

Unsupervised Discovery of Long-Term Spatiotemporal Periodic Workflows in Human Activities

Fan Yang¹, Quanting Xie³, Atsunori Moteki², Shoichi Masui²,
Shan Jiang², Kanji Uchino¹, Yonatan Bisk³, Graham Neubig³

¹ Fujitsu Research of America, USA ² Fujitsu Limited, Japan ³ Carnegie Mellon University, USA




Figure 1. **Periodic spatiotemporal activity in a yoga example.** While existing studies focus on short-term periods with simple structures and high-contrast patterns, we investigate a novel direction on long-term periods that involve complex workflows with low-contrast patterns.

Abstract

Periodic human activities with implicit workflows are common in manufacturing, sports, and daily life. While short-term periodic activities—characterized by simple structures and high-contrast patterns—have been widely studied, long-term periodic workflows with low-contrast patterns remain largely underexplored. To bridge this gap, we introduce the first benchmark comprising 580 multi-modal human activity sequences featuring long-term periodic workflows. The benchmark supports three evaluation tasks aligned with real-world applications: unsupervised periodic workflow detection, task completion tracking, and procedural anomaly detection. We also propose a lightweight, training-free baseline for modeling diverse periodic workflow patterns. Experiments show that: (i) our benchmark presents significant challenges to both unsupervised periodic detection methods and zero-shot approaches based on powerful large language models (LLMs); (ii) our baseline outperforms competing methods by a substantial margin in all evaluation tasks; and (iii) in real-world applications, our baseline demonstrates deployment advantages on par with traditional supervised workflow detection approaches, eliminating the need for annotation and retrain-

ing. Our project page is <https://sites.google.com/view/periodicworkflow>.

1. Introduction

The rhythmic repetition of activities, characterized by periodic spatiotemporal workflows, is ubiquitous in many aspects of human life and work, ranging from the structured routines of a factory production line to the cyclical movements of athletes in training [10, 11, 14, 31, 51, 55, 76]. Detecting and understanding these periodic workflows offers immense potential for optimizing working efficiency, enhancing training performance, improving motion synthesis, and facilitating knowledge transfer to new learners or robots [1, 45, 54, 62, 90].

Current methods for periodic activity detection generally focus on short-term periods with clear start and end points, often relying on supervised learning to detect high-contrast patterns between periods [16, 21, 23, 34, 47, 48, 53, 56, 58, 61, 75, 93, 94, 102]. However, many real-world activity periods are long-term, consisting of ordered activity tokens that form workflows, yet remain largely unexplored. For instance, in periodic workflows, identical action tokens may appear multiple times at different positions within a period, which complicates period boundary identification—



Figure 2. **Left:** Illustration of our benchmark. **TopRight:** Statistics of our dataset. **BottomRight:** Annotation of our benchmark. We collected 580 long-term workflows of periodic human activities, **characterized by a compact dataset yet encompassing a wide variety of real-world periodic tasks**, including factory production, exercise training, and shuttle routes.

an issue often overlooked in prior studies (Fig. 1, Sec. 2). Additionally, existing studies on human activity segmentation typically divide activities into multiple tokens based on a workflow. However, these studies often ignore the period detection, assuming a predefined period for segmentation [2, 9, 17, 19, 44, 52, 66, 67, 77, 78, 80, 87, 91, 95].

To bridge this gap, we introduce a pioneering benchmark for unsupervised detection of long-term periodic spatiotemporal workflows in human activities. Specifically, we collected 580 human activity sequences, **with a small size but a rich diversity to cover most real-world periodic activities**, such as factory production, exercise training, and shuttle routes. Our dataset covers rich modalities, including 3D body/hand pose sequences, 3D indoor trajectories, and 2D outdoor trajectories (Fig. 2). With our provided spatiotemporal features, we have designed three evaluation tasks. In the first task, we introduce an activity sequence consisting of more than three periods, requiring unsupervised identification of period boundaries and counting. The second task focuses on completion tracking, which involves using the extracted workflow to estimate the remaining portion of a new period when only part of that period is available. The final task utilizes the extracted workflow to localize procedural anomalies within a new period. These tasks correspond to real-world applications, such as assisting production line workers in optimizing schedules and reducing errors.

We provide a novel baseline method for our benchmark, starting with a spatiotemporal tokenization to tokenize all frames of an activity into a symbolic transcript using both

hard and soft tokens. The hard token for each frame is determined by the symbol of its nearest token feature centroid, while the soft token consists of a vector of normalized distances to all token feature centroids. We then form a 2D matrix from the soft tokens of all frames and analyze it in the frequency domain. We identify a rough periodic scale along the temporal axis and apply it to segment the entire transcript. Finally, we propose an innovative sequential mining algorithm to distill the common pattern within all segmented hard-tokenized transcripts. This approach yields unified periodic activity workflows with precisely defined boundaries. **Our method is fully unsupervised, eliminating the need for textual activity labels, which are often difficult to define for complex human activities** [46]. For example, whereas traditional action-to-text tokenization requires manual assignment of esoteric labels to actions (e.g., ‘pirouette’ in ballet), our scheme uses arbitrary symbols (e.g., ‘A’~‘Z’). This abstraction supports all downstream applications in our benchmark, without the need to explicitly know the activity names.

We summarize our contributions in three main areas:

- **Benchmark Innovation:** We introduce the first benchmark featuring long-term periodic spatiotemporal workflows of human activity, gathered from diverse real-world settings. This benchmark effectively mirrors the complexities encountered in practical applications and opens exciting research avenues in the field of human activity studies.
- **Algorithm Development:** We developed an innovative baseline to tackle challenges posed by our benchmark.

Our evaluation demonstrates its promising performance, establishing a robust methodology for future research endeavors and real-world applications.

- **Real-World Application:** We introduce three evaluation tasks to address real-world applications, including unsupervised period detection, task completion tracking, and procedural anomaly detection. We illustrate the practical value of our work through a successful deployment in a real-world factory production line.

2. Related Works

2.1. Short-term Periodic Activity Detection (S-PAD)

Existing research on periodic human activity detection primarily focuses on short-term repetitive actions [16, 21, 23, 34, 47, 48, 53, 56, 58, 61, 75, 93, 94, 102]. Methodologically, a common approach across these studies involves employing neural networks to encode spatiotemporal features and constructing a temporal self-similarity matrix (TSM) for periodicity detection. Furthermore, some investigations [34, 47, 48, 94] utilize transformers to capture temporal dependencies through supervised learning, while others [16, 61, 93] concentrate on learning periodic patterns without manual annotations or prior action knowledge. Beyond unimodal approaches, another line of research [102] integrates visual and auditory cues for multimodal period detection. In terms of datasets, studies such as [34, 75, 93, 94, 102] have introduced new datasets to advance research in this field. While these methods and datasets have proven effective for analyzing short repetitive actions, they have not addressed the complexities of long-term periodic activities characterized by embedded workflows. Our work bridges this critical gap by introducing both a novel benchmark and a robust baseline method specifically tailored for such challenging scenarios (Tab. 1).

