

# Understanding Emotional Body Expressions via Large Language Models

Haifeng Lu<sup>1,2</sup>, Jiuyi Chen<sup>3,4</sup>, Feng Liang<sup>1</sup>, Mingkui Tan<sup>3</sup>, Runhao Zeng<sup>1,5\*</sup>, Xiping Hu<sup>1,6\*</sup>

<sup>1</sup>Guangdong-Hong Kong-Macao Joint Laboratory for Emotional Intelligence and Pervasive Computing, Shenzhen MSU-BIT University,

<sup>2</sup>Lanzhou University,

<sup>3</sup>South China University of Technology,

<sup>4</sup>Peng Cheng Laboratory,

<sup>5</sup>Shenzhen University,

<sup>6</sup>Beijing Institute of Technology

luhf18@lzu.edu.cn, zengrh@smbu.edu.cn, huxp@smbu.edu.cn

## Abstract

Emotion recognition based on body movements is vital in human-computer interaction. However, existing emotion recognition methods predominantly focus on enhancing classification accuracy, often neglecting the provision of textual explanations to justify their classifications. In this paper, we propose an **Emotion-Action Interpreter** powered by **Large Language Model (EAI-LLM)**, which not only recognizes emotions but also generates textual explanations by treating 3D body movement data as unique input tokens within large language models (LLMs). Specifically, we propose a multi-granularity skeleton tokenizer designed for LLMs, which separately extracts spatio-temporal tokens and semantic tokens from the skeleton data. This approach allows LLMs to generate more nuanced classification descriptions while maintaining robust classification performance. Furthermore, we treat the skeleton sequence as a specific language and propose a unified skeleton token module. This module leverages the extensive background knowledge and language processing capabilities of LLMs to address the challenges of joint training on heterogeneous datasets, thereby significantly enhancing recognition accuracy on individual datasets. Experimental results demonstrate that our model achieves recognition accuracy comparable to existing methods. More importantly, with the support of background knowledge from LLMs, our model can generate detailed emotion descriptions based on classification results, even when trained on a limited amount of labeled skeleton data.

## 1 Introduction

In the field of human-computer interaction, a machine’s ability to understand human emotions directly impacts the user’s interaction experience (Fragopanagos and Taylor 2005; Mandryk, Atkins, and Inkpen 2006; Narayanan et al. 2020). Traditionally, the most common methods for emotion recognition involve analyzing facial expressions and vocal tones (Li and Deng 2022; El Ayadi, Kamel, and Karray 2011). However, accurately recognizing a user’s emotions becomes particularly challenging when the user is distant from the camera or when a microphone is unavailable.

Numerous studies have shown that full-body motion is a significant means of expressing emotion (Noroozi et al. 2018; Wang et al. 2024). Additionally, the large surface area of the human torso facilitates data collection from a distance, laying the foundation for more extensive emotion recognition scenarios (Wang et al. 2024). With recent advancements in motion capture technology and human pose estimation algorithms (Peng, Zheng, and Chen 2024; Zheng et al. 2023; Zhang 2012), the acquisition of 2D and 3D skeleton data has become more convenient. Notably, 3D skeleton data offer the advantage of being resistant to variations in viewpoint, lighting, and background clutter (Lu, Hu, and Hu 2023). As a result, emotion recognition based on 3D full-body skeleton data has garnered significant attention in recent years.

Most existing skeleton-based emotion recognition methods rely on handcrafted features (Fourati, Pelachaud, and Darmon 2019; Daoudi et al. 2017; Oğuz and Ertuğrul 2024). In addition, deep learning approaches such as those in (Beyan et al. 2023), ST-Gait++ (Lima et al. 2024), and BPM-GCN (Zhai et al. 2024) are also significant for learning emotion representations. However, these methods share some common limitations: (i) Traditional emotion recognition techniques primarily focus on improving classification accuracy and often lack the capability to provide textual explanations that support their classifications. (ii) Due to variations in data collection methods, 3D skeleton datasets are often heterogeneous, with differences in the number of joints and frame lengths. This heterogeneity impedes effective knowledge transfer between datasets and restricts the potential for improving classification accuracy within individual datasets.

**Motivation.** Inspired by the remarkable capabilities of large language models (LLMs) across various domains (Li et al. 2024; Qu, Cai, and Liu 2024), we are intrigued by the potential of utilizing LLMs as emotion recognizers that can not only classify emotions from skeleton data but also generate corresponding classification criteria. Several studies have shown that LLMs hold some features that are useful for emotion recognition. Specifically, these models are pre-trained on vast corpora that include descriptions of emotional actions (Qiu et al. 2024; Li et al. 2023a). Therefore, converting 3D full-body skeleton data into tokens recognizable by

\*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

LLMs, and utilizing the extensive background knowledge of LLMs to classify emotions and generate detailed explanations, would be highly valuable.

To address these challenges, we propose an **Emotion-Action Interpreter** powered by **Large Language Model (EAI-LLM)**, which is capable of simultaneously recognizing emotions from skeleton data and generating detailed emotion descriptions. First, we introduce a **Multi-Granularity Skeleton Tokenizer (MGST)** to enhance token diversity, enabling LLMs to produce more fine-grained emotion descriptions. To further improve performance, we implement a **Unified Skeleton Token (UST)** module that incorporates a spatio-temporal masking mechanism. This module not only addresses the challenges posed by heterogeneous datasets but also enhances recognition accuracy on individual datasets. Additionally, we pre-train the EAI-LLM using a combination of skeleton-language data and fine-tune it on prompt-based question-and-answer tasks. To convert skeleton features into tokens recognizable by LLMs, we also introduced a skeleton-text contrastive learning framework. This approach aligns the skeleton feature space with the semantic space using contrastive loss. Experimental results demonstrate that the recognition capability of EAI-LLM is comparable to existing methods, and it can generate detailed emotion descriptions from classification results by leveraging the background knowledge of LLMs.

