

MotionStreamer: Streaming Motion Generation via Diffusion-based Autoregressive Model in Causal Latent Space

Lixing Xiao¹ Shunlin Lu² Huaijin Pi³ Ke Fan⁴ Liang Pan³
 Yueer Zhou¹ Ziyong Feng⁵ Xiaowei Zhou¹ Sida Peng^{1†} Jingbo Wang⁶

¹Zhejiang University ²The Chinese University of Hong Kong (Shenzhen)

³The University of Hong Kong ⁴Shanghai Jiao Tong University

⁵DeepGlint

⁶Shanghai AI Laboratory

lixingxiao0@gmail.com, pengsida@zju.edu.cn

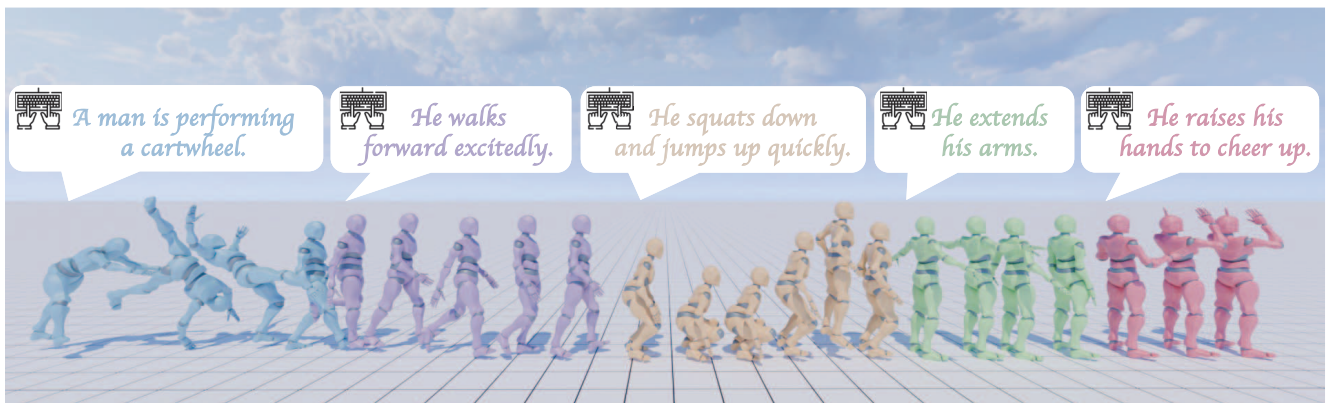


Figure 1. Visualization of streaming motion generation process. Texts are incrementally inputted and motions are generated online.

Abstract

This paper addresses the challenge of text-conditioned streaming motion generation, which requires us to predict the next-step human pose based on variable-length historical motions and incoming texts. Existing methods struggle to achieve streaming motion generation, e.g., diffusion models are constrained by pre-defined motion lengths, while GPT-based methods suffer from delayed response and error accumulation problem due to discretized non-causal tokenization. To solve these problems, we propose MotionStreamer, a novel framework that incorporates a continuous causal latent space into a probabilistic autoregressive model. The continuous latents mitigate information loss caused by discretization and effectively reduce error accumulation during long-term autoregressive generation. In addition, by establishing temporal causal dependencies between current and historical motion latents, our model fully utilizes the available information to achieve accurate online motion decoding. Experiments show that our method outperforms existing approaches while offering more applications, including multi-

round generation, long-term generation, and dynamic motion composition. Project Page: <https://zju3dv.github.io/MotionStreamer/>

1. Introduction

Streaming motion generation aims to incrementally synthesizing human motions while dynamically adapting to online text inputs and maintaining semantic coherence. Generating realistic and diverse human motions in a streaming manner is essential for various real-time applications, such as video games, animation, and robotics. Streaming motion generation presents a significant challenge due to two fundamental requirements. Firstly, the framework must incrementally process sequentially arriving textual inputs while maintaining online response. Secondly, the model should be able to continuously synthesize motion sequences that exhibit contextual consistency by effectively integrating historical information with incoming textual conditions, ensuring alignment between progressive text semantics and kinematic continuity across extended timelines.

Conventional diffusion-based motion generation models

[†] Corresponding author

[4, 12, 57] are constrained by their non-incremental generation paradigm with static text conditioning and fixed-length generation processes. This inherently limits their ability to dynamically evolving textual inputs in streaming scenarios. Other autoregressive motion generation frameworks [28, 68] are able to generate motions in a streaming manner. However, they have difficulties in achieving online response due to their non-causal tokenization architecture, which prevents partial token decoding until all sequence is generated. Real-time motion generation methods like DART [72] face a critical limitation in their reliance on fixed-window local motion primitives, which inherently restricts their capacity to model variable-length historical contexts and dynamically align with evolving textual inputs.

In this work, we propose a novel framework for streaming motion generation, named **MotionStreamer**. The visualization of the streaming generation process is illustrated in Fig. 1. Our core innovation is incorporating a diffusion head into an autoregressive model to predict the next motion latent, while introducing a causal motion compressor to enable online decoding in a streaming manner. Specifically, given an input text, we extract the textual feature, combine it with historical motion latents, and use an autoregressive model to generate a condition feature, which guides a diffusion model to generate the next motion latent. In contrast to previous methods [28, 68] that use vector quantization (VQ) based motion tokenizer and GPT architecture to generate discrete motion tokens, our continuous motion latents can avoid information loss of discrete tokens and accumulation of error during the streaming generation process, as demonstrated by our experimental results in Sec. 4.2.

A causal temporal AutoEncoder is then employed to convert motion latent into the next human pose. The causal network effectively establishes temporally causal dependencies between current and historical motion latents, allowing for online motion decoding. The key to achieving streaming generation is enabling the model to dynamically extract relevant information from a variable-length history to guide the next motion prediction. To enable the autoregressive model to self-terminate without a pre-defined sequence length, we additionally encode an “impossible pose” to get a reference end latent as the continuous stopping condition.

During experiments, we found that naive training of our model still suffers from error accumulation and cannot well support multi-round text input. To address these issues, we propose two training strategies: Two-forward training and Mixed training. Two-forward training strategy first generates motion latents using ground-truth, then replaces partial ground-truth latents with first-forward predictions for a hybrid second-forward, effectively mitigating the exposure bias inherent in autoregressive training while preserving parallel efficiency. Mixed training strategy unifies atomic (text, motion) pairs and contextual (text, history motion,

current motion) triplets in a single framework, enabling compositional semantics learning and generalization to unseen motion combinations.