2.2. Single-period Activity Segmentation (SAS)

Research in human activity segmentation aims to partition a single-period sequence of activities into distinct, bounded activity tokens that conform to specific workflows. We compare our work with prior SAS studies [4, 26, 28, 40, 49, 72, 79, 81] in Tabs. 1 and 2. Traditional SAS methods generally rely on labeled data to align activity features and tokens with semantic tags. While fully-supervised SAS [2, 44, 95] requires exhaustive annotations of boundaries and labels for all activity tokens, semi-supervised SAS [9, 17, 52, 66, 67, 77, 80, 87, 92] reduces annotation costs, as it requires only a limited number of annotated timestamps during training. Nevertheless, both approaches face a critical limitation: the need to annotate new activity semantic tags and retrain the model when applied to novel scenarios. This requirement poses a significant barrier to scalability, hindering the generalization of these methods

Benchmarks	w/ Periods	w/ Workflow	Learning
S-PAD (Sec. 2)	✓	✗	Supervised
SAS (Sec. 2)	✗	✓	Supervised / Unsupervised
Ours (Sec. 3)	✓	✓	Unsupervised

Table 1. Comparison of S-PAD, SAS, and our benchmarks.

Benchmark	No. multi-branch workflows	Unsupervised Tasks		
		Period	Completion	Anomaly
GTEA [26]	7	✗	-	-
Breakfast [40]	10	✗	-	-
50Salads [79]	<50	✗	-	-
Assembly101 [72]	362	✗	-	✓
Ours	580	✓	✓	✓

Table 2. Comparison of workflow analysis with SAS studies.

across diverse applications.

Unsupervised SAS [19, 42, 43, 50, 70, 71, 78, 89] overcomes this limitation by utilizing clustering-based features to represent action tokens, thereby circumventing the necessity for explicit mappings between activity features and textual semantic tags. Even without semantic tags, this approach can efficiently extract segmentation patterns (*i.e.*, workflows) from a single-period activity sequence. However, unsupervised SAS often assumes a single period is available for segmentation. This limits its applicability in scenarios where periodicity should be inferred. Our approach bridges this gap by jointly integrating period detection and unsupervised workflow detection.

2.3. Periodic Spatiotemporal Data Mining

Prior research in periodic spatiotemporal data mining, exemplified by studies such as [13, 38, 97–99], has predominantly focused on analyzing 2D spatiotemporal GPS trajectories, primarily within traffic contexts. These investigations typically examine sequences of locations visited at regular intervals over time. However, our work diverges significantly from this established paradigm by encompassing a wider spectrum of human activities, spanning diverse domains such as factory production line operations and exercise training regimens. Our dataset also offers greater modality diversity, incorporating 3D human body and hand pose sequences, along with 3D indoor and 2D outdoor trajectory data. Furthermore, conventional approaches frequently utilize trajectories restricted to fixed temporal boundaries, such as 24-hour cycles. In contrast, our dataset captures periodic human activities with flexible temporal boundaries. Moreover, many existing methodologies inadvertently mirror the principles of process mining [27, 83], where identical activity tokens positioned at different locations within a workflow can lead to self-loops in the graph, potentially losing long-term temporal dependencies in procedural analysis. Importantly, this field currently lacks a publicly accessible unified dataset and evaluation metrics.

To address this gap, we will release a comprehensive benchmark to promote further research in the field.

3. Our Benchmark

Dataset. As Fig. 2 shows, we collected 580 long-term workflows of periodic human activities by identifying real-world scenarios where such patterns naturally occur. For instance, while prior SAS datasets focus on single instances (*e.g.*, making one cup of coffee), our dataset captures periodic workflows (*e.g.*, repeatedly making multiple cups of coffee). Each activity sequence in our dataset is annotated with the number of periods, period boundaries, a unique workflow, and corresponding activity tokens (Note: the use of arbitrary tokens does not affect evaluation under our benchmark setting). To ensure sufficient temporal coverage, each sequence contains at least five and at most eight periods. For task 2 (completion tracking), we randomly truncate a portion of the second-last period; for task 3 (anomaly detection), we inject an anomaly into the final period. This design strikes a balance between complexity and tractability, allowing the study of realistic periodic workflows while keeping evaluation computationally manageable. To mitigate bias from various activity feature extraction models while protecting privacy, we provide processed spatiotemporal features—primarily motion trajectories—for periodic workflow analysis. These features are synthesized to reflect natural patterns of human activities and corresponding workflow structures, and they include accurate ground-truth labels. For practical applications, it is recommended to use off-the-shelf human body and hand pose estimators [63, 73, 74, 85, 101] to obtain 3D body and hand sequences in small-scale environments. Indoor human 3D trajectories can be captured using Radio Frequency Identification (RFID) [60, 96] and visual-inertial odometry (VIO) [7], while outdoor 2D trajectories are recorded via Global Navigation Satellite System (GNSS).

Evaluation Tasks. We define three evaluation tasks using the provided spatiotemporal features: **1.** detect period counts and boundaries on normal periods through unsupervised methods; **2.** based on workflows obtained in task 1, predict the remaining phase proportion of an ongoing period when only partial data are available; **3.** based on workflows obtained in task 1, localize procedural anomalies within new periods. Rather than directly assessing the workflow itself, we incorporate workflow flexibility by evaluating its effectiveness through downstream tasks 2 and 3. All three tasks are suitable for real-time livestream applications, with only task 1 requiring an initial observation period to infer the workflow structure.

Evaluation Metrics. For task 1, we adopt Mean Absolute Percentage Error (MAPE) to evaluate period counting accuracy. This approach aligns with previous evaluation protocols in S-PAD [21, 48, 56, 58, 75], explicitly using its cor-

rect name. Specifically, suppose we have N sequences, the estimated and ground truth (GT) period counts are denoted by $\{\hat{m}_1, \hat{m}_2, \dots, \hat{m}_N\}$ and $\{m_1, m_2, \dots, m_N, \forall m_i \geq 3\}$, respectively. The MAPE is defined as:

$$\text{MAPE} = \frac{1}{N} \sum_{i=1}^N \frac{|\hat{m}_i - m_i|}{m_i}. \quad (1)$$

To enhance long-term period assessment, we utilize the average score of Temporal Intersection over Union (tIoU) [6, 18, 84] to quantify boundary accuracy. Given that the number of estimated periods may differ from the GT, we apply Hungarian matching [41] to identify the optimal alignment between predictions and GTs, maximizing the average scores in each sequence. Similar approaches have been applied in previous unsupervised SAS evaluations [42, 43, 50, 70, 71, 89]. The average tIoU compares the start (s) and end (e) temporal points of the predicted and GT boundaries:

$$\text{tIoU}_{\text{period}} = \frac{1}{N} \sum_{i=1}^N \frac{1}{m_i} \sum_{j=1}^{m_i} \frac{\max(0, \min(e_{i,j}, \hat{e}_{i,j}) - \max(s_{i,j}, \hat{s}_{i,j}))}{\max(e_{i,j}, \hat{e}_{i,j}) - \min(s_{i,j}, \hat{s}_{i,j})}. \quad (2)$$