In summary, our main contributions include:

- We propose a novel 3D full-body skeleton-based emotion recognition model called EAI-LLM. To the best of our knowledge, EAI-LLM is the first approach to utilize LLMs as emotion recognizers while generating detailed emotion descriptions directly from skeleton data.
- We propose a Multi-Granularity Skeleton Tokenizer (MGST), which can enhance token diversity, with the goal of improving the model’s ability to infer emotion expressed through 3D skeleton movement sequences.
- We propose an Unified Skeleton Token (UST) module designed for joint training in heterogeneous datasets. It treats skeleton sequences as a specialized language, allowing heterogeneous datasets to be unified within a shared language space for training.

## 2 Related Work

### 2.1 Body Skeleton-based Emotion Recognition

Many studies have focused on using body movements, postures, and gestures to recognize emotion. Traditionally, emotion recognition from body movements has relied on handcrafted features (Piana et al. 2016; Daoudi et al. 2017; Fourati, Pelachaud, and Darmon 2019). For instance, (Glowinski et al. 2008) examined upper body motions, utilizing features such as the number of local maxima and the ratio between the maximum and the duration of the largest peaks to identify emotions. Similarly, (Fourati and Pelachaud 2015; Fourati, Pelachaud, and Darmon 2019) proposed a full-body motion encoding scheme that extracted over 110 motion features across three levels, anatomy, spatial direction, and posture/movement, to describe expres-

sive movements. Additionally, (Oğuz and Ertuğrul 2024) explored commonly used time, frequency, and statistical-based parameters, employing feature selection methods to select four features from each frame to construct a feature matrix for emotion recognition.

Recently, due to the impressive performance of neural networks (Zeng et al. 2024; You et al. 2024), deep learning models have dominated the field of emotion recognition from 3D skeleton movement sequences. (Ghaleb et al. 2021) proposed a classifier network based on Spatial-Temporal Graph Convolutional Networks (Yan, Xiong, and Lin 2018) to recognize emotions from body movements. In (Zhang et al. 2021), the authors introduced an attention-based stacked LSTM network to detect the relationship between emotions and movements. (Beyan et al. 2023) encoded various time intervals of 3D positional data into RGB images and then used Convolutional Neural Networks (CNNs) for classification. Similarly, (Wang et al. 2024) combined handcrafted features with CNN-learned image features and employed a linear classifier for emotion prediction. Unlike these previous models, which focus solely on classification, our method not only accurately identifies emotion from skeleton data but also provides interpretable explanations, significantly broadening the model’s application scenarios.

### 2.2 Multi-modal Large Language Models

Recently, LLMs have demonstrated impressive capabilities in generation and comprehension. Early models such as BERT (Devlin et al. 2018), GPT-2 (Radford et al. 2019), and T5 (Raffel et al. 2020), which were trained on web-scale text data, laid the groundwork for the popularity of large models. Subsequently, models with greater capacity, more parameters, and extensive training data, such as GPT-3 (Brown et al. 2020), Alpaca (Taori et al. 2023), PaLM (Chowdhery et al. 2023), Vicuna (Chiang et al. 2023), and LLaMA (Touvron et al. 2023), have been developed. Recent advancements, such as InstructGPT (Ouyang et al. 2022), have focused on aligning LLMs with human instructions and feedback. ChatGPT (OpenAI 2024), Anthropic (Anthropic 2024), and GPT-4 (OpenAI 2023) now interact with users and answer a broad range of diverse and complex questions.

Some studies have extended these models to other modalities. Flamingo (Alayrac et al. 2022) and BLIP-2 (Li et al. 2023b) aligned pre-trained vision encoders with language models using a cross-attention mechanism, establishing a foundation for subsequent research in multi-modal LLMs. More recently, the capabilities of LLMs to handle various tasks have been further explored. GPT-4 (OpenAI 2023) has demonstrated powerful visual understanding and reasoning abilities, while MiniGPT-4 (Zhu et al. 2023) and LLaVA (Liu et al. 2023) produced detailed and accurate image descriptions through a series of visual instruction-tuning datasets. VideoChat (Li et al. 2023c) and VideoChatGPT (Maaz et al. 2023) utilized ChatGPT to generate video instruction-tuning data, extending support to video modalities. Despite these advancements, there remains a gap in multi-modal LLMs capable of processing skeleton data. In this work, we propose a new framework that enables LLMs to handle skeleton-based body movement data.

