness of feedback between frames and action representations. To address this limitation, we propose Closed Loop Optimal Transport (CLOT), a novel OT-based framework with a multi-level cyclic feature learning mechanism. Leveraging its encoder-decoder architecture, CLOT learns pseudo-labels alongside frame and segment embeddings by solving two separate OT problems. It then refines both frame embeddings and pseudo-labels through cross-attention between the learned frame and segment embeddings, by integrating a third OT problem. Experimental results on four benchmark datasets demonstrate the benefits of cyclical learning for unsupervised action segmentation.¹

1. Introduction

Unsupervised action segmentation, the task of labeling each frame in a video or a set of videos with its action class without relying on a labeled training set, has been gaining increasing attention in recent years [12, 23, 28, 41, 45, 50]. This task is critical for numerous real-world applications, including action detection in sports events, video surveillance, and robotic manipulation [18, 53]. The classical approach to unsupervised action segmentation follows a multi-step pipeline, where frame embeddings are first learned and then clustered into action segments [22, 26, 46]. More recently, optimal transport (OT)-based methods have emerged, jointly learning action representations and clustering frames by leveraging OT theory to generate pseudo-labels for self-supervised training. [23, 45, 50].

¹<https://github.com/elenabbbuenob/CLOT>

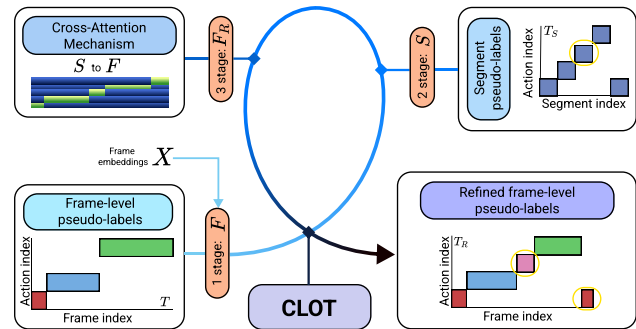


Figure 1. Given precomputed frame embeddings X as input, CLOT first solves an OT problem to generate pseudo-labels T alongside initial frame embeddings F . It then estimates segment embeddings S and pseudo-labels T_S by solving a second OT problem. Finally, a third OT problem is formulated to obtain refined frame embeddings F_R and pseudo-labels T_R , by leveraging cross-attention between the learned F and S , thereby closing the loop between frame and segment representations.

Among OT-based methods, ASOT [50] stands out as a particularly effective approach. Unlike other techniques, it does not impose assumptions about action ordering and achieves temporally consistent segmentation by solving a Gromov-Wasserstein OT problem between video frames and action labels. However, despite its strong performance, ASOT still struggles with accurately detecting action boundaries, particularly for short-duration segments, as it implicitly biases segment length through its prior assumptions. To address this, an alternative approach, HVQ [41], reformulates action segmentation as a hierarchical vector quantization problem, learning a codebook to represent action clusters. While this method improves the detection of short segments, it lacks effective feedback between representation learning and clustering, typical of OT-based formulations, resulting in poorer clustering assignments.

We argue that achieving high-quality, fine-grained segmentation requires *explicit feedback* between frame and segment representations to ensure that learned clusters align with true segment boundaries at the video level. To this end, we propose Closed Loop Optimal Transport (CLOT), a novel OT-based framework that introduces a multi-level cyclic learning mechanism to enhance segmentation quality.

Our approach, as illustrated in Fig. 1, initially generates pseudo-labels and jointly learns the frame embeddings. These embeddings are then used to produce segment embeddings. Both are optimized by solving two separate OT problems. To further improve segmentation consistency, CLOT refines the pseudo-labels through cross-attention between frame and segment representations, by integrating a third OT problem.

Our cyclic optimization framework employs a carefully designed encoder-decoder architecture (see Fig. 2), comprising an MLP-based encoder equipped with a feature dispatching mechanism, a parallel decoder, and a cross-attention-based mechanism. Additionally, we incorporate the Sliced Wasserstein (SW) distance to refine the initial cost matrix, leveraging its efficiency and robustness in high-dimensional representation spaces. Our contributions are as follows:

1. We propose a novel framework for unsupervised action segmentation that enforces segment-level consistency, achieving fine-grained segmentation and generalization across different videos.
2. We devised a novel encoder-decoder-based architecture that closes the loop between frame and segment embeddings by leveraging their cross-attention to refine frame embeddings.
3. We achieve very competitive performance at both video-level and activity-level over four benchmark datasets, namely Breakfast [21], YouTube INRIA [2], 50salads [42], and Desktop Assembly [23].
4. We provide extensive experiments, an ablation study, and qualitative results on four benchmark datasets, validating all components of the proposed architecture.

2. Related work

Fully supervised action segmentation methods remain the most reliable but require costly data annotations [4, 5, 15, 19, 27, 32, 51]. To improve scalability and practicality, research has increasingly shifted towards weakly-supervised [6, 9, 30, 31, 34, 40, 49, 52] and unsupervised approaches [7, 11, 12, 22, 23, 26, 28, 37, 38, 41, 45, 46, 50], aiming to reduce reliance on labeled data while maintaining competitive segmentation performance.

Unsupervised activity-level action segmentation. Unsupervised action segmentation traditionally follows a two-step pipeline: first, learning action representations in a self-supervised manner and then clustering the learned embeddings, typically assuming a predefined number of clusters. Classical methods strongly rely on temporal regularization to model the sequential nature of activities [22, 38]. This idea has been further refined through encoder-decoder architectures, incorporating either visual reconstruction losses (VTE) [46] or discriminative embedding losses (UDE) [43]

to enhance clustering performance. Other methods have framed the problem as a self-supervised learning task, where action prototypes are discovered via auxiliary classification objectives [12, 26]. For instance, CAD [12] introduced a framework that identifies action prototypes using an activity classification task, while ASAL [26] proposed a method that distinguishes between valid and invalid action orderings based on shuffled segment predictions.

Recently, Optimal Transport (OT) has emerged as a powerful tool for jointly learning action representations and pseudo-labels through self-training, enabling effective feedback between representation learning and clustering while directly optimizing for action segmentation. TOT [23] introduced a temporal OT formulation that generates pseudo-labels from predicted cluster assignments. However, this approach assumes fixed action ordering across all videos and uniform label assignment, contradicting the natural variability and long-tailed distribution of action labels. UFSA [45] addressed these limitations by incorporating frame- and segment-level cues from transcripts, allowing for action permutations within activities and non-uniform label assignments. However, it still requires prior knowledge of an estimated action order to infer the segmentation.

