

HERO: Human Reaction Generation from Videos

Chengjun Yu¹, Wei Zhai^{1,†}, Yuhang Yang¹, Yang Cao¹, Zheng-Jun Zha¹

¹ MoE Key Laboratory of Brain-inspired Intelligent Perception and Cognition, USTC

{yucj@mail., wzhai056@, yyuhang@mail., forrest@, zhazj@}ustc.edu.cn

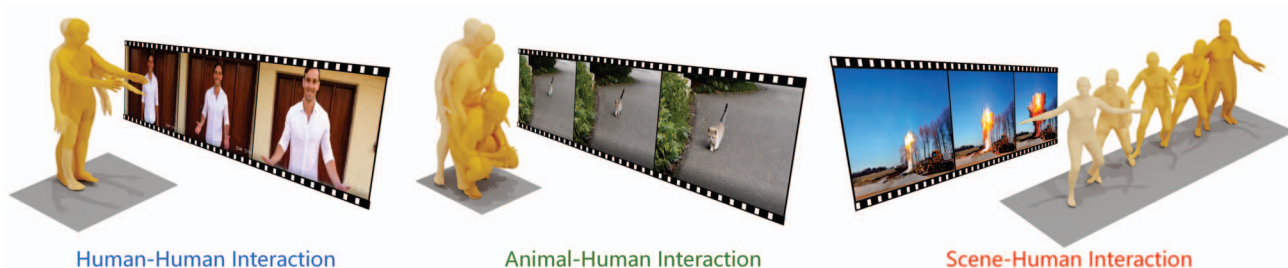


Figure 1. **Overview of our work.** We propose to generate 3D human reactions from RGB videos. To tackle this task, a simple yet effective framework, HERO, is presented. Furthermore, to facilitate research in this area, we introduce the ViMo dataset, which features a wide range of interaction categories covering three broad ones: **human-human interactions**, **animal-human interactions**, and **scene-human interactions**. For the human reactions visualized in the figure, the darker colors indicate the later in time.

Abstract

Human reaction generation represents a significant research domain for interactive AI, as humans constantly interact with their surroundings. Previous works focus mainly on synthesizing the reactive motion given a human motion sequence. This paradigm limits interaction categories to human-human interactions and ignores emotions that may influence reaction generation. In this work, we propose to generate 3D human reactions from RGB videos, which involves a wider range of interaction categories and naturally provides information about expressions that may reflect the subject’s emotions. To cope with this task, we present **HERO**, a simple yet effective framework for **H**uman **r**EAction **g**ene**R**ation from **v**ide**O**s. **HERO** considers both global and frame-level local representations of the video to extract the interaction intention, and then uses the extracted interaction intention to guide the synthesis of the reaction. Besides, local visual representations are continuously injected into the model to maximize the exploitation of the dynamic properties inherent in videos. Furthermore, the **ViMo** dataset containing paired **V**ideo-**M**otion data is collected to support the task. In addition to human-human interactions, these video-motion pairs also cover animal-human interactions and scene-human interactions. Extensive experiments demonstrate the superiority of our method.

1. Introduction

Human reaction generation refers to the technology that allows computers to synthesize realistic and natural human reactions based on input signals. It has a wide range of applications in interactive virtual and augmented reality, animation, games, human-robot interaction, and embodied AI, which promises to bring non-player characters (NPCs) to life and enable humanoid robots to understand and respond to social cues and environmental factors in a human-like manner [21, 59, 103]. Previous works [13, 21, 58, 59, 103] focus mainly on synthesizing the motions of reactors conditioned on the motions of actors. These motions are well-organized structured data that computers can easily process. However, intelligent systems such as humanoid robots often use raw visual signals as input. Therefore, human reaction generation from videos is worth exploring.

We note that using human motions as input introduces two limitations in the context of human reaction generation: restricted interaction categories and the absence of emotional information. On the one hand, it limits the interaction types to human-human interactions, overlooking other forms of interaction such as animal-human and scene-human interactions, thereby constraining the machine’s capacity to engage with the broader world. On the other hand, the motion data used in prior works do not reflect the emotional states of the characters. In fact, during human-human interactions, the actor’s emotions may influence the reactor’s decision-making process, particularly when the actor’s

[†]Corresponding Author.

Project page: <https://jackyu6.github.io/HERO>

motion cannot clearly convey interaction intention. Consider the following interaction: “*the actor walking towards the reactor*”. If the actor approaches the reactor with a neutral facial expression, the reactor may simply remain stationary, awaiting further actions from the actor. When the actor walks towards the reactor in an angry manner, the reactor may tend to respond by stepping back or moving aside. In contrast, if the actor approaches the reactor happily with a smile on their face, the reactor is more likely to move forward and may even initiate interactions such as handshakes or waves. The aforementioned example illustrates that the same action, when associated with different emotional information, could elicit distinct human reactions. Compared to human motion data, RGB videos can encompass a broader range of interaction categories, and naturally contain both the subjects’ actions and facial expressions that may reflect emotions (provided the face is visible) [86–88], without the need for expensive specialized equipment for capture.

In light of the significance to the application and the field discussed above, we propose to generate 3D human reactions from RGB videos. Unlike the previous task on human reaction generation [13, 21, 58, 59, 103] that requires to align the motion of the actor and the synthesized motion of the reactor in the same coordinate system, our aim is to generate a single reactive motion based on a given video, similar to some mainstream tasks of human motion generation, *e.g.*, text-driven human motion generation (text-to-motion generation) [12, 26, 27, 69, 90, 115, 116]. Nevertheless, a notable distinction exists between our task and text-to-motion generation, particularly with regard to the utilization of representations of the input signals. Specifically, in the text-to-motion paradigm, once the semantic representation of the input text is obtained, it is typically leveraged without further processing, directly serving as a guidance for motion generation. This implies that the model passively adheres to the textual semantics, with limited active comprehension or refinement of the encoded information. In contrast, our objective is not to have the model passively comply with video semantics to synthesize motion. Instead, we aim to enable the model to proactively extract the interaction intention from the visual representations, which is then exploited to guide the generation of the reaction.

The global representation of a video can, to some extent, reflect interaction intention. However, the interaction intention it represents is not ideal due to being derived from the average pooling of frame-level representations [45, 74, 96], which treats each frame equally. Notably, certain key-frames within a video are more revealing of interaction intention than other frames. For instance, the global semantics of a video might be described as: “*A person walks over; initially with no movement in the upper body, then spreads their arms.*” The frames associated with

“*spreading the arms*” reveal interaction intention of “*hugging*” more significantly than the preceding frames.

Motivated by the above observation, we propose a global-local representation cross-attention mechanism that simultaneously leverages global and local visual representations obtained from the video encoder [45] to extract interaction intention. Specifically, it assigns weights to each local frame-level representation on-the-fly based on the global representation of the input video and then performs weighted fusion of per-frame representations to obtain refined information on interaction intention. Next, we employ intention-conditioned self-attention to empower interaction intention to guide reaction generation. We further adopt motion-frame cross-attention to enable the model to obtain the details of the video content, fully leveraging the inherent dynamic properties of the video. We integrate the aforementioned techniques into a new framework, **HERO**, for human reaction generation from videos.

In addition, we collect the **ViMo** dataset, which contains video-motion pairs involving human-human, animal-human, and scene-human interactions, featuring diverse interactions of 32 subcategories. It serves as a test bed for model training and evaluation.

