

Human Motion Instruction Tuning

Lei Li^{1,2,*}, Sen Jia^{1,*}, Jianhao Wang³, Zhongyu Jiang¹, Feng Zhou⁴,
 Ju Dai⁵, Tianfang Zhang⁶, Zongkai Wu¹, Jenq-Neng Hwang¹

Abstract

This paper presents **LLaMo** (**Large Language and Human Motion Assistant**), a multimodal framework for human motion instruction tuning. In contrast to conventional instruction-tuning approaches that convert non-linguistic inputs, such as video or motion sequences, into language tokens, LLaMo retains motion in its native form for instruction tuning. This method preserves motion-specific details that are often diminished in tokenization, thereby improving the model’s ability to interpret complex human behaviors. By processing both video and motion data alongside textual inputs, LLaMo enables a flexible, human-centric analysis. Experimental evaluations across high-complexity domains, including human behaviors and professional activities, indicate that LLaMo effectively captures domain-specific knowledge, enhancing comprehension and prediction in motion-intensive scenarios. We hope LLaMo offers a foundation for future multimodal AI systems with broad applications, from sports analytics to behavioral prediction.

1. Introduction

Understanding human motion is a central challenge in multimodal AI, impacting numerous fields such as digital human, human-computer interaction, sports analytics, healthcare, and virtual human modeling [3–6]. Motion data, which captures skeletal movements and body dynamics, provides a structured, appearance-invariant representation of human actions, concentrating on the essential movement patterns while excluding irrelevant visual details [7]. This data has emerged as a privacy-conscious alternative to visual inputs, offering finer control over action recognition and facilitating analysis in applications ranging from digital human avatars to behavioral monitoring in healthcare [8–10]. As digital human representations and motion-centric

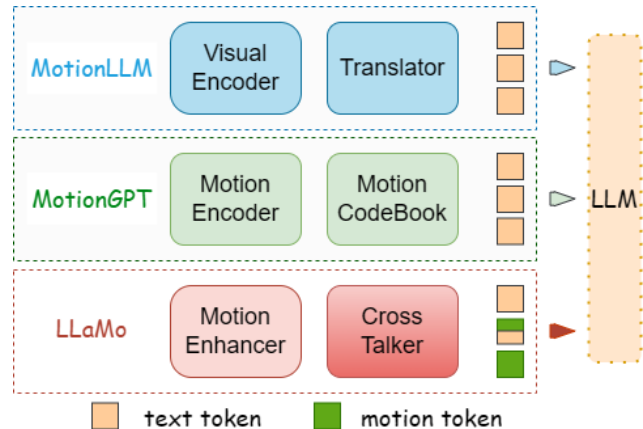


Figure 1. A comparison of MotionLLM [1], MotionGPT [2], and LLaMo highlights LLaMo’s motion-specific capabilities. Equipped with a Motion Enhancer and Cross Talker module to align motion and text, LLaMo supports both video and motion inputs, enabling text-aware, fine-grained motion analysis.

applications continue to expand, developing models that can leverage the unique characteristics of motion data becomes increasingly crucial for achieving robust, context-aware behavior analysis.

Recently, large language models (LLMs) have shown promise in motion analysis, leveraging their capacity to model temporal dependencies across multimodal data [1, 11]. State-of-the-art models, such as MotionGPT [12] and MotionLLM [1], treat motion data as a multimodal extension, encoding it as discrete tokens or translating it into language-like representations for processing by LLMs. These models have yielded notable results in motion understanding and recognition tasks, illustrating the adaptability of language models to non-linguistic data. However, their reliance on tokenization or textual conversion can limit the richness of the motion representation, abstracting away critical spatial-temporal details essential for high-fidelity human activity recognition.

Despite recent advancements, notable challenges remain in current approaches to motion analysis using LLMs. First, encoding motion data into language tokens, as implemented

* These authors contributed equally to this work.

† Corresponding Author. (lilei@di.ku.dk)

¹University of Washington ²University of Copenhagen ³Xi’an Jiaotong University ⁴North China University of Technology ⁵Peng Cheng Laboratory ⁶Tsinghua University

in models such as MotionGPT, relies on a quantization process that may obscure critical fine-grained motion information, potentially limiting the model’s capacity to accurately capture detailed spatial-temporal dynamics [13, 14]. The common approach of translating motion data into text tokens further contributes to this issue, often resulting in the loss of essential motion-specific nuances that are critical for in-depth behavior interpretation [11, 15]. This translation process poses considerable challenges in understanding the intricate details of human movement, impeding the model’s ability to achieve spatial and physical comprehension of human actions. Moreover, motion data inherently contains 3D structural information and insights into connected human behaviors, which are often difficult to fully capture when relying solely on language tokens, as this representation can overlook the underlying spatial dependencies crucial to modeling human motion accurately.

In addition, existing methods frequently treat video and motion data as isolated modalities, overlooking the potential benefits of a complementary, integrated approach [16]. Processing these data types independently neglects the rich environmental context and interaction cues provided by video inputs, which are crucial for accurately interpreting complex human behaviors [5, 8]. This separation limits the capacity for a holistic understanding of human motion, particularly in scenarios where context, spatial relations, and dynamic interactions play essential roles in behavior analysis. Furthermore, video-based approaches, while potentially more informative, are computationally intensive, posing significant constraints for large-scale or real-time applications, which limits their feasibility in practical motion analysis systems.

In response to these challenges, we propose LLaMo, a framework designed to integrate motion data as a distinct modality within LLMs without translating it into textual intermediaries. **LLaMo incorporates a Motion Estimator and an Enhancer to facilitate the direct input of motion or/and video data, thereby offering a unified, human-centric analysis pipeline.** The framework also introduces a **Cross Talker** module, which dynamically aligns the motion and text features, allowing for text-aware, fine-grained motion representations. This approach retains the original spatial-temporal characteristics of the motion data, providing a more precise understanding of human actions and enabling high-resolution behavior analysis in motion-intensive contexts [17, 18].