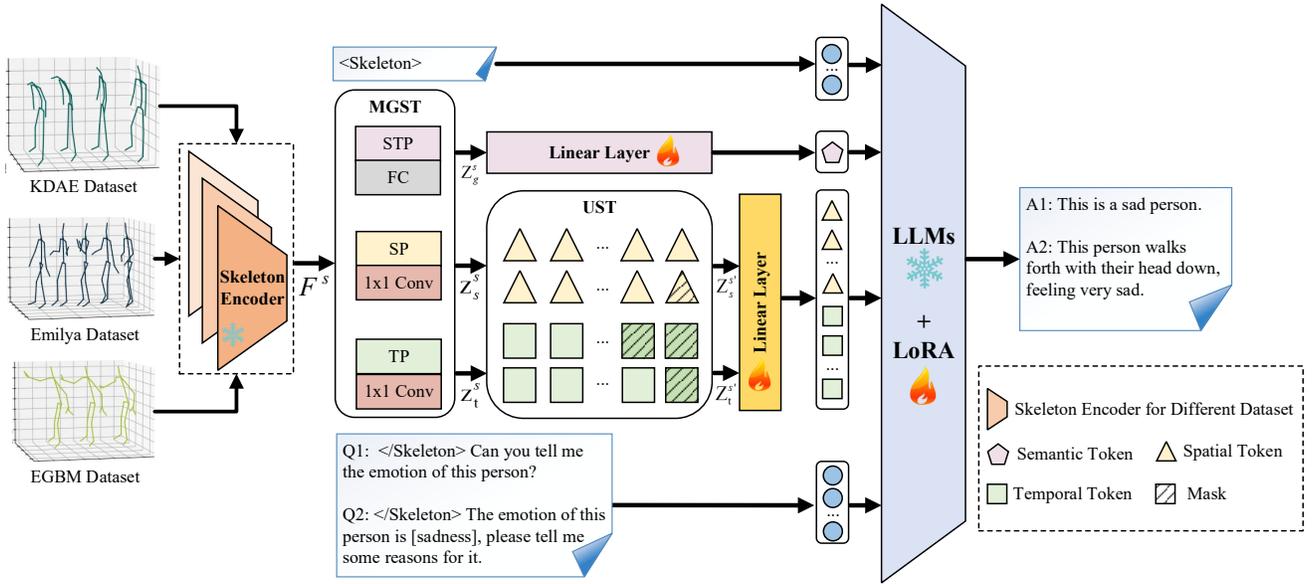


Figure 1: This work presents a novel approach for 3D full-body skeleton-based emotion recognition using fine-tuned LLMs, termed EAI-LLM. Unlike previous methods, EAI-LLM not only identifies emotions but also generates textual explanations by treating 3D body movement data as unique input tokens within the LLMs.

### 3 Method

**Problem definition.** Let  $X = \{X_t \in \mathbb{R}^{T*J*3}\}_{t=1}^T$  be a 3D skeleton sequence, where  $X_t$  denotes the frame at time  $t$  with  $J$  body joint. Our objective is to input a 3D skeleton movement sequence  $X$  into the EAI-LLM, which automatically encodes  $X$  into semantic space, outputs corresponding emotion labels for the body movement data and provides detailed emotion descriptions.

**Overview.** We propose an emotion recognition framework called EAI-LLM, as shown in Fig. 1. First, we introduce a multi-granularity skeleton tokenizer that extracts spatio-temporal tokens and semantic tokens distinctly. Second, we employ a masking mechanism to normalize tokens of varying lengths to a unified length, thereby enabling joint training. Third, we use a linear projection layer to bridge the gap between the skeleton encoder and LLMs, embedding the extracted multi-granularity skeleton tokens together with text prompts as input for LLMs. Finally, we pre-train EAI-LLM using a composite of skeleton-language data and fine-tune it on prompt-based question-and-answer tasks.

#### 3.1 Multi-Granularity Skeleton Tokenizer

Traditional skeleton-based emotion recognition methods typically use a skeleton encoder to extract features, denoted as  $F^s \in \mathbb{R}^{T*J*C}$ , where  $C$  represents the base channel. These features are then spatio-temporally compressed and passed through a fully connected (FC) layer to obtain classification results. However, the fundamental design of LLMs is to process text-based inputs, making it impossible to directly input skeleton sequence  $X$  into LLMs. To bridge this gap, we develop a skeleton tokenizer to convert skeleton data into a format compatible with the input structure of LLMs.

In this paper, we propose a multi-granularity skeleton tokenizer to enhance token diversity, thereby facilitating the generation of more detailed text descriptions. First, we perform spatio-temporal pooling (STP) on the feature  $F^s$ , followed by processing through an FC layer to obtain a semantic token  $z_g^s \in \mathbb{R}^C$ . This token encapsulates the global body motion encoding information. While the semantic token is advantageous for classification tasks, the spatio-temporal pooling process can result in information loss, which is detrimental in generation tasks.

To address this, we extract spatio-temporal tokens to enhance token diversity. Specifically, we employ a temporal pooling (TP) layer and a  $1 \times 1$  convolution layer to aggregate the skeleton features  $F^s$  across time. Similarly, we apply a spatial pooling (SP) layer and a  $1 \times 1$  convolution to average the spatial dynamics of all joints. The resulting spatio-temporal features are then flattened separately into unified representations, referred to as the spatial tokens  $z_s^s \in \mathbb{R}^{J*C}$  and temporal tokens  $z_t^s \in \mathbb{R}^{T*C}$ , respectively.

#### 3.2 Unified Skeleton Token Module

Due to variations in data collection methods across different datasets, the number of joints  $J$  and frame lengths  $T$  can differ significantly among datasets. Consequently, the dimensions of the spatio-temporal tokens extracted in Sec. 3.1 are inconsistent. Traditionally, classification tasks using skeleton data have addressed this heterogeneity by training separate models for each dataset. In this paper, we approach skeleton sequences as a specialized form of language, where differences in frame length and joint count are analogous to sentences of varying lengths. To unify these differences, we leverage the masking mechanism from natural language processing by applying a mask to all skeleton tokens, as shown

in Eq. (1). This approach allows us to merge all datasets into a larger, unified dataset in the language space for emotion representation learning.

$$\begin{aligned} z_s^{s'} &= z_s^s \odot M_s^s \\ z_t^{s'} &= z_t^s \odot M_t^s, \end{aligned} \quad (1)$$

where  $M_s^s$  and  $M_t^s$  are a  $1 \times L$  attention map that serves to retain the original tokens and nullify the padded elements.  $L$  represents the maximum length of all tokens.  $\odot$  denotes element-wise multiplication.

### 3.3 Skeleton-Aware Large Language Model

LLMs typically take sentences in human language as input instructions, but the skeleton sequences we use are not well-compatible with LLMs. To enable LLMs to recognize skeleton sequences and generate text descriptions, we apply Low-Rank Adaptation (LoRA) (Hu et al. 2021) to LLMs so that it better understands skeleton sequences while keeping the model’s pre-trained weights unchanged. Specifically, for each training sample composed of a skeleton sequence and its corresponding emotion label and emotion description, we perform the following steps:

1. We utilize a pre-trained GCN-based skeleton encoder to obtain the initial skeleton tokens ( $z_g^s, z_s^{s'}, z_t^{s'}$ ).
2. Prior work on multimodal LLMs (Liu et al. 2023) show that a learnable linear layer maps the visual feature space to the linguistic feature space, enabling LLMs to handle non-text data. Following this approach, we use a linear layer to transform the initial 768-dimensional skeleton token into a 4096-dimensional token, effectively bridging the gap between the skeleton encoder and LLMs.
3. We design two prompts that conform to the LLaMA-2 conversation template, as shown below.