To overcome this constraint, ASOT [50] introduced an OT-based method capable of producing temporally consistent segmentations without any prior assumptions about action order. This makes it particularly suited for both pseudo-labeling and decoding. However, ASOT enforces a strong structural prior, limiting its ability to detect short-duration actions, which are critical in many real-world applications. More recently, HVQ [41] tackled this issue by using a hierarchical vector quantization-based approach, significantly improving short-action detection. However, its learned codebook that represents action classes has limited action generalization capabilities compared to ASOT since it lacks explicit feedback between representation and clustering.

Unsupervised video-level action segmentation. Video-level action segmentation focuses on processing individual videos independently, without relying on predefined activity categories. Existing approaches can be broadly categorized into representation learning methods and clustering-based methods. Representation learning approaches aim to learn robust action features before applying a clustering algorithm. LSTM+AL [1] predicts future frames and assigns segmentation boundaries based on prediction errors, while TSA [7] introduced a contrastive learning framework that employs a triplet selection strategy. Clustering-based approaches directly segment videos using similarity metrics. While clustering has been underexplored in action segmentation, recent work on TW-FINCH [37] incorporates temporal proximity alongside semantic similarity for improved clustering. Similarly, ABD [14] detects action

boundaries by measuring adjacent frame similarities, and OTAS [28] enhances boundary detection by incorporating object-centric features.

Optimal Transport for Structured Prediction. Optimal Transport (OT) has become a key framework in machine learning for measuring distributional discrepancies, particularly in unsupervised clustering and representation learning [23, 45, 50]. Unlike traditional probability metrics, Wasserstein distance preserves the geometric structure of distributions, making it well-suited for structured prediction tasks [3, 25, 36]. However, its computational complexity limits its scalability. To address this, projection-based OT (POT) methods, such as Sliced Wasserstein (SW) distance, efficiently approximate OT by projecting high-dimensional distributions onto lower-dimensional subspaces. SW distance has been successfully applied to point-cloud processing, color transfer, Gaussian Mixture Model learning, and domain adaptation, making it a practical alternative for large-scale applications [10, 20, 24, 35, 48].

ASOT’s structural priors hinder short action detection and limit segmentation granularity, while its lack of feedback between frame and segment embeddings leads to sub-optimal clustering. Motivated by these challenges, we propose CLOT, a novel OT-based architecture that enhances segmentation performance by better capturing the underlying video structure through a multi-level cyclic feature learning mechanism.

3. Closed Loop Optimal Transport

Problem formulation. Given a dataset $\mathcal{V} := \{V^b\}_{b=1}^B$ consisting of B videos and, for each video V^b , an initial frame-level embeddings $X^b \in \mathbb{R}^{N \times D}$, where N is the number of video frames and D the dimension of embedding, our goal is to learn a set of K action cluster centroids, represented as $A := [a_1, a_2, \dots, a_K] \in \mathbb{R}^{K \times d}$, with $a_j \in \mathbb{R}^d$ corresponding to the centroid embedding of the j -th action.

CLOT learns these centroids in an unsupervised fashion, through a three-level cyclic architecture (see Fig. 2). Each level solves an OT problem that learns a soft assignment: 1) between frames and action clusters in the first level, 2) between segments and action clusters in the second level, and again 3) between frames and action clusters in the third level. In the latter level, the OT problem is based on refined frame-embeddings derived from the cross-attention between frames and segment embeddings learned in the previous levels. The solution of each OT is strongly rooted in ASOT but differs from it in the way the initial cost matrix is computed. In this section, we first detail our CLOT architecture and then explain how we solve the OT problem and train the proposed network.

3.1. CLOT architecture

The proposed architecture, illustrated in Fig. 2, consists of an MLP-based encoder with a feature dispatching mechanism and a parallel decoder. The encoder outputs a frame embedding F by solving a temporally consistent OT problem that operates on the cost matrix (frame-to-cluster) C . The frame embedding F is then decoded into a segment embedding S by the parallel decoder, which also solves a temporally consistent OT problem based on the cost matrix (segment-to-cluster) C_S . A cross-attention mechanism between the encoder output F and the decoder output S is used to refine the frame embeddings into F_R , while solving a projection-based OT problem that operates on the cost matrix (frame-to-segment-to-cluster) C_R . In the following, we detail each component.

Encoder with feature dispatching. The encoder consists of a multi-layer perceptron network (MLP) that processes the input video frames $X \in \mathbb{R}^{N \times D}$ into frame-level embedding $F \in \mathbb{R}^{N \times d}$, where N is the number of video frames and d the feature dimension. The proposed method employs an adaptive process that dispatches each frame within a cluster based on similarity, resulting in a more structured and meaningful representation of the overall structure and context within the cluster, close to [29]. The feature dispatching mechanism is driven by a learnable similarity function ϕ parameterized by α and β , defined as:

$$\phi(A, F) = \sigma \left(\beta + \alpha \cdot \frac{A \cdot F}{\|A\| \|F\|} \right), \quad (1)$$

where σ denotes the sigmoid activation function and A represents the cluster embeddings. Each frame is assigned to a cluster using an attention-based masking approach, selecting the most relevant cluster and updating its feature representation. For each video frame $f_i \in F$, the updated feature embedding $f'_i \in F'$, with $F' \in \mathbb{R}^{N \times d}$ is computed as:

$$f'_i = f_i + \frac{1}{K} \sum_{k=0}^K (\phi(A_k, f_i) * A_k) \quad (2)$$

where K is the total number of clusters and $i \in [1, \dots, N]$. This update ensures that feature representations are dynamically adjusted based on the learned similarities, leading to more structured and meaningful frame embeddings.

To simplify notation, the updated feature embedding F' is redefined as F , indicating that F now represents the encoder with the integrated feature-dispatching mechanism.

Parallel decoder. The second stage of our framework employs a parallel decoder, adapted from [17, 33] for the action anticipation task. This decoder is responsible for transforming the frame-level embeddings F into segment-level representations S in a parallel, structured and efficient manner.