We evaluate our approach on the HumanML3D [19] and BABEL [48] datasets, which are widely-used for text-to-motion benchmarks. Across these datasets, our method achieves state-of-the-art performance on both text-to-motion and long-term motion synthesis tasks. We also demonstrate the superiority of our method on abundant applications. Such streaming generation framework is suitable for online multi-round generation with progressive text inputs, long-term motion generation with multiple texts provided and dynamic motion composition where subsequent motions can be regenerated by altering textual conditions while preserving the initially generated prefix motion.

Overall, our contributions can be summarized as follows:

- We propose MotionStreamer, a novel framework combining a diffusion head with an autoregressive model to directly predict continuous motion latents, which enables streaming motion generation with incremental text inputs.
- We propose a causal motion compressor (Causal TAE) for continuous motion compression, which eliminates information loss from discrete quantization and establishes temporally causal latent dependencies to support streaming decoding and online response. We adopt Two-Forward training strategy to mitigate error accumulation in streaming generation scenarios.
- We demonstrate great performances of our approach on benchmark datasets. We also show various downstream applications, including online multi-round generation, long-term generation and dynamic motion composition.

2. Related Work

Text-conditioned motion generation. Text-conditioned motion generation aims to synthesize 3D human motions from natural language descriptions [19]. Various approaches [8–10, 25, 26, 32, 39, 40, 43–45, 61–63, 65, 67, 71] have explored diverse architectures and methodologies to enhance the naturalness and expressiveness for digital human and motion animation. Previous works [2, 19, 41] leverage VAE [29] to learn cross-modal mappings between text and motion spaces. Some works [3, 4, 12, 13, 16, 27, 30, 53, 57, 60, 64, 69] also apply diffusion models [24] to this task. Another line of works [18, 28, 35, 36, 68, 70, 73, 74] first discretize motions into discrete tokens and employ autoregressive models (e.g., GPT) for sequential token prediction. Furthermore, some approaches [17, 21, 38, 46, 47] adopt a BERT-style [14] bidirectional Transformer architecture [7] to reconstruct masked motion segments under text guidance. However, most existing works only focus on offline generation, where the entire motion sequence is generated at once. More recently, CAMDM [11] and AMDM

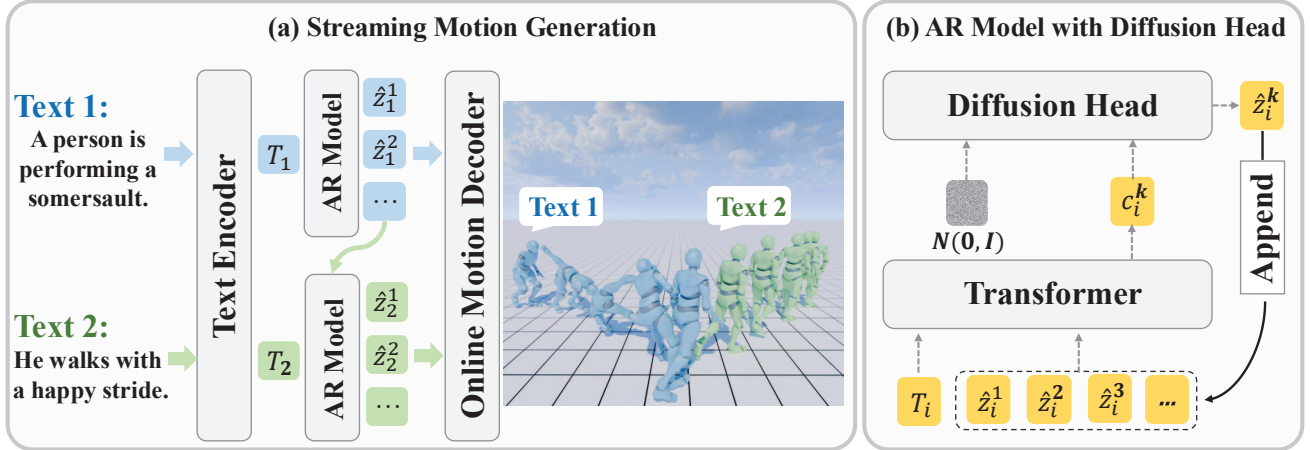


Figure 2. **Overview of MotionStreamer.** During inference, the AR model streamingly predicts next motion latents conditioned on the current text and previous motion latents. Each latent can be decoded into motion frames online as soon as it is generated.

[54] apply diffusion models into an autoregressive manner for real-time interactive character control. Ready-to-React [6] further explores this idea in two-character interaction. CLoSD [58] and DART [72] apply it for real-time text driven motion control. However, these methods are not strictly causal, as they rely on a fixed-length context window while our method could handle variable-length historical information and incrementally generate motions in a streaming manner.

Motion compression. Following the success of image generation [51], previous works [12, 68] first encode the raw motion sequences into a latent space and then generate motions within it. The most popular method is to use Vector Quantized Variational AutoEncoders (VQ-VAE) [59] for motion tokenization. TM2T [20] first introduces vector quantization for motion discretization. T2M-GPT [68] employs VQ-VAE to compress motion sequences into a discrete latent space and then uses a GPT for motion generation. MoMask [21] leverages Residual VQ-VAE (RVQ-VAE) [31] to progressively reduce quantization errors. In contrast, MLD [12] utilizes standard VAEs [29] to convert a motion sequence into an embedding and then use a diffusion model to generate the latent. However, existing methods require a whole motion sequence to be encoded and decoded, which is not suitable for streaming generation. In this paper, we propose a causal motion compression approach to achieve streaming motion generation with online response.

3. Method

We address the task of streaming motion generation with online response by introducing a novel framework, named MotionStreamer. The overview of the proposed framework is illustrated in Fig. 2. In section 3.1, we first introduce the problem formulation of streaming motion generation

and the motion representation used in this work. In section 3.2, we introduce a Causal Temporal AutoEncoder for continuous motion compression and online decoding. In section 3.3, we present a diffusion-based autoregressive motion generator and the streaming generation process.

3.1. Problem Formulation

Task definition. We first introduce the formulation of streaming motion generation. In contrast to previous text-to-motion generation [19] which is conditioned on a predefined fixed text prompt, we consider the case where a series of text prompts are given sequentially. Given a streaming sequence of text prompts $\{\mathcal{P}_i\}_{i=1}^M$, the goal is to generate a sequence of motion frames $\{x_j\}_{j=1}^N$ online, where \mathcal{P}_i is the i -th text prompt and x_j is the j -th frame pose.