Through extensive evaluations on benchmarks such as MoVid-Bench [1] and BABEL-QA [19], we demonstrate that LLaMo achieves state-of-the-art performance in motion-centric human activity recognition [11, 15]. These results underscore LLaMo’s potential to advance multimodal AI, providing a robust foundation for human behavior understanding. Our contributions can be summarized as

follows:

- We introduce LLaMo, which treats motion data as an independent modality within LLMs, preserving essential motion-specific details for robust analysis.
- We propose **Cross Talker**, a text-guided mechanism that dynamically focuses on and aggregates key motion frames features, optimizing computational efficiency and enabling the model to pick relevant motion features.
- LLaMo’s architecture supports the direct input of both raw motion and video data, providing a generalized framework for human-centric analysis across diverse applications, such as sports analytics, healthcare, and behavioral monitoring.

2. Related Work

2.1. Human-Centric Multimodal Representation

Multimodal representation learning is integral to human-centric analysis, especially for applications requiring spatial-temporal reasoning to decode complex behaviors [20–22]. Recent advances, such as Video-LLaVA, have made significant strides by embedding visual data from images and videos into a shared linguistic feature space, enabling sophisticated visual reasoning for behavioral analysis tasks [20, 23]. Despite these achievements, many current models are optimized primarily for static images or discrete video frames, limiting their effectiveness in sequential, dynamic scenarios where understanding progression and continuity in motion is crucial [9, 21, 24]. Privacy concerns further challenge the adoption of video-based methods in sensitive domains; thus, researchers are exploring motion data as a privacy-conscious alternative, enabling analysis focused on human actions without revealing identifiable visual information [25–27]. By combining visual and motion data, emerging multimodal frameworks hold promise for more comprehensive, privacy-aware human behavior analysis, leveraging the strengths of both modalities for adaptability and depth across applications.

2.2. Human Motion Understanding

Human motion understanding has traditionally relied on skeletal data represented as sequences of joint keypoints to capture movement while preserving user privacy [13, 28–30]. Early models like 2s-AGCN [31] and more recent transformer-based methods, such as MotionCLIP [2], have achieved success in tasks like activity recognition, captioning, and behavior analysis by mapping motion data to language tokens, effectively capturing structural aspects of movement. However, these methods often lack environmental context, which is crucial for nuanced interpretations in real-world applications, as similar movements can convey different meanings across varied scenarios [24, 32, 33]. To address this limitation, recent frameworks inte-

grate motion and visual data, allowing models to generalize across diverse, dynamic environments [34–36]. This fusion of skeletal data with video-provided context has shown promise in applications like sports and professional activity analysis, where contextual cues are essential. Additionally, skeletal-based data remains valuable in privacy-sensitive applications, such as healthcare and security, where it allows for robust analysis without compromising personal identity.

3. Methods

The LLaMo framework is composed of three primary modules to effectively process and integrate video, motion data and text. An overview of this pipeline is provided in Figure 2. The architecture comprises the following components:

1. **Multimodal Feature Extraction Module:** This module independently encodes the video or motion inputs, which makes LLaMo is a general human motion assistant with video input or motion input. Then the Motion features are enhanced by distilling valuable representations in video.
2. **Cross Talker:** Here, the enhanced motion features are aggregated in a text-guided manner and aligned with text within a shared semantic space, allowing the natural feature fusion between motion and text, which provides fine-grained inputs for LLM.
3. **Behavior Generation Module:** Using the aggregated and text-aligned motion features, along with the text representations, this module generates a contextually aware textual description of the observed human behavior.

3.1. Enhanced motion Feature Extraction Module

The enhanced motion feature extraction module consists of four primary components: the motion estimator, followed by the motion and video encoders, and the feature enhancement module. The motion estimator addresses scenarios where motion data may be unavailable by estimating motion information directly from video frames. The motion and video encoders then independently process the motion and video data, encoding them into their respective feature spaces. This approach allows the system to leverage complementary information from both modalities, creating a richer representation of human behavior.

Given a sequence of video frames $V = \{v_1, v_2, \dots, v_T\}$, the video encoder $f_v(\cdot)$ extracts a set of visual features:

$$F_V = f_v(V) = \{f_v(v_1), f_v(v_2), \dots, f_v(v_T)\}.$$

Simultaneously, the motion encoder $f_m(\cdot)$ processes the motion data $M = \{m_1, m_2, \dots, m_T\}$ to generate motion-specific features:

$$F_M = f_m(M) = \{f_m(m_1), f_m(m_2), \dots, f_m(m_T)\}.$$

Here, $F_V \in \mathbb{R}^{T \times H}$ and $F_M \in \mathbb{R}^{T \times H}$ represent the extracted feature sets for the video and motion data, respectively, where T is the sequence length and H denotes the feature dimension.

The feature enhancement module then aligns and fuses these modality-specific features. By using the motion features F_M as queries, this module selectively extracts semantically relevant information from the video features F_V . This alignment and fusion process allows the motion features to be enriched with contextual information from the video, producing refined, contextually enriched representations that enhance the model’s understanding of nuanced human behaviors. For the detailed design refer to the Appendix.

First, the video features and motion features will conduct self-attention operations respectively to get their augmented features F'_V and F'_M . This allows the video and motion to initially capture the dependencies on their respective time steps. Then, in order to enrich the motion data features with semantic information from the video, we subtly design a cross-attention block where augmented motion features as the query to extract valid semantic information in the video features. Finally, the selected video semantic representation and motion features are connected by residuals followed by a feed-forward neural network to obtain the enhanced motion features \tilde{F}_M .