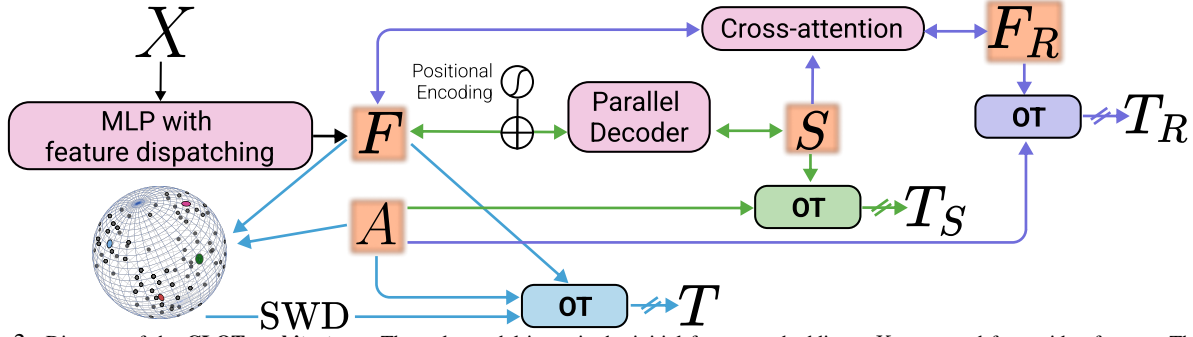


Figure 2. Diagram of the CLOT architecture. The only model input is the initial feature embeddings, X , extracted from video frames. The model consists of a multi-layer perceptron (MLP) encoder with a feature dispatching mechanism, a parallel decoder, and a cross-attention module. The magenta boxes indicate the architectural components of the model, whereas the orange boxes correspond to the learnable variables, which include the learnable action embeddings A as well as the learnable frame and segment embeddings. Arrows denote the computational/gradient flow, with crossed-out arrows ($\cancel{\rightarrow}$) indicating points where the gradient flow is stopped. Specifically, cyan arrows correspond to the first stage, green arrows to the second stage, and purple arrows to the third stage. The model leverages a closed-loop optimal transport mechanism to refine frame embeddings through cross-attention with segment embeddings, ensuring improved action segmentation accuracy. The estimated OT matrices (\mathbf{T} , \mathbf{T}_R , \mathbf{T}_S) act as pseudo-labels during training, and are computed at different stages by using the frame/segment features (F , S , and F_R) together with the action embeddings A , which are used to define the OT cost matrices.

Unlike autoregressive decoders like in [5, 16], which predict segments sequentially and suffer from error accumulation, our model adopts a parallel decoding strategy, where all action segments are predicted simultaneously.

The decoder is based on a query-based attention mechanism inspired by object detection architectures [8] and includes learnable queries $Q \in \mathbb{R}^{K' \times d_{dec}}$, where K' is the number of action queries and each query serves as a segment prototype. The query-based decoder also consists of multi-head cross-attention between the encoder and decoder, along with self-attention mechanisms, to obtain the final segment embedding $S \in \mathbb{R}^{K' \times d}$ with $K' \leq K$ denoting the number of detected segments.

Cross-attention for refinement. By incorporating the intrinsic information from the segment embeddings extracted for the decoder and the cluster centers, the method refines the feature embedding, enhancing the overall understanding of the video’s underlying structure. To integrate segment-level structure into frame embeddings, we apply a cross-attention mechanism

$$F_R = F + \text{softmax}\left(\frac{FS^\top}{\tau \cdot \sqrt{d}}\right)S, \quad (3)$$

where τ is a temperature parameter. This cross-attention-based alignment ensures that the feature frame representations are structured according to segment embeddings, helping to bring frame-level details to the temporal segmentation process.

3.2. Optimal Transport formulation

For the sake of completeness, we recall the OT unbalanced formulation proposed in [50], which combines the classical Kantorovich Optimal Transport (KOT) with a fused Gromov-Wasserstein Optimal Transport (GWOT).

KOT. The KOT problem [44] seeks for a transportation plan of the joint probability distribution \mathcal{T} and solves for the minimum cost of coupling $\mathbf{T}^* \in \mathcal{T}$ between the histograms μ and ν , where $\mu = \frac{1}{N}\mathbf{1}_N$ and $\nu = \frac{1}{K}\mathbf{1}_K$, being $\mathbf{1}_N$, and $\mathbf{1}_K$ uni-dimensional vectors of ones with N and K dimensions, respectively. The problem is formulated as:

$$\mathbf{T}^* = \min_{\mathbf{T} \in \mathcal{T}} \langle \mathbf{C}^k, \mathbf{T} \rangle, \quad (4)$$

$$\{\mathbf{T} \in \mathbb{R}_+^{N \times K} : \mathbf{T}\mathbf{1}_K = \mu, \mathbf{T}^\top \mathbf{1}_N = \nu\}, \quad (5)$$

where \mathcal{T} represents the transportation polytope and $\mathbf{C}^k \in \mathbb{R}_+^{N \times K}$ is the cost matrix. In the context of action segmentation, the variable $\mathbf{T} \in \mathbb{R}_+^{N \times K}$ can be interpreted as a soft assignment between two discrete sets, corresponding to frames and action indexes.

GW. The GM OT problem allows us to incorporate structural priors over the transport map, such as temporal consistency. The objective is defined as:

$$\mathcal{F}_{\text{GW}}(\mathbf{C}^v, \mathbf{C}^a, \mathbf{T}) := \sum_{i,j \in [n], l \in [m]} L(C_{ik}^v, C_{jl}^a) T_{ij} T_{kl}, \quad (6)$$

where $L : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ quantifies the discrepancies between the cost matrices, and $[n]$ and $[m]$ represent the set of video frames and action embedding, respectively.