Our main contributions can be summarized as follows:

- 1) We propose a new task aimed at generating 3D human reactions from RGB videos with interactivity. Due to the inherent nature of videos, we break through the limitations of previous works with regard to categories and emotions.
- 2) We present HERO, a simple yet effective framework that extracts interaction intention from the video, and utilize the interaction intention and the dynamic information of the video to synthesize plausible reactions.
- 3) We construct the ViMo dataset that contains video-motion pairs covering a wide range of interaction categories to support the proposed task setting. Extensive experiments demonstrate the feasibility of our task and the superiority of our methodology.

2. Related Work

Human Motion Generation. The goal of human motion generation is to synthesize human motion conditioned on various signals such as action labels [6, 9, 25, 62, 68, 101, 104], texts [12, 17, 18, 26, 27, 69, 90, 98, 115, 116], music [5, 22, 51, 52, 84, 92, 121], speech [3, 4, 30, 122] and sparse observations [2, 8, 16, 33, 39]. Different kinds of generative models are utilized in these works for effective modeling. While some works [97, 101, 104] adopt GAN [23] modeling for their methodologies, others [6, 9, 25, 26, 68, 69] utilize the VAE [46] framework to learn probabilistic mapping between conditions and motions. In addition, flow-based models [2, 75] are also employed. Diffusion models [31, 85] have recently been increasingly applied in human-related generation tasks [10, 53, 109]. Diffusion-based methods [14, 15, 34, 44, 48, 81, 92, 111, 120] generate mo-

tions through a reverse diffusion process from noise. For autoregressive-based methods [22, 27, 37, 115, 117, 119, 123], motions are first discretized as tokens through vector quantization [93] so that the Transformer can model the task as a next token prediction problem. Furthermore, some recent works [28, 35, 70, 72] propose to exploit masked motion modeling, which achieves real-time and high-fidelity motion generation thanks to the efficient inference process and powerful bidirectional attention. Therefore, we naturally adopt generative masked motion modeling in this work. However, unlike simply following the control signal to synthesize the motion, our task requires the model to extract the interaction intention from the video to generate a plausible human reaction.

Human Reaction Generation. Previous works [13, 21, 58, 59, 103] of human reaction generation mainly focus on human-human interactions, i.e., synthesizing the reactor’s motion conditioned on the actor’s motion. To achieve this goal, [13] proposes InterFormer, a Transformer network with temporal motion attention and spatial skeleton attention. ReGenNet [103] and ReMoS [21] adopt diffusion-based models to generate the full body motion of the reactor. [58] designs a method of reaction synthesis with or without objects utilizing affordance [82, 105–107, 118]. PhysReaction [59] introduces a forward dynamics-guided 4D imitation method to generate physically plausible reactions. However, none of the aforementioned works has considered animal-human and scene-human interactions, as well as emotional information of the actor. In contrast, we propose to generate human reaction from RGB videos, which cover human-human, animal-human, and scene-human interactions. In addition to the actor’s movement, the video could also provide information about the actor’s facial expression, which naturally reflects their emotion.

Human Interaction Dataset. There are many datasets on human interaction, as it is a topic that has long been of interest. Some datasets [29, 32, 54, 102, 110] provide motion data of human-human interactions. However, they lack facial expression information of participants. In other words, these data cannot reflect the emotions of the actors, as videos could. Other datasets [41, 64, 67, 78] supply videos of human-human interactions. However, most videos involving human-human interactions in these datasets contain the actor and the reactor simultaneously, i.e., they are not filmed from the perspective of the reactor. There are also datasets [19, 57, 94, 95, 112] that furnish both motion data and videos of human-human interactions. However, their videos suffer from the same issue mentioned above. [42, 43, 47, 66, 79, 80] do provide videos from the reactor’s view. However, they cover only a few interaction categories and are limited to human-human interactions. To support our work, we build a new dataset that contains video-motion pairs, covering human-human, animal-human, and scene-

human interactions.

3. Method

Given an RGB video V with interactivity, HERO generates a reactive 3D human motion M . Specifically, a video $V \in \mathbb{R}^{T \times H \times W \times 3}$ has a spatial resolution of $H \times W$ with T sampled frames, and a motion $M \in \mathbb{R}^{N \times D}$ is a sequence of N poses with D denoting the dimension of pose representations. Our framework consists of three modules: The video encoder (Sec. 3.1) is used to extract visual representations of the input video. The motion VQ-VAE (Sec. 3.2) learns a mapping between the raw motion and discrete code sequences. The reaction generation module (Sec. 3.3) extracts interaction intention from visual representations and uses it to guide the generation of code indices. The pipeline of HERO is illustrated in Fig. 2.

3.1. Video Encoder

To extract the representations of the input videos, we employ the video encoder of TC-CLIP [45], a paradigm for extending CLIP [73] to videos by encoding holistic video information through temporal contextualization. Given a video V , the video encoder obtains the per-frame representations $\mathbf{v}_l = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_T]$, which we call local visual representations. Finally, the global visual representation \mathbf{v}_g is produced by averaging local visual representations formulated as: $\mathbf{v}_g = \text{AvgPool}([\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_T])$, where $\mathbf{v}_g \in \mathbb{R}^{d_{vl}}$, $\mathbf{v}_i \in \mathbb{R}^{d_{vl}}$, $i \in \{1, 2, \dots, T\}$, d_{vl} is the dimension of the vision-language latent space.

3.2. Motion VQ-VAE

The functionality of the motion VQ-VAE is to learn discrete representations of motions for subsequent use by the generative module. It transforms the embedding \mathbf{z} output by the motion encoder \mathbf{E}_m into the entries of a learnable codebook $\mathbf{C} = \{\mathbf{c}_k\}_{k=1}^K$ containing K code entries with $\mathbf{c}_k \in \mathbb{R}^{d_c}$, where d_c is the dimension of the entries. Specifically, the 1D convolutional motion encoder \mathbf{E}_m first transforms the motion $M \in \mathbb{R}^{N \times D}$ into the latent embedding $\mathbf{z} = \mathbf{E}_m(M) \in \mathbb{R}^{n \times d_c}$ with a temporal downsampling ratio of $\alpha = n/N$. Next, we have the process of vector quantization \mathbf{Q} : Each vector in \mathbf{z} is substituted with its closest entry in the codebook \mathbf{C} , determined by computing the Euclidean distance. The indices of the sampled entries, i.e., motion tokens $\mathbf{m} \in \mathbb{R}^n$ serve as discrete representations of the input motion. Finally, the motion can be reconstructed by the motion decoder \mathbf{D}_m using the quantized vectors $\mathbf{q} = \mathbf{Q}(\mathbf{z}) \in \mathbb{R}^{n \times d_c}$ as inputs. Overall, the motion VQ-VAE is trained with the following loss function:

$$\mathcal{L}_{vq} = \mathcal{L}_{re} + \|sg[\mathbf{z}] - \mathbf{q}\|_2^2 + \beta \|\mathbf{z} - sg[\mathbf{q}]\|_2^2, \quad (1)$$

where $sg[\cdot]$ refers the stop-gradient operator, β is a hyperparameter. For the motion reconstruction, we use ℓ_1 smooth

