

# MCM: Multi-condition Motion Synthesis Framework

Zeyu Ling<sup>1†</sup>, Bo Han<sup>1†</sup>, Yongkang Wong<sup>2</sup>, Han Lin<sup>1</sup>, Mohan Kangankhali<sup>2</sup> and Weidong Geng<sup>1</sup>

<sup>1</sup>College of Computer Science and Technology, Zhejiang University

<sup>2</sup>School of Computing, National University of Singapore

{zeyuling, borishan815}@zju.edu.cn, yongkang.wong@nus.edu.sg, h0h972351@gmail.com, mohan@comp.nus.edu.sg, gengwd@zju.edu.cn

## Abstract

Conditional human motion synthesis (HMS) aims to generate human motion sequences that conform to specific conditions, which text and audio are the two predominant modalities employed as conditions. While existing research has primarily focused on single condition, the multi-condition human motion synthesis remains underexplored. In this study, we propose a multi-condition HMS framework, termed MCM, based on a dual-branch structure composed of a main branch and a control branch. This framework effectively extends the applicability of the diffusion model, which is initially predicated solely on textual conditions, to auditory conditions. This extension encompasses both music-to-dance and co-speech HMS while preserving the intrinsic quality of motion and the capabilities for semantic association inherent in the original model. Furthermore, we propose the implementation of a Transformer-based diffusion model, designated as MWNet, as the main branch. This model adeptly apprehends the spatial intricacies and inter-joint correlations inherent in motion sequences, facilitated by the integration of multi-wise self-attention modules. Extensive experiments show that our method achieves competitive results in single-condition and multi-condition HMS tasks.

## 1 Introduction

Human motion synthesis finds extensive applications in fields such as film production, game development, and simulation. However, traditional manual animation techniques are notably constrained in terms of efficiency. Following the advent of neural network-based generative models, a variety of these models, such as Variational Autoencoders (VAEs) [Kingma and Welling, 2013], Generative Adversarial Networks [Goodfellow *et al.*, 2014], Denoising Diffusion Probabilistic Models (DDPM) [Ho *et al.*, 2020] have been adapted and refined for the specific domain of human motion generation to achieve high-fidelity results.

Current conditional human motion generation, such as text-to-motion [Guo *et al.*, 2022a; Zhang *et al.*, 2024] and

music-to-dance [Siyao *et al.*, 2022; Tseng *et al.*, 2023] focuses on generating human motion from a singular modality condition. The integration of condition information from different modalities is an area that has not been extensively explored. The dual-modality-driven 3D motion generation task presents the following primary challenges:

- Current HMS datasets predominantly encompass unimodal conditions. For instance, HumanML3D [Guo *et al.*, 2022a] and AIST++ [Li *et al.*, 2021] only consider text modality and music modality, respectively.
- Text and audio, being temporally non-aligned modalities, pose significant challenges in ensuring both temporal and semantic correlation with 3D human motion spontaneously.
- Fine-tuning directly from an existing single-condition pre-trained model may necessitate structural adjustments to the existing model architecture on one hand, and, on the other hand, fine-tuning may lead to catastrophic forgetting of the pre-trained conditional association capabilities.

Typically, TM2D [Gong *et al.*, 2023a] and UDE [Zhou and Wang, 2023a] tokenize motion from HumanML3D [Guo *et al.*, 2022a] and AIST++ [Li *et al.*, 2021] using Motion VQ-VAE [Van Den Oord *et al.*, 2017]. In the generation phase, they use text and sound respectively to drive GPT [Hudson and Zitnick, 2021] to generate motion token sequences and then integrate these sequences by using weighted merging or replacement methods to achieve multi-condition driven. Employing the post-fusion approach circumvents the need to construct large-scale text-sound-motion data pairs; however, it harbors a critical drawback: *each motion token in the sequence is essentially generated under the influence of a single modality condition*. Tokens generated based on text do not conform to auditory conditions, while those generated from sound fail to align with semantic information.