The unbalanced OT formulation is as follows.

$$\min_{\mathbf{T} \in \mathcal{T}_p} \alpha \mathcal{F}_{\text{GW}}(\mathbf{C}^v, \mathbf{C}^a, \mathbf{T}) + (1 - \alpha) \mathcal{F}_{\text{KOT}}(\mathbf{C}^{\text{swd}}, \mathbf{T}) - \lambda KL(\mathbf{T}^\top \mathbf{1}_n \parallel \nu), \quad (7)$$

where $\alpha \in [0, 1]$ balance KOT and GW terms, and KL denotes the Kullback-Leibler divergence between \mathbf{T} and the

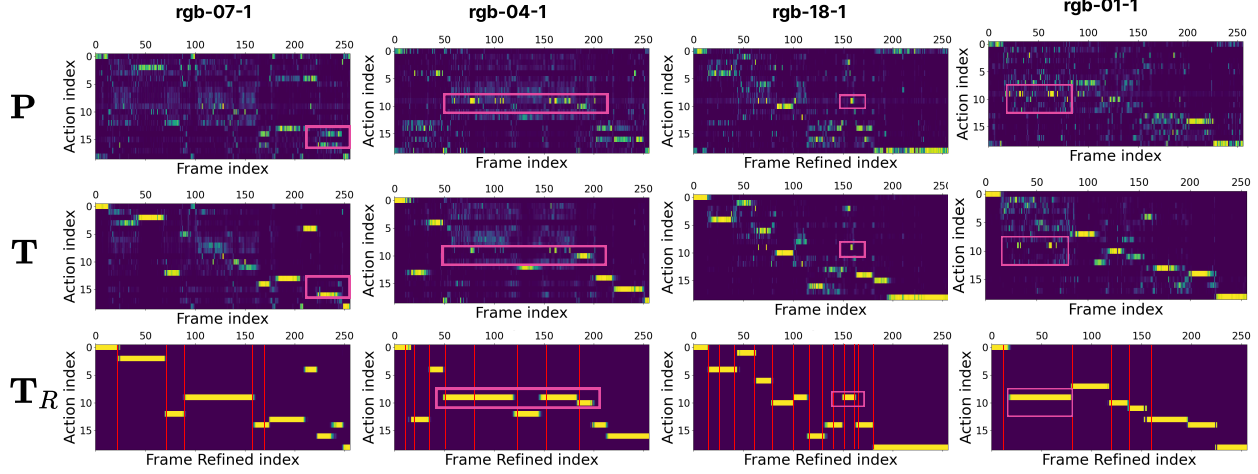


Figure 3. **Impact of cyclic multi-stage architecture.** Each column corresponds to different videos from the 50Salads EVAL dataset (rgb07-1, rgb04-1, rgb18-1, and rgb01-1). The pink boxes highlight instances where the model (\mathbf{T}_R) accurately detects the underlying segment or refines an existing one with greater precision, compared to using only the first stage (\mathbf{T}) or relying solely on the predicted frame-to-action assignment probabilities (\mathbf{P}). The red vertical lines denote segment boundaries. For this visualization, we use the subsampled features (256 per video) from the training set.

row-sum marginal distribution of \mathcal{T} with $\lambda > 0$. In particular, KL penalty on the column-sum marginals of the coupling matrix \mathcal{T} to ν , enabling a more flexible coupling solution.

Cost matrices. $\mathbf{C} := \{\mathbf{C}^k, \mathbf{C}^v, \mathbf{C}^a\}$ is the set of cost matrices for the KOT and GW subproblems. In our OT formulation, each of these matrices incorporates a set of submatrices: $\{\mathbf{C}^{sw}, \mathbf{C}^v, \mathbf{C}^a\}$ for frame embedding F and refined frame embedding F_R and $\{\mathbf{C}^k, \mathbf{C}^s, \mathbf{C}^a\}$ for segment embeddings S . They play a crucial role in defining the optimal transport problem at each stage since they encode structural priors over the transport map desirable for the video segmentation task (i.e., temporal consistency). Except for \mathbf{C}^{sw} , they are defined as in [50]. Specifically, the cost matrices $\mathbf{C}^v \in \mathbb{R}^{N \times N}$ and $\mathbf{C}^a \in \mathbb{R}^{K \times K}$ regulate the assignment of video frames and action categories, respectively by penalizing assignments of adjacent frames ($|i - k| \leq Nr, i \neq k$) to different action embedding ($j \neq l$) for two different assignments T_{ij} and T_{kl} . However, no penalty is applied to assignments outside temporal radius Nr or adjacent frames mapped to the same action ($j = l$). The last cost matrix \mathbf{C}^k is the visual component and is defined as $\mathbf{C}_{i,j}^k = \frac{x_i^T \cdot a_j}{\|x_i\| \|a_j\|} - \rho \cdot \mathbf{Z}_{ij}$ and \mathbf{Z} is the temporal prior used in [23, 50] defined as $\mathbf{Z}_{ij} = |1/N - j/K|$ for $\rho \geq 0$.

For \mathbf{C}^{sw} , we introduce the Sliced-Wasserstein (SW) distance as a complementary metric to the cosine distance, which has been successfully applied to many other structured prediction problems [10, 20, 24, 35, 48]. The p-Wasserstein distance between two probability measures μ

and ν in the probability space is defined as:

$$W_p(\mu, \nu) = \left(\inf_{\mathcal{T} \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} d^p(x, y) d\mathcal{T}(x, y) \right)^{\frac{1}{p}}, \quad (8)$$

where $\Pi(\mu, \nu)$ denotes the set of all couplings of μ and ν , i.e., joint distributions with marginals μ and ν . For $p \geq 1$, we estimate the discrete SWD distance empirically by discretizing the integral over the unit sphere \mathcal{S}^{d-1} in \mathbb{R}^d using a finite set of random projections $\{\theta_i\}_{i=1}^M$ between $x_i \in X$ and $a_j \in A$, leading to:

$$\text{SWD}_p(x_i, a_j) = \left(\frac{1}{M} \sum_{m=1}^M d(R_{\theta_m \#} x_i, R_{\theta_m \#} a_j) \right)^{\frac{1}{p}}, \quad (9)$$

for M randomly sampled θ , d quadratic loss and $\forall (i, j) \in N \times K$. In our method, we use $p = 1$, which means that it is equivalent to solving several one-dimensional optimal transport problems with closed-form solutions [36].