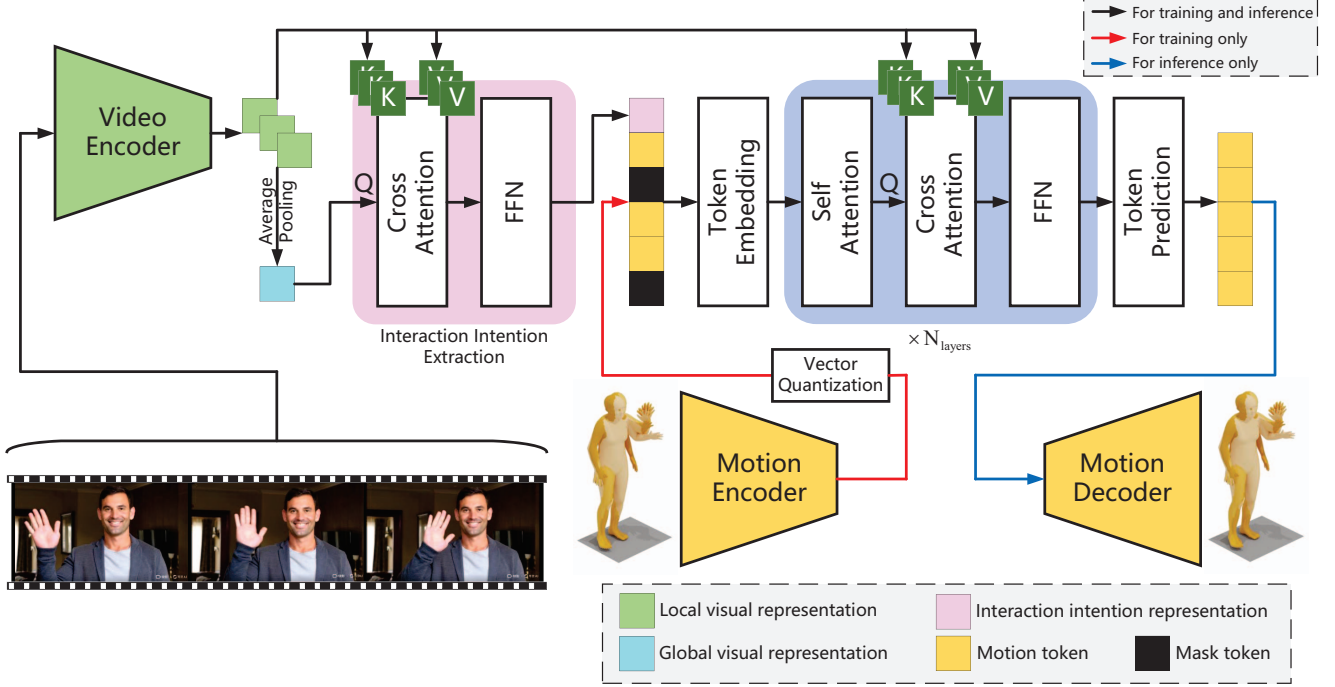


Figure 2. **The pipeline of HERO.** During training, the video and GT reactive motion are input into HERO. As for inference, only the video is provided. Note that we omit the residual motion refinement (See the end of Sec. 3.3) from the figure for clarity.

loss: $\mathcal{L}_{re} = \mathcal{L}_1^{smooth}(M, D_m(q))$. To improve codebook utilization, we apply the exponential moving average and codebook reset following [28, 71, 72, 115].

3.3. Reaction Generation Module

Masked Motion Modeling. Recently, generative masked motion modeling [28, 72] has achieved high-fidelity and real-time motion generation due to its powerful bidirectional attention and efficient inference process. Therefore, we naturally migrate masked motion modeling to our task. In our experiments, we use two learnable special tokens: [PAD] and [MASK]. The [PAD] token is utilized to pad shorter motion sequences, enabling the processing of batches containing multiple sequences with different lengths. The [MASK] represents input corruption. We first randomly replace a varying fraction of motion tokens \mathbf{m} with the [MASK] token to obtain a corrupted sequence $\tilde{\mathbf{m}}$. Then, a masked Transformer model is trained to predict the actual tokens in place of [MASK] given conditions c and $\tilde{\mathbf{m}}$. In practice, we linearly project c and $\tilde{\mathbf{m}}$ to the same latent dimension d_l . The training objective is to minimize the negative log-likelihood of target predictions:

$$\mathcal{L}_{mask} = \mathbb{E}_{\mathbf{m} \in \mathcal{D}} \left[\sum_{\forall \tilde{\mathbf{m}}_i = [\text{MASK}]} -\log p(\mathbf{m}_i | \tilde{\mathbf{m}}, c) \right]. \quad (2)$$

Unlike previous works [28, 72] using Transformer encoder, we implement the masked Transformer by stacking

N_{layers} Transformer decoder units to better exploit the dynamic information in the videos. We adopt a cosine function to schedule the masking ratio and use the remasking strategy following [28]. The inference process begins with an empty canvas, which is represented as a sequence where all tokens are [MASK]. Subsequently, the model iteratively predicts more tokens in parallel at each step.

Interaction Intention Extraction (IIE). The most different aspect from previous paradigms such as text-to-motion generation is that our goal is not to have the model passively follow control signals, but rather to enable it to proactively extract interaction intention from visual representations to guide reaction generation. In general, the global visual representation \mathbf{v}_g reflects the interaction intention to some extent, as it contains global semantic information about the video. In fact, some key-frames in a video could be more valuable than others for revealing the interaction intention. Therefore, instead of directly using the global visual representation \mathbf{v}_g derived from the average pooling of local visual representations \mathbf{v}_l as the interaction intention representation, we propose the global-local representation cross-attention to extract a refined interaction intention representation. Mathematically, our global-local representation cross-attention is expressed as:

$$Att_{gl} = \text{softmax} \left(\frac{\mathbf{Q}_g \mathbf{K}_l^T}{\sqrt{d_{vl}}} \right) \mathbf{V}_l, \quad (3)$$

where the query $\mathbf{Q}_g = \mathbf{v}_g \mathbf{W}_q \in \mathbb{R}^{d_{vl}}$, keys $\mathbf{K}_l = \mathbf{v}_l \mathbf{W}_k \in$

$\mathbb{R}^{T \times d_{vl}}$ and values $\mathbf{V}_l = \mathbf{v}_l \mathbf{W}_v \in \mathbb{R}^{T \times d_{vl}}$ are given by linear projections, d_{vl} is the dimension of the vision-language latent space. Thus, the model can proactively assign weights to each local visual representation on-the-fly based on the global visual representation of the video, and perform a weighted summation of local visual representations to produce an embedding Att_{gl} with refined information on interaction intention. Finally, we input Att_{gl} into an FFN to obtain the interaction intention representation \mathbf{e} .

Interaction Intention Injection. In order to inject interaction intention into the model so that it can guide the generation of reactions, we adopt intention-conditioned self-attention. Specifically, we first project the interaction intention representation \mathbf{e} and the corrupted motion token sequence $\tilde{\mathbf{m}}$ to the same latent dimension d_l , and concatenate them through the temporal dimension. Then we calculate the conditional self-attention between them. In this manner, the interaction intention is effectively integrated into the token sequence through the attention computation process, thereby facilitating the guidance of reaction synthesis by the interaction intention.