Motion Representation. Previous works [19, 21, 28, 57, 70] mainly uses the 263-dimensional pose representation [19] for motion generation. However, this representation requires an additional post-processing step [5], which is time-consuming and introduces rotation error [15] to be converted to SMPL [34] body parameters. To overcome this issue, we slightly modify it and directly use SMPL-based 6D rotation [34] as joint rotations. Similar to prior works on character control [55, 56], each pose x is represented by a 272-dimensional vector:

$$x = \{\dot{r}^x, \dot{r}^z, \dot{r}^a, j^p, j^v, j^r\}. \quad (1)$$

where we project the root on the XZ-plane (ground plane), $(\dot{r}^x, \dot{r}^z \in \mathbb{R})$ are root linear velocities on the XZ-plane, $\dot{r}^a \in \mathbb{R}^6$ denotes root angular velocity represented in 6D rotations, $j^p \in \mathbb{R}^{3K}$, $j^v \in \mathbb{R}^{3K}$, and $j^r \in \mathbb{R}^{6K}$ are local joint positions, local velocities, and local rotations relative to the root space, K is the number of joints. For SMPL [34] character, $K = 22$, then we get the $2 + 6 + 3 \times 22 +$

$3 \times 22 + 6 \times 22 = 272$ dimensions. This representation removes the post-processing step and we could directly use it for animating a SMPL character model.

3.2. Causal Temporal AutoEncoder

Streaming motion generation requires online motion decoding for dynamic text inputs. However, most existing works [21, 28, 68, 70] utilize temporal VQ-VAE to decode the whole sequence at once, where each frame inherently depends on past and future frames. Furthermore, the reliance on discrete tokenization induces quantization error accumulation across tokens, progressively degrading motion coherence in streaming generation scenarios. To address these issues, we introduce a Causal Temporal AutoEncoder (Causal TAE) to enable motion generation in a causal latent space.

Architecture. Causal TAE is designed to achieve continuous motion compression while explicitly modeling temporal dependencies and enforcing causal constraints for sequential motion representation. Fig. 3 shows the proposed Causal TAE network. We employ 1D causal convolution [66] for constructing temporal encoder \mathcal{E} and decoder \mathcal{D} to convert raw motion sequences into a causal latent space. The causality is guaranteed by a temporal padding scheme. Specifically, for a convolution layer with kernel size k_t , stride s_t and dilation rate d_t , we pad $(k_t - 1) \times d_t + (1 - s_t)$ frames at the beginning of the sequence. In this way, each frame only depends on the frames before it and the future frames are not involved in the computation. Moreover, explicitly modeling temporal causal structures in the latent space enables the model to learn temporal and causal dependencies inherent in the causally-related motion data.

Given a motion sequence $X = \{x_1, x_2, \dots, x_N\}$ with $x_t \in \mathbb{R}^D$, where N is the number of frames and D is the motion dimension ($D = 272$), we could obtain a set of temporal Gaussian distribution parameters $\{\mu_{1:N/l}, \sigma_{1:N/l}^2\}$ and preform reparameterization [29] to get continuous motion latent representation $Z = \{z_1, z_2, \dots, z_{N/l}\}$ with $z_i \in \mathbb{R}^{d_c}$, l represents the temporal downsampling rate of the Encoder \mathcal{E} . This architecture reconstructs motion frames while strictly preserving temporal causality across the sequence.

Training Objective. We use the same loss function as σ -VAE [52] to train the Causal TAE. In order to further enhance the reconstruction stability of the root joint, we add a root joint loss \mathcal{L}_{root} . The full loss function is defined as:

$$\mathcal{L} = \mathcal{L}_{recon} + D_{KL}(q(z|x)||p(z)) + \lambda \mathcal{L}_{root}. \quad (2)$$

where

$$\mathcal{L}_{recon} = \sum_{d=1}^D \sum_{i=1}^N \left(\frac{(x_{di} - \hat{x}_{di})^2}{2\sigma^{*2}} + \ln \sigma^* \right), \quad (3)$$

$$\mathcal{L}_{root} = \sum_{d=1}^{D_{root}} \sum_{i=1}^N \left(\frac{(x_{di} - \hat{x}_{di})^2}{2\sigma^{*2}} + \ln \sigma^* \right), \quad (4)$$

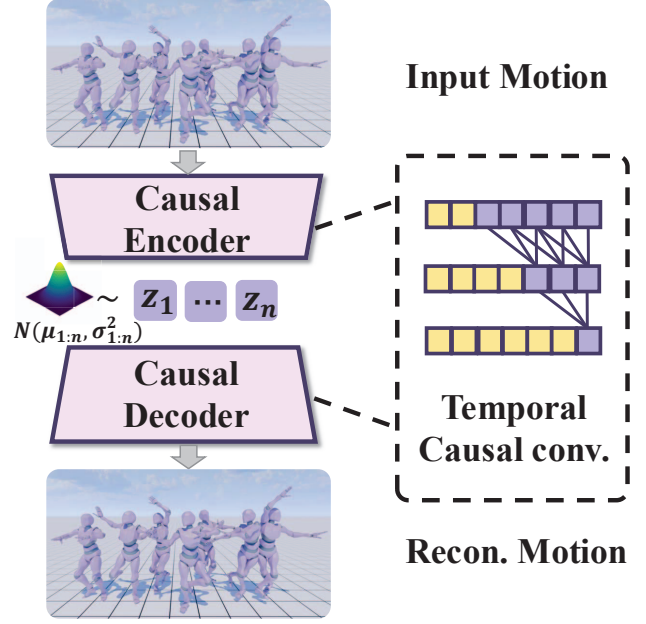


Figure 3. **Architecture of Causal TAE.** 1D temporal causal convolution is applied in both the encoder and decoder. Variables $z_{1:n}$ are sampled as continuous motion latent representations.

$$\sigma^{*2} = MSE(x, \hat{x}) = \frac{1}{DN} \sum_{d=1}^D \sum_{i=1}^N (x_{di} - \hat{x}_{di})^2, \quad (5)$$

$$D_{KL}(q(z|x)||p(z)) = \frac{1}{2} \sum_{d=1}^{d_c} \sum_{i'=1}^{N/l} (\mu_{di'}^2 + \sigma_{di'}^2 - \ln(\sigma_{di'}^2) - 1). \quad (6)$$

D , D_{root} and d_c represent the dimensions of the motion, root joint and the latent representation respectively. Specifically, $D = 272$, $D_{root} = 8$ (the first 8 dims relate to the root joint). $d_c = 16$ is the best choice in our experiments. N and l represent the number of frames and temporal downsampling rate, l is set to 4 here. x_{di} and \hat{x}_{di} are the ground-truth motion and the reconstructed motion at the i -th frame of the d -th dimension, σ^* is the analytic solution of the standard deviation [52]. D_{KL} represents the KL divergence, $q(z|x)$ is the distribution of latents given the motions, $p(z) = \mathcal{N}(0, I)$ is the prior distribution. $(\mu_{di'}, \sigma_{di'}^2)$ are the Gaussian distribution parameters at the i' -th latent of the d -th latent dimension, which derive from the Causal TAE Encoder \mathcal{E} . λ is the balancing hyperparameter.