Specifically, the cost matrix $\mathbf{C}^{sw} \in \mathbb{R}_+^{N \times K}$ captures the visual similarity between frames and actions, defined as

$$\mathbf{C}_{i,j}^{sw} = 1 + \text{SWD}(x_i, a_j) - \mathbf{C}_{i,j}^k \quad (10)$$

3.3. CLOT learning pipeline

We learn the parameters of our architecture in an unsupervised manner. The training objective minimizes a cross-entropy loss between predicted action probabilities P and their corresponding soft pseudo-labels T derived by solving an OT problem. See Fig. 3 to visualize the impact of our cyclic multi-stage training. Specifically, we integrate three key components: frame embeddings (\mathbf{F}) generated by the

MLP, segment embeddings generated by the decoder (\mathbf{S}), and refined frame representations (\mathbf{F}_R) obtained by using the cross-attention between \mathbf{F} and \mathbf{S} . Each of these components is associated with soft pseudo-label $\mathbf{T} \in \mathbb{R}^{K \times N}$, $\mathbf{T}_S \in \mathbb{R}^{K \times K'}$ and $\mathbf{T}_R \in \mathbb{R}^{K \times N}$, respectively. For a given video b , we define the predicted frame-to-action assignment probabilities as

$$P_{ij}^b = \text{softmax}(\mathbf{F}^b \mathbf{A}^\top)_{ij} / \tau, \quad (11)$$

where $\tau > 0$ is a temperature scaling parameter. The frame-to-action learning loss is then formulated as

$$\mathcal{L}(\mathbf{T}, \mathbf{P}) = -\frac{1}{B} \sum_{b=1}^B \sum_{i=1}^N \sum_{j=1}^K T_{ij}^b \log P_{ij}^b, \quad (12)$$

Similarly, we compute the segment-to-action assignment probabilities \mathbf{P}_S and the frame-to-segment-to-action assignment probabilities \mathbf{P}_R . Their corresponding loss functions are defined analogously. The final training objective is a sum of these three losses

$$\mathcal{L}_{\text{train}} = \mathcal{L}(\mathbf{T}, \mathbf{P}) + \mathcal{L}(\mathbf{T}_S, \mathbf{P}_S) + \mathcal{L}(\mathbf{T}_R, \mathbf{P}_R) \quad (13)$$

4. Experimental Results

4.1. Experimental setting

Datasets and Features. We evaluate our approach using four widely used video datasets, each with pre-extracted feature representations for fair comparison with prior work [12, 22, 23, 41, 45, 50]:

- **Breakfast (BF)** [21] comprises approximately 1,700 videos depicting individuals preparing various breakfast items. Each video belongs to one of 10 activity categories, covering 48 distinct actions such as *cracking an egg* or *pouring flour*. The videos range in duration from 30 seconds to several minutes. We use Fisher vector features extracted from IDT [47].
- **YouTube Instructions (YTI)** [2] consists of 150 instructional videos spanning 5 activity categories. These videos, averaging around 2 minutes in length, primarily focus on tutorials, with a significant portion of frames ($\approx 75\%$) containing background content. We use a concatenation of Histogram of Optical Flow (HOF) descriptors and features extracted from the VGG16-conv5 layer [39].
- **50 Salads (FS)** [42] consists of 50 videos, totaling 4.5 hours, in which actors perform cooking activities. The dataset provides two levels of granularity: Mid (19 action classes) and Eval (12 action classes), where some Mid-level actions are aggregated into single actions in the Eval level. We extract Fisher vector features from IDT [47].
- **Desktop Assembly (DA)** [23] includes 76 videos, each approximately 1.5 minutes long, depicting actors assembling items in a predefined order. The task involves 22 sequentially performed actions. We use the features provided by [45].

Metrics. We follow established evaluation protocols for unsupervised action segmentation, consistent with prior studies [13]. To align learned action clusters with ground truth labels, we apply Hungarian Matching at the video level [7, 14, 37], where matching is performed within each video individually, and at the activity level [22, 23, 41, 45, 46, 50], where matching is performed across all videos of the same activity category.

To assess action segmentation quality, we employ the following metrics: (i) **Mean over Frames (MoF)**, which represents the percentage of correctly predicted frames and is sensitive to class imbalance, favoring dominant classes; (ii) **F1 Score**, the harmonic mean of precision and recall, as proposed in [22], which complements MoF by treating large segments as a whole and mitigating the effects of class imbalance; and (iii) **Mean Intersection over Union (mIoU)**, which computes the average intersection-over-union across all classes, explicitly addressing the impact of class imbalance.

Implementation details The encoder employs an MLP with a single hidden layer, while the decoder follows a parallel architecture. Dropout rates, projection dimensions, the number of predicted segments K' per video and other dataset-specific hyperparameters are adjusted for each dataset, with full details provided in the supplementary material. Optimization is performed using the Adam optimizer with a learning rate of 10^{-3} and weight decay of 10^{-4} . Action embeddings are initialized through k-means clustering. Training is conducted with 256 frames per video, sampled from uniformly distributed intervals. The number of action clusters K is defined based on the ground truth action count for each dataset, ensuring consistency with prior work [22, 23, 41, 45, 46, 50].

4.2. State-of-the-Art Comparisons

Traditionally, unsupervised action segmentation at the video level and the activity level have been treated as separate challenges. Video-level segmentation is generally considered the less complex of the two, as it does not require generalizing action representations across multiple videos. In contrast, activity-level segmentation demands the ability to maintain accurate action boundary detection while ensuring generalization across different instances of the same activity. This trade-off has historically limited performance improvements across both settings. CLOT establishes a new state-of-the-art in both video-level and activity-level segmentation. It consistently outperforms previous approaches on almost all benchmark datasets. Below, we analyze our results in detail, comparing them to existing methods and interpreting the architectural contributions that drive our superior performance. Fig 4 contains qualitative examples where actions are seen in different orders across videos within a dataset. More results are in the supp. mat.