Dynamic Information Exploitation (DIE). Videos inherently possess dynamic attributes that can be well utilized by the model. Intuitively, we linearly project local visual representations to the latent dimension d_l , and compute the motion-frame cross-attention to maximize the exploitation of available dynamic information in the video. The calculation of the cross-attention is formulated as follows:

$$Att_{mf} = \text{softmax} \left(\frac{\mathbf{Q}_m \mathbf{K}_f^T}{\sqrt{d_l}} \right) \mathbf{V}_f, \quad (4)$$

where $\mathbf{Q}_m \in \mathbb{R}^{(n+1) \times d_l}$ is derived from linear projection of the output of the previous conditional self-attention, $\mathbf{K}_f \in \mathbb{R}^{T \times d_l}$ and $\mathbf{V}_f \in \mathbb{R}^{T \times d_l}$ are linear projections of the local visual representations in the latent space. Through this process, the model can take advantage of the inherent dynamic properties of the video to browse the detailed content in the video, and learn the importance of each key-frame in the video for reaction generation.

Residual Motion Refinement. We utilize residual motion refinement as in [28, 71] to further improve the quality of the generated motion. In the motion tokenization stage, a residual VQ-VAE (RVQ-VAE) [113] discretizes the raw motion sequence into multiple code sequences to minimize information loss during the process of quantization. Specifically, each code sequence is produced by a distinct quantization layer, with each quantization layer responsible for encoding the quantization error (namely residual) remaining from the preceding one. In the motion generation stage, the masked Transformer is trained to generate the most informative code sequence (a.k.a. base-layer motion tokens) from the first quantization layer. Using base-layer motion tokens as input, another residual Transformer that has a sim-




Figure 3. **ViMo dataset** contains 32 subcategorized interactions, each belonging to one of three broad categories: **human-human interactions**, **animal-human interactions**, and **scene-human interactions**. Among them, human-human interactions cover daily socializing, sports, physical confrontations, and others.

ilar architecture to the masked Transformer predicts the remaining code sequences (a.k.a. residual-layer tokens). Finally, all the code sequences add up to form one sequence used by the decoder for motion reconstruction.

4. Dataset

We introduce the ViMo dataset, which contains paired video-motion data to support our task setting. Fig. 3 provides some information on the categories.

Video Collection. We mainly acquire RGB videos in two ways: one is to select from publicly available datasets and the other is to utilize video generation models. Specifically, we carefully select 1,335 clips from datasets including UCF101 [89], JPL-Interaction [79], JPL-Interaction-e [80], IDEA400 [56], HMDB51 [50], Ego-Exo4D [24], Harmony4D [43], Ego-Humans [42], etc. The videos we collect adhere to the condition that humans could react to the content if they were seeing it in real life. Long videos are cropped into clips containing only one atomic action. These videos span many categories of human-human interactions, such as daily socializing and sports. However, videos containing certain contents, such as “a tiger roaring towards the camera” and “a car driving fast towards the camera” are not easy to find in these datasets. Therefore, we employ several high-performance text-to-video generation models [1, 7, 20, 49, 65, 83, 91, 108, 114] to gener-

ate the desired videos and manually filter 1390 high-fidelity ones from them. In addition, to further enrich the diversity and quantity of videos, while obtaining videos of characters with specified actions, we use the following approach: The motions of the actors in Inter-X [102] are first rendered as RGB videos from the viewpoint of the reactor. The videos are then converted into a realistic style using Gen-3 Alpha [77], a model with the capability of video-to-video generation. Finally, 775 desired videos are selected from them. In summary, applying the methods described above, we collect a total of 3,500 videos featuring humans, animals, and scenes that allow human interaction with the depicted contents. These videos contain 375.61 k frames, with a total duration of 3.83 hours, which cover 32 subcategories.

Motion Collection and Processing. The collection of motions is similar to that of videos. On the one hand, we select the corresponding motion data of human reactions from datasets such as Inter-X [102], Harmony4D [43] and Ego-Humans [42] based on the collected videos. On the other hand, we employ MoMask [28] to generate the desired motion data from text prompts. A manual post-processing step ensues to filter out abnormal motions. We obtain a total of 2,000 motions in a ratio of 3:1 using the above two approaches. The pose representations of all motions are converted to the format used in HumanML3D [26]. We scale motions to 20 FPS and crop them into the most semantically informative ones, each with a maximum duration of 10 seconds. For data normalization, motions are then retargeted to a default human skeletal template and properly rotated to face the Z+ direction initially as in [26]. As a result, the motions contain a total of 193.61 k frames, equivalent to 2.69 hours, in which the average motion length is 4.84 seconds.

Video-motion Pairs. Since all of our motions are hand-picked human reactions corresponding to the content of the collected videos, we can easily pair them up. That is, for each video in our dataset, we can find one or even more corresponding plausible reactions from our motion set. We manually pair each video with a motion to form our base dataset that contains 3,500 video-motion pairs. We also investigate pairing each video with several motions for training in the Sup. Mat. These paired data cover a wide range of interaction categories, with three broad ones: human-human, animal-human, and scene-human interactions. Among them, human-human interactions consist of daily socializing, sports, physical confrontations, and others. More details on ViMo can be found in the Sup. Mat.

5. Experiments

5.1. Experimental Settings

Dataset Partitioning. For each subcategory of video-motion pairs, we divide them into the training set and the test set with a ratio of 4:1, resulting in a base dataset (ViMo-

Method	FID↓	Diversity→	MModality↑
Real	-	7.954 \pm 0.074	-
MDM [90]	1.688 \pm 0.030	7.385 \pm 0.088	2.117\pm0.064
MLD [12]	1.565 \pm 0.041	7.431 \pm 0.090	2.102 \pm 0.062
T2M-GPT [115]	1.154 \pm 0.038	7.721 \pm 0.081	1.936 \pm 0.032
BAMM [71]	0.930 \pm 0.031	7.619 \pm 0.055	1.885 \pm 0.047
MoMask [28]	0.856 \pm 0.015	7.394 \pm 0.056	1.567 \pm 0.043
HERO	0.427\pm0.014	7.801\pm0.061	1.614 \pm 0.040

Table 1. **Quantitative evaluation on the ViMo test set.** \pm indicates 95% confidence interval, and \rightarrow means the closer to the real motions the better. **Bold** face indicates the best result.




Figure 4. **User study results.** The higher the scores, the better.

base) with 2800 pairs for training and 700 pairs for testing. **Evaluation Metrics.** We employ metrics that are widely adopted in prior works [25, 26, 103] on human motion generation, including: Frechet Inception Distance (FID), diversity, and multimodality.

Implementation Details. During training and testing HERO, the parameters of the video encoder are frozen. We initialize the RVQ-VAE with the weights obtained by pretraining on HumanML3D [26] and then fine-tune it on ViMo for 10 epochs with a batch size of 256. For the Transformers, we train them for 200 epochs with a batch size of 64. Please refer to the Sup. Mat. for more details.

5.2. Comparison to State-of-the-art Approaches

As no prior research has addressed human reaction generation from videos, we reimplement several existing state-of-the-art methods of human-centric generative models as baselines. These methods range from diffusion-based models [12, 90], autoregressive model [115], to masked generative models [28, 71]. For the sake of fairness, all of their signal encoders are replaced with the video encoder of TC-CLIP [45], and all of their VAE/VQ-VAEs are initialized with the weights obtained by pretraining on HumanML3D.