Causal TAE offers distinct technical advantages for motion compression. Its causal property inherently supports online decoding without requiring access to future frames, which is critical for streaming generation. With the employment of continuous token representation, it bypasses the discretization bottleneck of existing VQ-based methods.

3.3. MotionStreamer

In this section, we present MotionStreamer, a streaming generation pipeline based on a causally-structured latent space. In order to handle coherence between arriving text inputs and historical generated motions, we hypothesize that current motion should only be conditioned on previous motion and current text. As illustrated in Fig. 2, MotionStreamer comprises a pre-trained text encoder, a diffusion-based autoregressive model, and the online motion decoder (the learned Causal TAE decoder).

Training. Each training sample can be represented as: $S_i = (T_i, C_i, Z_i)$, where $T_i \in \mathbb{R}^{1 \times d_t}$ is the text embedding obtained via a pre-trained language model (e.g. T5-XXL [50]), $C_i \in \mathbb{R}^{k \times d_c}$ and $Z_i \in \mathbb{R}^{n \times d_c}$ are the previous motion latents and current motion latents encoded by the learned Causal TAE, where k, n, d_t and d_c denote the lengths of previous motion latents, current motion latents, the text embedding dimension and the latent dimension respectively. We concatenate them along temporal axis to form a sequence $S_i = [T_i, C_i, Z_i]$. We employ a diffusion-based autoregressive Transformer to predict motion latents. The latent sequence S is first processed by the Transformer and a causal mask is applied to ensure the temporal causality [68]. After the Transformer processing, we obtain the intermediate latents $\{c_i^1, c_i^2, \dots, c_i^n\}$, which serve as the condition for the diffusion head (a small MLP) to predict motion latents $\{\hat{z}_i^1, \hat{z}_i^2, \dots, End_i\}$. End_i is the reference end latent inserted at the end of a sequence as the continuous stopping condition, which we will elaborate later. Following [24, 33], the loss function is defined as:

$$\mathcal{L} = \mathbb{E}_{\epsilon, t} [|\epsilon - \epsilon_\theta(Z_t | t, C_i, T_i)|^2]. \quad (7)$$

where t denotes the timestep of noise schedule. We employ QK normalization (i.e., normalize both queries and keys) [22] before self-attention layer to enhance training stability. **Two-Forward strategy.** We observe that using teacher-forcing [1] directly during training often leads to error accumulation in the autoregressive generation process. To this end, we propose a Two-Forward strategy that progressively introduces the test-time distribution during training. Specifically, after the first forward pass, we replace a subset of ground-truth motion latents with their generated counterparts, creating a mixture of real and generated motion latents. This hybrid input is then used in the second forward pass, where gradients are backpropagated. We employ a cosine scheduler to control the proportion of replaced motion latents. More details are provided in Sec.A of the appendix.

Mixed training. The datasets contain two types of training samples, so we set C_i to Null if there is no historical motion in the dataset. We find that this simple strategy enables a seamless transition between two consecutive motions.

Continuous stopping condition. Streaming generation

requires automatically determining the generation length for each text prompt. Previous method [37] uses a binary classifier to determine whether to stop generation, which suffers from a strong class imbalance. In contrast, we introduce an ‘‘impossible pose’’ prior (i.e., all-zero vectors $\mathbf{0} \in \mathbb{R}^D$) as the stopping condition and use the causal TAE to convert it into the latent space. The encoded latent serves as the reference end latent. The generation should stop when the distance between the currently generated latent and the reference end latent is less than a threshold. Therefore, MotionStreamer is able to stop generation automatically and enables online and multi-round generation.

Inference. During inference, given a stream of text prompts $\{\mathcal{P}_i\}_{i=1}^M$, the first text embedding T_1 is first fed into the autoregressive motion generator to generate the first predicted motion latent sequence $\hat{Z}_1 = \{\hat{z}_1^1, \hat{z}_1^2, \dots, \hat{z}_1^{n_1}\}$. As soon as a motion latent is predicted, it can be immediately processed by the online motion decoder (i.e., the learned Causal TAE decoder) to get the output motion frames, benefiting from its causal property. If the distance between the currently predicted motion latent and the reference end latent is lower than a threshold, the generation process of this prompt stops. Then, we replace T_1 with T_2 as the current text embedding. The already generated motion latent sequence \hat{Z}_1 is appended to the end of the second text embedding, forming the contextual latents used as input for the next autoregressive step. We then generate the second predicted motion latent sequence $\hat{Z}_2 = \{\hat{z}_2^1, \hat{z}_2^2, \dots, \hat{z}_2^{n_2}\}$. Next, we replace T_2 with future text embedding, removes \hat{Z}_1 from the condition latents and uses \hat{Z}_2 as the historical motion latents. Therefore, the third sequence could be predicted. This streaming generation process is repeated until the entire motion sequence $\{\hat{Z}_i\}_{i=1}^N$ is generated, ensuring online response during streaming generation process.

4. Experiment

4.1. Experimental Setup

Dataset. We evaluate the proposed MotionStreamer on HumanML3D [19] and BABEL [48] datasets, with the original train and test splits. The HumanML3D dataset integrates motion sequences with three distinct textual descriptions. The BABEL dataset provides frame-level textual descriptions with explicit inter-segment transition labels. Unlike recent methods [4, 53] that use different motion representations for both HumanML3D and BABEL datasets, we employ the 272-dimensional motion representation as mentioned in Sec. 3.1 for both datasets. All motion sequences are uniformly resampled to 30 FPS.