To address the above challenges, we propose a novel end-to-end framework **MCM** (Multi-Condition Motion synthesis), which is tailored for multi-condition driven 3D human motion synthesis. MCM adopts a dual-branch structure, comprising the main branch and the control branch. Initially, the main branch leverages an arbitrary pre-trained text2motion DDPM network like MotionDiffuse [Zhang *et al.*, 2024] and MDM [Tevet *et al.*, 2022], ensuring the motion quality and semantic coherence during multi-condition motion synthesis as shown in Figure 1. Subsequently, the control branch ini-

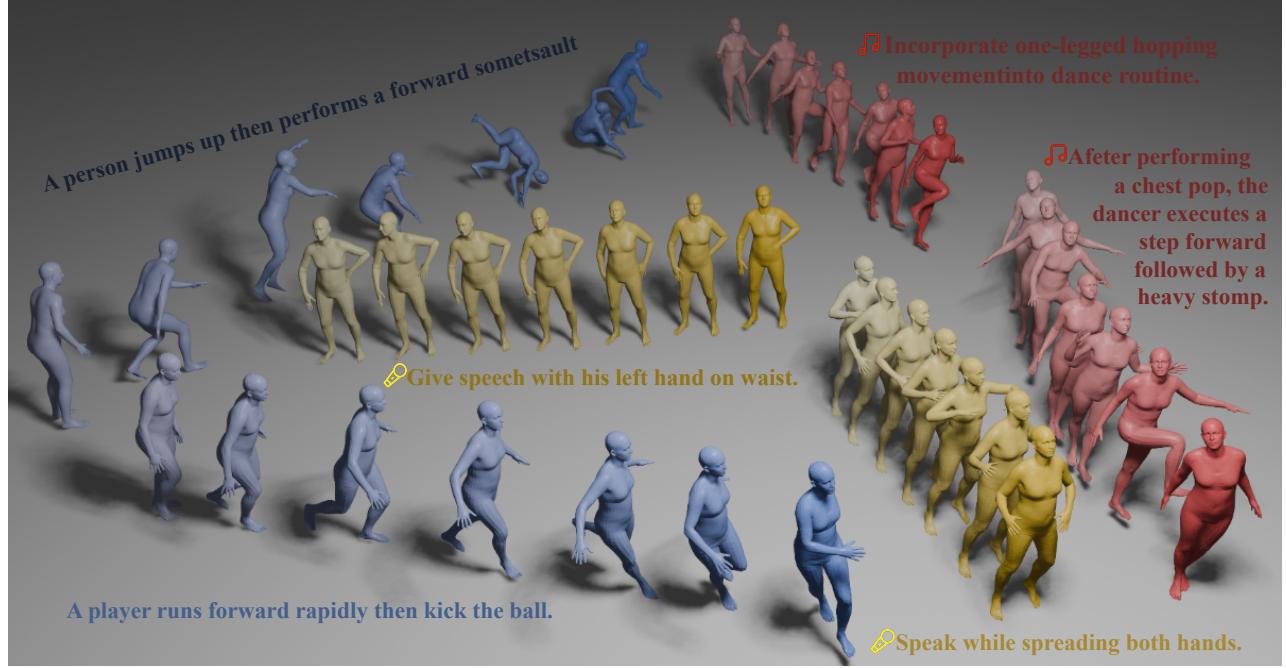


Figure 1: Samples of Multi-Condition Motion (MCM) synthesis, where human motions can be generated across various scenarios (*i.e.*, **text-to-motion**, **dance**, and **co-speech** motion synthesis) under multiple conditions. The generated motion not only conforms to the rhythm of the music and speech but also exhibits consistency with the textual descriptions of dance and gestures.

tializes its parameters, mirroring the structure of the main branch, and assumes the responsibility of modifying the motion in accordance with the auditory conditions. During the training of tasks conditioned on sound, MCM freezes the parameters of the main branch and activates those of the control branch. This approach is designed to preserve the motion quality and semantic correlation capabilities of the main branch without compromise. This methodology obviates the need for collecting a text-audio-motion dataset, yet enables both audio and text conditions to simultaneously influence every part of the motion sequence.

Regarding the structural design of the main branch, existing text2motion models [Zhang *et al.*, 2024; Tevet *et al.*, 2022] employ self-attention along the temporal dimension to grasp the sequential associations of motion sequences and perform cross-attention with textual descriptions to capture semantic information. However, when considering motion sequences, it is imperative to recognize that the channel dimension holds valuable spatial information and inter-joint relationships within the human body, aspects that have often been underappreciated. In this work, we propose a Transformer-based DDPM network, termed **MWNet**, that incorporates a **Multi-Wise** self-attention mechanism (channel-wise self-attention) to better learn spatial information.