As reported in Tab. 1, HERO outperforms the baselines on FID and diversity, suggesting that it learns the reaction patterns well. Although multimodality is important, [28] highlights its role as a secondary measure which should be




Figure 5. **Visual comparisons** between the different methods given three distinct videos from ViMo test set.

Method	FID↓	Diversity→	MModality↑
Real	-	7.954 ± 0.074	-
w/o IIE	0.535 ± 0.017	7.671 ± 0.056	1.626 ± 0.042
w/o DIE	0.521 ± 0.011	7.769 ± 0.068	1.443 ± 0.052
HERO	0.427 ± 0.014	7.801 ± 0.061	1.614 ± 0.040

Table 2. **Quantitative ablation studies.** w/o means without. IIE and DIE represent Interaction Intention Extraction and Dynamic Information Exploitation, respectively.




Figure 6. **Qualitative ablation studies** given the same video. A plausible reaction would be to defend against the shooting.

evaluated alongside primary performance metrics such as FID. The qualitative comparisons in Fig. 5 show that HERO is more capable of generating plausible reactions than other methods.

User Study. We ask participants to rate (1 to 5, the higher the better) the motions they observe in terms of the quality of the reactive motion (motion quality), and the plausibility of the reaction given the video (reaction plausibility) as in [21]. For each method, 50 reactions are generated using the same video pool from the ViMo test set. We report the average scores from the feedback of 40 participants in Fig. 4, excluding the feedback that did not pass our validation checks. In both metrics surveyed, HERO-generated reactions receive the highest scores, except for GT.

5.3. Ablation Study

We can observe from Tab. 2 that both interaction intention extraction (IIE) and dynamic information exploitation (DIE) effectively bring the generated reactions closer to the real distribution. Fig. 6 illustrates that the absence of IIE may produce an implausible reaction, while the absence of DIE leads to a decrease in motion quality.

5.4. Performance Analysis

Different Broad Categories of Interactions. Fig. 5 show the reactions synthesized by HERO in different broad categories of interactions: human-human, animal-human and




Figure 7. **Visualized cases of generating on the Unseen set.** Note that the data of the subcategories to which each video in the figure belongs are not utilized when training HERO on the **Seen** set.



Figure 8. **Visualized cases of “walking towards” with different emotions.** The left side of each row in the figure is the input video and the right side is the reaction generated by HERO.

scene-human interactions. The visualization results validate the effectiveness of HERO in different broad categories of interactions. The quantitative results are presented in the Sup. Mat.

The Same Action with Different Emotions. As demonstrated in Fig. 8, HERO can generate different plausible reactions depending on the subjects’ different emotions, even if the subjects’ actions are the same.

Generalization Ability. We further divide the 3500 pairs in the base dataset into two parts, **Seen** and **Unseen**. The **Seen** set covers 26 subcategories of interactions, and the **Unseen**

set contains the remaining 6 subcategories. To verify the generalization ability of HERO, we train it on the **Seen** set and use it to generate reactions on the **Unseen** set. The visualized cases in Fig. 7 show that HERO does have a certain generalization ability for unseen video categories. More experimental results can be found in the Sup. Mat.

6. Conclusion

In this paper, we introduce a new task focused on generating 3D human reactions from RGB videos with interactivity. With the collected ViMo dataset containing video-motion pairs, we train HERO, a model that extracts interaction intention from input video and uses it to guide the generation of the plausible human reaction. Extensive experiments demonstrate the effectiveness and superiority of HERO. We hope that this work can contribute new insights into 3D human motion generation involving interaction.

Limitations and Future Work. Currently, the categories and volume of video-motion pairs still need to be expanded, and our framework is as simple as possible to verify the feasibility of the task. In the future, we plan to integrate the bidirectional autoregressive approach [71] into our framework so that the model can predict the length of the reaction by itself. Also, we can incorporate physical constraints [111] into the pipeline to eliminate artifacts such as floating and foot sliding, often observed with human motion generation models.

Acknowledgments. This work is supported by the National Natural Science Foundation of China (NSFC) under Grants 62306295, 62225207 and 62436008.

References

- [1] AISphere. Pixverse-v3. <https://app.pixverse.ai>, 2024.11. 5, 1
- [2] Sadegh Aliakbarian, Pashmina Cameron, Federica Bogo, Andrew Fitzgibbon, and Thomas J Cashman. Flag: Flow-based 3d avatar generation from sparse observations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13253–13262, 2022. 2
- [3] Tenglong Ao, Qingzhe Gao, Yuke Lou, Baoquan Chen, and Libin Liu. Rhythmic gesticulator: Rhythm-aware co-speech gesture synthesis with hierarchical neural embeddings. *ACM Transactions on Graphics (TOG)*, 41(6):1–19, 2022. 2
- [4] Tenglong Ao, Zeyi Zhang, and Libin Liu. Gesturediffuclip: Gesture diffusion model with clip latents. *ACM Transactions on Graphics (TOG)*, 42(4):1–18, 2023. 2
- [5] Andreas Aristidou, Anastasios Yiannakidis, Kfir Aberman, Daniel Cohen-Or, Ariel Shamir, and Yiorgos Chrysanthou. Rhythm is a dancer: Music-driven motion synthesis with global structure. *IEEE Transactions on Visualization and Computer Graphics*, 29(8):3519–3534, 2022. 2
- [6] Nikos Athanasiou, Mathis Petrovich, Michael J Black, and Gül Varol. Teach: Temporal action composition for 3d humans. In *2022 International Conference on 3D Vision (3DV)*, pages 414–423. IEEE, 2022. 2
- [7] ByteDance. Jimeng video. <https://jimeng.jiaying.com>, 2024.06. 5, 1
- [8] Angela Castillo, Maria Escobar, Guillaume Jeanneret, Albert Pumarola, Pablo Arbeláez, Ali Thabet, and Artsiom Sanakoyeu. Bodiffusion: Diffusing sparse observations for full-body human motion synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4221–4231, 2023. 2
- [9] Pablo Cervantes, Yusuke Sekikawa, Ikuro Sato, and Koichi Shinoda. Implicit neural representations for variable length human motion generation. In *European Conference on Computer Vision*, pages 356–372. Springer, 2022. 2
- [10] Junhao Chen, Mingjin Chen, Jianjin Xu, Xiang Li, Juntong Dong, Mingze Sun, Puhua Jiang, Hongxiang Li, Yuhang Yang, Hao Zhao, et al. Dancetogether! identity-preserving multi-person interactive video generation. *arXiv preprint arXiv:2505.18078*, 2025. 2
- [11] Ling-Hao Chen, Shunlin Lu, Ailing Zeng, Hao Zhang, Benyou Wang, Ruimao Zhang, and Lei Zhang. Motion-llm: Understanding human behaviors from human motions and videos. *arXiv preprint arXiv:2405.20340*, 2024. 4
- [12] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18000–18010, 2023. 2, 6, 1
- [13] Baptiste Chopin, Hao Tang, Naima Otherdout, Mohamed Daoudi, and Nicu Sebe. Interaction transformer for human reaction generation. *IEEE Transactions on Multimedia*, 25: 8842–8854, 2023. 1, 2, 3
- [14] Rishabh Dabral, Muhammad Hamza Mughal, Vladislav Golyanik, and Christian Theobalt. Mofusion: A framework for denoising-diffusion-based motion synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9760–9770, 2023. 2
- [15] Wenxun Dai, Ling-Hao Chen, Jingbo Wang, Jinpeng Liu, Bo Dai, and Yansong Tang. Motionlcm: Real-time controllable motion generation via latent consistency model. In *European Conference on Computer Vision*, pages 390–408. Springer, 2024. 2
- [16] Yuming Du, Robin Kips, Albert Pumarola, Sebastian Starke, Ali Thabet, and Artsiom Sanakoyeu. Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 481–490, 2023. 2
- [17] Ke Fan, Junshu Tang, Weijian Cao, Ran Yi, Moran Li, Jingyu Gong, Jiangning Zhang, Yabiao Wang, Chengjie Wang, and Lizhuang Ma. Freemotion: A unified framework for number-free text-to-motion synthesis. In *European Conference on Computer Vision*, pages 93–109. Springer, 2024. 2
- [18] Ke Fan, Shunlin Lu, Minyue Dai, Runyi Yu, Lixing Xiao, Zhiyang Dou, Juntong Dong, Lizhuang Ma, and Jingbo Wang. Go to zero: Towards zero-shot motion generation with million-scale data. *arXiv preprint arXiv:2507.07095*, 2025. 2
- [19] Mihai Fieraru, Mihai Zanfir, Elisabeta Oneata, Alin-Ionut Popa, Vlad Olaru, and Cristian Sminchisescu. Three-dimensional reconstruction of human interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7214–7223, 2020. 3
- [20] Genmo. Mochi 1. <https://www.genmo.ai/play>, 2024.11. 5, 1
- [21] Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. Remos: 3d motion-conditioned reaction synthesis for two-person interactions. In *European Conference on Computer Vision (ECCV)*, page 3, 2024. 1, 2, 3, 7
- [22] Kehong Gong, Dongze Lian, Heng Chang, Chuan Guo, Zihang Jiang, Xinxin Zuo, Michael Bi Mi, and Xinchao Wang. Tm2d: Bimodality driven 3d dance generation via music-text integration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9942–9952, 2023. 2, 3
- [23] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [24] Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, et al. Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19383–19400, 2024. 5