Evaluation Metrics. We adopt the metrics from [19] for evaluation, including: (1) Frchet Inception Distance (FID) [23], indicating the distribution distance between the generated and real motion; (2) Mean Per Joint Position Error

Methods	FID ↓	R@1 ↑	R@2 ↑	R@3 ↑	MM-D ↓	Div →
Real motion	0.002	0.702	0.864	0.914	15.151	27.492
MDM [57]	23.454	0.523	0.692	0.764	17.423	26.325
MLD [12]	18.236	0.546	0.730	0.792	16.638	26.352
T2M-GPT [68]	12.475	0.606	0.774	0.838	16.812	27.275
MotionGPT [28]	14.375	0.456	0.598	0.628	17.892	27.114
MoMask [21]	<u>12.232</u>	<u>0.621</u>	<u>0.784</u>	<u>0.846</u>	16.138	27.127
AttT2M [73]	15.428	0.592	0.765	0.834	15.726	26.674
Ours	11.790	0.631	0.802	0.859	<u>16.081</u>	27.284

Table 1. **Comparison with baseline text-to-motion generation methods** on HumanML3D [19] test set. MM-D and Div denote Multimodal Distance and Diversity respectively.

(MPJPE), the average distance between the predicted and ground-truth joint positions, measuring the reconstruction quality; (3) R-Precision (Top-1, Top-2, and Top-3 accuracy), the accuracy of the top-k retrieved motions; (4) Multimodal Distance (MM-Dist), the average Euclidean distances between the generated motion feature and its text feature. (5) Diversity, the average Euclidean distances of the randomly sampled 300 motion pairs, measuring the diversity of motions. (6) Peak Jerk (PJ) [4], the maximum value throughout the transition motion over all joints. (7) Area Under the Jerk (AUJ) [4], the area under the jerk curve. Both PJ and AUJ measures the smoothness of motions.

4.2. Quantitative Results

Comparison on Text-to-Motion Generation. We trained an evaluator based on TMR [42] to evaluate the quality of the generated motions. We use the processed 272-dimensional motion data from HumanML3D [19] train set for text-to-motion model training. All baselines are trained from scratch following their original implementations. As shown in Tab. 1, our method receives better performance across multiple metrics on HumanML3D [19] test set.

Comparison on Long-Term Motion Generation. We adopt a Mix Training Strategy for streaming long-term generation training. Specifically, we create training samples by pairing adjacent subsequences from long motion sequences in BABEL. Additionally, we incorporate text-motion pairs from the HumanML3D dataset for mix training. The comparison results on BABEL [48] dataset are demonstrated in Tab. 2. Following FlowMDM [4], the motion transition length is set to 30 frames. We modified T2M-GPT to support streaming generation (marked as T2M-GPT*) and also adapted our model to use VQ for discretization (marked as VQ-LLaMA). Experimental results show that neither the existing long-term generation baseline nor the discrete autoregressive model performs as well as our streaming generation approach in the continuous latent space.

Comparison on First-frame Latency. As streaming gen-

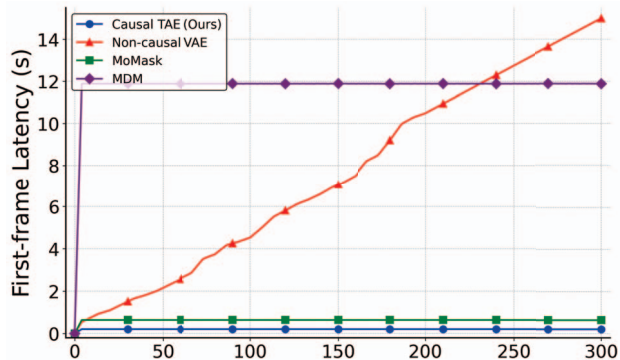


Figure 4. **Comparison on the First-frame Latency** of different methods. Horizontal axis: the number of generated frames. Vertical axis: the time required to produce the first output frame.

eration requires the model to generate motion progressively and respond online. Therefore, we adopt the First-frame Latency to evaluate the efficiency of different methods. First-frame Latency refers to the time taken by the model to produce its first predicted frame, serving as a key metric for evaluating online response ability. Experimental results in Fig. 4 show that our proposed Causal TAE achieves the lowest First-frame Latency, benefiting from the causal property of motion latents, which can be decoded immediately after generation. In contrast, non-causal VAE must wait until the entire sequence is generated before decoding, causing First-frame Latency to increase as the number of generated frames grows. Another fixed-length generation methods like [21, 57] exhibit higher First-frame Latency as they process the entire sequence at once rather than generating frames progressively in a streaming manner.

4.3. Qualitative Results

Figure 5 shows the qualitative results of our method compared with T2M-GPT [68], MoMask [21], AttT2M [73], and FlowMDM [4]. For the text-to-motion generation, we observe that VQ-based methods have difficulty in generating motions that are accurate and aligned with the textual description. In the case of “a man jumps on one leg,” T2M-GPT [68] and AttT2M [73] generate a motion where the person jumps with both legs instead. MoMask [21] employs residual vector quantization (RVQ) to reduce quantization errors but still suffers from fine-grained motion details loss. Specifically, the generated motion starts with a one-leg jump but later switches to two-leg jumps or alternating legs, along with noticeable sliding artifacts. However, our method can generate motions that are more accurate with more details preserved as we use a continuous latent space without discretization process.

For long-term motion generation, we compare with FlowMDM [4] with a stream of prompts: [“a man walks forward with arms swinging,” “then he jumps up,” “he

Methods	Subsequence				Transition			
	R@3 \uparrow	FID \downarrow	Div \rightarrow	MM-D \downarrow	FID \downarrow	Div \rightarrow	PJ \rightarrow	AUJ \downarrow
GT	0.634	0.000	24.907	17.543	0.000	21.472	0.03	0.00
DoubleTake [53]	0.452	23.937	22.732	21.783	51.232	18.892	0.48	1.83
FlowMDM [4]	<u>0.492</u>	<u>18.736</u>	<u>23.847</u>	<u>20.253</u>	<u>34.721</u>	<u>20.293</u>	<u>0.06</u>	0.51
T2M-GPT* [68]	0.364	39.482	24.317	20.692	43.823	20.797	0.12	1.43
VQ-LLaMA	0.383	24.342	19.329	38.285	36.293	19.932	0.08	1.20
Ours	0.568	15.743	23.546	15.397	32.888	19.986	0.04	<u>0.90</u>

Table 2. Comparison with long-term motion generation methods on BABEL [48] dataset.