Methods	Breakfast			YTI			50Salads (Mid)			50Salads (Eval)			DA		
	MoF	F1	mIoU	MoF	F1	mIoU	MoF	F1	mIoU	MoF	F1	mIoU	MoF	F1	mIoU
CTE [22]	41.8	26.4	-	39.0	28.3	-	30.2	-	-	35.5	-	-	47.6	44.9	-
VTE [46]	48.1	-	-	-	29.9	-	24.2	-	-	30.6	-	-	-	-	-
UDE [43]	47.4	31.9	-	43.8	29.6	-	-	-	-	42.2	34.4	-	-	-	-
ASAL [26]	52.5	37.9	-	44.9	32.1	-	34.4	-	-	39.2	-	-	-	-	-
TOT [23]	47.5	31.0	-	40.6	30.0	-	31.8	-	-	47.4	42.8	-	56.3	51.7	-
TOT+ [23]	39.0	30.3	-	45.3	32.9	-	34.3	-	-	44.5	48.2	-	58.1	53.4	-
UFSA [45]	52.1	38.0	-	49.6	32.4	-	36.7	30.4	-	55.8	50.3	-	65.4	63.0	-
ASOT [50]	<u>56.1</u>	38.3	18.6	<u>52.9</u>	<u>35.1</u>	24.7	<u>46.2</u>	<u>37.4</u>	<u>24.9</u>	<u>59.3</u>	<u>53.6</u>	<u>30.1</u>	70.4	<u>68.0</u>	<u>45.9</u>
HVQ [41]	54.4	<u>39.7</u>	-	50.3	<u>35.1</u>	-	-	-	-	-	-	-	-	-	-
CLOT (Ours)	60.1	40.1	<u>18.5</u>	54.4	36.7	<u>23.4</u>	50.6	46.6	31.4	59.4	63.2	38.8	<u>68.8</u>	72.6	48.1

Table 1. Comparisons of action segmentation performance obtained by applying the [Hungarian matching](#) at the [activity-level](#) on the Breakfast [21], Youtube Instr. [2], 50Salads [42] and Desktop Assembly [23] benchmarks. The highest accuracy is indicated in **bold**, and the second highest is underlined.

Methods	Breakfast			YTI			50Salads (Mid)			50Salads (Eval)			DA		
	MoF	F1	mIoU	MoF	F1	mIoU	MoF	F1	mIoU	MoF	F1	mIoU	MoF	F1	mIoU
TWF* [37]	62.7	49.8	42.3	56.7	48.2	-	66.8	<u>56.4</u>	48.7	<u>71.7</u>	-	-	73.3	67.7	57.7
ABD* [14]	64.0	52.3	-	67.2	49.2	-	<u>71.8</u>	-	-	71.2	-	-	-	-	-
OTAS* [28]	67.9	-	-	65.7	-	-	72.4	-	-	73.5	-	-	-	-	-
TSA* [7] (kmeans)	63.7	58.0	53.3	59.7	55.3	-	-	-	-	-	-	-	-	-	-
ASOT [50]	63.3	53.5	35.9	71.2	63.3	47.8	64.3	51.1	33.4	64.5	<u>58.9</u>	<u>33.0</u>	73.4	68.0	47.6
CLOT (Ours)	<u>66.3</u>	<u>55.9</u>	<u>37.1</u>	<u>69.3</u>	<u>60.8</u>	48.2	69.4	63.8	<u>45.0</u>	64.6	69.7	42.5	73.5	75.2	<u>52.4</u>

Table 2. Comparisons on the Breakfast [21], Youtube Instr. [2], 50Salads [42] and Desktop Assembly [23] datasets computed by applying [Hungarian matching](#) per [video](#). * denotes whether the method has a training stage on target videos. The highest accuracy is indicated in **bold**, and the second highest is underlined.

	Breakfast			YTI			50Salads (Mid)			50Salads (Eval)			DA		
	MoF	F1	mIoU	MoF	F1	mIoU	MoF	F1	mIoU	MoF	F1	mIoU	MoF	F1	mIoU
CLOT	60.1	40.1	18.5	54.4	36.7	23.4	50.6	46.6	31.4	59.4	63.2	38.8	68.8	72.6	48.1
w/o SWD	59.8	39.8	18.0	53.8	34.8	20.7	47.9	38.5	25.0	58.1	52.7	25.4	58.5	62.3	38.5
w/o FD	51.9	35.9	16.2	51.9	34.6	22.2	45.0	31.2	24.9	51.7	51.9	26.2	63.6	68.3	44.3
w/o Decoder	58.0	38.9	18.2	51.5	35.8	23.3	44.9	37.5	26.0	57.7	52.7	26.3	66.2	68.4	46.5
w/o Refinement	59.7	39.6	18.3	51.9	35.3	20.9	47.1	37.8	25.3	58.2	52.6	24.5	68.2	70.6	47.6

Table 3. Ablation Study on the four datasets: Breakfast [21], YTI [2], 50Salads [42] and Desktop Assembly [23]

Activity-level. CLOT surpasses all existing methods, achieving a significant performance gain. Specifically, our approach improves over the previous best-performing method by an average of +2.54% in MoF, +5.08% in F1 Score, and +3.2% in mIoU across five benchmark datasets.

Our F1 results consistently outperform the state of the art, with a particularly large margin in 50Salads and DA. This suggests that our action embeddings are more descriptive than the cluster embeddings from ASOT [50] or the vector quantization from HVQ [41]. However, our mIoU scores are marginally lower than those of ASOT on the BF and YTI datasets. The mIoU metric assigns equal weight to all classes, meaning that infrequent actions disproportionately affect the final score. If certain actions are never de-

tected due to their rare occurrence, mIoU is more affected than MoF or F1 Score. This explains why our F1 scores remain high even in datasets where mIoU is slightly lower. Nonetheless, in the remaining datasets, we demonstrate a substantial improvement in mIoU ($\sim 5.8\%$ on average).

CLOT lies in its parallel decoder architecture, which ensures that segmentation errors do not propagate across an entire sequence. Unlike autoregressive models, where errors accumulate over time, our parallel decoder prevents error propagation, ensuring that a single misclassified segment does not negatively impact the entire sequence. This allows the model to effectively avoid segmentation drift while maintaining high-quality action boundaries (comparisons in supp. mat).




Figure 4. **Qualitative results.** We display the ground-truth (GT), the results of CLOT (Ours), and ASOT [50]. Results from different datasets and activities are displayed for comparison.

Video-level CLOT and ASOT [50] learns a single model across all videos of an activity. In contrast, * methods are trained and optimized independently per video and are tailored to each test instance, often achieving higher performance by design, without generalization across videos. CLOT outperforms ASOT [50] across all datasets except for YTI, while often achieving F1 and mIoU scores comparable to methods specifically designed for video-level segmentation. This improvement suggests that CLOT dynamically refines action representations, allowing it to adapt more effectively to varying segment structures, which is particularly important at the video-level.