3.2. Cross Talker Module

Given the enhanced motion features \tilde{F}_M from the feature enhancement module, directly feeding the entire sequence of motion frames into the Large Language Model (LLM) incurs significant computational costs and may lead to sub-optimal performance. This limitation arises from the self-attention mechanism in transformers, which has a computational complexity of $\mathcal{O}(L^2)$, where L denotes the sequence length. When incorporating both motion and text, the sequence length becomes $L = L_T + T$, where L_T is the number of text tokens and T is the number of motion frames. As T increases, the quadratic growth in computational cost becomes prohibitive, leading to complexity proportional to $(L_T + T)^2$.

Furthermore, computing attention over long sequences can dilute attention weights, making it challenging for the model to focus on essential parts of the sequence and effectively capture long-range dependencies. Furthermore, the motion features are not well aligned with text either. To address these issues, we propose the Cross Talker Module, which selectively distills key motion frames—referred to as *viewpoint frames*—based on their relevance to the text context, as well as aligning the motion features with the text. An overview of the Cross Talker Module is shown in Figure 3.

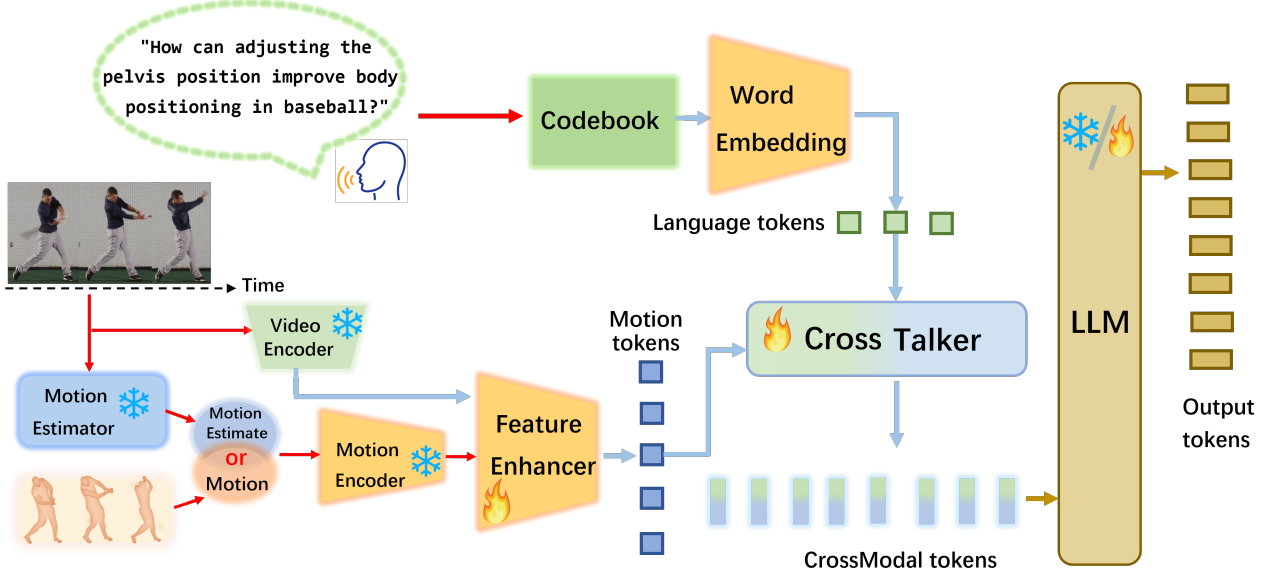


Figure 2. Overview of the LLaMo framework. It includes three main modules: (1) Multimodal Feature Extraction for encoding video and motion data; (2) Cross Talker for aligning and fusing motion and text features; and (3) Behavior Generation Module to produce text descriptions of human behavior based on integrated features.

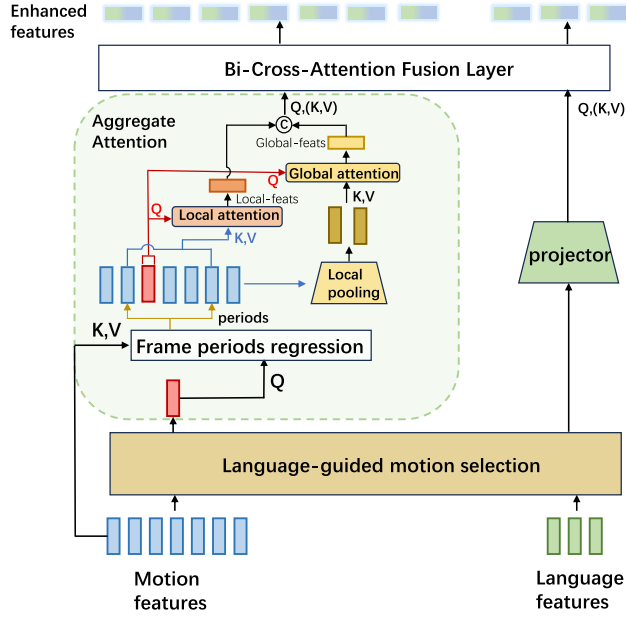


Figure 3. Overview of the Cross Talker Module, which selects key frames based on text guidance and fuses them with text features for enhanced analysis.

Language-Guided Frame Selection Given the enhanced motion features $\tilde{F}_M \in \mathbb{R}^{T \times H}$ and text embeddings $F_T \in \mathbb{R}^{L_T \times H}$, we aim to identify K viewpoint frames that are most relevant to the textual context. This selection is achieved through a cross-attention operation, where text

features act as queries and motion features as keys and values:

$$A = \text{Softmax} \left(\frac{F_T W_Q (\tilde{F}_M W_K)^T}{\sqrt{d}} \right), \quad (1)$$

where $W_Q, W_K \in \mathbb{R}^{H \times d}$ are learnable projection matrices, and d is the dimensionality of the queries and keys. The resulting attention matrix $A \in \mathbb{R}^{L_T \times T}$ contains attention weights, with each element $A_{i,j}$ indicating the relevance between the i -th text token and the j -th motion frame.

