

AMD: Autoregressive Motion Diffusion

Bo Han¹, Hao Peng², Minjing Dong³, Yi Ren¹, Yixuan Shen⁴, Chang Xu^{3*}

¹College of Computer Science and Technology, Zhejiang University

²Unity China

³School of Computer Science, Faculty of Engineering, The University of Sydney

⁴National University of Singapore

borishan815@gmail.com, caspian.peng@unity.cn, mdon0736@uni.sydney.edu.au, rayeren613@gmail.com, yshe0148@gmail.com, c.xu@sydney.edu.au

Abstract

Human motion generation aims to produce plausible human motion sequences according to various conditional inputs, such as text or audio. Despite the feasibility of existing methods in generating motion based on short prompts and simple motion patterns, they encounter difficulties when dealing with long prompts or complex motions. The challenges are two-fold: 1) the scarcity of human motion-captured data for long prompts and complex motions. 2) the high diversity of human motions in the temporal domain and the substantial divergence of distributions from conditional modalities, leading to a many-to-many mapping problem when generating motion with complex and long texts. In this work, we address these gaps by 1) elaborating the first dataset pairing long textual descriptions and 3D complex motions (HumanLong3D), and 2) proposing an autoregressive motion diffusion model (AMD). Specifically, AMD integrates the text prompt at the current timestep with the text prompt and action sequences at the previous timestep as conditional information to predict the current action sequences in an autoregressive iterative manner. Furthermore, we present its generalization for X-to-Motion with “No Modality Left Behind”, enabling the generation of high-definition and high-fidelity human motions based on user-defined modality input.

Introduction

Human motion generation is a crucial task in computer animation and has applications in various fields including gaming, robots, and film. Traditionally, new motion is accessed through motion capture in the gaming industry, which can be costly. As a result, automatically generating motion from textual descriptions or audio signals can be more time-efficient and cost-effective. Related research work is currently flourishing, exploring human motion generation from different modalities (Tevet et al. 2022b; Zhang et al. 2022; Tseng, Castellon, and Liu 2022; Li et al. 2021).

Current text-based conditional human motion synthesis approaches have demonstrated plausible mapping from text to motion (Petrovich, Black, and Varol 2022; Tevet et al. 2022b; Zhang et al. 2022; Guo et al. 2022; Zhang et al. 2023). They are mainly divided into three categories:

Latent space strategy (Petrovich, Black, and Varol 2022; Ahuja and Morency 2019; Tevet et al. 2022a): This is typically done by separately learning a motion Variational AutoEncoder (VAE) (Kingma and Welling 2013) and a text encoder, and then constraining them to a compatible latent space using the Kullback-Leibler (KL) divergence loss. However, since the distributions of natural language and human motion are vastly different, forcibly aligning these two simple Gaussian distributions can result in misalignments and diminished generative diversity.

Diffusion-based approach (Tevet et al. 2022b; Zhang et al. 2022; Xin et al. 2023): diffusion models (Ho and Salimans 2022; Song et al. 2020) have recently attracted significant attention and have shown remarkable breakthroughs in various areas such as video (Luo et al. 2023), image (Ramesh et al. 2022), and 3D point cloud generation (Han, Liu, and Shen 2023), etc. Current motion generation methods based on diffusion models (Tevet et al. 2022b; Zhang et al. 2022; Xin et al. 2023) have achieved exceptional results using different denoising strategies.

Typically, MDM (Tevet et al. 2022b) proposes a motion diffusion model on raw motion data to learn the relationship between motion and text conditions. However, these models tend to only generate single motions or contain several motion sequences and are often inefficient for complex long texts. Autoregressive method (Gopalakrishnan et al. 2019; Pavllo et al. 2020; Athanasiou et al. 2022): they can process varying motion lengths, tackling the issue of fixed motion duration. However, their single-step generation methods often rely on traditional VAE models (Kingma and Welling 2013), which are less effective than diffusion models. Despite the progress made by existing methods, text-based conditional human motion generation remains a challenging task for several reasons:

- **Lack of motion-captured data:** At present, there are few widely used text-to-motion datasets (Plappert, Mandery, and Asfour 2016; Guo et al. 2022; Punnakkal et al. 2021), which mostly contain simple motions and are deficient in long text prompts, i.e., “he is flying kick with his left leg”.
- **Weak correlation:** Due to the differing distributions of language and human motion, resulting in a multiple mapping problem (Tevet et al. 2022b). This issue is further exacerbated when generating long text-based human motions.

*Corresponding author

To address the aforementioned limitations and challenges, we propose Autoregressive Motion Diffusion model (AMD) that can generate motion sequences with complex long content, variable duration, and multiple modalities. It leverages the generative capabilities of the diffusion model and the temporal modeling strengths of the autoregressive model. Considering the high dimensionality of complex long motion sequences, in order to better capture the dependencies between texts and motions in long sequences, AMD combines the text description at the current timestep with the text description and motion information at the previous timestep as conditional information to predict the motion sequence at the current timestep. AMD continuously employs the diffusion method to synthesize the corresponding motion sequence from the previous timestep and finally can generate the motion sequences of all texts. Besides, we explicitly design several geometric constraints to encourage matching physical realism in including height, joint position, joint rotation, joint velocity, and sliding loss. To address the scarcity of human motion-captured data for long prompts and complex motions, we have developed HumanLong3D - the first dataset to pair long textual descriptions with complex 3D human motions, i.e., "A person is doing martial art action raising knees and stretching feet, and then the person performs step forward with his right foot". The dataset comprises 158,179 textual descriptions and 43,696 3D human motions. It encompasses a broad spectrum of complex motion types. Importantly, it features annotations for motion coherence. In addition, we have also developed the HumanMusic dataset to evaluate the generation effect across different modalities. This dataset pairs 137,136 motions with corresponding audio data. They all follow the format of the HumanML3D dataset (Guo et al. 2022). The codes for AMD and demos can be found in the Supplementary Materials.