In summary, our core contributions are as follows:

- We introduce a unified MCM framework for 3D human motion synthesis based on multiple conditions. Remarkably, without necessitating structural reconfiguration of the network, MCM extends the capabilities of DDPM-based methods to sound-conditional inputs.

- We propose a Transformer-based architecture **MWNet**, enriched with a multi-wise attention mechanism, which leverages spatial information within motions to achieve better motion generation quality and comparable semantic matching capacity.
- Extensive experiments show that our method demonstrates competitive performance in the single-condition-driven tasks (text2motion and music2dance) and multi-condition-driven tasks. Furthermore, we also present an ablation analysis to elucidate the contribution of each component, enhancing understanding of their individual and collective impact on the system’s performance.

## 2 Related Work

### 2.1 Text-to-motion

Text-to-motion converts descriptive text into motion sequences using various generative models. VAE-based methods [Guo *et al.*, 2022a; Guo *et al.*, 2022b; Petrovich *et al.*, 2022] learn motion distribution in latent space but are limited by VAE’s Gaussian posterior estimation, resulting in sub-optimal results. Recently, diffusion models have shown impressive image synthesis capabilities, prompting increased research into their application for motion generation.

MotionDiffuse [Zhang *et al.*, 2024] and MDM [Tevet *et al.*, 2022] employ DDPM for human motion synthesis. MotionDiffuse features a linear self-attention mechanism, whereas MDM uses a vanilla Transformer Encoder and predicts the original motion sequence at each time step, instead of step-specific noise. MLD [Chen *et al.*, 2023] utilizes

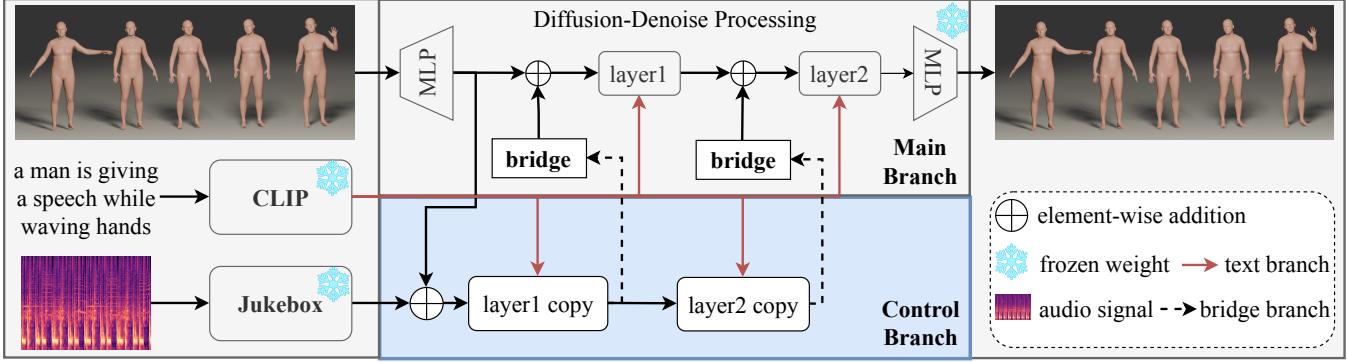


Figure 2: Overview of the simplified two-layer MCM framework. MCM employs a dual-branch structure consisting of the main branch and the control branch. The layer-wise outputs from the control branch are connected to the main branch via bridge modules, which are fully connected layers or 1d-convolutions with parameters initialized to zero. The output of each bridge module is directly added to the input feature vector of corresponding layers in the main branch. The condition encoders encompass several pre-trained feature extractors for different modal conditions.

a latent vector DDPM for forward noising and reverse denoising in motion latent space. MAA [Azadi *et al.*, 2023] enhances performance on out-of-distribution data by pre-training a diffusion model on a large dataset of (text, static pseudo-pose) pairs. ReMoDiffuse [Zhang *et al.*, 2023c] introduces a retrieval-enhanced DDPM, significantly improving in-distribution text-to-motion capabilities.

With advancements in Large Language Models, autoregressive generative models are now used to link two different modalities. TM2T [Guo *et al.*, 2022b], T2MGPT [Zhang *et al.*, 2023a], and MotionGPT [Jiang *et al.*, 2023] approach text-to-motion by using VQ-VAE [Van Den Oord *et al.*, 2017] to discretize motion sequences and employ GRU [Chung *et al.*, 2014], GPT [Vaswani *et al.*, 2017; Radford *et al.*, 2019], or T5 [Raffel *et al.*, 2020] models to learn the correlations between motion and text tokens. HumanTomato [Lu *et al.*, 2023] proposes using hierarchical VQ-VAEs and GPT for the generation of full-body motion sequences.