To determine the importance of each motion frame, we aggregate attention weights across all text tokens using max pooling:

$$s_j = \max_{i=1, \dots, L_T} A_{i,j}, \quad j = 1, \dots, T. \quad (2)$$

The resulting scores s_j represent the significance of each frame with respect to the text context. We then select the top K frames with the highest relevance scores as viewpoint frames, reducing the effective sequence length from T to K . This reduction in sequence length lowers the self-attention complexity from $\mathcal{O}((L_T + T)^2)$ to $\mathcal{O}((L_T + K)^2)$, which is computationally efficient when $K \ll T$.

Adaptive Contextual Feature Aggregation For each selected viewpoint frame, we enrich its representation by aggregating both local and global contexts. An adaptive receptive field size r_k is predicted for each viewpoint frame k using a receptive field regression module. The module performs a cross-attention operation where each viewpoint

frame serves as a query, and the unselected motion features serve as keys and values, followed by a sigmoid activation to estimate r_k . Based on r_k , we define a local window W_k around frame k :

$$W_k = \{j \mid |j - k| \leq r_k \times T\}. \quad (3)$$

Within this window, we apply local attention to capture fine-grained details:

$$F_{\text{local}}(k) = \tilde{F}_M(k) + \text{Attention}(\tilde{F}_M(k), \tilde{F}_M(W_k), \tilde{F}_M(W_k)), \quad (4)$$

where $\tilde{F}_M(W_k)$ denotes the motion features within the local window around frame k .

To incorporate global context, we partition the full motion sequence into N segments and compute segment-level features via average pooling:

$$F_M^{\text{seg}} = \{\bar{F}_M(n)\}_{n=1}^N, \quad (5)$$

where the segment size is controlled by the hyperparameter S_n . We then apply global attention:

$$F_{\text{global}}(k) = F_{\text{local}}(k) + \text{Attention}(F_{\text{local}}(k), F_M^{\text{seg}}, F_M^{\text{seg}}). \quad (6)$$

The final representation for each viewpoint frame is obtained by concatenating local and global features:

$$F_M(k) = [F_{\text{local}}(k); F_{\text{global}}(k)] \in \mathbb{R}^{2H}, \quad (7)$$

allowing the model to capture both fine-grained motion details and broader contextual information.

Bidirectional Cross-Modal Fusion To integrate the enhanced motion features with the text embeddings, we employ a bidirectional cross-attention mechanism. First, we update the motion features by performing a cross-attention operation, where the motion features act as queries and the text features serve as keys and values, and vice versa for the text features. The updated features are then passed through feed-forward networks, and we concatenate the motion and text representations to form the input for the next stage:

$$F_{\text{fusion}} = [F_T; \{F_M(k)\}_{k \in K}]. \quad (8)$$

This bidirectional fusion enhances the interaction between motion and text, enabling a coherent, contextually grounded understanding of human behavior. By selectively focusing on key motion frames and integrating enriched motion and text features, the Cross Talker Module reduces computational costs and improves the model's ability to capture complex human actions.

3.3. Behavior Generation Module

The Behavior Generation Module leverages the fused features F_{fusion} to generate a textual description Y that encapsulates the observed human behavior. This module utilizes a language model $h(\cdot)$ to transform F_{fusion} into a semantically coherent and contextually relevant output sequence:

$$Y = h(F_{\text{fusion}}) = \{y_1, y_2, \dots, y_L\}, \quad (9)$$

where y_t denotes the token generated at the t -th time step, and L is the length of the generated description. The probability of generating each token y_t is given by:

$$p(y_t \mid y_{1:t-1}, F_{\text{fusion}}) = \text{Softmax}(W_o h_t), \quad (10)$$

where h_t represents the hidden state of the language model at time t , and W_o is an output projection matrix. The hidden state h_t is derived from the model's internal mechanisms, which integrate attention over the input features and previously generated tokens. This autoregressive approach enables LLaMo to produce detailed and contextually aligned descriptions that incorporate nuanced information from both motion and text modalities.

By iteratively generating tokens based on the fused features and preceding context, the Behavior Generation Module constructs a meaningful and comprehensive narrative of human activity, grounded in the multimodal input data.

3.4. Training Objective

To train LLaMo, we finetune a pretrained language model using supervised learning, guiding it to generate accurate and contextually relevant descriptions based on the input motion, video, and text features. Both the motion and video encoders remain frozen during training, focusing the optimization on the language model's alignment with multimodal inputs. The training objective minimizes the discrepancy between the generated descriptions and the ground-truth annotations in the dataset.

Given a training set of N samples, each with video frames $V^{(i)}$, motion data $M^{(i)}$, and corresponding ground-truth textual descriptions $\hat{Y}^{(i)} = \{\hat{y}_1^{(i)}, \hat{y}_2^{(i)}, \dots, \hat{y}_L^{(i)}\}$, the overall training objective is to minimize the negative log-likelihood of the ground-truth tokens across all samples:

$$L = -\frac{1}{N} \sum_{i=1}^N \sum_{t=1}^{L^{(i)}} \log p(\hat{y}_t^{(i)} \mid \hat{y}_{1:t-1}^{(i)}, F_{\text{fusion}}^{(i)}), \quad (11)$$

where $F_{\text{fusion}}^{(i)}$ represents the fused input features for the i -th sample. The term $p(\hat{y}_t^{(i)} \mid \hat{y}_{1:t-1}^{(i)}, F_{\text{fusion}}^{(i)})$ denotes the probability of the ground-truth token $\hat{y}_t^{(i)}$ given the previous tokens $\hat{y}_{1:t-1}^{(i)}$ and the fused features.

This objective ensures that the generated descriptions are syntactically and semantically aligned with the target

outputs, capturing the detailed aspects of human behavior present in the video and motion inputs.

4. Experiments

To evaluate LLaMo, we performed experiments across datasets representing complex human behaviors and tasks, including professional sports and cyclic action repetition counts.