In summary, our contributions include:

- We propose a novel continuous autoregressive diffusion model that combines state-of-the-art performance for generating complex and variable motions on long texts.
- We construct two large-scale cross-modal 3D human motion datasets HumanLong3D and HumanMusic, which could serve as the benchmark datasets.
- Our proposed AMD achieves impressive performances on the HumanML3D, HumanLong3D, AIST++, and HumanMusic datasets, which highlights its ability to generate high-fidelity motion given different modality inputs.

Background

Human motion generation has been an active area of research for many years (Badler, Phillips, and Webber 1993). Early work in this field focused on unconditional motion generation (Rose, Cohen, and Bodenheimer 1998; Mukai and Kuriyama 2005; Ikemoto, Arikan, and Forsyth 2009), with some studies attempting to predict future motion based on an initial pose or starting motion sequence (O’rourke and Badler 1980; Gavrilu 1999). Statistical models such as Principal Component Analysis (PCA) (Ormonet et al. 2005) and Motion Graphs (Min and Chai 2012) were commonly used

for these generative tasks. The development of deep learning has led to the emergence of an increasing number of sophisticated generative architectures (Kingma and Welling 2013; Vaswani et al. 2017; Goodfellow et al. 2020; Kingma and Dhariwal 2018; Ho, Jain, and Abbeel 2020). These advanced generative models have encouraged researchers to explore conditional motion generation. Conditional human motion generation can be modulated by a variety of signals that describe the motion, with high-level guidance provided through various means such as action classes (Petrovich, Black, and Varol 2022), audio (Aristidou et al. 2022), and natural language (Ahuja and Morency 2019; Petrovich, Black, and Varol 2022).

Text-to-Motion

Due to the language descriptors are the most user-friendly and convenient. Text-to-motion has been driving and dominating research frontiers. In recent years, the leading approach for the Text-to-Motion task is to learn a shared latent space for language and motion. JL2P (Ahuja and Morency 2019) learns from the KIT-ML dataset (Plappert, Mandery, and Asfour 2016) with an auto-encoder, limiting one-to-one mapping from text to motion. TEMOS (Ahuja and Morency 2019) and T2M (Guo et al. 2022) propose using a VAE (Kingma and Welling 2013) to map a text prompt into a normal distribution in latent space. Recently, MotionCLIP (Tevet et al. 2022a) has leveraged the shared text-image latent space learned by CLIP to expand text-to-motion beyond data limitations and enable latent space editing. However, due to the inconsistency of the two data distributions of natural language and human motion, it is very difficult to align them in the shared latent space. Diffusion Generative Models (Sohl-Dickstein et al. 2015) achieve significant success in the image synthesis domain, such as Imagen (Saharia et al. 2022), DALL2 (Ramesh et al. 2022) and Stable Diffusion (Rombach et al. 2022). Inspired by their works, most recent methods (Tevet et al. 2022b; Zhang et al. 2022; Xin et al. 2023) leverage diffusion models for human motion synthesis. MotionDiffuse (Zhang et al. 2022) is the first work to generate human motion that corresponds to text utilizing a diffusion model. Recently, MDM (Tevet et al. 2022b) has been proposed, which operates on raw motion data to learn the relationship between motion and input conditions. Inspired by Stable Diffusion (Rombach et al. 2022), MLD (Xin et al. 2023) implements the human motion diffusion process in the latent space. Despite their ability to produce exceptional results, these models are typically limited to short text descriptions and simple motions. Additionally, several works (Gopalakrishnan et al. 2019; Pavllo et al. 2020; Athanasiou et al. 2022) have been developed based on the concept of autoregression, which can generate human actions of any length. ARDMs (Hoogetboom et al. 2022) combines the order-agnostic autoregressive model and the discrete diffusion model, which eliminates the need for causal masking of model representations and enables fast training, allowing it to scale favorably to high-dimensional data. Consequently, for long text prompts, we combine the advantages of the diffusion model in generating motion for short text descriptions with the concept of autoregression to achieve

superior human motion results for continuous long text.

Motion Datasets

Common forms of description for human motion data are 2D keypoints, 3D keypoints, and statistical model parameters (Yu et al. 2020; Cai et al. 2022). For the text-conditioned motion generation task, KIT (Plappert, Mandery, and Asfour 2016) is the first 3D human motion dataset with matching text annotations for each motion sequence. HumanML3D (Guo et al. 2022) provides more textual annotation for some motions of AMASS (Mahmood et al. 2019). They are also our focus in the text-to-motion task. Babel (Punnakkal et al. 2021) also collects motions from AMASS (Mahmood et al. 2019) and provides action and behavior annotations, it annotates each frame of the action sequence, thereby dividing compound actions into simple action groups. In this paper, we use the HumanML3D dataset to evaluate the proposed methods for simple motions and short prompts. In addition, we collected and labeled pairs of complex motion data and text prompts (HumanLong3D). More importantly, we provided temporal motion-coherence information to support long text-to-motion generation tasks.

Audio-to-Motion

Generating natural and realistic human motion from audio is also a challenging problem. Many early approaches follow a motion retrieval paradigm (Fan, Xu, and Geng 2011; Lee, Lee, and Park 2013). A traditional approach to motion synthesis involves constructing motion maps. New motions are synthesized by combining different motion segments and optimizing transition costs along graph paths (Safonova and Hodgins 2007). More recent approaches employ RNN (Tang, Jia, and Mao 2018; Alemi, Françoise, and Pasquier 2017; Huang et al. 2020), GANs (Lee et al. 2019; Sun et al. 2020), Transformer (Li et al. 2022, 2021; Siyao et al. 2022), and CNN (Holden, Saito, and Komura 2016) models to map the given music to a joint sequence of the continuous human pose space directly. Such methods would regress to nonstandard poses that are beyond the dancing subspace during inference. In contrast, our proposed method does not produce the phenomenon of limb drift.