## 2.2 Music-to-dance

The music-to-dance task involves generating dance movements synchronized with music beats. DanceNet [Zhuang *et al.*, 2022] uses an LSTM-based auto-regressive method to create dance movements. DeepDance [Sun *et al.*, 2020] employs Generative Adversarial Networks for human motion generation. FACT [Li *et al.*, 2021] introduces the Full-Attention Cross-modal Transformer that utilizes self-attention among music, dance movements, and combined music-dance tokens. Bailando [Siyao *et al.*, 2022] utilizes a VQ-VAE encoding for separate upper and lower body parts, integrating them into full-body dance sequences using GPT. EDGE [Tseng *et al.*, 2023] employs a Transformer-based DDPM for precise music-rhythm matching and introduces a contact consistency loss to address physical errors in foot movements. EnchantDance [Han *et al.*, 2023] develops a robust dance latent space to train a dance diffusion model.

## 2.3 Multi-condition HMS

The aforementioned method has a limitation in that it can only accept a single modality condition as input, which is insufficient to meet the requirements of many scenarios. Some methods have made efforts to integrate multiple conditions and accommodate various scenarios. TriModal [Yoon *et al.*, 2020] and CaMN [Liu *et al.*, 2022] generate co-speech motion based on speech and spoken text, with CaMN integrating additional conditions such as facial expressions. MoFusion [Dabral *et al.*, 2023] for the first time, implements a unified UNet-based DDPM architecture, which, through two different types of conditional encoders, can effectively function in both text-to-motion and music-to-dance tasks. UDE [Zhou and Wang, 2023b] and TM2D [Gong *et al.*, 2023b] both adopt a model structure based on VQ-VAE and GPT and propose a similar approach to integrate music and text conditions. They generate motion token sequences conditioned on sound and text respectively, and then either weight or concatenate the generated logits to obtain a smooth motion sequence that blends motion tokens conditioned on both text and audio.

## 3 Method

### 3.1 Problem Definition

The objective of the multi-condition human motion synthesis task is to generate a motion sequence  $X \in \mathbb{R}^{T \times D}$  under a set of constraint conditions  $C$ .  $X$  is an array of  $x_i$ , where  $i \in \{1, 2, \dots, T\}$ , and  $T$  denotes the number of frames. Each  $x_i \in \mathbb{R}^D$  represents the  $D$ -dimensional pose state vector at the  $i$ -th frame.  $c_j \in C$  could be textual description, speech voice, or background music.

### 3.2 MCM Framework

As depicted in Figure 2, MCM framework employs a dual-branch architecture, comprising a pre-trained main branch with frozen parameters and a trainable control branch. The main branch is composed of any neural network grounded in DDPM, such as MotionDiffuse or MDM. In this study, we have devised and implemented a DDPM network named

MCM as the main branch of MCM, which will be comprehensively detailed in Section 3.3. The control branch mirrors the structural framework of the main branch and is initialized utilizing the parameters derived from the main branch.

Drawing inspiration from ControlNet [Zhang *et al.*, 2023b], we implement a strategy of optimizing the main branch and control branch independently during distinct stages of the training process. We pre-train the main branch on the text-to-motion task, with the objective of this phase being to endow the MCM with foundational motion quality and semantic association capabilities. In the training phase of audio-to-motion tasks, all parameters, excluding those assigned to the control branch and bridge module, are maintained in a fixed state. This strategy is implemented to guarantee the retention of the main branch’s generative quality and its capabilities in semantic association.

For the input to the Control branch, we perform element-wise addition of the Jukebox features with the motion latent vector to incorporate audio information into MCM. The output of each layer of the Control branch is added to the input of the main branch through the bridge module, thereby causing a slight offset in the output of the main branch under the control of audio conditions. By performing zero-initialization on the bridge module, we ensure that the initial output of MCM is identical to that of the main branch and gradually adjust the parameters according to the audio conditions during the iterative process.

### 3.3 MWNet Architecture

As previously elucidated, the quality of the main branch exerts a substantial impact on the final quality and semantic association capacities of the motions generated by MCM. Consequently, we delved deeper into the structural aspects of the main branch in the context of the text-to-motion task and proposed the MWNet model.