4.1. Implementation Details

Training Datasets: LLaMo was trained using a combination of bimodal video-motion datasets and unimodal video or motion datasets to enable a robust understanding of both modalities. For video-motion data, we utilized MoVid from MotionLLM [1] and our custom Swing dataset containing 20,000 annotated videos of baseball and golf swings with corresponding motion capture data and expert Q&A instructions. Motion-only datasets included HumanML3D [37] and KIT-ML [3], while Mo-RepCount, a filtered subset of RepCount [38], provided high-quality repetition count videos. This multimodal data configuration ensures that LLaMo effectively learns from both single and combined modalities.

Evaluation Datasets: LLaMo was evaluated on MoVid-Bench [1] to assess both video and motion understanding capabilities, with additional tests on BABEL-QA [19] to evaluate motion-based question answering. For specific tasks, we used Mo-RepCount to gauge motion details and Swing to evaluate professional sports guidance.

Evaluation Metrics: For MoVid-Bench, we follow previous LLM evaluation metrics [39] [22] [40] on accuracy and scores, which contains body-part awareness, sequentially, direction analysis, reasoning ability, and hallucination, respectively. BABEL-QA used the metric followed by [19], Mo-RepCount applied metrics such as OBO, MAE, OBZ, and RMSE, and Swing performance was assessed on reasonableness, coherence, pertinence, and adaptability using GPT-4 evaluations, scoring from 0 to 5.

4.2. Results

Evaluation on Motion Understanding in MoVid-Bench. We evaluated LLaMo’s motion understanding capabilities on the MoVid-Bench, focusing on five critical aspects. Both accuracy and score metrics were used to comprehensively assess LLaMo’s performance. Table 1 provides a comparison with several baselines, including GPT-3.5 [41], MotionGPT [2], and MotionLLM [1]. The GPT-3.5 baseline, limited to textual data processing, struggles significantly with motion interpretation, yielding low performance. MotionGPT, while capable of incorporating motion data, is pri-

marily code motion into a codebook, which limits its reasoning ability on coiled motion data and leaves it vulnerable to hallucinations. Although MotionLLM provides improvements in understanding motion sequences, it still conducts translation on motion data, resulting in low accuracy on seq. and Hall.

LLaMo, in contrast, achieves superior results across nearly all metrics on MoVid-Bench-Motion, attaining the highest accuracy in most categories. This performance is primarily attributed to its preservation of critical motion-specific nuances, as LLaMo processes motion data as a distinct modality without requiring any conversion. LLaMo effectively captures the intricate details of 3D sequences associated with human behaviors. Its Feature Enhancer, combined with the Cross Talker module, enables simultaneous input of video and motion data, allowing LLaMo to selectively focus on motion frames most relevant to the text context. This design enhances the model’s ability to capture complex patterns of human motion and improves its reasoning on motion-intensive tasks. Overall, LLaMo’s results on MoVid-Bench demonstrate its effectiveness in advancing human behavior analysis, highlighting its potential for broader applications in general motion understanding.

Evaluation on Video Understanding in MoVid-Bench.

We evaluated LLaMo on the video component of the MoVid-Bench dataset to assess its capability for comprehensive video understanding. As presented in Table 1, LLaMo achieves strong performance across most metrics, demonstrating LLaMo’s extraordinary comprehension ability on videos even without motion input.

Compared to other models like GPT-3.5 [41], VideoLLAVA [39], and MotionLLM [1], while LLaMo slightly trails behind MotionLLM in body-part awareness, LLaMo shows significant improvements in reasoning ability and overall scores. This indicates that our model excels in understanding complex video content and making accurate inferences about human behaviors.

The rewarding results can be attributed to our model’s powerful ability to estimate motion data accurately in video and align it seamlessly with video inputs. This approach retains the motion capture of motion data and the auxiliary information provided by video, enriching the understanding of dynamic scenes.

Evaluation on BABEL-QA. We evaluated LLaMo on the BABEL-QA dataset to assess its capability to handle complex motion-based queries. As shown in Table 2, LLaMo achieves the highest overall score of 0.458, outperforming all baselines. It excels particularly in the *Action* category with a score of 0.525, demonstrating superior action recognition capabilities. Additionally, it achieves top scores in temporal reasoning tasks (*Before* and *After*), indicating a