Our Approach

In this section, we first introduce the problem formulation for long text-to-motion. To enable adaptive motion generation for different text descriptions, we propose the inclusion of a motion duration prediction network to approximate the duration. To generate human motions that correspond to continuous long text descriptions, we establish an autoregressive motion diffusion model.

Problem Description

To generate complex motion sequences with long-term text prompts, we propose to feed multiple text prompts in order. Given N text prompts $S^{1:N} = \{S^1, S^2, \dots, S^N\}$, the model is required to generate N motion segments $X^{1:N} = \{X^1, X^2, \dots, X^N\}$ consistent with the text descriptions, where N denotes the number of motion segments involved

in the entire motion sequence. Each motion segment is defined as $X^i = \{x^1, x^2, \dots, x^{F^i}\}$, where F^i is the total number of frames of the motion segment X^i and x^j denotes the 3D human body pose representation of the j -th frame. It is imperative that each generated motion segment and the corresponding number of motion frames adhere to the specifications outlined in the text prompt. Additionally, a seamless transition from X^{i-1} to X^i is crucial for the generation of high-fidelity motion.

Overall Framework

It is important to note that daily human motions encompass not only simple, single motions but also complex, prolonged motions that more accurately reflect real-life scenarios. Specifically, given a series of semantic prompts $S^{1:N}$, a series of randomly sampled temporal motion sequences $X_T^{1:N} \sim \mathcal{N}(0, I)$ obeying the standard normal distribution, and a maximum noise scale $T \in \mathbb{N}$ where each semantic prompt S^i describes a single and distinct motion. Our goal is to generate noise-free temporal motion sequences $X_0^{1:N}$, which are guided by the semantic prompts, with smooth transitions between adjacent motions X_0^{i-1} and X_0^i . The overall process is illustrated in Figure 1, and each pair of blue and green blocks represents each step of the AMD model. $S^{1:N}$ employs the model iteratively to synthesize motion $X_0^{1:N}$. The blue block represents the context encoder and the green block is the motion diffusion module.

Motion Duration Prediction Network Given a semantic prompt S^i , the duration of X^i may vary. For instance, in the HumanML3D dataset, the prompt "a man kicks something or someone with his left leg" corresponds to 116 frames, while the prompt "a person squats down then jumps" corresponds to 35 frames. Consequently, we propose predicting the motion duration in order to generate motions with adaptive length. Similar to T2M, we use probability density estimation to determine the number of frames required for the motion synthesis based on text prompts. Due to the diversity of the motion duration, it is more reasonable to model the mapping problem as a density estimation problem than directly regressing the specific value. By utilizing the semantic prompt S^i as input for the motion duration prediction network, a probability density estimation is conducted on the discrete group encompassing all possible motion durations $L = \{L_{min}, L_{min} + 1, \dots, L_{max}\}$. The loss function of the network is designed as the cross-entropy loss of multi-classification.

Context Encoder It includes the motion duration prediction network E_D , the semantic conditional encoder E_S , and the motion linear layer. The CLIP model (Radford et al. 2021) is utilized as the semantic conditional encoder. Given that our primary focus is on long text-to-motion generation, it is necessary to consider timing-related information associated with long texts. To this end, we encode the previous motion X_0^{i-1} by the motion linear layer to obtain z_m^{i-1} and encode semantic information S^{i-1} by the semantic conditional encoder E_S to obtain z_c^{i-1} . These are then concatenated to form the final prior condition feature z_{past}^{i-1} . Simultaneously,

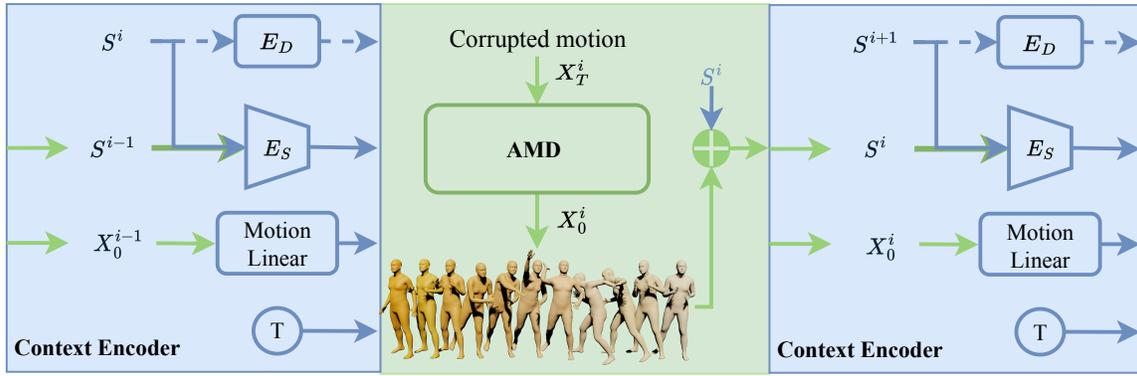


Figure 1: Overview of the Autoregressive Motion Diffusion model. Given the current text prompt S^i , the last text prompt S^{i-1} , and motion X_0^{i-1} (green arrow), we first encode the context information (blue block). Then, we feed the input conditions and corrupted motion X_T^i to AMD Module (Figure 2) to generate the original cleaned motion X_0^i . Afterward, we send the current text prompt S^i and motion X_0^i to the next time step. Iteratively, we can obtain motion sequences for long text prompts.

the current semantic information is input into the motion duration prediction network and semantic conditional encoder to obtain F^i and z_c^i , respectively. In order to avoid overfitting, we perform a random mask on the semantic conditional information z_c^i . For the corrupted motion X_t^i , the same motion linear layer is utilized to obtain the encoded information z_m^i . We feed the diffusion time scale t to a Multi-layer Perceptron (MLP) to obtain the time embedding z_t . The final condition information z is defined as follows:

$$z = C(C(C(z_m^{i-1}, z_c^{i-1}) + RM(z_c^i), z_t), z_m^i, PE(F^i))) \quad (1)$$

where C represents the concatenation operation, RM denotes random mask, and position embedding refers to position embedding. It is important to note that during training, we utilize the actual motion duration present in the dataset, whereas, during the inference phase, the predicted duration information is used.