As delineated in Section 3.1, the channel dimension of the motion sequence encompasses the entire spatial information pertinent to the motion sequence. Nevertheless, current DDPM-based motion generation models, including Motion-Diffuse and MDM, predominantly concentrate on time-wise self-attention and cross-attention mechanisms. These are employed to model the correlations at the temporal level and the semantic level between motions and their respective conditions. In the spatial dimension, information is currently modeled exclusively through linear layers and activation functions. We posit that this approach is inadequate for capturing the complex spatial relationships present in intricate motion sequences, such as the correlations between joints. Consequently, we advocate for the use of channel-wise self-attention [Ding *et al.*, 2022] and introduce the Multi-Wise attention mechanism, specifically designed to model these critical aspects of spatial information. MWNet is composed of layers arranged in the configuration depicted in Figure 3 (a).

Similar to Stable Diffusion [Rombach *et al.*, 2022] and GLIDE [Nichol *et al.*, 2021], we use FiLM [Perez *et al.*, 2018] blocks to furnish timestamp information to MWNet after every attention or Feed Forward Network (FFN) modules. In the FiLM module, the weights and biases for the affine transformation performed on the motion latent vector are de-

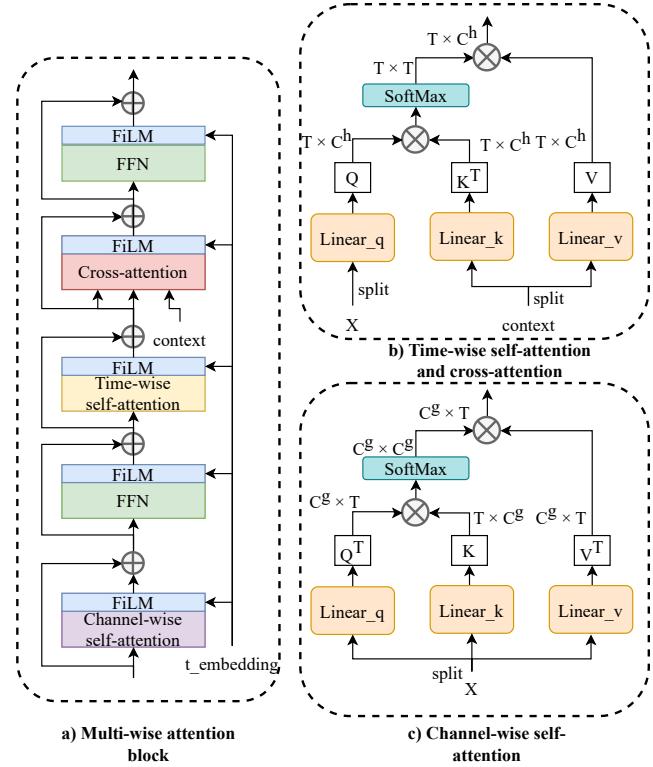


Figure 3: Model architecture for a multi-wise attention block. It incorporates three distinct types of attention modules, which are employed alternately. The symbols “+” and “ $\times$ ” separately represent feature addition and multiplication operations.  $T$  symbolizes the length of the input sequence, while  $C^g$  and  $C^h$  signify the number of channels for the matrices  $Q$ ,  $K$ , and  $V$  after. The split operation means splitting the channels into  $g$  groups or  $h$  heads. Context represents text condition for cross-attention and is exactly equal to  $X$  for time-wise self-attention.

rived from the timestep mapping, as shown in the following:

$$FiLM(x, \epsilon_t) = x + LN(x \odot (W_1 + I)\epsilon_t) + W_2\epsilon_t \quad (1)$$

$LN$  denotes layer normalization module [Ba *et al.*, 2016].  $W_1$ , and  $W_2$  are two projection matrices.  $I$  denotes a matrix wherein all elements are 1, possessing a shape congruent with that of  $x$ .  $\odot$  denotes the element-wise multiplication.