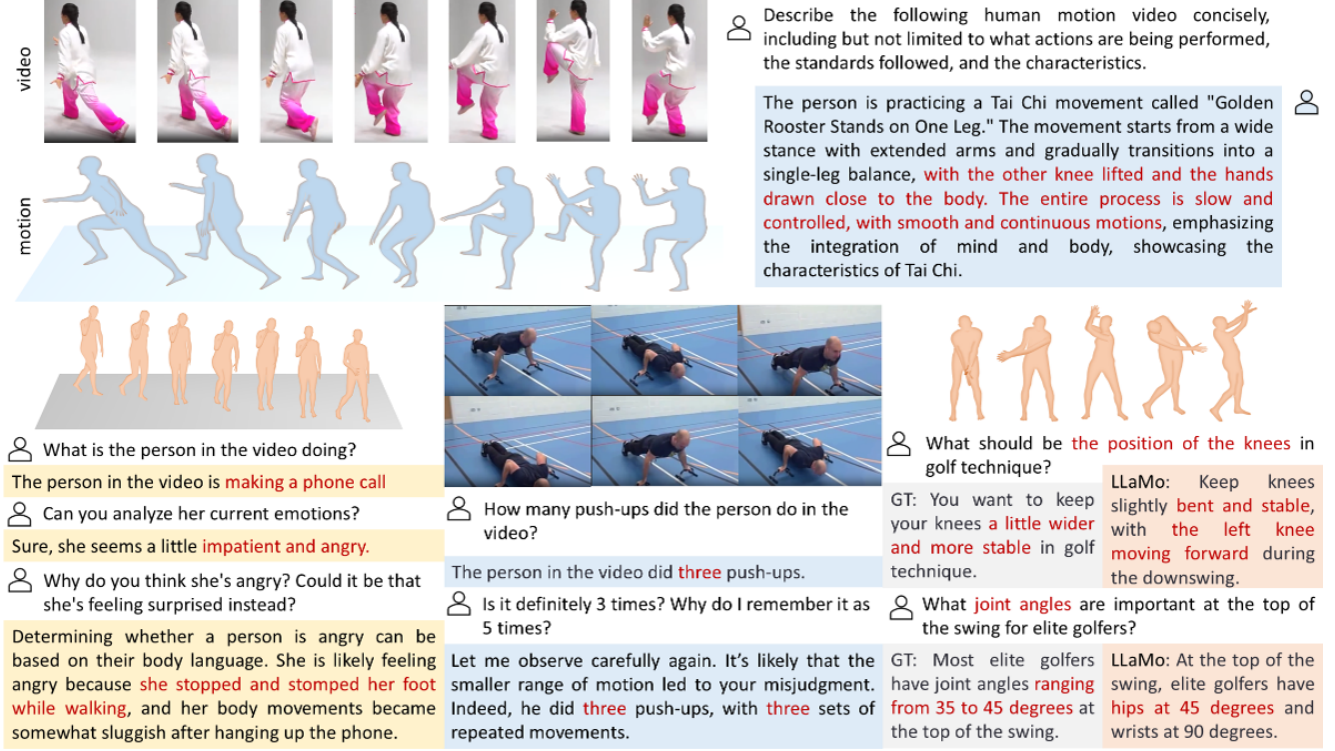


Figure 4. Example outputs from LLaMo across human activities and professional sports, showcasing its reasoning capabilities and domain-specific knowledge in motion-intensive scenarios.

Table 1. Expected Comparison on the MoVid-Bench. The top part of the table presents motion-related results, and the bottom part presents video-related results. Higher accuracy and score values indicate better performance.

| MoVid-Bench-Motion | Body. | | Seq. | | Dir. | | Rea. | | Hall. | | All | |
|--------------------|--------------|-------------|--------------|-------------|--------------|-------------|--------------|-------------|--------------|-------------|--------------|-------------|
| | Acc. | Score | Acc. | Score | Acc. | Score | Acc. | Score | Acc. | Score | Acc. | Score |
| GT | 100.00 | 5.00 | 100.00 | 5.00 | 100.00 | 5.00 | 100.00 | 5.00 | 100.00 | 5.00 | 100.00 | 5.00 |
| GPT-3.5 [41] | 24.51 | 2.04 | 30.41 | 2.25 | 27.14 | 2.19 | 39.19 | 2.64 | 58.33 | 3.22 | 31.33 | 2.31 |
| MotionGPT [2] | 31.22 | 3.98 | 42.69 | 3.16 | 44.29 | 3.50 | 35.81 | 3.06 | 16.66 | 2.25 | 36.86 | 3.11 |
| MotionLLM [1] | 50.49 | 3.55 | 36.84 | 3.14 | 58.57 | 3.76 | 52.70 | 3.58 | 55.56 | 3.39 | 49.50 | 3.49 |
| LLaMo | 59.30 | 4.01 | 44.01 | 3.12 | 60.91 | 3.99 | 58.21 | 3.64 | 61.17 | 3.53 | 55.32 | 3.67 |
| MoVid-Bench-Video | Body. | | Seq. | | Dir. | | Rea. | | Hull. | | All | |
| | Acc. | Score | Acc. | Score | Acc. | Score | Acc. | Score | Acc. | Score | Acc. | Score |
| GT | 100.00 | 5.00 | 100.00 | 5.00 | 100.00 | 5.00 | 100.00 | 5.00 | 100.00 | 5.00 | 100.00 | 5.00 |
| GPT-3.5 [41] | 2.40 | 1.23 | 1.39 | 1.00 | 4.65 | 1.09 | 5.41 | 1.65 | 0.00 | 0.94 | 3.03 | 1.26 |
| Video-LLAVA [39] | 33.53 | 2.76 | 25.46 | 2.72 | 41.86 | 2.84 | 52.97 | 3.28 | 58.83 | 1.89 | 42.53 | 2.70 |
| MotionLLM [1] | 34.13 | 2.93 | 32.87 | 2.92 | 44.18 | 3.14 | 63.20 | 3.55 | 70.59 | 2.30 | 49.00 | 2.97 |
| LLaMo | 33.83 | 2.85 | 36.01 | 3.11 | 45.50 | 3.32 | 67.59 | 3.73 | 72.81 | 2.25 | 52.33 | 3.10 |

strong understanding of temporal relationships in motion sequences. This capability benefits from the direct integration of motion data as an independent modality and the effective text-aware motion frames extract mechanism in LLaMo, preserving crucial motion-specific nuances, as well as leading to more accurate and contextually rich interpretations.

Professional Sports Analysis. We further evaluate LLaMo on the challenging swing-dataset, termed as

profession-swing, details refer to Appendix, where the ground truth (GT) consists of answers provided by professional coaches, serving as expert benchmarks. The evaluation focuses on four key indicators: **Reasonableness**: Logical soundness and plausibility of the responses. **Coherence**: Consistency and logical flow within the responses. **Pertinence**: Relevance of the answers to the specific questions. **Adaptability**: Ability to adjust responses based on the stage on the athlete.