AMD Module The network architecture of the AMD module is depicted in Figure 2. The denoising process (gray) and the diffusion process (yellow) span a total of T timesteps, where T represents the pre-defined maximum time scale. We directly predict the original cleaning motion clips during the denoising process, while the diffusion process operates in the opposite direction. The single step of the diffusion process is essentially the transfer process from X_{t-1}^i to X_t^i , as defined in the following:

$$q(X_t^i | X_{t-1}^i) = \mathcal{N}(X_t^i; \sqrt{1 - \beta_t} X_{t-1}^i, \beta_t I), \quad (2)$$

Where β_t is pre-defined to regulate the magnitude of noise addition. The transition probability, denoted as Equation 3, from X_0^i to X_t^i can be derived using Equation 2 in conjunction with the multiplication formula for Gaussian distribution, where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$.

$$q(X_t^i | X_0^i) = \prod_{t=1}^t q(X_t^i | X_{t-1}^i) = \mathcal{N}(X_t^i; \sqrt{\bar{\alpha}_t} \hat{X}_0^i, (1 - \bar{\alpha}_t) I). \quad (3)$$

The single step of the denoising process is essentially the transfer process from X_t^i to X_{t-1}^i , the transfer strategy requires a network with parameters θ to learn the sampling distribution as:

$$\begin{aligned} p_\theta(X_{t-1}^i | X_t^i, z) &= q(X_{t-1}^i | \hat{X}_0^i(X_t^i, t, z)) \\ &= \mathcal{N}(X_{t-1}^i; \sqrt{\bar{\alpha}_t} \hat{X}_0^i(X_t^i, t, z), (1 - \bar{\alpha}_t) I), \end{aligned} \quad (4)$$

where $\hat{X}_0^i(X_t^i, t, z)$ represents the neural network with parameter θ , which takes in X_t^i , t , and conditional information z as input.

After predicting \hat{X}_0 through X_t using the inverse diffusion network, it is necessary to perform the forward diffusion process according to the calculation method outlined in Equation 3 to obtain the noise motion clips X_{t-1} of the subsequent noise scale, thus completing the single-step inverse diffusion process. In summary, the algorithm randomly samples the noise motion segment X_T with the largest noise scale from the standard normal distribution $\mathcal{N}(0, I)$ and iteratively executes the single-step inverse diffusion model. Each prediction approximately coarse cleaned action sequence \hat{X}_0 before forward diffusion is performed to obtain the action segment of the next noise scale. This process continues until the noise scale reaches 0 and returns to the original cleaned action sequence X_0 .

Explicit Constraints In the AMD Module, the original motion sequence is explicitly predicted. To enhance physical realism, we design multiple geometric constraints, including L_h , L_p , L_r , L_v , and L_f . The loss function of each part of human motion is defined as the L2 loss between the predicted values and the ground truth.

Among them, L_h represents the height loss. L_p represents the joint position loss. L_r represents the joint rotation loss. L_v represents the joint velocity loss. L_f represents the sliding foot loss. Finally, the loss function is defined as:

$$\mathcal{L}_{train} = \lambda_h \mathcal{L}_h + \lambda_p \mathcal{L}_p + \lambda_r \mathcal{L}_r + \lambda_v \mathcal{L}_v + \lambda_f \mathcal{L}_f \quad (5)$$

where λ denotes the coefficients to balance the loss terms.

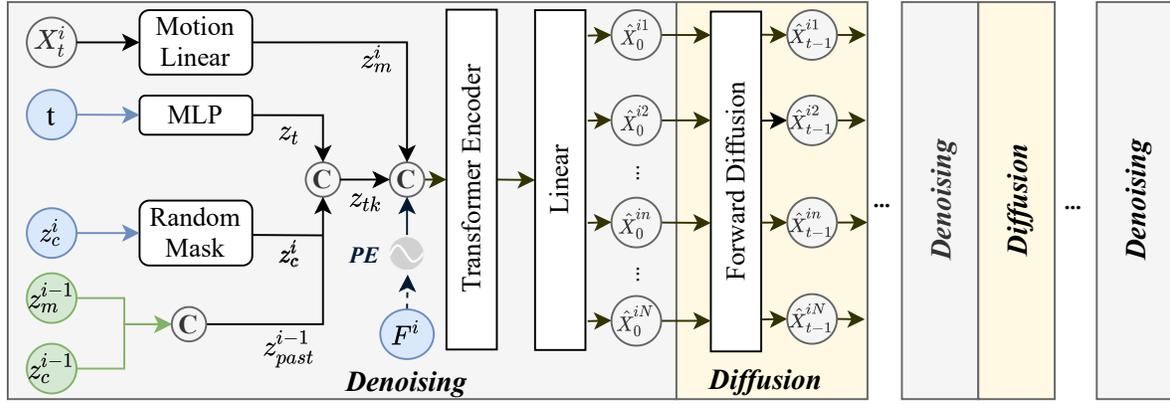


Figure 2: AMD Module. The gray blocks denote the denoising process, while the yellow blocks represent the diffusion process. Within the AMD module, they appear in pairs T times (with the exception of the last one).

With the proposed AMD, we are able to generate motion sequences according to ordered text prompts iteratively. Specifically, we commence from the first prompt S^1 and utilize the AMD to synthesize the corresponding clean motion sequence X_0^1 . The remaining high-fidelity motion sequences $X_0^{2:N}$ can be synthesized using prior condition information as well as $S^{2:N}$. Ultimately, a coherent motion sequence of any length can be synthesized.