**Emotion Recognition** #Human: <Skeleton> <SkeletonFeature> </Skeleton> Can you tell me the emotion of this person? #Assistant:

**Emotion Description** #Human: <Skeleton> <SkeletonFeature> </Skeleton> The emotion of this person is [shame], please tell me some reasons for it. #Assistant:

In these prompts, <SkeletonFeature >represents a skeleton token derived in step 2, and the word within square brackets represents an emotion label.

4. We concatenate the skeleton tokens with the prompt tokens and input them together into LLMs.
5. We use Eq. (2) to constrain the similarity between the ground truth tokens  $t_g$  and the predicted tokens  $t_p$ , adjusting EAI-LLM through LoRA.

$$\mathcal{L}_{LoRA} = \mathcal{L}_{ce}(t_p, t_g), \quad (2)$$

where  $\mathcal{L}_{ce}(\cdot, \cdot)$  is the cross-entropy loss.

### 3.4 Skeleton-Language Alignment

Skeleton and text data exist in distinct feature spaces, making it challenging for LLMs to interpret unaligned skeleton features. To address this, we employ contrastive learning to align the skeleton feature space with the linguistic feature space, thus improving the compatibility of the skeleton tokens with LLMs.

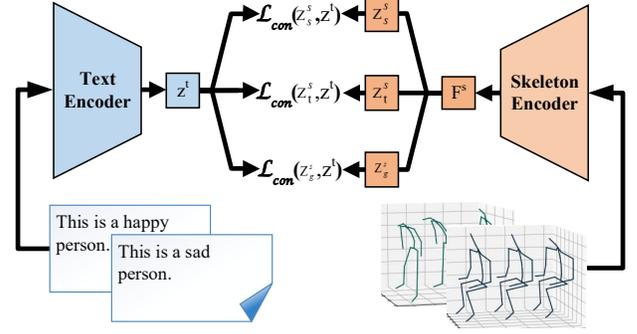


Figure 2: Diagram of skeleton-language alignment.

**Skeleton Encoder** Our skeleton encoder is designed in a manner similar to CTR-GCN (Chen et al. 2021), enabling the extraction of skeleton features  $F^s$ . Following pooling and dimensional transformation, we derive a C-dimensional feature vector  $z^s$ . The backbone of the encoder is flexible and can be substituted with other GCN-based networks.

**Text Encoder** We employ a pre-trained CLIP model (specifically ViT-L/14) (Radford et al. 2021) as the text encoder to extract a feature vector  $z^t$  representing the text description. Drawing inspiration from ActionCLIP (Wang et al. 2023), we manually design the text descriptions in the format: “This is a [happy] person.”, where the word within the brackets changes according to the skeleton label.

**Semantic Alignment** Contrastive learning excels at unifying data from different modalities into a shared feature space, facilitating cross-modal knowledge transfer (Radford et al. 2021; Xiang et al. 2023). Consequently, we employ contrastive learning to align the skeleton feature space with the linguistic feature space, enabling the compatibility of these tokens with LLMs.

To bring the pairwise skeleton representation  $z^s$  and text description label representation  $z^t$  closer together, we calculate the skeleton-to-text and text-to-skeleton similarity scores as specified in Eq. (3).

$$\begin{aligned} p_i^{s2t}(z_i^s) &= \frac{\exp(\cos(z_i^s, z_i^t) / \tau)}{\sum_{j=1}^N \exp(\cos(z_i^s, z_j^t) / \tau)} \\ p_i^{t2s}(z_i^t) &= \frac{\exp(\cos(z_i^t, z_i^s) / \tau)}{\sum_{j=1}^N \exp(\cos(z_i^t, z_j^s) / \tau)}, \end{aligned} \quad (3)$$

where  $\cos$  represents cosine similarity,  $\tau$  is the temperature hyper-parameter,  $N$  represents the batch size. Since the number of skeleton sequences greatly exceeds the number of emotion labels, there will be multiple skeleton sequences with the same label within the same batch. Consequently, using cross-entropy to calculate the similarity between  $p_i^{s2t}$  and  $p_i^{t2s}$  is not suitable. Therefore, we redefine the skeleton-text contrastive loss using the Kullback–Leibler (KL) divergence, as shown in Eq. (4).

$$\begin{aligned} \mathcal{L}_{con}(z^s, z^t) &= \frac{1}{2} \mathbb{E}_{(z^s, z^t) \sim \mathcal{D}} [\text{KL}(p^{s2t}(z^s), \hat{y}) \\ &\quad + \text{KL}(p^{t2s}(z^t), \hat{y})], \end{aligned} \quad (4)$$

where  $\mathcal{D}$  is the entire dataset. The ground-truth  $\hat{y}$  is defined as 1 for positive pairs and 0 for negative pairs.

In this paper, we optimize the semantic tokens and spatio-temporal tokens using skeleton-language alignment, making them more easily recognized by LLMs. The loss functions can be represented by Eq. (5) and Eq. (6), respectively.

$$\mathcal{L}_{st} = \mathcal{L}_{ce}(z_g^s, y) + \frac{1}{2}(\mathcal{L}_{con}(z_s^s, z^t) + \mathcal{L}_{con}(z_t^s, z^t)), \quad (5)$$

$$\mathcal{L}_{se} = \mathcal{L}_{ce}(z_g^s, y) + \mathcal{L}_{con}(z_g^s, z^t), \quad (6)$$

where  $y$  is the one-hot presentation of the emotion label.

## 4 Experiment

### 4.1 Datasets

**Emilya** The EMotional body expression In dALLY Actions dataset (Fourati and Pelachaud 2016) comprises 8,206 samples. Eleven actors performed eight emotions—Neutral, Joy, Anger, Panic, Fear, Anxiety, Sadness, and Shame—across seven daily actions, including sitting, walking, and lifting. Data were recorded using the Xsens MVN system, capturing 28 3D joints at a frequency of 120 Hz.