As presented in Table 3, LLaMo significantly outper-

Table 2. Comparison on BABEL-QA dataset. Higher scores indicate better performance.

| Model | Pred. type | Overall ↑ | Action ↑ | Direction ↑ | Body Part ↑ | Before ↑ | After ↑ | Other ↑ |
|-------------------|------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| MotionCLIP-M [41] | cls. | 0.430 | 0.485 | 0.361 | 0.272 | 0.372 | 0.321 | 0.404 |
| MotionCLIP-R [41] | cls. | 0.420 | 0.489 | 0.310 | 0.250 | 0.398 | 0.314 | 0.387 |
| MotionLLM [41] | gen. | 0.436 | 0.517 | 0.354 | 0.154 | 0.427 | 0.368 | 0.529 |
| LLaMo | gen. | 0.458 | 0.525 | 0.398 | 0.224 | 0.443 | 0.392 | 0.518 |

Table 3. Performance on the profession-swing dataset across four indicators. Higher accuracy and score values indicate better performance.

| Model | Reasonableness | | Coherence | | Pertinence | | Adaptability | | All | |
|---------------|----------------|-------------|--------------|-------------|--------------|-------------|--------------|-------------|--------------|-------------|
| | Acc | Score | Acc | Score | Acc | Score | Acc | Score | Acc | Score |
| GT | 100 | 5 | 100 | 5 | 100 | 5 | 100 | 5 | 100 | 5 |
| GPT-3.5 [41] | 4.51 | 0.83 | 4.20 | 0.74 | 3.65 | 0.66 | 1.71 | 0.22 | 3.42 | 0.61 |
| MotionGPT [2] | 10.53 | 1.25 | 15.42 | 1.66 | 12.64 | 1.1 | 14.79 | 1.35 | 14.35 | 1.4 |
| MotionLLM [1] | 12.33 | 1.51 | 19.20 | 1.87 | 17.98 | 1.79 | 10.20 | 1.22 | 16.53 | 1.57 |
| LLaMo | 21.10 | 2.11 | 27.10 | 2.71 | 31.81 | 3.12 | 20.22 | 1.98 | 24.80 | 2.48 |

forms both GPT-3.5 [41] and MotionLLM [1] across all metrics. While MotionLLM shows improvements over GPT-3.5, with higher accuracy and scores, it still lags behind LLaMo. Specifically, LLaMo achieves an overall accuracy of 24.80

The superior performance of LLaMo can be attributed to the design of the Cross Talker module, which enables language-guided frame selection to capture complex motion relationships inherent in professional swings. This allows LLaMo to effectively model intricate motion patterns and generate responses that are more reasonable, coherent, pertinent, and adaptable, closely aligning with the expert-level insights provided by professional coaches.

In contrast, while MotionLLM demonstrates certain strengths over MotionGPT—such as better pertinence and coherence. It lacks the advanced motion understanding capabilities of LLaMo, lagging behind LLaMo distinctly.

Table 4. Motion and video details capture evaluation on Mo-RepCount

| Model | OBO | MAE | OBZ | RMSE |
|---------------|--------------|--------------|--------------|-------------|
| EScounts [42] | 0.397 | 0.291 | 0.198 | 5.58 |
| PoseRAC [43] | 0.382 | 0.312 | 0.204 | 5.95 |
| TransRAC [38] | 0.276 | 0.444 | 0.105 | 8.56 |
| RepNet [44] | 0.009 | / | / | / |
| LLaMo | 0.389 | 0.324 | 0.222 | 6.15 |

Evaluation on Mo-RepCount. We evaluated LLaMo on the Mo-RepCount dataset to assess its ability to capture fine-grained motion details and complex temporal features required for accurate repetition counting. For a fair comparison, we trained all state-of-the-art models, including RepNet [44], TransRAC [38], PoseRAC [43] and ES-

counts [42], on the Mo-RepCount dataset without using any additional data. As presented in Table 4, LLaMo outperforms the RepNet and TransRAC across almost all metrics. Since both of them are models built for loop counting tasks, the effect of our model on this task is convincing, which shows that LLaMo is also competent in capturing the details of the motion and videos.

While LLaMo did not outperform the specialized counting models EScounts and poseRAC on many metrics, it excelled on the OBZ metric, demonstrating an unmatched ability to capture fine-grained motion details and spatio-temporal relationships. These results confirm LLaMo’s strong potential in motion and video analysis, underscoring its suitability for detail-oriented tasks and applications.

5. Conclusion

We introduce LLaMo, a novel framework that integrates motion data directly into large language models, thereby enhancing multimodal understanding. By unifying text, video, and motion data in an architecture fine-tuned for human motion instruction, LLaMo achieves notable progress in behavior comprehension, motion-focused captioning, and interactive Q&A, capturing the fine-grained motion details vital for precise, context-aware insights. Experimental results, especially in professional sports analysis, demonstrate LLaMo’s capacity to interpret complex human motion with accuracy. Future work will extend LLaMo to a broad range of domains where detailed motion analysis and multimodal understanding are essential. The main challenges lie in refining cross-modal integration techniques and improving efficiency for real-time applications, thus broadening LLaMo’s potential in human-centered AI. We hope our work will spur further exploration into more advanced human-centric multimodal models.

Acknowledgement

This work was supported in part by the Pioneer Centre for AI, DNRF grant number P1.