Experiments

Datasets and Evaluation Metrics

HumanLong3D We collected motion data using motion capture equipment and online sources and annotated each motion sequence with various semantic labels to create the HumanLong3D dataset. The data format of the HumanLong3D dataset is consistent with that of HumanML3D, and it additionally includes coherence information for motion sequences to support temporal motion generation tasks.

HumanML3D The dataset involves the textual re-annotation of motion capture data from the AMASS (Mahmood et al. 2019) and HumanAct12 (Guo et al. 2020), comprising 14,616 motions annotated with 44,970 textual descriptions. The comparison of datasets is shown in Table 1.

HumanMusic We collected dance videos from online sources and extracted the pose parameters of the dancers in the videos, converting the motion data into the HumanML3D format. The frame rate of each video was normalized to 20 FPS and the sampling rate of the accompanying music was standardized to 10240Hz. In total, we obtained 137,136 paired dance and music data samples, with each dance sample consisting of 200 frames. Music features were extracted using the publicly available audio processing toolbox Librosa (Jin et al. 2017).

AIST++ This dataset (Li et al. 2021) comprises 992 high-quality 3D pose sequences in SMPL format (Loper et al. 2015), captured at 60 FPS, with 952 sequences designated for training and 40 for evaluation. We followed the approach of Bailando (Siyao et al. 2022) to partition the dataset.

Evaluation Metrics For text-to-motion evaluation, we employ metrics consistent with existing methods (Tevet

Dataset	Motion	Textual descriptions	Duration
KIT-ML	3911	6248	10.33h
HumanML3D	14616	44970	28.59h
HumanLong3D	43696	158179	85.87h

Table 1: Text-to-motion dataset description

et al. 2022b; Zhang et al. 2022). Specifically, (a) Fréchet Inception Distance (FID) is used as the primary metric to evaluate the feature distributions between generated and real motions in feature space (Guo et al. 2022), and (b) R-Precision (top 3) calculates the top 3 matching accuracy between text and motion in feature space. (c) *MultiModal Dist* calculates the distance between motions and texts. (d) *Diversity* measures variance through features. (e) *MultiModality* assesses the diversity of generated motions for the same text. For music-to-dance evaluation, we employ metrics consistent with existing methods (Siyao et al. 2022).

Implement Details

Motion Representation Our motion representation adopts the same format as HumanML3D, i.e., $X \in \mathbb{R}^{263 \times F}$. Each frame of motion is 263-dimensional data, including the position, linear velocity, angular velocity, joint space rotation of three-dimensional human joints, and label information for judging whether the foot joints are still. Since images are often represented as $I \in \mathbb{R}^{W \times H \times C}$.

Motion Duration Prediction Network L_{min} is set to 10 and L_{max} is 50, each unit increment corresponds to 4 motion frames, i.e., 0.2s motion duration, so the duration prediction range covers the lower bound of 2s and the upper bound of 9.8s of the data samples. The motion duration prediction network is pretrained, with the motion duration prediction network being used only during inference.

AMD Module We set the maximum noise scale T to be 1000, the coefficient $\beta_{1:T}$ is set to a linear increment from 10^{-4} to 0.02, latent vector dimensions are 512, the number of layers of the motion encoder is 6, and the number of heads of the multi-head attention mechanism is set to 6, the