For task 2, we employ Mean Absolute Error (MAE) to compare the estimated remaining proportion value and the GT:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |\hat{z}_i - z_i|, \quad (3)$$

where $\hat{z}_i$ and z_i are the estimated and GT remaining phase proportions, respectively. For task 3, we currently only assign one anomalous section per test period. Thus, unlike task 1, we do not apply Hungarian matching but instead, rely on tIoU to evaluate the overlap between predicted anomalous regions and GT in each sequence [30], as follows:

$$\text{tIoU}_{\text{anomaly}} = \frac{1}{N} \sum_{i=1}^N \frac{\max(0, \min(e_i, \hat{e}_i) - \max(s_i, \hat{s}_i))}{\max(e_i, \hat{e}_i) - \min(s_i, \hat{s}_i)}. \quad (4)$$

4. Our Baseline Method

Although human activity occurs continuously in spatiotemporal space, its features can be discretized into tokenized representations. Ideally, activity tokens might be represented as textual semantic tags [65, 104]; however, despite advances in vision foundation models [3, 8, 82, 100], it is challenging to fully capture nuanced spatial-temporal relationships, intricate coordination, and procedural complexities that transcend verbal descriptions [46]. For instance, text alone cannot precisely convey the dynamic motions of two hands. To effectively manage activity tokens without additional model training, an alternative approach is to represent them using a finite set of symbols [32, 33, 46], with




Figure 3. **Our baseline method.** In Step 1, we construct an activity transcript using soft tokens to determine the initial period window size (w). In Step 2, we apply our sequential mining algorithm to extract the workflow and identify the period boundaries.

each symbol corresponding to a specific activity spatiotemporal state. Consequently, a transcript of human activity can be tokenized into a symbolic transcript (*i.e.*, a string). Given a human activity transcript of length T with frame-wise spatiotemporal features $\{\mathbf{x}_t \in \mathbb{R}^{n \times 1}\}_{t=0}^{T-1}$, we perform spatiotemporal clustering to organize these features into K clusters. While K can be automatically determined using non-parametric clustering methods [68], we simply use K-means clustering [39] for convenience. Empirically, a setting K between 6 and 14 is sufficient to represent most workflows aimed at our task completion tracking and anomaly detection.

After the clustering, let $\mathbf{C}_k$ be the center of the k -th cluster $k = 0, 1, \dots, K-1$. We then tokenize $\{\mathbf{x}_t\}_{t=0}^{T-1}$ into hard and soft tokens, denoted as y_t^{hard} and y_t^{soft} , respectively, as follows:

$$y_t^{\text{hard}} = \arg \min_k \|\mathbf{x}_t - \mathbf{C}_k\|; \quad (5)$$

$$y_t^{\text{soft}} = \{d_{tk}\}_{k=0}^{K-1}, d_{tk} = \frac{\exp(-\|\mathbf{x}_t - \mathbf{C}_k\|)}{\sum_{j=0}^{K-1} \exp(-\|\mathbf{x}_t - \mathbf{C}_j\|)}. \quad (6)$$

The hard-encoded transcript $\{y^{\text{hard}}\}_{t=0}^{T-1}$ forms a 1D array, while the soft-encoded transcript $\{y^{\text{soft}}\}_{t=0}^{T-1}$ forms a 2D matrix. Next, we attempt to obtain a coarse period window size from our encoded activity transcript. As shown in Fig. 1, the period cannot be detected by simply searching for peak signals within the activity transcript. Alternatively,

Fast Fourier Transform (FFT) [20] and AutoCorrelation Function (ACF) [12] are two prevalent approaches for detecting periods in sequential data. FFT converts a sequence into the frequency domain and predicts the period using the inverse of the frequency with the highest power. ACF assesses the similarity between a signal and its time-shifted versions, revealing repeating patterns through peaks at lag values corresponding to the period. Both FFT and ACF are proposed to work on 1D data [12, 24, 51], corresponding to our $\{y^{\text{hard}}\}_{t=0}^{T-1}$. However, in highly overlapping spatiotemporal trajectories, boundary ambiguities inherited in $\{y^{\text{hard}}\}_{t=0}^{T-1}$ can introduce noise and impede period estimation. To reduce ambiguities, we use $\{y^{\text{soft}}\}_{t=0}^{T-1}$ to incorporate the global spatiotemporal context, represented by K features at each time step t . Next, we accommodate 2D FFT with a context marginalization to estimate the period window size w in the temporal domain, which is formulated as follows:

$$\mathcal{F}(u, v) = \sum_{t=0}^{T-1} \sum_{k=0}^{K-1} y_{k,t}^{\text{soft}} \cdot e^{-j2\pi(\frac{u}{K} + \frac{v}{T})}, \quad (7)$$

where (u, v) are the frequency indices. Let $|\cdot|$ represent the magnitude spectrum, the frequency components along the temporal axis and the corresponding magnitude are:

$$f(v) = \frac{v}{T}, \quad v = 0, 1, \dots, T-1; \quad (8)$$

$$\text{Mag}(f(v)) = \sum_{u=0}^{K-1} |\mathcal{F}(u, v)|. \quad (9)$$

Since the contextual periodic signal is present in K dimensions, marginalization amplifies the periodic component in $\text{Mag}(f(v))$. We then select the top-3 dominant frequencies $\{f_{v,\max}^{(0)}, f_{v,\max}^{(1)}, f_{v,\max}^{(2)}\}$ and their corresponding period window sizes $\{w_{\max}^{(0)}, w_{\max}^{(1)}, w_{\max}^{(2)}\}$ by using:

$$[f_{v,\max}^{(0)}, f_{v,\max}^{(1)}, f_{v,\max}^{(2)}] = \text{argsort}_{f(v)}(\text{Mag}(f(v)))[-3:], \quad (10)$$

$$\text{where } w_{\max}^{(j)} = \frac{1}{f_{v,\max}^{(j)}}, \quad f_{v,\max}^{(j)} \neq 0, \quad j = 0, 1, 2.$$

The obtained window sizes could be further refined. We first use the window candidates to partition $\{y^{\text{soft}}\}_{t=0}^{T-1}$ into segments, denoted as $\{\{y^{\text{soft}}\}_{t=0}^{w_{\max}^{(j)}}, \{y^{\text{soft}}\}_{t=w_{\max}^{(j)}}^{2w_{\max}^{(j)}}, \dots\}$. Next, we compute the Dynamic Time Warping (DTW) [37, 69] distances between these segments to rerank the period window sizes, yielding $\{w_{\text{rerank}}^{(0)}, w_{\text{rerank}}^{(1)}, w_{\text{rerank}}^{(2)}\}$. Finally, we select $w_{\text{rerank}}^{(0)}$ as the initial period window size.

Although $w_{\text{rerank}}^{(0)}$ has been obtained, the period boundaries are still unknown. Some prior works assume a uniform period scale and directly use the estimated window to segment periods. However, we argue that period window sizes can vary. Therefore, we employ the FFT-derived scale as an initial estimate and subsequently extract the workflow to pinpoint precise period boundaries (Fig. 3).