**KDAE** The Kinematic Dataset of Actors Expressing Emotions (Zhang et al. 2020) was collected using a portable motion capture system that tracked 72 body markers at a frequency of 125 Hz. This dataset includes 1,402 recordings of seven emotions: Happiness, Sadness, Neutral, Anger, Disgust, Fear, and Surprise, performed by 22 semi-professional actors. For our analysis, we excluded hand joints and focused on 24 full-body joints.

**EGBM** The Emotional Gestures and Body Movements Corpus (Sapiński et al. 2019) contains 560 samples recorded by a Kinect V2 camera at 30 Hz. Sixteen Polish professional actors, evenly split between men and women aged 25 to 64, expressed seven emotions: Happiness, Sadness, Neutral, Anger, Disgust, Fear, and Surprise. Each emotion is represented by 80 samples, with 3D positions for 25 joints.

**Emotion Description** The aforementioned datasets only include emotional labels without detailed descriptions of emotional actions. To fully leverage the generative capabilities of LLMs, we manually annotated a subset of new data with fine-grained descriptions of emotional actions. For annotation, we referred to the emotional action descriptions provided in (Noroozi et al. 2018). In total, 174 samples from the Emilya dataset and 105 samples from the KDAE dataset were labeled with these detailed descriptions.

### 4.2 Implementation Details

All experiments are conducted on NVIDIA 4×A100 GPUs and are implemented using the PyTorch framework (Paszke et al. 2019). We use LLaMA-7B (Touvron et al. 2023) as the base LLMs. For data preprocessing, we followed the scheme outlined in CTR-GCN (Chen et al. 2021), adjusting each sample to 64 frames by padding shorter sequences with previous frames and downsampling longer sequences.

Dataset	Spatial Tokens	Temporal Tokens	Accuracy
Emilya	✓		85.08
	✓	✓	<b>86.36</b>
KDAE	✓		62.99
	✓	✓	<b>63.35</b>
EGBM	✓		53.21
	✓	✓	<b>55.96</b>

Table 1: Classification results using spatial tokens and temporal tokens.

Token Type	Dataset	Separate	Joint
Semantic	Emilya	80.63	<b>85.44</b>
	KDAE	67.97	<b>71.17</b>
	EGBM	61.47	<b>66.97</b>
Spatio-temporal	Emilya	86.36	<b>86.42</b>
	KDAE	<b>63.35</b>	62.28
	EGBM	55.96	<b>63.30</b>

Table 2: Classification results for various training strategies.

**Training** For emotion description, the model is trained for 10,000 steps using labeled datasets. The LoRA parameters are configured with a rank of 64, an alpha of 16, and a dropout rate of 0.05. The global batch size is 16, and the maximum learning rate is 1e-5. For emotion recognition, the model undergoes 800,000 training steps, with the global batch size increased to 64. The LoRA parameters and maximum learning rate remain the same as in the emotion description stage. During skeleton encoder pre-training, we optimize the model for 200 epochs using the stochastic gradient descent (SGD) optimizer with an initial learning rate of 0.1 and a batch size of 64. The learning rate is reduced by a factor of 10 at epochs 100, 150, and 175. A warm-up strategy is applied for the first 5 epochs.

**Evaluation Protocols** We randomly split the dataset into training and testing sets at a 4:1 ratio. In emotion recognition, we extract the emotion labels from the generated sentences and compare them with the ground truth labels to compute accuracy. It is important to note that if a sentence contains multiple emotion labels or no emotion label, it will be marked as 'Error'. However, if the labels are synonyms or different grammatical forms of the same label, the recognition is deemed successful. For emotion description, we use three commonly employed metrics in natural language processing: Rouge, BLEU, and METEOR.

		Emotion Description Task				Emotion Recognition Task			
Token Type	Order	Accuracy	Rouge	BLEU	METEOR	Accuracy	Rouge-1	BLEU	METEOR
Semantic	D→R	N/A	0.1939	0.1301	0.2632	81.30	0.2018	0.1177	0.2148
Semantic	R→D	82.04		N/A		44.98	0.3138	0.2488	0.3408
Spatio-temporal	D→R	N/A	0.2774	0.2239	0.2941	81.40	0.2065	0.0401	0.0735
Spatio-temporal	R→D	81.45		N/A		19.59	0.2714	0.2149	0.2849

Table 3: Comparison of classification and emotion description results under different training orders. D→R indicates that the emotion description task is completed first, followed by the recognition task, while R→D indicates the reverse order.

### 4.3 Comparisons for Emotion Recognition

**Evaluation on Spatio-temporal Tokens** We evaluate the classification results utilizing spatial and temporal tokens both separately and in combination, as presented in Tab. 1. The results indicate that the separate classification of spatial and temporal tokens leads to lower accuracy when compared to their combined application across all three datasets. The integration of spatial and temporal tokens provides the model with more comprehensive information, thereby validating the effectiveness of our proposed multi-granularity skeleton tokenizer.

**Evaluation of Different Training Strategies** We evaluate the effectiveness of separate training strategy and joint training strategies with UST module by using skeleton tokens at different granularities. The results, presented in Tab. 2, show that using semantic tokens with joint training significantly improved classification accuracy across all three datasets, with an average increase of 4.5%. Specifically, when employing spatio-temporal tokens, the accuracy on the EGBM dataset increased by 7.34%. However, improvements are less pronounced for other datasets, with a slight decrease in accuracy observed on the KDAE dataset. These results suggest that joint training with semantic tokens enhances knowledge transfer between datasets, leading to better classification accuracy for individual datasets.

To further investigate the impact of UST module on recognition results, we plot the confusion matrix for the Emilya dataset, which has the largest number of samples. As shown in Fig. 3, the classification accuracy for anxiety, shame, and pride has significantly improved after employing the joint training strategy. Notably, these three emotions are absent in the KDAE and EGBM datasets. Compared to separate training, the joint training strategy has enhanced the model’s recognition of emotions common to all three datasets, thereby positively influencing the recognition of emotion labels unique to the Emilya dataset.