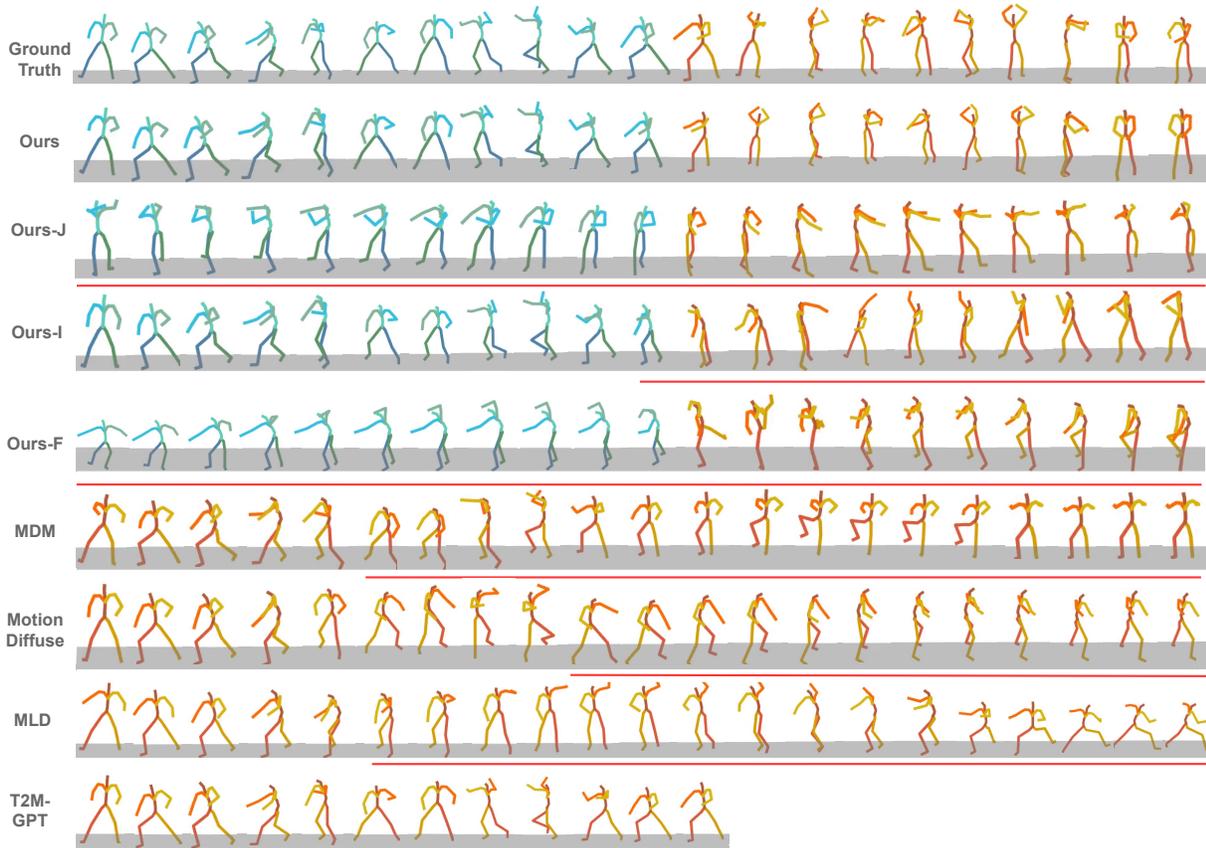


Figure 3: Result for compound motion synthesis (blue: “there is a man doing left smash right cover.” yellow: “then he steps forward and turn around”). The part (red line) indicates a discrepancy between the generated motion and the ground truth.

learning rate is fixed at 10^{-4} , the number of training steps is 200000, and we use AdamW optimizer.

Other Settings The output dimension of the motion linear layer and the latent vector dimension of the AMD module are both 512. The semantic conditional encoder adopts the CLIP-ViT-B/32 checkpoint. During inference, the semantic prompt S^i is input into the motion duration prediction network E_D to obtain the estimated value F^i of the motion sequence duration, which is used to determine the timing dimension for motion sequence sampling.

Comparisons on Compound Motion

We compare compound motion generation with SOTA methods. Since the HumanML3D dataset does not contain motion coherence information, we conducted this experiment only on the HumanLong3D dataset, and we divided the dataset into training, test, and validation sets using a ratio of 0.85:0.10:0.05. Additionally, we designed three benchmarks based on TEACH (Athanasios et al. 2022): 1) Joint prediction (ours-J): The long semantic prompt $S^{i-1:i}$ formed by the combination of two coherent prompts are used as the input of the diffusion model, and a coherent time-series motion sequence $X_0^{i-1:i}$ is obtained by direct joint prediction. 2) Linear interpolation (ours-I): This method interpolates the

results of two independent motion synthesis. 3) Motion filling (ours-F): Similar to linear interpolation, two independent motion synthesis are required to obtain X_0^{i-1} and X_0^i , and the time window is set to 10% of the motion sequence duration. All frame data except for the time window are fixed, and the frame data within the time window are filled with random normal distribution noise. The coherent motion sequence is then restored through the denoising process.

As shown in Table 2, Among the five evaluation metrics, AMD achieved top 3 performance in four of them, with FID and Diversity, the primary metrics for motion generation quality, ranking first, demonstrating its superiority in the long text-to-compound motion generation task. Notably, AMD outperformed other methods by a significant margin in the FID metric. While the “Ours-J” scheme, despite having the highest Multimodality index, performed poorly in terms of FID, indicating its inability to generate reasonable human movements. In cases where synthesis quality is extremely low, high diversity in Multimodality may indicate that the synthesized actions are chaotic and irregular.

As illustrated in Figure 3, compared to the ground truth, AMD keeps with the highest degree of similarity, while MDM, MotionDiffuse, and MLD all exhibited varying degrees of limb stiffness. Although T2M-GPT achieves results

Method	R-Precision(Top3) \uparrow	FID \downarrow	MultiModal Dist \downarrow	Diversity \rightarrow	MultiModality \uparrow
Ours-J	0.120 \pm .006	12.679 \pm .063	9.991 \pm .047	4.683 \pm .117	3.646 \pm .285 \diamond
Ours-I	0.142 \pm .003	0.827 \pm .036*	7.989 \pm .056	4.329 \pm .065	2.249 \pm .206
Ours-F	0.132 \pm .008	0.576 \pm .042 \dagger	7.937 \pm .023	4.312 \pm .052	2.257 \pm .269
MDM (Tevet et al. 2022b)	0.096 \pm .005	27.348 \pm .349	8.203 \pm .039	0.781 \pm .040	0.547 \pm .037
MotionDiffuse (Zhang et al. 2022)	0.156 \pm .002*	6.860 \pm .113	8.783 \pm .028	4.529 \pm .076 \dagger	2.409 \pm .204
T2M-GPT (Zhang et al. 2023)	0.159 \pm .005 \diamond	1.249 \pm .026	7.653 \pm .019 \diamond	4.895 \pm .047	3.093 \pm .104 \dagger
MLD (Xin et al. 2023)	0.144 \pm .002	3.843 \pm .058	7.847 \pm .011*	4.365 \pm .033*	2.831 \pm .072*
Ours	0.158 \pm .006 \dagger	0.225 \pm .013 \diamond	7.745 \pm .029 \dagger	4.515 \pm .135 \diamond	1.242 \pm .118
GT	0.162 \pm .005	0.003 \pm .001	7.119 \pm .013	4.456 \pm .073	-

Table 2: Compound motion generation evaluation on HumanLong3D Dataset. For each metric, we repeat the evaluation 20 times (except MultiModality runs 5 times). \diamond , \dagger , and * indicate the first, the second, and the third best result.

Method	Motion Quality		Motion Diversity		BAS \uparrow
	FID $_k$	FID $_g^\dagger$	Div $_k$	Div $_g^\dagger$	
DanceNet	69.18	25.59	2.86	2.85	0.1430
FACT	35.35	22.11	5.94	6.18	0.2209
Bailando	28.16	9.62	7.83	6.34	0.2332
Ours	30.28	16.11	6.75	6.29	0.2302

Table 3: Music-to-dance evaluation on AIST++ Dataset.