Algorithm 1 Multiple Transcript Alignment (MTA)

```
1: function MTA(transcripts, gap_penalty = -1)
2:   F, P  $\leftarrow$  InitializeMatrix(transcripts)  $\triangleright$  Alg. 2
3:   dims  $\leftarrow$  F.shape
4:   for all pos  $\in$  product(*[range(1, d) for d in dims]) do
5:     max_score, best_moves  $\leftarrow$   $-\infty$ , []
6:     for all neighbor in GetNeighbors(pos, dims) do  $\triangleright$  Alg. 2
7:       chars  $\leftarrow$  ['.' if  $c = p$  else transcripts[i][c - 1] for
         i, (c, p) in enumerate(zip(pos, neighbor))]
8:       score  $\leftarrow$  F[neighbor] + ScoreMatch(chars)  $\triangleright$  Alg. 2
9:       if score > max_score then
10:        max_score  $\leftarrow$  score, best_moves  $\leftarrow$  [neighbor]
11:       else if score == max_score then
12:        best_moves.append(neighbor)
13:       F[pos], P[pos]  $\leftarrow$  max_score, best_moves
14:   aligned  $\leftarrow$  [[] for _ in transcripts]
15:   current_pos  $\leftarrow$  tuple(d - 1 for d in dims)
16:   while any(current_pos) and P[current_pos] do
17:     prev_pos  $\leftarrow$  P[current_pos][0]
18:     for all i, (c, p)  $\in$  enumerate(zip(current_pos, prev_pos)) do
19:       aligned[i].append(transcripts[i][c - 1] if  $c \neq p$  else '-')
20:     current_pos  $\leftarrow$  prev_pos
21:   return [('', join(seq))[: -1] for seq in aligned]
```

Algorithm 2 Helper Functions

```
1: function InitializeMatrix(transcripts)  $\triangleright$  Create DP score and pointer
   matrix
2:   dims  $\leftarrow$  [len(seq) + 1 for seq in transcripts]
3:   F, P  $\leftarrow$  zeros(dims), empty(dims)
4:   for all idx in ndindex(*dims) do P[idx]  $\leftarrow$  []
5:   for all idx, dim in enumerate(dims) do
6:     F[tuple([slice(None) if i == idx else 0 for i in range(len(dims))])]
7:      $\leftarrow$  linspace(0, -len(transcripts)  $\times$  dim, dim)
8:   return F, P
9: function GetNeighbors(current_pos, dims)  $\triangleright$  List valid positions in
   the DP matrix
10:  neighbors  $\leftarrow$  []
11:  for all i  $\in$  [0, 2len(dims)) do
12:    for all j, pos  $\in$  enumerate(current_pos) do
13:      if i & (1 << j) and pos > 0 then
14:        neighbor.append(pos - 1)
15:      else neighbor.append(pos)
16:    if len(neighbor) == len(dims) then
17:      neighbors.append(tuple(neighbor))
18:  return neighbors[1:]
19: function ScoreMatch(chars)  $\triangleright$  Compute alignment score for given
   transcripts
20:  if '.' in chars then return -count([c for c in chars if c == '.'])
21:  else return  $\sum_{i,j} (\text{chars}[i] == \text{chars}[j]) - \text{len}(\text{chars})$ 
```

Specifically, we segment the entire transcript into multiple sections based on the estimated time window and incorporating a 20% buffer at both sides to account for temporal variations. We then adapt a modified Needleman–Wunsch (NW) algorithm [59] to detect the workflow shared across all segments. The original NW algorithm optimizes sequence alignment by minimizing gaps (*i.e.*, '-') and maximizing matches between two sequences while considering all possible alignments. We extend the NW algorithm into

a Multiple Transcript Alignment (MTA) algorithm, which jointly optimizes multiple sequence alignments through dynamic programming (DP). The detailed methodology is outlined in Algs. 1 & 2. The core idea is to fill a multi-dimensional DP matrix **F**, where each cell **F**[pos] represents the optimal score to align prefixes of the input sequences ending at pos:

$$\mathbf{F}[\text{pos}] = \max_{\text{prev} \in \text{GetNeighbors}(\text{pos}, \text{dims})} \left(\mathbf{F}[\text{prev}] + \text{ScoreMatch} \left(\begin{bmatrix} \text{'.'} & \text{if } \text{pos}_i = \text{prev}_i \\ \text{transcripts}[\text{z}][\text{pos}_i - 1] & \text{if } \text{pos}_i \neq \text{prev}_i \end{bmatrix}_i \right) \right) \quad (11)$$

Due to the inclusion of buffers, redundant transcripts could be added. Thus, we implement a bidirectional traversal strategy to identify the first recurring motif (inter repetition) as the definitive boundary point. Additionally, gaps ('.') occurring before the start symbol or after the end symbol are also considered as boundary points. Elements situated beyond these demarcated boundaries are systematically eliminated to improve alignment accuracy. Moreover, we model inner variation through a multi-branch workflow to capture the complexity of divergent transcriptional pathways.

We perform period boundary detection, completion tracking, and anomaly detection by aligning input stream tokens with workflow tokens. Specifically, we initialize two pointers: one for the workflow and one for the input stream. The workflow pointer loops within the workflow and is initialized at its end, while the input stream pointer starts at the beginning of the stream. At each time step, we compare the tokens pointed to in both sequences. If the current input token matches the workflow's start token while the workflow pointer is at its end, this indicates that a new period has begun and the previous one has ended. As the stream progresses, we advance both pointers in synchronization, enabling continuous, online processing of all three tasks.

5. Experiments

We conducted a series of experiments using the metrics introduced in Sec. 3 and the activity tokenization method introduced in Sec. 4. For task 1, we follow the common unsupervised evaluation protocols and apply all samples as the test set, allowing the model to learn the underlying patterns and perform estimations without any training data. The workflows extracted from task 1 are then used for tasks 2 and 3. We evaluate our baseline against traditional unsupervised mining methods, supervised PAD and SAS approaches, and LLMs. Our results are presented in Tabs. 3, 5, and 6, and Figs. 4, 5, 6, and 7. The best and second-best results are marked in purple and orange, respectively. Below, we summarize our key findings from the results.