- [25] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2021–2029, 2020. 2, 6
- [26] 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. 2, 6, 1, 3
- [27] 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 *European Conference on Computer Vision*, pages 580–597. Springer, 2022. 2, 3
- [28] Chuan Guo, Yuxuan Mu, Muhammad Gohar Javed, Sen Wang, and Li Cheng. Momask: Generative masked modeling of 3d human motions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1900–1910, 2024. 3, 4, 5, 6, 1
- [29] Wen Guo, Xiaoyu Bie, Xavier Alameda-Pineda, and Francesc Moreno-Noguer. Multi-person extreme motion prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13053–13064, 2022. 3
- [30] Ikhsanul Habibie, Mohamed Elgharib, Kripasindhu Sarkar, Ahsan Abdullah, Simbarashe Nyatsanga, Michael Neff, and Christian Theobalt. A motion matching-based framework for controllable gesture synthesis from speech. In *ACM SIGGRAPH 2022 conference proceedings*, pages 1–9, 2022. 2
- [31] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2
- [32] Tao Hu, Xinyan Zhu, Wei Guo, and Kehua Su. Efficient interaction recognition through positive action representation. *Mathematical Problems in Engineering*, 2013(1):795360, 2013. 3
- [33] Yinghao Huang, Manuel Kaufmann, Emre Aksan, Michael J Black, Otmar Hilliges, and Gerard Pons-Moll. Deep inertial poser: Learning to reconstruct human pose from sparse inertial measurements in real time. *ACM Transactions on Graphics (TOG)*, 37(6):1–15, 2018. 2
- [34] Yiheng Huang, Hui Yang, Chuanchen Luo, Yuxi Wang, Shibiao Xu, Zhaoxiang Zhang, Man Zhang, and Junran Peng. Stablemofusion: Towards robust and efficient diffusion-based motion generation framework. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 224–232, 2024. 2
- [35] Muhammad Gohar Javed, Xingyu Li, et al. Intermask: 3d human interaction generation via collaborative masked modelling. In *The Thirteenth International Conference on Learning Representations*, 2025. 3
- [36] Eric Jensen. *Brain-Based Learning: Teaching the Way Students Really Learn*. Turning Point Publishing, Del Mar, CA, 1996. 4
- [37] Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. Motiongpt: Human motion as a foreign language. *Advances in Neural Information Processing Systems*, 36:20067–20079, 2023. 3
- [38] Biao Jiang, Xin Chen, Chi Zhang, Fukun Yin, Zhuoyuan Li, Gang Yu, and Jiayuan Fan. Motionchain: Conversational motion controllers via multimodal prompts. In *European Conference on Computer Vision*, pages 54–74. Springer, 2024. 4
- [39] Jiayi Jiang, Paul Strel, Huajian Qiu, Andreas Fender, Larissa Laich, Patrick Snape, and Christian Holz. Avatarposer: Articulated full-body pose tracking from sparse motion sensing. In *European conference on computer vision*, pages 443–460. Springer, 2022. 2
- [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. 4
- [41] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 3
- [42] Rawal Khirodkar, Aayush Bansal, Lingni Ma, Richard Newcombe, Minh Vo, and Kris Kitani. Ego-humans: An ego-centric 3d multi-human benchmark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19807–19819, 2023. 3, 5, 6, 1
- [43] Rawal Khirodkar, Jyun-Ting Song, Jinkun Cao, Zhengyi Luo, and Kris M Kitani. Harmony4d: A video dataset for in-the-wild close human interactions. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. 3, 5, 6, 1
- [44] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. Flame: Free-form language-based motion synthesis & editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8255–8263, 2023. 2
- [45] Minji Kim, Dongyoon Han, Taekyung Kim, and Bohyung Han. Leveraging temporal contextualization for video action recognition. In *European Conference on Computer Vision*, pages 74–91. Springer, 2024. 2, 3, 6
- [46] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [47] Woo-Ri Ko, Minsu Jang, Jaeyeon Lee, and Jaehong Kim. Air-act2act: Human–human interaction dataset for teaching non-verbal social behaviors to robots. *The International Journal of Robotics Research*, 40(4-5):691–697, 2021. 3
- [48] Hanyang Kong, Kehong Gong, Dongze Lian, Michael Bi Mi, and Xinchao Wang. Priority-centric human motion generation in discrete latent space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14806–14816, 2023. 2
- [49] Kuaishou. Kling ai. <https://klingai.kuaishou.com/>, 2024.06. 5, 1
- [50] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *2011 Interna-*