### 4.4 Comparisons for Emotion Description

In this section, we compare the impact of semantic tokens and spatio-temporal tokens on emotion description and recognition results. Our model is designed to handle both emotion description and emotion recognition tasks simultaneously, so we also report the effects of different training orders on the results. Note that all results are obtained un-

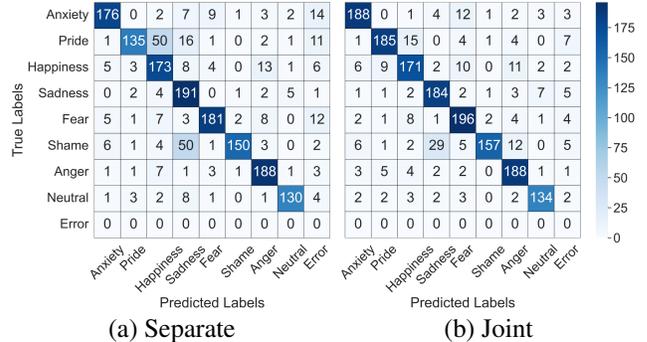


Figure 3: Confusion matrices for Emilya dataset using different training strategies.

der joint training strategy, and both accuracy and generation metrics are averaged across the three datasets.

As shown in Tab. 3, we can draw three conclusions: (i) **In the emotion description task, spatio-temporal tokens outperform semantic tokens, whereas in the emotion recognition task, the opposite is true.** This advantage is likely because spatio-temporal tokens retain richer spatial and temporal information, which enhances the detail and quality of the generated descriptions. Semantic tokens are more concise and contribute better to classification. (ii) **The training order of different tasks significantly affects the recognition results.** When training the emotion description and emotion recognition tasks separately, good results can be achieved regardless of which tokens are used. However, when the emotion recognition task is completed first, and then the emotion description task is fine-tuned, catastrophic forgetting occurs. This results in a significant drop in average recognition accuracy by 37.06% and 61.86% for semantic and spatio-temporal tokens, respectively. This issue may arise from conflicting demands within the model for both emotion recognition and description capabilities. EAI-LLM imposes stringent constraints on recognition outcomes, expecting results to be framed in specific sentences, such as ”This is a happy person.”. These constraints often conflict with the diverse requirements of generation tasks, leading to reduced recognition accuracy when generation tasks are performed after recognition tasks. (iii) **Semantic tokens offer a balanced trade-off between recognition ac-**

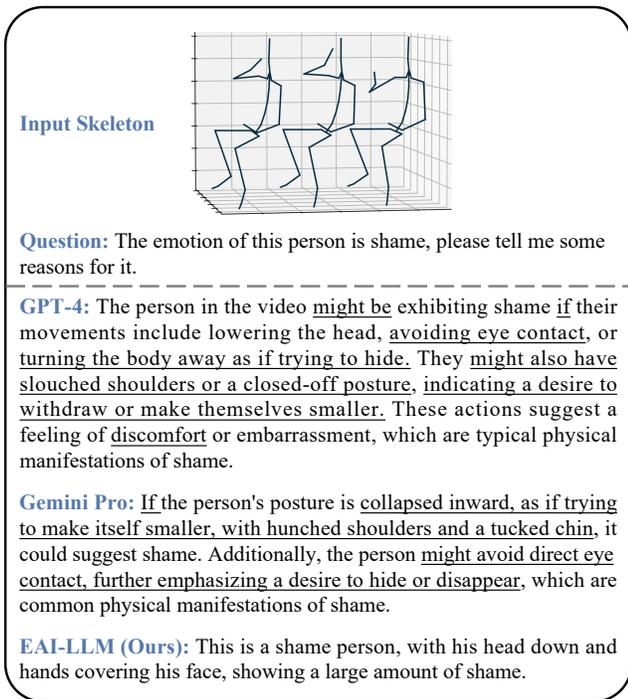


Figure 4: Examples for emotion description capabilities of EAI-LLM. Underlined text indicates that the description is unrelated to the input sequences.

**curacy and description capability.** Although semantic tokens show lower accuracy compare to spatio-temporal tokens in emotion recognition tasks, their performance in subsequent emotion description tasks does not decline significantly, thereby maintaining relatively good accuracy.

#### 4.5 Comparison with the State-of-the-Art

Due to the limited reporting on skeleton-based emotion recognition and the absence of a standardized evaluation method, we have chosen to compare our approach with existing methods used in skeleton-based action recognition, which are analogous to emotion recognition. To ensure a fair comparison, we re-implemented all methods and standardized the data preprocessing procedures. As shown in Tab. 4, our method demonstrates accuracy comparable to existing approaches on the KDAE and EGBM datasets, with the best accuracy reaching 71.17% and 66.97%, respectively. However, on the Emilya dataset, the accuracy only reached 85.44%, still lagging behind other methods.

During the annotation process of the emotion description dataset, we convert the skeleton sequences into videos and then labeled the emotional action features by viewing these videos. Detailed examples are provided in the supplementary materials. Since current multi-modal large models do not support direct input of the original skeleton sequences, we instead input the videos derived from these skeleton sequences into the multi-modal LLMs, using the same prompts as EAI-LLM. As shown in Tab. 5, when using the same prompt words, the metrics produced by our method surpass

	Emilya	KDAE	EGBM	D
AGCN (Shi et al. 2019)	88.92	56.58	22.94	×
CTR-GCN (Chen et al. 2021)	<b>89.77</b>	<u>70.46</u>	63.30	×
GAP (Xiang et al. 2023)	<u>89.16</u>	67.26	<u>66.06</u>	×
<b>EAI-LLM (Ours)</b>	85.44	<b>71.17</b>	<b>66.97</b>	✓

Table 4: Comparisons of emotion recognition capabilities with the state-of-the-art methods. Bold and underline indicate the best and the second best result. D represents emotion description ability from 3D skeleton sequences.

	Rouge	BLEU	METEOR	R
GPT-4	0.1282	0.0686	0.2006	×
Gemini 1.5 Pro	0.1007	0.0596	0.1823	×
<b>EAI-LLM (Ours)</b>	<b>0.2018</b>	<b>0.1177</b>	<b>0.2148</b>	✓

Table 5: Comparison of emotion description capabilities with mainstream multimodal LLMs. R represents emotion recognition ability from 3D skeleton sequences.

those generated by popular multi-modal LLMs.