4.3. Ablation Study

To assess the impact of each component of our proposed framework, we conducted an ablation study on four benchmark datasets by systematically removing key elements from our architecture and measuring the resulting performance degradation. Further ablation experiments and sensitivity analysis are provided in the supp. mat. The results are reported in Tab 3. Our full model, CLOT, achieves the highest performance across all datasets.

Removing SWD leads to a decline in performance across all datasets, particularly in terms of mIoU. This suggests that SWD plays a critical role in enhancing the initial cost matrix. The performance drop is evident in the 50Salads and DA datasets. By projecting data onto one-dimensional

subspaces, SWD helps the visual information to reduce the error in videos where actions have a strong bias between actions.

Removing the feature dispatching (FD) mechanism results in a noticeable decrease in MoF and F1. This suggests to us FD prioritises informative frames and filters out irrelevant and noise frames (no visual information or background). Without the Decoder, Mof and mIoU drop significantly, particularly in 50S and YTI, suggesting that segment-level embeddings are essential for short-duration actions. The decoder improves structured segmentation and enables a more global representation of actions, complementing frame-level embeddings.

Eliminating the refinement stage results in consistent performance degradation, particularly in F1 and mIoU. This suggests the importance of incorporating segment-level structure to refine frame embeddings and action boundaries.

The results demonstrate that each of the introduced components adds value to the final segmentation results, even if the magnitude of the impact may vary significantly depending on the dataset. This study confirms that the integration of these components mitigates reliance on strong structural priors, filters noisy frames, and enables a more comprehensive action representation.

5. Conclusion

We presented CLOT, a novel Optimal Transport-based framework for unsupervised action segmentation that introduces a multi-level cyclic learning mechanism to refine frame and segment embeddings. Unlike previous methods, CLOT closes the loop between representation learning and clustering, ensuring a structured and adaptive segmentation process. The key architectural components of CLOT are the enhancement of the initial cost matrix through the Sliced Wasserstein Distance, which captures finer inter-frame relationships, a feature dispatching mechanism that filters out noisy or background frames, a parallel decoder that provides segment-level estimations, and a cross-attention mechanism that allows refining frame-level embedding. Altogether, these components ensure high-quality segmentation accuracy across dataset variability. Our experiments on four benchmark datasets confirm the effectiveness of CLOT, consistently surpassing existing methods in both video-level and activity-level segmentation.

Limitations. While CLOT achieves state-of-the-art performance, action mislabeling remains the main open challenge, highlighting the need for improvements in pseudo-label refinement and segment alignment. We believe CLOT represents a significant step forward in OT-based action segmentation, paving the way for advances in self-supervised learning and structured prediction tasks.

Acknowledgements

This work was supported by grant PRE2020-094714, the project PID2019-110977GA-I00 and the project PID2023-151351NB-I00, funded by Ministerio de Ciencia e Innovación (MCIN)/ Agencia Estatal de Investigación (AEI) /10.13039/501100011033, by European Social Fund (ESF) Investing in your future and by ERDF, UE. It was also supported by an Alexander von Humboldt (AvH) fellowship for experienced researchers funded by the AvH Foundation.

References

- [1] Sathyanarayanan N. Aakur and Sudeep Sarkar. A perceptual prediction framework for self supervised event segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [2] Jean-Baptiste Alayrac, Piotr Bojanowski, Nishant Agrawal, Josef Sivic, Ivan Laptev, and Simon Lacoste-Julien. Unsupervised learning from narrated instruction videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4575–4583, 2016. 2, 6, 7
- [3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning (ICML)*, pages 214–223. PMLR, 2017. 3
- [4] Emad Bahrami, Gianpiero Francesca, and Juergen Gall. How much temporal long-term context is needed for action segmentation? In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2
- [5] Nadine Behrmann, S Alireza Golestaneh, Zico Kolter, Juergen Gall, and Mehdi Noroozi. Unified fully and timestamp supervised temporal action segmentation via sequence to sequence translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2, 4
- [6] Elena Bueno-Benito and Mariella Dimiccoli. 2by2: Weakly-supervised learning for global action segmentation. In *International Conference on Pattern Recognition (ICPR)*, 2024. 2
- [7] Elena Bueno-Benito, Biel Tura, and Mariella Dimiccoli. Leveraging triplet loss for unsupervised action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2023. 2, 6, 7
- [8] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 213–229. Springer, 2020. 4
- [9] Xiaobin Chang, Frederick Tung, and Greg Mori. Learning discriminative prototypes with dynamic time warping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [10] Ishan Deshpande, Yuan-Ting Hu, Ruoyu Sun, Ayis Pyrros, Nasir Siddiqui, Sanmi Koyejo, Zhizhen Zhao, David Forsyth, and Alexander G. Schwing. Max-sliced wasserstein distance and its use for gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3, 5
- [11] Catarina Dias and Mariella Dimiccoli. Learning event representations by encoding the temporal context. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 2
- [12] Guodong Ding and Angela Yao. Temporal action segmentation with high-level complex activity labels. *IEEE Trans. Multimedia*, 2023. 1, 2, 6
- [13] Guodong Ding, Fadime Sener, and Angela Yao. Temporal action segmentation: An analysis of modern techniques. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 6
- [14] Zexing Du, Xue Wang, Guoqing Zhou, and Qing Wang. Fast and unsupervised action boundary detection for action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 6, 7
- [15] Yazan Abu Farha and Juergen Gall. Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [16] Rohit Girdhar and Kristen Grauman. Anticipative video transformer. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13505–13515, 2021. 4
- [17] Dayoung Gong, Joonseok Lee, Manjin Kim, Seong Jong Ha, and Minsu Cho. Future transformer for long-term action anticipation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3052–3061, 2022. 3
- [18] Yuchen He, Zeqing Yuan, Yihong Wu, Liqi Cheng, Dazhen Deng, and Yingcai Wu. Vistec: Video modeling for sports technique recognition and tactical analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024. 1
- [19] Feng Huang, Xiao-Diao Chen, Wen Wu, and Weiyin Ma. Not all samples are equal: Boosting action segmentation via selective incremental learning. *Engineering Applications of Artificial Intelligence*, 147:110334, 2025. 2
- [20] Soheil Kolouri, Kimia Nadjahi, Umut Simsekli, Roland Badeau, and Gustavo Rohde. Generalized sliced wasserstein distances. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019. 3, 5
- [21] Hilde Kuehne, Ali Arslan, and Thomas Serre. 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)*, 2014. 2, 6, 7
- [22] Anna Kukleva, Hilde Kuehne, Fadime Sener, and Juergen Gall. Unsupervised learning of action classes with continuous temporal embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 6, 7
- [23] Sateesh Kumar, Sanjay Haresh, Awais Ahmed, Andrey Konin, M. Zeeshan Zia, and Quoc-Huy Tran. Unsupervised action segmentation by joint representation learning and on-line clustering. In *Proceedings of the IEEE/CVF Conference*