- tional conference on computer vision, pages 2556–2563. IEEE, 2011. 5
- [51] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. Dancing to music. *Advances in neural information processing systems*, 32, 2019. 2
- [52] Buyu Li, Yongchi Zhao, Shi Zhelun, and Lu Sheng. Danceformer: Music conditioned 3d dance generation with parametric motion transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1272–1279, 2022. 2
- [53] Hongxiang Li, Yaowei Li, Yuhang Yang, Junjie Cao, Zhihong Zhu, Xuxin Cheng, and Long Chen. Dispose: Disentangling pose guidance for controllable human image animation. *arXiv preprint arXiv:2412.09349*, 2024. 2
- [54] Han Liang, Wenqian Zhang, Wenxuan Li, Jingyi Yu, and Lan Xu. Intergen: Diffusion-based multi-human motion generation under complex interactions. *International Journal of Computer Vision*, pages 1–21, 2024. 3, 1
- [55] Bin Lin, Yang Ye, Bin Zhu, Jiaxi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5971–5984, 2024. 4
- [56] Jing Lin, Ailing Zeng, Shunlin Lu, Yuanhao Cai, Ruimao Zhang, Haoqian Wang, and Lei Zhang. Motion-x: A large-scale 3d expressive whole-body human motion dataset. *Advances in Neural Information Processing Systems*, 36: 25268–25280, 2023. 5
- [57] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C Kot. Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2684–2701, 2019. 3
- [58] Yunze Liu, Changxi Chen, and Li Yi. Interactive humanoid: Online full-body motion reaction synthesis with social affordance canonicalization and forecasting. *arXiv preprint arXiv:2312.08983*, 2023. 1, 2, 3
- [59] Yunze Liu, Changxi Chen, Chenjing Ding, and Li Yi. Phys-reaction: Physically plausible real-time humanoid reaction synthesis via forward dynamics guided 4d imitation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 3771–3780, 2024. 1, 2, 3
- [60] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 3
- [61] Shunlin Lu, Ling-Hao Chen, Ailing Zeng, Jing Lin, Ruimao Zhang, Lei Zhang, and Heung-Yeung Shum. Human-tomato: Text-aligned whole-body motion generation. In *International Conference on Machine Learning*, pages 32939–32977. PMLR, 2024. 4
- [62] Thomas Lucas, Fabien Baradel, Philippe Weinzaepfel, and Grégory Rogez. Posegpt: Quantization-based 3d human motion generation and forecasting. In *European Conference on Computer Vision*, pages 417–435. Springer, 2022. 2
- [63] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12585–12602, 2024. 4
- [64] Marcin Marszalek, Ivan Laptev, and Cordelia Schmid. Actions in context. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2929–2936. IEEE, 2009. 3
- [65] MiniMax. Hailuo ai. <https://hailuoai.com/video>, 2024.09. 5, 1
- [66] Evonne Ng, Donglai Xiang, Hanbyul Joo, and Kristen Grauman. You2me: Inferring body pose in egocentric video via first and second person interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9890–9900, 2020. 3
- [67] Alonso Patron-Perez, Marcin Marszalek, Andrew Zisserman, and Ian Reid. High five: Recognising human interactions in tv shows. In *BMVC*, page 33, 2010. 3
- [68] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10985–10995, 2021. 2
- [69] Mathis Petrovich, Michael J Black, and Gül Varol. Temos: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision*, pages 480–497. Springer, 2022. 2
- [70] Ekkasit Pinyoanuntapong, Muhammad Usama Saleem, Korrave Karunratanakul, Pu Wang, Hongfei Xue, Chen Chen, Chuan Guo, Junli Cao, Jian Ren, and Sergey Tulyakov. Controlmm: Controllable masked motion generation. *arXiv preprint arXiv:2410.10780*, 2024. 3
- [71] Ekkasit Pinyoanuntapong, Muhammad Usama Saleem, Pu Wang, Minwoo Lee, Srijan Das, and Chen Chen. Bamm: bidirectional autoregressive motion model. In *European Conference on Computer Vision*, pages 172–190. Springer, 2024. 4, 5, 6, 8, 1
- [72] Ekkasit Pinyoanuntapong, Pu Wang, Minwoo Lee, and Chen Chen. Mmm: Generative masked motion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1546–1555, 2024. 3, 4
- [73] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3
- [74] Hanoona Rasheed, Muhammad Uzair Khattak, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Fine-tuned clip models are efficient video learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6545–6554, 2023. 2
- [75] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015. 2
- [76] Lawrence D. Rosenblum. *See What I'm Saying: The Extraordinary Powers of Our Five Senses*. W. W. Norton & Company, New York, 2010. 4

- [77] Runway. Gen-3 alpha. <https://runwayml.com>, 2024.06. 6, 1
- [78] Michael S Ryoo and JK Aggarwal. Ut-interaction dataset, icpr contest on semantic description of human activities (sdha). In *IEEE International Conference on Pattern Recognition Workshops*, page 4, 2010. 3
- [79] Michael S Ryoo and Larry Matthies. First-person activity recognition: What are they doing to me? In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2730–2737, 2013. 3, 5
- [80] Michael S Ryoo, Thomas J Fuchs, Lu Xia, Jake K Aggarwal, and Larry Matthies. Robot-centric activity prediction from first-person videos: What will they do to me? In *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction*, pages 295–302, 2015. 3, 5
- [81] Yoni Shafir, Guy Tevet, Roy Kapon, and Amit Haim Bermano. Human motion diffusion as a generative prior. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [82] Yawen Shao, Wei Zhai, Yuhang Yang, Hongchen Luo, Yang Cao, and Zheng-Jun Zha. Great: Geometry-intention collaborative inference for open-vocabulary 3d object affordance grounding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17326–17336, 2025. 3
- [83] ShengShu-AI. Vidu 1.5. <https://www.vidu.studio/>, 2024.11. 5, 1
- [84] Li Siyao, Weijiang Yu, Tianpei Gu, Chunze Lin, Quan Wang, Chen Qian, Chen Change Loy, and Ziwei Liu. Bailando: 3d dance generation by actor-critic gpt with choreographic memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11050–11059, 2022. 2
- [85] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015. 2
- [86] Siyang Song, Micol Spitale, Cheng Luo, German Barquero, Cristina Palmero, Sergio Escalera, Michel Valstar, Tobias Baur, Fabien Ringeval, Elisabeth Andre, et al. React2023: The first multiple appropriate facial reaction generation challenge. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 9620–9624, 2023. 2
- [87] Siyang Song, Micol Spitale, Cheng Luo, Cristina Palmero, German Barquero, Hengde Zhu, Sergio Escalera, Michel Valstar, Tobias Baur, Fabien Ringeval, et al. React 2024: the second multiple appropriate facial reaction generation challenge. In *2024 IEEE 18th International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–5. IEEE, 2024.
- [88] Siyang Song, Micol Spitale, Xiangyu Kong, Hengde Zhu, Cheng Luo, Cristina Palmero, German Barquero, Sergio Escalera, Michel Valstar, Mohamed Daoudi, et al. React 2025: the third multiple appropriate facial reaction generation challenge. *arXiv preprint arXiv:2505.17223*, 2025. 2
- [89] K Soomro. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 5
- [90] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 6, 1
- [91] Ali Tongyi. Wanxiang video. <https://tongyi.aliyun.com/wanxiang/videoCreation>, 2024.09. 5, 1
- [92] Jonathan Tseng, Rodrigo Castellon, and Karen Liu. Edge: Editable dance generation from music. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 448–458, 2023. 2
- [93] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 3
- [94] NP Van der Aa, Xinghan Luo, Geert-Jan Giezeman, Robby T Tan, and Remco C Veltkamp. Umpm benchmark: A multi-person dataset with synchronized video and motion capture data for evaluation of articulated human motion and interaction. In *2011 IEEE international conference on computer vision workshops (ICCV Workshops)*, pages 1264–1269. IEEE, 2011. 3
- [95] Coert Van Gemeren, Ronald Poppe, and Remco C Veltkamp. Spatio-temporal detection of fine-grained dyadic human interactions. In *Human Behavior Understanding: 7th International Workshop, HBU 2016, Amsterdam, The Netherlands, October 16, 2016, Proceedings 7*, pages 116–133. Springer, 2016. 3
- [96] Mengmeng Wang, Jiazheng Xing, and Yong Liu. Action-clip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*, 2021. 2, 3
- [97] Zhenyi Wang, Ping Yu, Yang Zhao, Ruiyi Zhang, Yufan Zhou, Junsong Yuan, and Changyou Chen. Learning diverse stochastic human-action generators by learning smooth latent transitions. In *Proceedings of the AAAI conference on artificial intelligence*, pages 12281–12288, 2020. 2
- [98] Lixing Xiao, Shunlin Lu, Huaijin Pi, Ke Fan, Liang Pan, Yueer Zhou, Ziyong Feng, Xiaowei Zhou, Sida Peng, and Jingbo Wang. Motionstreamer: Streaming motion generation via diffusion-based autoregressive model in causal latent space. *arXiv preprint arXiv:2503.15451*, 2025. 2
- [99] Xianghui Xie, Bharat Lal Bhatnagar, Jan Eric Lenssen, and Gerard Pons-Moll. Template free reconstruction of human-object interaction with procedural interaction generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10003–10015, 2024. 1
- [100] Xianghui Xie, Jan Eric Lenssen, and Gerard Pons-Moll. Intertrack: Tracking human object interaction without object templates. In *International Conference on 3D Vision*. IEEE, 2025. 1
- [101] Liang Xu, Ziyang Song, Dongliang Wang, Jing Su, Zhicheng Fang, Chenjing Ding, Weihao Gan, Yichao Yan, Xin Jin, Xiaokang Yang, et al. Actformer: A gan-based