We present the visualization results in Fig. 4. The emotion descriptions generated by EAI-LLM exhibit significant consistency with the input sequences. In contrast, while both GPT-4 (OpenAI 2023) and Gemini 1.5 Pro (Research 2024) can also generate emotion descriptions from video, these descriptions largely reflect the models' background knowledge and do not align as well with the input sequences.

**Limitations** Since EAI-LLM is built upon LLMs, it inherits the LLM's limitations, such as hallucination, which can lead to ambiguous outputs. For instance, when given an input labeled as "Anxiety", EAI-LLM might generate: "This person is expressing anxiety or fear." The output contains multiple emotional labels and cannot be clearly categorized under a specific emotional label, thereby resulting in a recognition failure.

## 5 Conclusion

In this paper, we have introduced a novel emotion recognition model, EAI-LLM, designed to identify emotions from 3D full-body skeleton sequences and generate detailed emotion descriptions. We have used LLaMA as the language decoder and fine-tuned it on prompt-based question-answering tasks. Specifically, we have extracted skeleton tokens at various granularities to enable our model to produce more nuanced emotion descriptions. Additionally, the unified skeleton token module has significantly enhanced the accuracy of emotion recognition on individual datasets. Experimental results have demonstrated that our model effectively recognizes emotions from skeleton data and generates detailed emotion descriptions, even with limited data annotations.

## Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant 62072190 and 62202311, and in part by the Shenzhen Natural Science Foundation (the Stable Support Plan Program) under Grant 20220809180405001, and in part by the Innovation Team Project of Guangdong Province of China 2024KCXTD017, and in part by the Excellent Science and Technology Creative Talent Training Program of Shenzhen Municipality under Grant RCBS20221008093224017, and in part by the Guangdong Basic and Applied Basic Research Foundation under Grants 2023A1515011512, and in part by the Key Scientific Research Project of the Department of Education of Guangdong Province 2024ZDZX3012.

## References

- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *NeurIPS*, 35: 23716–23736.
- Anthropic. 2024. Claude: An AI Assistant by Anthropic. <https://www.anthropic.com/product>. Accessed: 2024-08-01.
- Beyan, C.; Karumuri, S.; Volpe, G.; Camurri, A.; and Niewiadomski, R. 2023. Modeling Multiple Temporal Scales of Full-Body Movements for Emotion Classification. *IEEE Transactions on Affective Computing*, 14(2): 1070–1081.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *NeurIPS*, 33: 1877–1901.
- Chen, Y.; Zhang, Z.; Yuan, C.; Li, B.; Deng, Y.; and Hu, W. 2021. Channel-wise topology refinement graph convolution for skeleton-based action recognition. In *ICCV*, 13359–13368.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; Zhang, H.; Zheng, L.; Zhuang, S.; Zhuang, Y.; Gonzalez, J. E.; Stoica, I.; and Xing, E. P. 2023. Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%\* ChatGPT Quality.
- Chowdhery, A.; Narang, S.; Devlin, J.; Bosma, M.; Mishra, G.; Roberts, A.; Barham, P.; Chung, H. W.; Sutton, C.; Gehrmann, S.; et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240): 1–113.
- Daoudi, M.; Berretti, S.; Pala, P.; Delevoe, Y.; and Del Bimbo, A. 2017. Emotion recognition by body movement representation on the manifold of symmetric positive definite matrices. In *ICIAP*, 550–560.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv:1810.04805*.
- El Ayadi, M.; Kamel, M. S.; and Karray, F. 2011. Survey on speech emotion recognition: Features, classification schemes, and databases. *Pattern recognition*, 44(3): 572–587.
- Fourati, N.; and Pelachaud, C. 2015. Multi-level classification of emotional body expression. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, volume 1, 1–8.
- Fourati, N.; and Pelachaud, C. 2016. Perception of emotions and body movement in the emilya database. *IEEE Transactions on Affective Computing*, 9(1): 90–101.
- Fourati, N.; Pelachaud, C.; and Darmon, P. 2019. Contribution of temporal and multi-level body cues to emotion classification. In *2019 8th International Conference on Affective Computing and Intelligent Interaction*, 116–122.
- Fragopanagos, N.; and Taylor, J. G. 2005. Emotion recognition in human-computer interaction. *Neural Networks*, 18(4): 389–405.
- Ghaleb, E.; Mertens, A.; Asteriadis, S.; and Weiss, G. 2021. Skeleton-based explainable bodily expressed emotion recognition through graph convolutional networks. In *FG*, 1–8.
- Glowinski, D.; Camurri, A.; Volpe, G.; Dael, N.; and Scherer, K. 2008. Technique for automatic emotion recognition by body gesture analysis. In *CVPR Workshops*, 1–6.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv:2106.09685*.
- Li, C.; Wang, J.; Zhang, Y.; Zhu, K.; Hou, W.; Lian, J.; Luo, F.; Yang, Q.; and Xie, X. 2023a. Large language models understand and can be enhanced by emotional stimuli. *arXiv:2307.11760*.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023b. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 19730–19742.
- Li, K.; He, Y.; Wang, Y.; Li, Y.; Wang, W.; Luo, P.; Wang, Y.; Wang, L.; and Qiao, Y. 2023c. Videochat: Chat-centric video understanding. *arXiv:2305.06355*.
- Li, K.; Wang, Y.; He, Y.; Li, Y.; Wang, Y.; Liu, Y.; Wang, Z.; Xu, J.; Chen, G.; Luo, P.; et al. 2024. Mvbench: A comprehensive multi-modal video understanding benchmark. In *CVPR*, 22195–22206.
- Li, S.; and Deng, W. 2022. Deep Facial Expression Recognition: A Survey. *IEEE Transactions on Affective Computing*, 13(3): 1195–1215.
- Lima, M. L.; De Lima Costa, W.; Martínez, E. T.; and Teichrieb, V. 2024. ST-Gait++: Leveraging spatio-temporal convolutions for gait-based emotion recognition on videos. In *CVPR*, 302–310.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual Instruction Tuning. In *NeurIPS*, volume 36, 34892–34916.
- Lu, H.; Hu, X.; and Hu, B. 2023. See Your Emotion from Gait Using Unlabeled Skeleton Data. In *AAAI*, volume 37, 1826–1834.
- Maaz, M.; Rasheed, H.; Khan, S.; and Khan, F. S. 2023. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv:2306.05424*.