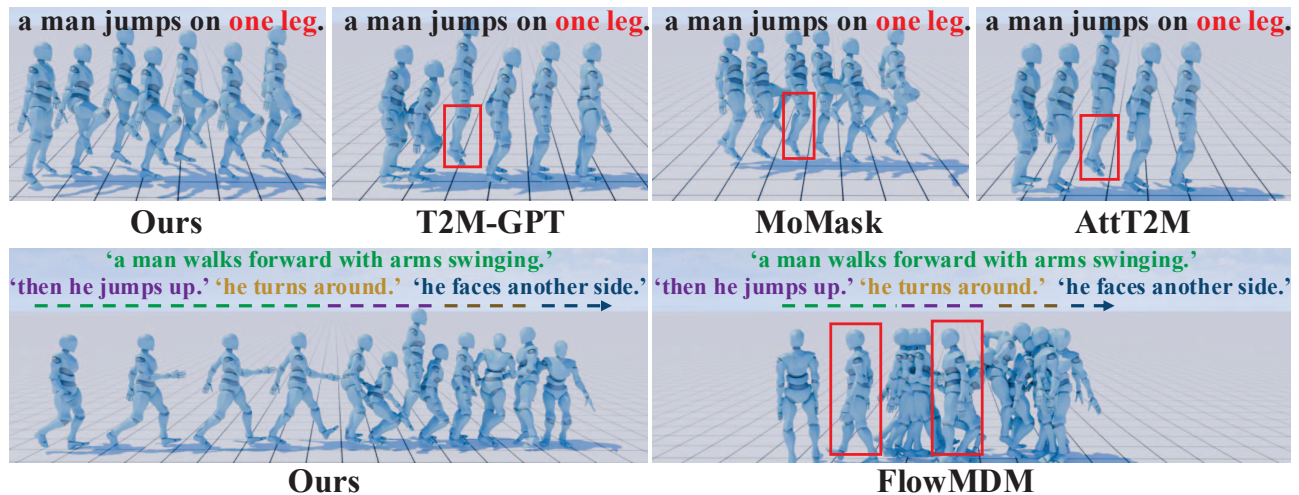


Figure 5. Visualization results between our method and some baseline methods [4, 21, 68, 73]. The first row shows text-to-motion generation results, the second row shows long-term generation results and the third row shows the application of dynamic motion composition.

turns around.”, “he faces another side.”]. The visualization results show that FlowMDM fails to generate the initial “walking” motion, instead producing in-place stepping. However, we can generate more coherent and natural long-term motions streamingly as our model has the ability to dynamically extract relevant information from variable-length motion histories. Please refer to the supplementary videos in our project page for more dynamic visualizations.

4.4. Ablation Study

Architecture of the Causal TAE. We comprehensively evaluate the reconstruction performance and the corresponding generation quality of different Causal TAE architectures on the HumanML3D [19] test set, as shown in Tab. 3. We replace the motion compression stage with VQ-VAE [68] to discretize the motions, while keeping the second-stage model architecture identical to ours. We also experimented with replacing Causal TAE with a non-causal temporal VAE and a standard temporal AE without vector quantization. The results show that our continuous repre-

sentsation avoids the VQ process, effectively reducing information loss and minimizing quantization error, thus performs better than the VQ-VAE baseline. The non-causal VAE performs worse than Causal TAE in both reconstruction and generation, as Causal TAE inherently models the causal structure of motion data during compression. This causal latent space is better suited for autoregressive generation, aligning naturally with the causal masking process. While AE achieves the best reconstruction quality by learning a near-identity mapping, its generation performance is significantly worse. This highlights the crucial role of latent space representation in determining the effectiveness of subsequent motion generation.

We provide a more detailed ablation on the latent dimension and hidden size of Causal TAE, as shown in Sec.B in the appendix. Notably, we observe that a larger latent dimension results in less compression rate, improving reconstruction quality. However, this comes at the cost of poorer generation performance, as insufficient compression and ineffective latent space representation makes it harder for the



Figure 6. **Dynamic motion composition.** Our model supports composition of multiple motions with different textual descriptions while maintaining previous motions unchanged.

Methods	Reconstruction		Generation			
	FID ↓	MPJPE ↓	FID ↓	R@3 ↑	MM-D ↓	Div →
Real motion	-	-	0.002	0.914	15.151	27.492
VQ-VAE	5.173	63.9	<u>13.226</u>	<u>0.824</u>	<u>16.746</u>	<u>27.024</u>
AE	0.001	1.7	43.828	0.463	22.040	27.382
VAE	2.092	26.2	19.902	0.735	17.926	<u>27.312</u>
Ours	<u>0.661</u>	<u>22.9</u>	11.790	0.859	16.081	27.284

Table 3. **Ablation Study of different motion compressors** on HumanML3D [19] test set. MPJPE is measured in millimeters.

model to learn meaningful motion generation. Meanwhile, the hidden size determines the model’s capacity, requiring a careful balance between compression rate and hidden size to ensure high reconstruction quality while enhancing generation performance. Ablation on the hyperparameter λ is provided in Sec.A of the appendix.

Design choices of AR Model. We analyze the impact of different design choices of the AR model, as shown in Tab. 4. The results show that the QK normalization and Two-forward strategy are effective. We also remove the diffusion head and use MSE loss for autoregressive training, which leads to a significant drop in generation quality. Moreover, we find that using T5-XXL [50] improves the generation performance compared with the CLIP [49] tokenizer. We found that applying a binary classifier to predict whether to stop generation, as in [37], fails to learn the correct stopping condition. As a result, we did not evaluate the model without the proposed continuous stopping condition.

4.5. Applications

MotionStreamer offers various applications, including multi-round generation, long-term generation, and dynamic motion composition. (1) **Multi-round generation** requires iteratively generating motion in response to sequential or interactive textual inputs. Our model can process incre-

AR Design choices	FID ↓	R@3 ↑	MM-D ↓	Div →
Real motion	0.002	0.914	15.151	27.492
w/o QK Norm	<u>12.324</u>	0.836	16.524	27.126
w/o Two-Forward	12.944	<u>0.842</u>	<u>16.432</u>	27.282
w/o Diffusion Head	59.296	0.360	22.864	27.325
CLIP	14.234	0.791	17.534	<u>27.316</u>
Ours	11.790	0.859	16.081	27.284

Table 4. **Analysis of design choices of the AR model** on HumanML3D [19] test set. CLIP indicates the use of CLIP model [49] as the text encoder to extract text features.

mental text inputs, respond online, and autonomously determine when to stop generation. (2) **Long-term generation** requires smoothly generating long motion sequences in response to sequential textual inputs. (3) **Dynamic motion composition** refers to the capability of seamlessly integrating diverse motion sequences while preserving the consistency of previously generated motion prefix, as shown in Fig. 6. Our Causal TAE enables this application, which eliminates the need for full-sequence re-decoding during the generation of subsequent motion latents, requiring only the decoding of the newly generated motion latents.