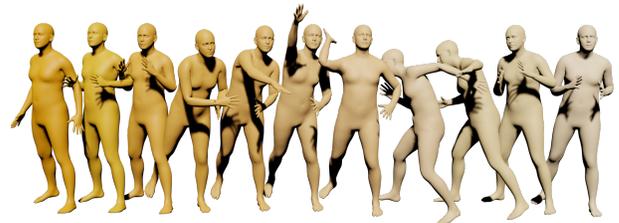
comparable to the ground truth in the first half of motion generation, its performance deteriorates in longer text-to-motion generation tasks. This is due to its premature prediction of the terminator, resulting in a lack of corresponding motion sequences for the second half of the text. For T2M-GPT, we also tried to separate the long texts into short texts and generated single motion on the short texts individually. T2M-GPT performs well in single motion generation tasks but struggles with compound motion generation tasks. Additionally, when generating long text-to-compound motion, dividing the long text and generating it separately often results in unnatural transitional motion clips.

Comparisons on Single Motion We also single motion generation experiments on HumanML3D and HumanLong3D Dataset. For single motion generation, our conditional information includes the estimated motion duration and semantic information but excludes prior motion and semantic information. The visualization results are shown in Figure 4. It can be seen that AMD is capable of generating corresponding motion in response to text prompts containing a single motion while achieving smooth transitions.

Comparisons on Music-to-Dance We conducted comparative experiments with SOTA methods, including DanceNet (Zhuang et al. 2022), Dancerevolution (Huang et al. 2020), FACT (Li et al. 2021), and Bailando (Siyao et al. 2022), on the public dataset AIST++. We employed the same data partitioning strategy as the aforementioned prior works, and we converted the data in AIST++ into HumanML3D format. The quantitative results are presented in Table 3. As can be observed from the table, our method is second only to Bailando in various performance metrics, particularly in the measurement of music and dance consistency



(a) the person picks an object up off the floor with their left hand



(b) a person throws an object with his right hand.

Figure 4: Visualization on HumanML3D Dataset.

indicators BAS. Bailando employs a customized reinforcement learning module to enhance the BAS index. In contrast, our method does not incorporate any enhancements yet still achieves comparable performance to Bailando. These results demonstrate that our method not only excels in the text-to-motion task but also exhibits strong generalization capabilities to other human motion generation tasks.

Conclusion

In this paper, we present the HumanLong3D - the first dataset that pairs complex motions with long textual descriptions to address the scarcity of such data. Given the suboptimal performance of current motion generation methods on long text descriptions, we introduce a novel network architecture AMD, which combines autoregressive and diffusion models to effectively capture the information contained in long texts. Furthermore, we extend our approach to incorporate audio conditional input and construct a large-scale music-dance dataset - HumanMusic can serve as a benchmark in the field of music-to-dance.

Acknowledgments

This work was supported in part by the Australian Research Council under Project DP210101859 and China Scholarship Council (CSC) Grant.

References

- Ahuja, C.; and Morency, L.-P. 2019. Language2pose: Natural language grounded pose forecasting. In *2019 International Conference on 3D Vision (3DV)*, 719–728. IEEE.
- Alemi, O.; François, J.; and Pasquier, P. 2017. GrooveNet: Real-time music-driven dance movement generation using artificial neural networks. *networks*, 8(17): 26.
- Aristidou, A.; Yiannakidis, A.; Aberman, K.; Cohen-Or, D.; Shamir, A.; and Chrysanthou, Y. 2022. Rhythm is a Dancer: Music-Driven Motion Synthesis with Global Structure. *IEEE Transactions on Visualization & Computer Graphics*, (01): 1–1.
- Athanasioiu, N.; Petrovich, M.; Black, M. J.; and Varol, G. 2022. TEACH: Temporal Action Compositions for 3D Humans. In *International Conference on 3D Vision (3DV)*.
- Badler, N. I.; Phillips, C. B.; and Webber, B. L. 1993. *Simulating humans: computer graphics animation and control*. Oxford University Press.
- Cai, Z.; Ren, D.; Zeng, A.; Lin, Z.; Yu, T.; Wang, W.; Fan, X.; Gao, Y.; Yu, Y.; Pan, L.; et al. 2022. Humman: Multimodal 4d human dataset for versatile sensing and modeling. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, 557–577. Springer.
- Fan, R.; Xu, S.; and Geng, W. 2011. Example-based automatic music-driven conventional dance motion synthesis. *IEEE transactions on visualization and computer graphics*, 18(3): 501–515.
- Gavrila, D. M. 1999. The visual analysis of human movement: A survey. *Computer vision and image understanding*, 73(1): 82–98.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Gopalakrishnan, A.; Mali, A.; Kifer, D.; Giles, L.; and Ororbia, A. G. 2019. A neural temporal model for human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12116–12125.
- Guo, C.; Zou, S.; Zuo, X.; Wang, S.; Ji, W.; Li, X.; and Cheng, L. 2022. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5152–5161.
- Guo, C.; Zuo, X.; Wang, S.; Zou, S.; Sun, Q.; Deng, A.; Gong, M.; and Cheng, L. 2020. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, 2021–2029.
- Han, B.; Liu, Y.; and Shen, Y. 2023. Zero3D: Semantic-Driven Multi-Category 3D Shape Generation. *arXiv preprint arXiv:2301.13591*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Holden, D.; Saito, J.; and Komura, T. 2016. A deep learning framework for character motion synthesis and editing. *ACM Transactions on Graphics (TOG)*, 35(4): 1–11.
- Hoogeboom, E.; Gritsenko, A. A.; Bastings, J.; Poole, B.; van den Berg, R.; and Salimans, T. 2022. Autoregressive Diffusion Models. In *ICLR 2022*. OpenReview.net.
- Huang, R.; Hu, H.; Wu, W.; Sawada, K.; Zhang, M.; and Jiang, D. 2020. Dance revolution: Long-term dance generation with music via curriculum learning. *arXiv preprint arXiv:2006.06119*.
- Ikemoto, L.; Arikian, O.; and Forsyth, D. 2009. Generalizing motion edits with gaussian processes. *ACM Transactions on Graphics (TOG)*, 28(1): 1–12.
- Jin, Y.; Zhang, J.; Li, M.; Tian, Y.; Zhu, H.; and Fang, Z. 2017. Towards the automatic anime characters creation with generative adversarial networks. *arXiv preprint arXiv:1708.05509*.
- Kingma, D. P.; and Dhariwal, P. 2018. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Lee, H.-Y.; Yang, X.; Liu, M.-Y.; Wang, T.-C.; Lu, Y.-D.; Yang, M.-H.; and Kautz, J. 2019. Dancing to music. *Advances in neural information processing systems*, 32.
- Lee, M.; Lee, K.; and Park, J. 2013. Music similarity-based approach to generating dance motion sequence. *Multimedia tools and applications*, 62: 895–912.
- Li, B.; Zhao, Y.; Zhelun, S.; and Sheng, L. 2022. Danceformer: Music conditioned 3d dance generation with parametric motion transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 1272–1279.
- Li, R.; Yang, S.; Ross, D. A.; and Kanazawa, A. 2021. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13401–13412.
- Loper, M.; Mahmood, N.; Romero, J.; Pons-Moll, G.; and Black, M. J. 2015. SMPL: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6): 1–16.
- Luo, Z.; Chen, D.; Zhang, Y.; Huang, Y.; Wang, L.; Shen, Y.; Zhao, D.; Zhou, J.; and Tan, T. 2023. VideoFusion: Decomposed Diffusion Models for High-Quality Video Generation. *arXiv e-prints*, arXiv–2303.
- Mahmood, N.; Ghorbani, N.; Troje, N. F.; Pons-Moll, G.; and Black, M. J. 2019. AMASS: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5442–5451.