References

- [1] Ling-Hao Chen, Shunlin Lu, Ailing Zeng, Hao Zhang, Benyou Wang, Ruimao Zhang, and Lei Zhang. Motionllm: Understanding human behaviors from human motions and videos. *arXiv preprint arXiv:2405.20340*, 2024. 1, 2, 6, 7, 8
- [2] Ling-Hao Chen, Jiawei Zhang, Wen Liu, Gang Yu, and Tao Chen. Motiongpt: Human motion as a foreign language. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2024. 1, 2, 6, 7, 8
- [3] Matthias Plappert, Christian Mandery, and Tamim Asfour. The kit motion-language dataset. *Big data*, 4(4):236–252, 2016. 1, 6
- [4] Yan Zhang, Michael J Black, and Siyu Tang. We are more than our joints: Predicting how 3d bodies move. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [5] Fangzhou Hong, Liang Pan, Zhongang Cai, and Ziwei Liu. Versatile multi-modal pre-training for human-centric perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16156–16166, 2022. 2
- [6] Haoxuan Qu, Yujun Cai, and Jun Liu. Llm are good action recognizers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18395–18406, 2024. 1
- [7] Shan Chen, Jiale Zhou, and Lei Li. Dense point clouds matter: Dust-gs for scene reconstruction from sparse viewpoints. *arXiv preprint arXiv:2409.08613*, 2024. 1
- [8] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 1, 2
- [9] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [10] Li Lei, Jia Sen, Wang Jianhao, An Zhaochong, Li Jiaang, Hwang Jenq-Neng, and Belongie Serge. Chatmotion: A multimodal multi-agent for human motion analysis. *arXiv preprint arXiv:2502.18180*, 2025. 1
- [11] Mark Endo, Joy Hsu, Jiaman Li, and Jiajun Wu. Motion question answering via modular motion programs. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2023. 1, 2
- [12] Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. Motiongpt: Human motion as a foreign language. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 1
- [13] Yaya Shi, Haiyang Xu, Chunfeng Yuan, Bing Li, Weiming Hu, and Zheng-Jun Zha. Learning video-text aligned representations for video captioning. *IEEE Transactions on Multimedia (TMM)*, 2023. 2
- [14] Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 580–597, 2022. 2
- [15] Kunchang Li, Yali Wang, Yinan He, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [16] Zhaochong An, Guolei Sun, Yun Liu, Runjia Li, Min Wu, Ming-Ming Cheng, Ender Konukoglu, and Serge Belongie. Multimodality helps few-shot 3d point cloud semantic segmentation. *arXiv preprint arXiv:2410.22489*, 2024. 2
- [17] Bobby Bodenheimer, Chuck Rose, Seth Rosenthal, and John Pella. The process of motion capture: Dealing with the data. In *Proceedings of the Eurographics Workshop on Computer Animation and Simulation (EG Workshop)*, pages 3–18, 1997. 2
- [18] Xiangyu Peng, Yuxuan Zhang, Xiaolong Wang, and Ziwei Liu. Text-driven human motion generation with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [19] Mark Endo, Joy Hsu, Jiaman Li, and Jiajun Wu. Motion question answering via modular motion programs. In *International Conference on Machine Learning*, pages 9312–9328. PMLR, 2023. 2, 6
- [20] Jie Lin, Wei Zhang, and Zhe Chen. Videollm: Language models for video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2405–2415, 2023. 2
- [21] Xu Ning, Hongwei Li, and Yu Zhao. Videobench: A benchmark for large-scale video understanding. *arXiv preprint arXiv:2304.12345*, 2023. 2
- [22] KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023. 2, 6
- [23] Lei Li. Segment any building. In *Computer Graphics International Conference*, pages 155–166. Springer, 2023. 2
- [24] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*, 2023. 2
- [25] Enxin Song, Guanhong Zhang, and Haoyang Zhou. Adaptive multimodal learning for behavior analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [26] Yunhua Yang, Ziwang Zhao, and Yiming Xie. Recognizing human behaviors with skeletal data in llm-based frameworks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [27] Lei Li, Tianfang Zhang, Zhongfeng Kang, and Xikun Jiang. Mask-fpan: Semi-supervised face parsing in the wild with de-occlusion and uv gan. *Computers & Graphics*, 116:185–193, 2023. 2

- [28] Matthias Plappert, Christian Mandery, and Tamim Asfour. Learning a bidirectional mapping between human whole-body motion and natural language using deep recurrent neural networks. *Robotics and Autonomous Systems (RAS)*, 109: 13–26, 2018. [2](#)
- [29] Lei Li. Cpsseg: Finer-grained image semantic segmentation via chain-of-thought language prompting. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 513–522, 2024.
- [30] Yunhua Yang, Liang Zhang, and Hui Li. Understanding human behaviors from skeletal data: A review of datasets and methods. *arXiv preprint arXiv:2310.12998*, 2023. [2](#)
- [31] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12026–12035, 2019. [2](#)
- [32] Enxin Song, Wenhao Chai, and Yucheng Zhang. Fine-grained spatial-temporal motion understanding in complex video environments. *arXiv preprint arXiv:2310.08639*, 2023. [2](#)
- [33] Zhaochong An, Guolei Sun, Zongwei Wu, Hao Tang, and Luc Van Gool. Temporal-aware hierarchical mask classification for video semantic segmentation. *arXiv preprint arXiv:2309.08020*, 2023. [2](#)
- [34] Xianghong Liu, Haoxuan Wang, Zhiwei Zhang, Huan Wu, Baoquan Chen, and Xiaoguang Han. Category-agnostic pose estimation for point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [3](#)
- [35] Xiaolong He, Yi Zhang, Chen Chen, and Kyoung Mu Lee. Activitynet++: A large-scale benchmark for video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [36] Lei Li. Hierarchical edge aware learning for 3d point cloud. In *Computer Graphics International Conference*, pages 81–92. Springer, 2023. [3](#)
- [37] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5152–5161, 2022. [6](#)
- [38] 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. [6](#), [8](#)
- [39] Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning unified visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*, 2023. [6](#), [7](#)
- [40] Peng Jin, Ryuichi Takanobu, Wancai Zhang, Xiaochun Cao, and Li Yuan. Chat-univi: Unified visual representation empowers large language models with image and video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13700–13710, 2024. [6](#)
- [41] OpenAI. Gpt-3.5: Generative pre-trained transformer 3.5, 2023. <https://platform.openai.com/docs/models/gpt-3-5>. [6](#), [7](#), [8](#)
- [42] Saptarshi Sinha, Alexandros Stergiou, and Dima Damen. Every shot counts: Using exemplars for repetition counting in videos. *arXiv preprint arXiv:2403.18074*, 2024. [8](#)
- [43] Ziyu Yao, Xuxin Cheng, and Yuexian Zou. Poserac: Pose saliency transformer for repetitive action counting. *arXiv preprint arXiv:2303.08450*, 2023. [8](#)
- [44] Bastian Wandt and Bodo Rosenhahn. Repnet: Weakly supervised training of an adversarial reprojection network for 3d human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7782–7791, 2019. [8](#)