5. Conclusion

We present MotionStreamer, a novel framework for streaming motion generation that integrates diffusion-based autoregressive model to directly predict causal motion latents. By introducing Causal TAE, MotionStreamer supports online response with progressive textual inputs. We propose a Two-Forward training strategy to mitigate cumulative errors in streaming generation process. Our method outperforms baseline approaches, demonstrating its competitiveness in motion generation while providing greater flexibility. It can be applied to multi-round motion generation, long-term motion generation, and dynamic motion composition.

Acknowledgement

This work was partially supported by the National Key R&D Program of China (No. 2024YFB2809102), NSFC (No. 62402427, NO. U24B20154), Zhejiang Provincial Natural Science Foundation of China (No. LR25F020003), DeepGlint, Zhejiang University Education Foundation Qizhen Scholar Foundation, and Information Technology Center and State Key Lab of CAD&CG, Zhejiang University, the National Key R&D Program of China (2022ZD0160201), and Shanghai Artificial Intelligence Laboratory.

References

- [1] Kushal Arora, Layla El Asri, Hareesh Bahuleyan, and Jackie Chi Kit Cheung. Why exposure bias matters: An imitation learning perspective of error accumulation in language generation. *arXiv preprint arXiv:2204.01171*, 2022. 5
- [2] Nikos Athanasiou, Mathis Petrovich, Michael J. Black, and Gül Varol. Teach: Temporal action compositions for 3d humans. In *International Conference on 3D Vision (3DV)*, 2022. 2
- [3] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. Multidiffusion: Fusing diffusion paths for controlled image generation. 2023. 2
- [4] German Barquero, Sergio Escalera, and Cristina Palmero. Seamless human motion composition with blended positional encodings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 457–469, 2024. 2, 5, 6, 7
- [5] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *Computer Vision – ECCV 2016*. Springer International Publishing, 2016. 3
- [6] Zhi Cen, Huaijin Pi, Sida Peng, Qing Shuai, Yujun Shen, Hujun Bao, Xiaowei Zhou, and Ruizhen Hu. Ready-to-react: Online reaction policy for two-character interaction generation. In *ICLR*, 2025. 3
- [7] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11315–11325, 2022. 2
- [8] Junhao Chen, Mingjin Chen, Jianjin Xu, Xiang Li, Junting 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
- [9] Ling-Hao Chen, Jiawei Zhang, Yewen Li, Yiren Pang, Xiaobo Xia, and Tongliang Liu. Humanmac: Masked motion completion for human motion prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9544–9555, 2023.
- [10] Ling-Hao Chen, Shunlin Lu, Ailing Zeng, Hao Zhang, Benyou Wang, Ruimao Zhang, and Lei Zhang. Motionllm: Understanding human behaviors from human motions and videos. *arXiv preprint arXiv:2405.20340*, 2024. 2
- [11] Rui Chen, Mingyi Shi, Shaoli Huang, Ping Tan, Taku Komura, and Xuelin Chen. Taming diffusion probabilistic models for character control. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–10, 2024. 2
- [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, 3, 6
- [13] 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
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186, 2019. 2
- [15] GitHub discussion. rotation discussion. <https://github.com/EricGuo5513/HumanML3D/issues/26>, 2023. 2023-03-04. 3
- [16] 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
- [17] Ke Fan, Jiangning Zhang, Ran Yi, Jingyu Gong, Yabiao Wang, Yating Wang, Xin Tan, Chengjie Wang, and Lizhuang Ma. Textual decomposition then sub-motion-space scattering for open-vocabulary motion generation. *arXiv preprint arXiv:2411.04079*, 2024. 2
- [18] Ke Fan, Shunlin Lu, Minyue Dai, Runyi Yu, Lixing Xiao, Zhiyang Dou, Junting 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] 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 (CVPR)*, pages 5152–5161, 2022. 2, 3, 5, 6, 7, 8
- [20] 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. 3
- [21] 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. 2, 3, 4, 6, 7
- [22] Alex Henry, Prudhvi Raj Dachapally, Shubham Pawar, and Yuxuan Chen. Query-key normalization for transformers. *arXiv preprint arXiv:2010.04245*, 2020. 5
- [23] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a

- two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 5
- [24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 5
- [25] Bin Ji, Ye Pan, Yichao Yan, Ruizhao Chen, and Xiaokang Yang. Stylevr: Stylizing character animations with normalizing flows. *IEEE Transactions on Visualization and Computer Graphics*, 30(7):4183–4196, 2023. 2
- [26] Bin Ji, Ye Pan, Zhimeng Liu, Shuai Tan, Xiaogang Jin, and Xiaokang Yang. Pomp: Physics-consistent motion generative model through phase manifolds. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22690–22701, 2025. 2
- [27] Bin Ji, Ye Pan, Zhimeng Liu, Shuai Tan, and Xiaokang Yang. Sport: From zero-shot prompts to real-time motion generation. *IEEE Transactions on Visualization and Computer Graphics*, 2025. 2
- [28] 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. 2, 3, 4, 6
- [29] Diederik P Kingma, Max Welling, et al. Auto-encoding variational bayes, 2013. 2, 3, 4
- [30] 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
- [31] Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11523–11532, 2022. 3
- [32] 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
- [33] Tianhong Li, Yonglong Tian, He Li, Mingyang Deng, and Kaiming He. Autoregressive image generation without vector quantization. *arXiv preprint arXiv:2406.11838*, 2024. 5
- [34] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866. 2023. 3
- [35] Shunlin Lu, Ling-Hao Chen, Ailing Zeng, Jing Lin, Ruimao Zhang, Lei Zhang, and Heung-Yeung Shum. Humantomato: Text-aligned whole-body motion generation. *arXiv preprint arXiv:2310.12978*, 2023. 2
- [36] Shunlin Lu, Jingbo Wang, Zeyu Lu, Ling-Hao Chen, Wenxun Dai, Juntong Dong, Zhiyang Dou, Bo Dai, and Ruimao Zhang. Scamo: Exploring the scaling law in autoregressive motion generation model. *arXiv preprint arXiv:2412.14559*, 2024. 2
- [37] Lingwei Meng, Long Zhou, Shujie Liu, Sanyuan Chen, Bing Han, Shujie Hu, Yanqing Liu, Jinyu Li, Sheng Zhao, Xixin Wu, et al. Autoregressive speech synthesis without vector quantization. *arXiv preprint arXiv:2407.08551*, 2024. 5, 8
- [38] Zichong Meng, Yiming Xie, Xiaogang Peng, Zeyu Han, and Huaizu Jiang. Rethinking diffusion for text-driven human motion generation. *arXiv preprint arXiv:2411.16575*, 2024. 2
- [39] Zichong Meng, Zeyu Han, Xiaogang Peng, Yiming Xie, and Huaizu Jiang. Absolute coordinates make motion generation easy. *arXiv preprint arXiv:2505.19377*, 2025. 2
- [40] Liang Pan, Zeshi Yang, Zhiyang Dou, Wenjia Wang, Buzhen Huang, Bo Dai, Taku Komura, and Jingbo Wang. Tokenhsi: Unified synthesis of physical human-scene interactions through task tokenization. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5379–5391, 2025. 2
- [41] 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
- [42] Mathis Petrovich, Michael J Black, and Gül Varol. Tmr: Text-to-motion retrieval using contrastive 3d human motion synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9488–9497, 2023. 6
- [43] Huaijin Pi, Sida Peng, Minghui Yang, Xiaowei Zhou, and Hujun Bao. Hierarchical generation of human-object interactions with diffusion probabilistic models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15061–15073, 2023. 2
- [44] Huaijin Pi, Ruoxi Guo, Zehong Shen, Qing Shuai, Zechen Hu, Zhumei Wang, Yajiao Dong, Ruizhen Hu, Taku Komura, Sida Peng, et al. Motion-2-to-3: Leveraging 2d motion data to boost 3d motion generation. *arXiv preprint arXiv:2412.13111*, 2024.
- [45] Huaijin Pi, Zhi Cen, Zhiyang Dou, and Taku Komura. Coda: Coordinated diffusion noise optimization for whole-body manipulation of articulated objects. *arXiv preprint arXiv:2505.21437*, 2025. 2
- [46] 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. 2
- [47] 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. 2
- [48] Abhinanda R. Punakkal, Arjun Chandrasekaran, Nikos Athanasiou, Alejandra Quiros-Ramirez, and Michael J. Black. BABEL: Bodies, action and behavior with english labels. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 722–731, 2021. 2, 5, 6, 7
- [49] 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. 8