- Min, J.; and Chai, J. 2012. Motion graphs++ a compact generative model for semantic motion analysis and synthesis. *ACM Transactions on Graphics (TOG)*, 31(6): 1–12.
- Mukai, T.; and Kuriyama, S. 2005. Geostatistical motion interpolation. In *ACM SIGGRAPH 2005 Papers*, 1062–1070.
- Ormoneit, D.; Black, M. J.; Hastie, T.; and Kjellström, H. 2005. Representing cyclic human motion using functional analysis. *Image and Vision Computing*, 23(14): 1264–1276.
- O’rourke, J.; and Badler, N. I. 1980. Model-based image analysis of human motion using constraint propagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (6): 522–536.
- Pavlo, D.; Feichtenhofer, C.; Auli, M.; and Grangier, D. 2020. Modeling human motion with quaternion-based neural networks. *International Journal of Computer Vision*, 128: 855–872.
- Petrovich, M.; Black, M. J.; and Varol, G. 2022. TEMOS: Generating diverse human motions from textual descriptions. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, 480–497. Springer.
- Plappert, M.; Mandery, C.; and Asfour, T. 2016. The KIT motion-language dataset. *Big data*, 4(4): 236–252.
- Punnakkal, A. R.; Chandrasekaran, A.; Athanasiou, N.; Quiros-Ramirez, A.; and Black, M. J. 2021. BABEL: Bodies, action and behavior with english labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 722–731.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.
- Rose, C.; Cohen, M. F.; and Bodenheimer, B. 1998. Verbs and adverbs: Multidimensional motion interpolation. *IEEE Computer Graphics and Applications*, 18(5): 32–40.
- Safonova, A.; and Hodgins, J. K. 2007. Construction and optimal search of interpolated motion graphs. In *ACM SIGGRAPH 2007 papers*, 106–es.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35: 36479–36494.
- Siyao, L.; Yu, W.; Gu, T.; Lin, C.; Wang, Q.; Qian, C.; Loy, C. C.; and Liu, Z. 2022. Bailando: 3d dance generation by actor-critic gpt with choreographic memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11050–11059.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, 2256–2265. PMLR.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Sun, G.; Wong, Y.; Cheng, Z.; Kankanhalli, M. S.; Geng, W.; and Li, X. 2020. DeepDance: music-to-dance motion choreography with adversarial learning. *IEEE Transactions on Multimedia*, 23: 497–509.
- Tang, T.; Jia, J.; and Mao, H. 2018. Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis. In *Proceedings of the 26th ACM international conference on Multimedia*, 1598–1606.
- Tevet, G.; Gordon, B.; Hertz, A.; Bermano, A. H.; and Cohen-Or, D. 2022a. Motionclip: Exposing human motion generation to clip space. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, 358–374. Springer.
- Tevet, G.; Raab, S.; Gordon, B.; Shafir, Y.; Cohen-Or, D.; and Bermano, A. H. 2022b. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*.
- Tseng, J.; Castellon, R.; and Liu, C. K. 2022. EDGE: Editable Dance Generation From Music. *arXiv preprint arXiv:2211.10658*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Xin, C.; Jiang, B.; Liu, W.; Huang, Z.; Fu, B.; Chen, T.; Yu, J.; and Yu, G. 2023. Executing your Commands via Motion Diffusion in Latent Space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yu, Z.; Yoon, J. S.; Lee, I. K.; Venkatesh, P.; Park, J.; Yu, J.; and Park, H. S. 2020. Humbi: A large multiview dataset of human body expressions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2990–3000.
- Zhang, J.; Zhang, Y.; Cun, X.; Huang, S.; Zhang, Y.; Zhao, H.; Lu, H.; and Shen, X. 2023. T2M-GPT: Generating Human Motion from Textual Descriptions with Discrete Representations. *arXiv preprint arXiv:2301.06052*.
- Zhang, M.; Cai, Z.; Pan, L.; Hong, F.; Guo, X.; Yang, L.; and Liu, Z. 2022. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*.
- Zhuang, W.; Wang, C.; Chai, J.; Wang, Y.; Shao, M.; and Xia, S. 2022. Music2dance: Dancenet for music-driven dance generation. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 18(2): 1–21.