Our tokenization enables LLM-based periodic analysis. Given our tokenized transcripts (Fig. 3), we apply

	Task 1 (PD)		Task 2 (CT)		Task 3 (AD)	
	MAPE $\downarrow$ tloU _{period} $\uparrow$	Runtime $\downarrow$	MAPE $\downarrow$ tloU _{period} $\uparrow$	Runtime $\downarrow$	tloU _{anomaly} $\uparrow$	Runtime $\downarrow$
Unsupervised sequential periodic detection (applicable to task 1: Period Detection (PD))						
ClaSP [25]	1.284	0.331	0.4X	-	-	-
CFDAutoPeriod [64]	8.396	0.146	6X	-	-	-
1D FFT [36]	2.746	0.399	0.6X	-	-	-
ACF [12]	2.432	0.439	4X	-	-	-
Unsupervised sequential anomaly detection (applicable to task 3: Anomaly Detection (AD))						
TadGAN [5, 29]	-	-	-	-	0.066	>100X
AER [5, 86]	-	-	-	-	0.068	>100X
AnomalTrans [88]	-	-	-	-	0.085	>100X
Unsupervised sequential reasoning (applicable to all tasks: PD, Completion Tracking (CT), AD)						
GPT-4o [3, 35]	0.038	0.882	$\approx 5X$	0.388	$\approx 5X$	$\approx 5X$
Gemini-2.5 Pro [15]	0.030	0.915	$\approx 10X$	0.286	$\approx 10X$	$\approx 10X$
Our Baseline	0.025	0.937	1X	0.106	1X	1X

(a) Comparison of unsupervised methods on our 3 benchmark tasks. Workflows are extracted from task 1 and used for tasks 2 and 3.

Token for w	Buffer Size	Task 1 (PD)	Task 2 (CT)	Task 3 (AD)
Initialization		MAPE $\downarrow$ tloU _{period} $\uparrow$	MAPE $\downarrow$	tloU _{anomaly} $\uparrow$
w/ window size re-ranking				
$\{y_{\text{hard}}\}_{i=0}^{T-1}$	10%	0.149	0.842	0.164
$\{y_{\text{hard}}\}_{i=0}^{T-1}$	20%	0.147	0.854	0.160
$\{y_{\text{hard}}\}_{i=0}^{T-1}$	40%	0.150	0.833	0.167
$\{y_{\text{soft}}\}_{i=0}^{T-1}$	10%	0.026	0.930	0.108
$\{y_{\text{soft}}\}_{i=0}^{T-1}$	20%	0.025	0.937	0.106
$\{y_{\text{soft}}\}_{i=0}^{T-1}$	40%	0.036	0.928	0.124
w/o window size re-ranking				
$\{y_{\text{hard}}\}_{i=0}^{T-1}$	20%	1.392	0.375	0.623
$\{y_{\text{soft}}\}_{i=0}^{T-1}$	20%	0.932	0.539	0.342

(b) Ablation studies on our baseline (re-ranking & token & buffer). Our standard baseline parameters are rendered in green.

Table 3. **Evaluations on our benchmark.** We apply a constant $K = 10$ for tokenization and use the identical transcripts for all methods as inputs. Runtime on each task of our benchmark is normalized to our baseline (1X).




Figure 4. **Ablation studies with various K .** Larger values improve performance at the expense of higher computational cost.




Figure 5. **An example of workflows generated for the same activity sequence with different values of K .** The symbol ‘.’ represents a skipped token. As K increases, the workflows become more detailed and extended.

both GPT-4o[3, 35] and Gemini-2.5 [15] to analyze our benchmark, using the task descriptions directly as prompts (Tab. 3a). While they are comparable to our baseline method on task 1, they struggle with tasks 2 and 3. One possible explanation is that when provided with hundreds of repetitive, low-contrast tokens, LLMs can identify the repetitive patterns but may have difficulty extracting the long-term workflow used for tasks 2 and 3.

Traditional unsupervised sequential mining methods struggle on our benchmark. Unlike previous datasets used for periodicity and anomaly detection studies, our dataset features low-contrast periodic patterns and complex workflows that significantly challenge existing approaches (Tab. 3a). For instance, in period counting, the MAPE scores of unsupervised sequential periodic detection methods we evaluated [12, 25, 36, 64] are over 50 times higher than that of our baseline method. Similarly, in anomaly detection, the representative methods we assessed [29, 86, 88]

perform nearly at the level of random guessing. Therefore, our benchmark presents a novel challenge to the research community, encouraging the development of new solutions.

Ablation studies on baseline parameters. We demonstrate the effectiveness of employing our initialized window re-ranking, soft tokenization, and 20% buffer size for period window estimation, as detailed in Tab. 3b. We also explore the impact of varying the parameter K in Fig. 4. Changes in K influence the representation of the extracted workflows. In complicated scenarios, a larger K leads to fine-grained tokens and may remove ambiguities with a smaller K , thereby improving performance on our benchmark. However, this improvement comes with increased computational cost. Fig. 5 illustrates the extracted workflows for various K values. As K increases, the workflows become more detailed and extended. Notably, the extracted workflows exhibit multi-branch structures that effectively capture diverse approaches within the activity while




Figure 6. Illustration of applying our baseline to analyze production workflows (see supp. video for details).

accounting for spatiotemporal variations and noise.

Our baseline outperforms supervised SAS methods on task 3.

Although existing SAS methods cannot handle periodic workflow detection in task 1, they can be adapted to task 3 by training on our annotated task 1 data to estimate

	tIoU _{anomaly} (Task 3)↑
Supervised (w/ task 1 labeled data)	
ASFormer [95]	0.294
FACT [57]	0.312
Unsupervised (w/ task 1 unlabeled data)	
HVQ [78]	0.087
Our Baseline	0.423

Table 4. Comparison of SAS methods for task 3.

workflows and detect anomalies. We compare our baseline with both supervised and unsupervised SAS methods in Tab. 4. Our baseline outperforms supervised methods. The unsupervised SAS method HVQ [78], which extracts patterns from unlabeled periodic data without explicitly modeling periodicity, fails to generalize to task 3. As the SAS benchmarks do not involve periodicity, which is beyond the scope of our study, we do not evaluate our baseline on them.

Our baseline outperforms on the S-PAD benchmark, but supervised S-PAD methods underperform on our benchmark. Tab. 5 illustrates that our baseline method, enhanced with advanced activity tokenization and sequential modeling, effectively detects periodic patterns in an unsupervised manner. This approach surpasses several supervised methods [22, 34, 75, 94] on the traditional S-PAD benchmark RepCount [34]. In contrast, representative conventional S-PAD methods [75, 94] perform poorly on our benchmark—even when provided with supervision—highlighting their limitations in modeling long-term periodic patterns. These results demonstrate that our method is capable of effectively handling both short- and long-term periodic structures.

Our baseline is efficient in real-world deployments. We validate the effectiveness of our approach by deploying it




Figure 7. Comparison of deployment costs. **Top:** Deployment for one livestream. **Bottom:** Deployment for two different livestreams.

	MAPE (Task 1) ↓
Supervised	
TransRAC [34]	0.443
RepNet [22]	0.331
SkimFocus [103]	0.249
PoseRAC [94]	0.236
ESCounts [75]	0.213
Unsupervised	
Our Baseline	0.156

Table 5. Our baseline on the S-PAD benchmark.

	MAPE (Task 1) ↓
Supervised	
PoseRAC [94]	0.738
ESCounts [75]	0.673
Unsupervised	
Our Baseline	0.023

Table 6. S-PAD methods on our benchmark. 380 samples with video inputs are used for task 1 evaluation. We use 4-fold cross-validation for supervised methods.

on a factory production line (see Fig. 6). Our method accurately identifies the production workflow patterns from raw periodic video streams, providing actionable insights for tracking the remaining phase proportion of an ongoing cycle and detecting procedural anomalies. We also conceptually analyze the deployment costs of our method relative to supervised alternatives (see Fig. 7). Overall, our method supports seamless deployment across various production lines, even as procedures and schedules change frequently, demonstrating both its scalability and efficiency in real-world applications.