- on *Computer Vision and Pattern Recognition (CVPR)*, pages 20174–20185, 2022. 1, 2, 3, 5, 6, 7
- [24] Chen-Yu Lee, Tanmay Batra, Mohammad Haris Baig, and Daniel Ulbricht. Sliced wasserstein discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10285–10295, 2019. 3, 5
- [25] Kwonjoon Lee, Weijian Xu, Fan Fan, and Zhuowen Tu. Wasserstein introspective neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3702–3711, 2018. 3
- [26] Jun Li and Sinisa Todorovic. Action shuffle alternating learning for unsupervised action segmentation. 2021. 1, 2, 7
- [27] Shi-Jie Li, Yazan AbuFarha, Yun Liu, Ming-Ming Cheng, and Juergen Gall. Ms-tcn++: Multi-stage temporal convolutional network for action segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 2
- [28] Yuerong Li, Zhengrong Xue, and Huazhe Xu. Otas: Unsupervised boundary detection for object-centric temporal action segmentation. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2024. 1, 2, 3, 7
- [29] James C Liang, Yiming Cui, Qifan Wang, Tong Geng, Wenguan Wang, and Dongfang Liu. Clusterformer: Clustering as a universal visual learner. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 3
- [30] Zijia Lu and Ehsan Elhamifar. Weakly-supervised action segmentation and alignment via transcript-aware union-of-subspaces learning. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [31] Zijia Lu and Ehsan Elhamifar. Set-supervised action learning in procedural task videos via pairwise order consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19871–19881, 2022. 2
- [32] Zijia Lu and Ehsan Elhamifar. Fact: Frame-action cross-attention temporal modeling for efficient action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [33] Megha Nawhal, Akash Abdu Jyothi, and Greg Mori. Rethinking learning approaches for long-term action anticipation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 558–576. Springer, 2022. 3
- [34] Yan Bin Ng and Basura Fernando. Weakly supervised action segmentation with effective use of attention and self-attention. *Computer Vision and Image Understanding*, 213: 103298, 2021. 2
- [35] Khai Nguyen and Nhat Ho. Energy-based sliced wasserstein distance. *Advances in Neural Information Processing Systems (NeurIPS)*, 36:18046–18075, 2023. 3, 5
- [36] Julien Rabin, Gabriel Peyré, Julie Delon, and Marc Bernt. Wasserstein barycenter and its application to texture mixing. In *Scale Space and Variational Methods in Computer Vision: Third International Conference, SSVI*, pages 435–446. Springer, 2012. 3, 5
- [37] M. Saquib Sarfraz, Naila Murray, Vivek Sharma, Ali Diba, Luc Van Gool, and Rainer Stiefelwagen. Temporally-weighted hierarchical clustering for unsupervised action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 6, 7
- [38] Fadime Sener and Angela Yao. Unsupervised learning and segmentation of complex activities from video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6
- [40] Yaser Souri, Mohsen Fayyaz, Luca Minciullo, Gianpiero Francesca, and Juergen Gall. Fast weakly supervised action segmentation using mutual consistency. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2
- [41] Federico Spurio, Emad Bahrami, Gianpiero Francesca, and Juergen Gall. Hierarchical vector quantization for unsupervised action segmentation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024. 1, 2, 6, 7
- [42] Sebastian Stein and Stephen J McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 729–738, 2013. 2, 6, 7
- [43] Sirmam Swetha, Hilde Kuehne, Yogesh S Rawat, and Mubarak Shah. Unsupervised discriminative embedding for sub-action learning in complex activities. 2021. 2, 7
- [44] Matthew Thorpe. Introduction to optimal transport. *Notes of Course at University of Cambridge*, 2018. 4
- [45] Quoc-Huy Tran, Ahmed Mehmood, Muhammad Ahmed, Muhammad Naufil, Andrey Konin, and M. Zeeshan Zia. Permutation-aware activity segmentation via unsupervised frame-to-segment alignment. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2024. 1, 2, 3, 6, 7
- [46] Rosaura G. VidalMata, Walter J. Scheirer, Anna Kukleva, David D. Cox, and Hilde Kuehne. Joint visual-temporal embedding for unsupervised learning of actions in untrimmed sequences. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1237–1246, 2021. 1, 2, 6, 7
- [47] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2013. 6
- [48] Jiqing Wu, Zhiwu Huang, Dinesh Acharya, Wen Li, Janine Thoma, Danda Pani Paudel, and Luc Van Gool. Sliced wasserstein generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3713–3722, 2019. 3, 5
- [49] Angchi Xu and Wei-Shi Zheng. Efficient and effective weakly-supervised action segmentation via action-transition-aware boundary alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [50] Ming Xu and Stephen Gould. Temporally consistent unbalanced optimal transport for unsupervised action segmentation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2, 3, 4, 5, 6, 7, 8

- [51] Fangqiu Yi, Hongyu Wen, and Tingting Jiang. Asformer: Transformer for action segmentation. In *British Machine Vision Conference (BMVC)*, 2021. [2](#)
- [52] Runzhong Zhang, Suchen Wang, Yueqi Duan, Yansong Tang, Yue Zhang, and Yap-Peng Tan. Hoi-aware adaptive network for weakly-supervised action segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023. [2](#)
- [53] Ido Zuckerman, Nicole Werner, Jonathan Kouchly, Emma Huston, Shannon DiMarco, Paul DiMusto, and Shlomi Laufer. Depth over rgb: automatic evaluation of open surgery skills using depth camera. *International Journal of Computer Assisted Radiology and Surgery*, pages 1–9, 2024. [1](#)