With projection weights  $W^Q$ ,  $W^K$ ,  $W^V$ ,  $X$  are projected to  $Q = XW^Q$ ,  $K = XW^K$ ,  $V = XW^V$  and split into  $N_h$  heads or  $N_g$  groups. We denote  $Q_i$ ,  $K_i$ , and  $V_i$  for each head or group. Time-wise self-attention can be denoted as follows:

$$SA_T(Q_i, K_i, V_i) = \text{Softmax}\left(\frac{Q_i K_i^T}{\sqrt{C_h}}\right) V_i \quad (2)$$

$$SA_T(Q, K, V) = \{SA_T(Q_i, K_i, V_i)\}_{i=0}^{N_h} \quad (3)$$

whereas, channel-wise self-attention can be denoted as:

$$SA_C(Q_i, K_i, V_i) = \left( \text{Softmax}\left(\frac{Q_i^T K_i}{\sqrt{C_g}}\right) V_i^T \right)^T \quad (4)$$

$$SA_C(Q, K, V) = \{SA_C(Q_i, K_i, V_i)\}_{i=0}^{N_g} \quad (5)$$

where  $C_h$  and  $C_g$  denote the channel numbers for head and group, respectively.

## 4 Experiments

### 4.1 Dataset

We evaluate the proposed method using the HumanML3D [Guo *et al.*, 2022a], AIST++ [Li *et al.*, 2021], and BEAT [Liu *et al.*, 2022] dataset. In all experiments, we process the motion of all involved datasets to the same format with a 22-joint skeleton (the first 22 joints of the SMPL skeleton) and 20 FPS. Subsequently, following HumanML3D, we use a 263-dimension representation  $x = \text{concat}(\dot{r}^a, \dot{r}^x, \dot{r}^z, r^y, j^p, j^v, j^r, c^f)$  to represents motions at every frame.  $\dot{r}^a \in \mathbb{R}$  is root angular velocity along the Y-axis;  $\dot{r}^x, \dot{r}^z \in \mathbb{R}$  are root linear velocities on XZ-plane;  $r^y$  is root height;  $j^p, j^v \in \mathbb{R}^{3j}$  and  $j^r \in \mathbb{R}^{6j}$  are the local joints positions, velocities, and rotations in root space, with  $j$  denoting the number of joints;  $c^f \in \mathbb{R}^4$  is binary features obtained by thresholding the heel and toe joint velocities to emphasize the foot ground contacts.

We also processed AIST++ dataset and BEAT dataset to generate pseudo text descriptions, thereby forming audio-text-motion sample pairs for training in the control stage.

### 4.2 Evaluation Metrics

For text2motion evaluation, we employ the established metrics [Guo *et al.*, 2022a] to assess the generated motion sequences across various dimensions. (1) **Motion quality**: Frechet Inception Distance (FID) [Onuma *et al.*, 2008] evaluates the dissimilarity between two distributions by calculating the difference between feature vector (extracted by feature extractor) statistical measures (mean and covariance). (2) **Diversity**: The Diversity and MultiModality metrics respectively assess the degree of variation in the generated motions with different text inputs and with the same text input. (3) **Semantic matching**: The motion-retrieval precision (R Precision) evaluates the accuracy of matching between texts and motions using Top 1/2/3 retrieval accuracy. Multi-modal Distance (MM Dist) measures the distance between motions and text features extracted by the feature extractor trained with contrastive learning.

For music2dance evaluation, we adopt the commonly used evaluation metrics following [Dabral *et al.*, 2023]. (1) **FID**: utilizing kinetic and geometry features implemented within fairmotion [Gopinath and Won, 2020]. (2) **Diverisy**: it computes the average pairwise Euclidean distance of the kinetic and geometry features of the motions generated from music in the validation and test set. (3) **Beat Alignment Score (BAS)**: a metric that quantifies the congruence between kinematic beats and musical beats. Kinematic beats correspond to the local minima of kinetic velocity within a motion sequence, signifying points where motion momentarily halts.

### 4.3 Implementation Details

We conduct training of MCMs utilizing distinct DDPM-like main branch architecture, including MotionDiffuse [Zhang *et al.*, 2024], MDM [Tevet *et al.*, 2022], and our MWNet. The conditioning inputs from diverse modalities are pre-processed through the employment of pre-trained condition encoders. We use CLIP [Radford *et al.*, 2021] to extract features from text conditions. Similar to EDGE, we use the prior layer

of Jukebox [Dhariwal *et al.*, 2020] to extract features from all audio conditions(music and human voice) and downsample them to the same frame rate as the motion samples (20 FPS). Regarding the diffusion model, we set the number of diffusion steps at 1000, while the variances  $\beta_t$  follow a linear progression from 0.0001 to 0.02. We employ the Adam optimizer for training the model, employing a learning rate of 0.0002 throughout both training phases. In concordance with MDM, we adopt the strategy of predicting  $x_{start}$  as an alternative to the prediction of noise. This approach is aimed at achieving an enhanced quality of motion.