6. Conclusion

We introduce a pioneering benchmark and baseline for the unsupervised detection of long-term periodic spatiotemporal workflows in human activities. We show its practical applications through three evaluation tasks: unsupervised period detection, task completion tracking, and anomaly detection. Our baseline, enhanced with advanced activity tokenization and sequential modeling, effectively detects periodic patterns in an unsupervised manner. We believe our work lays the foundation for future research and applications in long-term periodic human activity analysis.

References

- [1] Ali Abedi, Paritosh Bisht, Riddhi Chatterjee, Rachit Agrawal, Vyom Sharma, Dinesh Babu Jayagopi, and Shehroz S Khan. Rehabilitation exercise repetition segmentation and counting using skeletal body joints. In *2023 20th Conference on Robots and Vision (CRV)*, pages 288–295. IEEE, 2023. 1
- [2] Yazan Abu Farha and Juergen Gall. MS-TCN: Multi-stage temporal convolutional network for action segmentation. In *CVPR*, 2019. 2, 3
- [3] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 4, 7
- [4] 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 conference on computer vision and pattern recognition*, pages 4575–4583, 2016. 3
- [5] Sarah Alnegheimish, Dongyu Liu, Carles Sala, Laure Berti-Equille, and Kalyan Veeramachaneni. Sintel: A machine learning framework to extract insights from signals. In *Proceedings of the 2022 International Conference on Management of Data*, page 1855–1865. Association for Computing Machinery, 2022. 7
- [6] Humam Alwassel, Fabian Caba Heilbron, and Bernard Ghanem. Action search: Spotting actions in videos and its application to temporal action localization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 251–266, 2018. 4
- [7] Apple Inc. Arkit. developer.apple.com/augmented-reality/arkit/, 2017. Accessed: 2025. 4
- [8] Muhammad Awais, Muzammal Naseer, Salman Khan, Rao Muhammad Anwer, Hisham Cholakkal, Mubarak Shah, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Foundation models defining a new era in vision: a survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 4
- [9] Nadine Behrmann, S Alireza Golestaneh, Zico Kolter, Jürgen Gall, and Mehdi Noroozi. Unified fully and timestamp supervised temporal action segmentation via sequence to sequence translation. In *European Conference on Computer Vision*, pages 52–68. Springer, 2022. 2, 3
- [10] Tudor O Bompá and Carlo Buzzichelli. *Periodization: theory and methodology of training*. Human kinetics, 2019. 1
- [11] Borja Bordel, Ramón Alcarria, and Tomás Robles. Recognizing human activities in industry 4.0 scenarios through an analysis-modeling-recognition algorithm and context labels. *Integrated Computer-Aided Engineering*, 29(1):83–103, 2022. 1
- [12] Tim Breitenbach, Bartosz Wilkusz, Lauritz Rasbach, and Patrick Jahnke. On a method for detecting periods and repeating patterns in time series data with autocorrelation and function approximation. *Pattern Recognition*, 138:109355, 2023. 5, 7
- [13] Huiping Cao, Nikos Mamoulis, and David W Cheung. Discovery of periodic patterns in spatiotemporal sequences. *IEEE Transactions on Knowledge and Data Engineering*, 19(4):453–467, 2007. 3
- [14] Sheng-Hsien Cheng, Muhammad Atif Sarwar, Yousef-Awwad Daraghmi, Tsi-Uí Ik, and Yih-Lang Li. Periodic physical activity information segmentation, counting and recognition from video. *IEEE Access*, 11:23019–23031, 2023. 1
- [15] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Naveen Sachdeva, Inderjit Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. 7
- [16] Matteo Destro and Michael Gygli. Cyclecl: Self-supervised learning for periodic videos. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2861–2870, 2024. 1, 3
- [17] Guodong Ding and Angela Yao. Leveraging action affinity and continuity for semi-supervised temporal action segmentation. In *European Conference on Computer Vision*, pages 17–32. Springer, 2022. 2, 3
- [18] Guodong Ding, Fadime Sener, and Angela Yao. Temporal action segmentation: An analysis of modern techniques. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 4
- [19] 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*, pages 3323–3332, 2022. 2, 3
- [20] Pierre Duhamel and Martin Vetterli. Fast fourier transforms: a tutorial review and a state of the art. *Signal processing*, 19(4):259–299, 1990. 5
- [21] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. Counting out time: Class agnostic video repetition counting in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10387–10396, 2020. 1, 3, 4
- [22] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. A short note on evaluating repnet for temporal repetition counting in videos. *arXiv preprint arXiv:2411.08878*, 2024. 8
- [23] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, and Andrew Zisserman. Ovr: A dataset for open vocabulary temporal repetition counting in videos. *arXiv preprint arXiv:2407.17085*, 2024. 1, 3
- [24] Mohamed G Elfeky, Walid G Aref, and Ahmed K Elmagarmid. Periodicity detection in time series databases. *IEEE Transactions on Knowledge and Data Engineering*, 17(7): 875–887, 2005. 5
- [25] Arik Ermshaus, Patrick Schäfer, and Ulf Leser. Clasp: parameter-free time series segmentation. *Data Mining and Knowledge Discovery*, 37(3):1262–1300, 2023. 7
- [26] Alireza Fathi, Xiaofeng Ren, and James M Rehg. Learning to recognize objects in egocentric activities. In *CVPR 2011*, pages 3281–3288. IEEE, 2011. 3