- [50] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020. 5, 8
- [51] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [52] Oleh Rybkin, Kostas Daniilidis, and Sergey Levine. Simple and effective vae training with calibrated decoders. In *International conference on machine learning*, pages 9179–9189. PMLR, 2021. 4
- [53] Yonatan Shafir, Guy Tevet, Roy Kapon, and Amit H Bermano. Human motion diffusion as a generative prior. *arXiv preprint arXiv:2303.01418*, 2023. 2, 5, 7
- [54] Yi Shi, Jingbo Wang, Xuekun Jiang, Bingkun Lin, Bo Dai, and Xue Bin Peng. Interactive character control with autoregressive motion diffusion models. *ACM Transactions on Graphics (TOG)*, 43(4):1–14, 2024. 3
- [55] Yi Shi, Jingbo Wang, Xuekun Jiang, Bingkun Lin, Bo Dai, and Xue Bin Peng. Interactive character control with autoregressive motion diffusion models. *ACM Transactions on Graphics (TOG)*, 43(4):1–14, 2024. 3
- [56] Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. Neural state machine for character-scene interactions. *ACM Transactions on Graphics*, 38(6):178, 2019. 3
- [57] 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, 3, 6
- [58] Guy Tevet, Sigal Raab, Setareh Cohan, Daniele Reda, Zhengyi Luo, Xue Bin Peng, Amit H Bermano, and Michiel van de Panne. Clossd: Closing the loop between simulation and diffusion for multi-task character control. *arXiv preprint arXiv:2410.03441*, 2024. 3
- [59] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 3
- [60] Yin Wang, Zhiying Leng, Frederick WB Li, Shun-Cheng Wu, and Xiaohui Liang. Fg-t2m: Fine-grained text-driven human motion generation via diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22035–22044, 2023. 2
- [61] Yabiao Wang, Shuo Wang, Jiangning Zhang, Ke Fan, Jiafu Wu, Zhucun Xue, and Yong Liu. Timotion: Temporal and interactive framework for efficient human-human motion generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 7169–7178, 2025. 2
- [62] Shuyang Xu, Zhiyang Dou, Mingyi Shi, Liang Pan, Leo Ho, Jingbo Wang, Yuan Liu, Cheng Lin, Yuexin Ma, Wenping Wang, et al. Mospa: Human motion generation driven by spatial audio. *arXiv preprint arXiv:2507.11949*, 2025.
- [63] 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
- [64] Zhao Yang, Bing Su, and Ji-Rong Wen. Synthesizing long-term human motions with diffusion models via coherent sampling. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3954–3964, 2023. 2
- [65] Chengjun Yu, Wei Zhai, Yuhang Yang, Yang Cao, and Zheng-Jun Zha. Hero: Human reaction generation from videos. *arXiv preprint arXiv:2503.08270*, 2025. 2
- [66] Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Vighnesh Birodkar, Agrim Gupta, Xiuye Gu, et al. Language model beats diffusion–tokenizer is key to visual generation. *arXiv preprint arXiv:2310.05737*, 2023. 4
- [67] Runyi Yu, Yinhuai Wang, Qihan Zhao, Hok Wai Tsui, Jingbo Wang, Ping Tan, and Qifeng Chen. Skillmimic-v2: Learning robust and generalizable interaction skills from sparse and noisy demonstrations. *arXiv preprint arXiv:2505.02094*, 2025. 2
- [68] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Shaoli Huang, Yong Zhang, Hongwei Zhao, Hongtao Lu, and Xi Shen. T2m-gpt: Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 3, 4, 5, 6, 7
- [69] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiandiffuse: Text-driven human motion generation with diffusion model. *IEEE transactions on pattern analysis and machine intelligence*, 46(6):4115–4128, 2024. 2
- [70] 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. 2, 3, 4
- [71] Yuhong Zhang, Guanlin Wu, Ling-Hao Chen, Zhuokai Zhao, Jing Lin, Xiaoke Jiang, Jiamin Wu, Zhuoheng Li, Hao Frank Yang, Haoqian Wang, et al. Humanmm: Global human motion recovery from multi-shot videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1973–1983, 2025. 2
- [72] Kaifeng Zhao, Gen Li, and Siyu Tang. Dart: A diffusion-based autoregressive motion model for real-time text-driven motion control. *arXiv preprint arXiv:2410.05260*, 2024. 2, 3
- [73] Chongyang Zhong, Lei Hu, Zihao Zhang, and Shihong Xia. Attt2m: 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. 2, 6, 7
- [74] Zixiang Zhou, Yu Wan, and Baoyuan Wang. Avatargpt: All-in-one framework for motion understanding planning generation and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1357–1366, 2024. 2