- Mandryk, R. L.; Atkins, M. S.; and Inkpen, K. M. 2006. A continuous and objective evaluation of emotional experience with interactive play environments. In *SIGCHI*, 1027–1036.
- Narayanan, V.; Manoghar, B. M.; Dorbala, V. S.; Manocha, D.; and Bera, A. 2020. Proximo: Gait-based emotion learning and multi-view proxemic fusion for socially-aware robot navigation. In *IROS*, 8200–8207.
- Noroozi, F.; Corneanu, C. A.; Kamińska, D.; Sapiński, T.; Escalera, S.; and Anbarjafari, G. 2018. Survey on emotional body gesture recognition. *IEEE Transactions on Affective Computing*, 12(2): 505–523.
- Oğuz, A.; and Ertuğrul, Ö. F. 2024. Emotion recognition by skeleton-based spatial and temporal analysis. *Expert Systems with Applications*, 238: 121981.
- OpenAI. 2023. GPT-4: Technical Report. <https://cdn.openai.com/papers/gpt-4.pdf>. Accessed: 2024-08-15.
- OpenAI. 2024. ChatGPT: An AI Language Model. <https://chat.openai.com/>. Accessed: 2024-08-15.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*, volume 35, 27730–27744.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, volume 32.
- Peng, Q.; Zheng, C.; and Chen, C. 2024. A Dual-Augmentor Framework for Domain Generalization in 3D Human Pose Estimation. In *CVPR*, 2240–2249.
- Piana, S.; Staglianò, A.; Odone, F.; and Camurri, A. 2016. Adaptive body gesture representation for automatic emotion recognition. *ACM Transactions on Interactive Intelligent Systems*, 6(1): 1–31.
- Qiu, F.; Zhang, W.; Liu, C.; Li, L.; Du, H.; Guo, T.; and Yu, X. 2024. Language-guided Multi-modal Emotional Mimicry Intensity Estimation. In *CVPR*, 4742–4751.
- Qu, H.; Cai, Y.; and Liu, J. 2024. LLMs are good action recognizers. In *CVPR*, 18395–18406.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Research, G. 2024. Gemini 1.5: Unlocking multi-modal understanding across millions of tokens of context. *arXiv:2403.05530*.
- Sapiński, T.; Kamińska, D.; Pelikant, A.; Ozcinar, C.; Avots, E.; and Anbarjafari, G. 2019. Multimodal database of emotional speech, video and gestures. In *ICPR*, 153–163.
- Shi, L.; Zhang, Y.; Cheng, J.; and Lu, H. 2019. Two-Stream Adaptive Graph Convolutional Networks for Skeleton-Based Action Recognition. In *CVPR*, 12026–12035.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Stanford Alpaca: An Instruction-following LLaMA model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv:2302.13971*.
- Wang, M.; Xing, J.; Mei, J.; Liu, Y.; and Jiang, Y. 2023. Actionclip: Adapting language-image pretrained models for video action recognition. *IEEE Transactions on Neural Networks and Learning Systems*, 1–13.
- Wang, T.; Liu, S.; He, F.; Dai, W.; Du, M.; Ke, Y.; and Ming, D. 2024. Emotion Recognition From Full-Body Motion Using Multiscale Spatio-Temporal Network. *IEEE Transactions on Affective Computing*, 15(3): 898–912.
- Xiang, W.; Li, C.; Zhou, Y.; Wang, B.; and Zhang, L. 2023. Generative action description prompts for skeleton-based action recognition. In *ICCV*, 10276–10285.
- Yan, S.; Xiong, Y.; and Lin, D. 2018. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, volume 32, 7444–7452.
- You, Z.; Wen, Z.; Chen, Y.; Li, X.; Zeng, R.; Wang, Y.; and Tan, M. 2024. Toward Long Video Understanding via Fine-detailed Video Story Generation. *IEEE Transactions on Circuits and Systems for Video Technology*, 1–1.
- Zeng, R.; Zhuo, Y.; Li, J.; Yang, Y.; Wu, H.; Chen, Q.; Hu, X.; and Leung, V. C. 2024. Improving Video Moment Retrieval by Auxiliary Moment-Query Pairs with Hyper-Interaction. *IEEE Transactions on Circuits and Systems for Video Technology*, 1–1.
- Zhai, Y.; Jia, G.; Lai, Y.-K.; Zhang, J.; Yang, J.; and Tao, D. 2024. Looking into Gait for Perceiving Emotions via Bilateral Posture and Movement Graph Convolutional Networks. *IEEE Transactions on Affective Computing*, 1–15.
- Zhang, H.; Yi, P.; Liu, R.; and Zhou, D. 2021. Emotion recognition from body movements with as-lstm. In *ICVR*, 26–32.
- Zhang, M.; Yu, L.; Zhang, K.; Du, B.; Zhan, B.; Chen, S.; Jiang, X.; Guo, S.; Zhao, J.; Wang, Y.; et al. 2020. Kinematic dataset of actors expressing emotions. *Scientific data*, 7(1): 292.
- Zhang, Z. 2012. Microsoft kinect sensor and its effect. *IEEE multimedia*, 19(2): 4–10.
- Zheng, C.; Wu, W.; Chen, C.; Yang, T.; Zhu, S.; Shen, J.; Kehtarnavaz, N.; and Shah, M. 2023. Deep learning-based human pose estimation: A survey. *ACM Computing Surveys*.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv:2304.10592*.