- [27] Philippe Fournier-Viger, Tin Truong Chi, Youxi Wu, Jun-Feng Qu, Jerry Chun-Wei Lin, and Zhitian Li. Finding periodic patterns in multiple sequences. *Periodic Pattern Mining: Theory, Algorithms, and Applications*, pages 81–103, 2021. 3
- [28] Yixin Gao, S Swaroop Vedula, Carol E Reiley, Narges Ahmadi, Balakrishnan Varadarajan, Henry C Lin, Lingling Tao, Luca Zappella, Benjamin Béjar, David D Yuh, et al. Jhu-isi gesture and skill assessment working set (jigsaws): A surgical activity dataset for human motion modeling. In *MICCAI workshop: M2cai*, page 3, 2014. 3
- [29] Alexander Geiger, Dongyu Liu, Sarah Alnegheimish, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. Tadgan: Time series anomaly detection using generative adversarial networks. In *2020 IEEE International Conference on Big Data (IEEE BigData)*, pages 33–43. IEEE, 2020. 7
- [30] Reza Ghoddoosian, Isht Dwivedi, Nakul Agarwal, and Behzad Dariush. Weakly-supervised action segmentation and unseen error detection in anomalous instructional videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10128–10138, 2023. 4
- [31] Abhirup Ghosh, Christopher Lucas, and Rik Sarkar. Finding periodic discrete events in noisy streams. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 627–636, 2017. 1
- [32] Gutemberg Guerra-Filho and Yiannis Aloimonos. A language for human action. *Computer*, 40(5):42–51, 2007. 4
- [33] Yiwei Guo, Zhihan Li, Hankun Wang, Bohan Li, Chongtian Shao, Hanglei Zhang, Chenpeng Du, Xie Chen, Shujie Liu, and Kai Yu. Recent advances in discrete speech tokens: A review. *arXiv preprint arXiv:2502.06490*, 2025. 4
- [34] Huazhang Hu, Sixun Dong, Yiqun Zhao, Dongze Lian, Zhengxin Li, and Shenghua Gao. Transrac: Encoding multi-scale temporal correlation with transformers for repetitive action counting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19013–19022, 2022. 1, 3, 8
- [35] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 7
- [36] RJ Hyndman. *Forecasting: principles and practice*. OTexts, 2018. 7
- [37] Fumitada Itakura. Minimum prediction residual principle applied to speech recognition. *IEEE Transactions on acoustics, speech, and signal processing*, 23(1):67–72, 2003. 5
- [38] Tanvi Jindal, Prasanna Giridhar, Lu-An Tang, Jun Li, and Jiawei Han. Spatiotemporal periodical pattern mining in traffic data. In *Proceedings of the 2nd ACM SIGKDD international workshop on urban computing*, pages 1–8, 2013. 3
- [39] Trupti M Kodinariya, Prashant R Makwana, et al. Review on determining number of cluster in k-means clustering. *International Journal*, 1(6):90–95, 2013. 5
- [40] 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 conference on computer vision and pattern recognition*, pages 780–787, 2014. 3
- [41] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 4
- [42] Anna Kukleva, Hilde Kuehne, Fadime Sener, and Jurgen Gall. Unsupervised learning of action classes with continuous temporal embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12066–12074, 2019. 3, 4
- [43] Sateesh Kumar, Sanjay Hareesh, Awais Ahmed, Andrey Konin, M Zeeshan Zia, and Quoc-Huy Tran. Unsupervised action segmentation by joint representation learning and online clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20174–20185, 2022. 3, 4
- [44] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 156–165, 2017. 2, 3
- [45] Chan-Su Lee and Ahmed Elgammal. Human motion synthesis by motion manifold learning and motion primitive segmentation. In *International Conference on Articulated Motion and Deformable Objects*, pages 464–473. Springer, 2006. 1
- [46] Kenneth Li, Xiao Sun, Zhirong Wu, Fangyun Wei, and Stephen Lin. Towards tokenized human dynamics representation. *arXiv preprint arXiv:2111.11433*, 2021. 2, 4
- [47] Kun Li, Xinge Peng, Dan Guo, Xun Yang, and Meng Wang. Repetitive action counting with hybrid temporal relation modeling. *IEEE Transactions on Multimedia*, 2025. 1, 3
- [48] Xinjie Li and Huijuan Xu. Repetitive action counting with motion feature learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6499–6508, 2024. 1, 3, 4
- [49] Yin Li, Zhefan Ye, and James M Rehg. Delving into ego-centric actions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 287–295, 2015. 3
- [50] Yuerong Li, Zhengrong Xue, and Huazhe Xu. Otas: Unsupervised boundary detection for object-centric temporal action segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6437–6446, 2024. 3, 4
- [51] Zhenhui Li, Bolin Ding, Jiawei Han, Roland Kays, and Peter Nye. Mining periodic behaviors for moving objects. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1099–1108, 2010. 1, 5
- [52] Zhe Li, Yazan Abu Farha, and Jurgen Gall. Temporal action segmentation from timestamp supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8365–8374, 2021. 2, 3
- [53] Zishi Li, Xiaoxuan Ma, Qiuyan Shang, Wentao Zhu, Hai Ci, Yu Qiao, and Yizhou Wang. Efficient action counting with dynamic queries. *arXiv preprint arXiv:2403.01543*, 2024. 1, 3

- [54] Jonathan Feng-Shun Lin, Vladimir Joukov, and Dana Kulić. Classification-based segmentation for rehabilitation exercise monitoring. *Journal of rehabilitation and assistive technologies engineering*, 5:2055668318761523, 2018. 1
- [55] Daniel S Lorenz, Michael P Reiman, and John C Walker. Periodization: current review and suggested implementation for athletic rehabilitation. *Sports Health*, 2(6):509–518, 2010. 1
- [56] Jiada Lu, WeiWei Zhou, Xiang Qian, Dongze Lian, Yanyu Xu, Weifeng Wang, Lina Cao, and Shenghua Gao. Fcarac: First cycle annotated repetitive action counting. *arXiv preprint arXiv:2406.12178*, 2024. 1, 3, 4
- [57] 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*, pages 18175–18185, 2024. 8
- [58] Yanan Luo, Jinhui Yi, Yazan Abu Farha, Moritz Wolter, and Juergen Gall. Rethinking temporal self-similarity for repetitive action counting. In *2024 IEEE International Conference on Image Processing (ICIP)*, pages 2187–2193, 2024. 1, 3, 4
- [59] Saul B Needleman and Christian D Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of molecular biology*, 48(3):443–453, 1970. 6
- [60] Lionel M Ni, Yunhao Liu, Yiu Cho Lau, and Abhishek P Patil. Landmarc: Indoor location sensing using active rfid. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications, 2003.(PerCom 2003)*., pages 407–415. IEEE, 2003. 4
- [61] Costas Panagiotakis, Giorgos Karvounas, and Antonis Argyros. Unsupervised detection of periodic segments in videos. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 923–927. IEEE, 2018. 1, 3
- [62] Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees, Chelsea Finn, and Sergey Levine. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025. 1
- [63] Rolandos Alexandros Potamias, Jinglei Zhang, Jiankang Deng, and Stefanos Zafeiriou. Wilor: End-to-end 3d hand localization and reconstruction in-the-wild. *arXiv preprint arXiv:2409.12259*, 2024. 4
- [64] Tom Puech, Matthieu Boussard, Anthony D’Amato, and Gaëtan Millerand. A fully automated periodicity detection in time series. In *Advanced Analytics and Learning on Temporal Data: 4th ECML PKDD Workshop, AALTD 2019, Würzburg, Germany, September 20, 2019, Revised Selected Papers 4*, pages 43–54. Springer, 2020. 7
- [65] Haoxuan Qu, Yujun Cai, and Jun Liu. Llms are good action recognizers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18395–18406, 2024. 4
- [66] Rahul Rahaman, Dipika Singhania, Alexandre Thiery, and Angela Yao. A generalized and robust framework for timestamp supervision in temporal action segmentation. In *European Conference on Computer Vision*, pages 279–296. Springer, 2022. 2, 3
- [67] Alexander Richard, Hilde Kuehne, Ahsan Iqbal, and Juergen Gall. NeuralNetwork-Viterbi: A framework for weakly supervised video learning. In *CVPR*, 2018. 2, 3
- [68] Stephen J Roberts. Parametric and non-parametric unsupervised cluster analysis. *Pattern Recognition*, 30(2):261–272, 1997. 5
- [69] Stan Salvador and Philip Chan. Toward accurate dynamic time warping in linear time and space. *Intelligent Data Analysis*, 11(5):561–580, 2007. 5
- [70] Saquib Sarfraz, Naila Murray, Vivek Sharma, Ali Diba, Luc Van Gool, and Rainer Stiefelhagen. Temporally-weighted hierarchical clustering for unsupervised action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11225–11234, 2021. 3, 4
- [71] Fadime Sener and Angela Yao. Unsupervised learning and segmentation of complex activities from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8368–8376, 2018. 3, 4
- [72] Fadime Sener, Dibyaadip Chatterjee, Daniel Shelepov, Kun He, Dipika Singhania, Robert Wang, and Angela Yao. Assembly101: A large-scale multi-view video dataset for understanding procedural activities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21096–21106, 2022. 3
- [73] Zehong Shen, Huaijin Pi, Yan Xia, Zhi Cen, Sida Peng, Zechen Hu, Hujun Bao, Ruizhen Hu, and Xiaowei Zhou. World-grounded human motion recovery via gravity-view coordinates. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–11, 2024. 4
- [74] Soyong Shin, Juyong Kim, Eni Halilaj, and Michael J Black. Wham: Reconstructing world-grounded humans with accurate 3d motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2070–2080, 2024. 4
- [75] Saptarshi Sinha, Alexandros Stergiou, and Dima Damen. Every shot counts: Using exemplars for repetition counting in videos. pages 384–402, 2024. 1, 3, 4, 8
- [76] Xiaomin Song, Qingsong Wen, Yan Li, and Liang Sun. Robust time series dissimilarity measure for outlier detection and periodicity detection. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 4510–4514, 2022. 1
- [77] 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*, 44(10):6196–6208, 2022. 2, 3
- [78] Federico Spurio, Emad Bahrami, Gianpiero Francesca, and Juergen Gall. Hierarchical vector quantization for unsupervised action segmentation. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2025. 2, 3, 8
- [79] 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. 3