- transformer towards general action-conditioned 3d human motion generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2228–2238, 2023. 2
- [102] Liang Xu, Xintao Lv, Yichao Yan, Xin Jin, Shuwen Wu, Congsheng Xu, Yifan Liu, Yizhou Zhou, Fengyun Rao, Xingdong Sheng, et al. Inter-x: Towards versatile human-human interaction analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22260–22271, 2024. 3, 6, 1
- [103] Liang Xu, Yizhou Zhou, Yichao Yan, Xin Jin, Wenhan Zhu, Fengyun Rao, Xiaokang Yang, and Wenjun Zeng. Regenet: Towards human action-reaction synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1759–1769, 2024. 1, 2, 3, 6
- [104] Sijie Yan, Zhizhong Li, Yuanjun Xiong, Huahan Yan, and Dahua Lin. Convolutional sequence generation for skeleton-based action synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4394–4402, 2019. 2
- [105] Yuhang Yang, Wei Zhai, Hongchen Luo, Yang Cao, Jiebo Luo, and Zheng-Jun Zha. Grounding 3d object affordance from 2d interactions in images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10905–10915, 2023. 3
- [106] Yuhang Yang, Wei Zhai, Hongchen Luo, Yang Cao, and Zheng-Jun Zha. Lemon: Learning 3d human-object interaction relation from 2d images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16284–16295, 2024.
- [107] Yuhang Yang, Wei Zhai, Chengfeng Wang, Chengjun Yu, Yang Cao, and Zheng-Jun Zha. Egochoir: Capturing 3d human-object interaction regions from egocentric views. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 3
- [108] Yuhang Yang, Ke Fan, Shangkun Sun, Hongxiang Li, Ailing Zeng, FeiLin Han, Wei Zhai, Wei Liu, Yang Cao, and Zheng-Jun Zha. Videogen-eval: Agent-based system for video generation evaluation. *arXiv preprint arXiv:2503.23452*, 2025. 5
- [109] Yuhang Yang, Fengqi Liu, Yixing Lu, Qin Zhao, Pingyu Wu, Wei Zhai, Ran Yi, Yang Cao, Lizhuang Ma, Zheng-Jun Zha, et al. Sigman: Scaling 3d human gaussian generation with millions of assets. *arXiv preprint arXiv:2504.06982*, 2025. 2
- [110] Yifei Yin, Chen Guo, Manuel Kaufmann, Juan Jose Zarate, Jie Song, and Otmar Hilliges. Hi4d: 4d instance segmentation of close human interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17016–17027, 2023. 3
- [111] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16010–16021, 2023. 2, 8
- [112] Kiwon Yun, Jean Honorio, Debaleena Chattopadhyay, Tamara L Berg, and Dimitris Samaras. Two-person interaction detection using body-pose features and multiple instance learning. In *2012 IEEE computer society conference on computer vision and pattern recognition workshops*, pages 28–35. IEEE, 2012. 3
- [113] Neil Zeghidour, Alejandro Luebs, Ahmed Omran, Jan Skoglund, and Marco Tagliasacchi. Soundstream: An end-to-end neural audio codec. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:495–507, 2021. 5
- [114] Ailing Zeng, Yuhang Yang, Weidong Chen, and Wei Liu. The dawn of video generation: Preliminary explorations with sora-like models. *arXiv preprint arXiv:2410.05227*, 2024. 5
- [115] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Yong Zhang, Hongwei Zhao, Hongtao Lu, Xi Shen, and Ying Shan. Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14730–14740, 2023. 2, 3, 4, 6, 1
- [116] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *IEEE transactions on pattern analysis and machine intelligence*, 46(6):4115–4128, 2024. 2
- [117] Yaqi Zhang, Di Huang, Bin Liu, Shixiang Tang, Yan Lu, Lu Chen, Lei Bai, Qi Chu, Nenghai Yu, and Wanli Ouyang. Motiongpt: Finetuned llms are general-purpose motion generators. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7368–7376, 2024. 3
- [118] Zichen Zhang, Hongchen Luo, Wei Zhai, Yang Cao, and Yu Kang. Pear: Phrase-based hand-object interaction anticipation. *Science China Information Sciences*, 68(5):1–20, 2025. 3
- [119] Chongyang Zhong, Lei Hu, Zihao Zhang, and Shihong Xia. Att2m: Text-driven human motion generation with multi-perspective attention mechanism. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 509–519, 2023. 3
- [120] Wenyang Zhou, Zhiyang Dou, Zeyu Cao, Zhouyingcheng Liao, Jingbo Wang, Wenjia Wang, Yuan Liu, Taku Komura, Wenping Wang, and Lingjie Liu. Emdm: Efficient motion diffusion model for fast and high-quality motion generation. In *European Conference on Computer Vision*, pages 18–38. Springer, 2024. 2
- [121] Zixiang Zhou and Baoyuan Wang. Ude: A unified driving engine for human motion generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5632–5641, 2023. 2
- [122] Lingting Zhu, Xian Liu, Xuanyu Liu, Rui Qian, Ziwei Liu, and Lequan Yu. Taming diffusion models for audio-driven co-speech gesture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10544–10553, 2023. 2
- [123] Qiran Zou, Shangyuan Yuan, Shian Du, Yu Wang, Chang Liu, Yi Xu, Jie Chen, and Xiangyang Ji. Parco: Part-coordinating text-to-motion synthesis. In *European Conference on Computer Vision*, pages 126–143. Springer, 2024. 3