- [80] Yuhao Su and Ehsan Elhamifar. Two-stage active learning for efficient temporal action segmentation. In *European Conference on Computer Vision*, page 161–183, 2024. 2, 3
- [81] Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou. Coin: A large-scale dataset for comprehensive instructional video analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1207–1216, 2019. 3
- [82] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 4
- [83] Wil Van Der Aalst. Process mining. *Communications of the ACM*, 55(8):76–83, 2012. 3
- [84] Binglu Wang, Yongqiang Zhao, Le Yang, Teng Long, and Xuelong Li. Temporal action localization in the deep learning era: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 4
- [85] Yufu Wang, Ziyun Wang, Lingjie Liu, and Kostas Daniilidis. Tram: Global trajectory and motion of 3d humans from in-the-wild videos. In *European Conference on Computer Vision*, pages 467–487. Springer, 2025. 4
- [86] Lawrence Wong, Dongyu Liu, Laure Berti-Equille, Sarah Alnegheimish, and Kalyan Veeramachaneni. Aer: Auto-encoder with regression for time series anomaly detection. In *2022 IEEE International Conference on Big Data (IEEE BigData)*, pages 1152–1161. IEEE, 2022. 7
- [87] 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*, pages 18253–18262, 2024. 2, 3
- [88] Jiehui Xu, Haixu Wu, Jianmin Wang, and Mingsheng Long. Anomaly transformer: Time series anomaly detection with association discrepancy. In *International Conference on Learning Representations*, 2022. 7
- [89] Ming Xu and Stephen Gould. Temporally consistent unbalanced optimal transport for unsupervised action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14618–14627, 2024. 3, 4
- [90] Chenguang Yang, Chao Zeng, Yang Cong, Ning Wang, and Min Wang. A learning framework of adaptive manipulative skills from human to robot. *IEEE Transactions on Industrial Informatics*, 15(2):1153–1161, 2018. 1
- [91] Fan Yang, Yang Wu, Sakriani Sakti, and Satoshi Nakamura. Make skeleton-based action recognition model smaller, faster and better. In *Proceedings of the 1st ACM International Conference on Multimedia in Asia*, pages 1–6, 2019. 2
- [92] Fan Yang, Shigeyuki Odashima, Shoichi Masui, and Shan Jiang. Is weakly-supervised action segmentation ready for human-robot interaction? no, let’s improve it with action-union learning. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9800–9807. IEEE, 2023. 3
- [93] Yuzhe Yang, Xin Liu, Jiang Wu, Silviu Borac, Dina Katabi, Ming-Zher Poh, and Daniel McDuff. Simper: Simple self-supervised learning of periodic targets. In *International Conference on Learning Representations*, 2023. 1, 3
- [94] Ziyu Yao, Xuxin Cheng, and Yuexian Zou. Poserac: Pose saliency transformer for repetitive action counting. *arXiv preprint arXiv:2303.08450*, 2023. 1, 3, 8
- [95] Fangqiu Yi, Hongyu Wen, and Tingting Jiang. ASFormer: Transformer for action segmentation. In *BMVC*, 2021. 2, 3, 8
- [96] Faheem Zafari, Athanasios Gkelias, and Kin K Leung. A survey of indoor localization systems and technologies. *IEEE Communications Surveys & Tutorials*, 21(3):2568–2599, 2019. 4
- [97] Yue Zeng, Hongchao Qin, Rong-Hua Li, Kai Wang, Guoren Wang, and Xuemin Lin. Mining quasi-periodic communities in temporal network. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, pages 2476–2488. IEEE, 2024. 3
- [98] Dongzhi Zhang. *Periodic pattern mining from spatio-temporal trajectory data*. PhD thesis, James Cook University, 2018.
- [99] Dongzhi Zhang, Kyungmi Lee, and Ickjai Lee. Periodic pattern mining for spatio-temporal trajectories: a survey. In *2015 10th International Conference on Intelligent Systems and Knowledge Engineering (ISKE)*, pages 306–313. IEEE, 2015. 3
- [100] Jingyi Zhang, Jiaying Huang, Sheng Jin, and Shijian Lu. Vision-language models for vision tasks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 4
- [101] Jinglei Zhang, Jiankang Deng, Chao Ma, and Rolandos Alexandros Potamias. Hawor: World-space hand motion reconstruction from egocentric videos. *arXiv preprint arXiv:2501.02973*, 2025. 4
- [102] Yunhua Zhang, Ling Shao, and Cees GM Snoek. Repetitive activity counting by sight and sound. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14070–14079, 2021. 1, 3
- [103] Zhengqi Zhao, Xiaohu Huang, Hao Zhou, Kun Yao, Errui Ding, Jingdong Wang, Xinggang Wang, Wenyu Liu, and Bin Feng. Skim then focus: Integrating contextual and fine-grained views for repetitive action counting. *arXiv preprint arXiv:2406.08814*, 2024. 8
- [104] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. Motionbert: A unified perspective on learning human motion representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15085–15099, 2023. 4