### 4.4 Music-to-Dance Evaluation

During the training phase on the AIST++ dataset, we processed the training data into motion segments with a maximum length of 196 frames, equivalent to 9.8 seconds. In the evaluation stage, prior studies were assessed using the complete dance motion sequences from the validation and test sets of AIST++ dataset, maintaining the original frame rate of 60 FPS. To facilitate an equitable comparison with these studies, we linearly interpolated the generated motion sequences, initially at 20 FPS, to upscale them to 60 FPS. Table 4 presents the evaluation outcomes of MCM on the AIST++ dataset. When utilizing MCM, three distinct DDPM methodologies, previously trained on the text2motion dataset, all demonstrated performance on par with state-of-the-art (SOTA) methods in the music2dance task. As demonstrated in Table 4, our approach achieved the best  $FID_k$ , while UDE achieved the best  $FID_g$ , and the diversity of the dance movements generated by each method was found to be remarkably similar. Furthermore, our methodology, while preserving a comparable level of dance motion quality and diversity, attains SOTA metrics in beat alignment exclusively via the elementary element-wise integration of Jukebox music features. Among various implementations utilizing the MCM framework, MWNet + MCM demonstrates relatively superior dance motion quality and beat alignment. Conversely, MotionDiffuse + MCM excels in achieving enhanced diversity in its outputs.

### 4.5 Text-to-motion Evaluation

We train and evaluate our main branch model MWNet on HumanML3D dataset. The quantitative results are shown in Table 1. Overall, we achieved the best motion quality and next to best semantic relevance among all methods. Moreover, it surpasses all non-DDPM methods. Owing to its retrieval-based mechanism, ReMoDiffuse remains the model with the most robust semantic correlation capabilities. Compared to the TM2D and MoFusion approach, which similarly accommodates the integration of auditory conditions, our method exhibits a distinct advantage in both semantic relevance and the quality of motion. A robust capability for semantic association lays the groundwork for maintaining strong semantic relational abilities in multi-condition HMS scenarios. Our method exhibits moderate performance on the MultiModality metric. The underlying causes of this phenomenon were explored in our ablation studies.

Methods	R Precision $\uparrow$			FID $\downarrow$	MultiModal Dist $\downarrow$	Diversity $\rightarrow$	MultiModality $\uparrow$
	Top 1	Top 2	Top 3				
Real motions	0.511 $\pm$ .003	0.703 $\pm$ .003	0.797 $\pm$ .002	0.002 $\pm$ .000	2.974 $\pm$ .008	9.503 $\pm$ .065	-
T2M	0.457 $\pm$ .002	0.639 $\pm$ .003	0.740 $\pm$ .003	1.067 $\pm$ .002	3.340 $\pm$ .008	9.188 $\pm$ .002	2.090 $\pm$ .083
TM2T	0.424 $\pm$ .003	0.618 $\pm$ .003	0.729 $\pm$ .002	1.501 $\pm$ .017	3.467 $\pm$ .011	8.589 $\pm$ .076	2.424 $\pm$ .093
TEMOS	0.424 $\pm$ .002	0.612 $\pm$ .002	0.722 $\pm$ .002	3.734 $\pm$ .028	3.703 $\pm$ .008	8.973 $\pm$ .071	0.368 $\pm$ .018
TM2D	-	-	0.319	1.021	4.098	<b>9.513</b>	<b>4.139</b>
T2MGPT	0.491 $\pm$ .003	0.680 $\pm$ .003	0.775 $\pm$ .002	0.116 $\pm$ .004	3.118 $\pm$ .011	9.761 $\pm$ .081	1.856 $\pm$ .011
MotionGPT	0.492 $\pm$ .003	0.681 $\pm$ .003	0.778 $\pm$ .002	0.232 $\pm$ .008	3.096 $\pm$ .008	<u>9.528</u> $\pm$ .071	2.008 $\pm$ .084
AttT2M	0.499 $\pm$ .003	0.690 $\pm$ .002	0.786 $\pm$ .002	0.112 $\pm$ .006	3.038 $\pm$ .007	9.700 $\pm$ .090	2.452 $\pm$ .051
MoFusion	-	-	0.492	-	-	8.82	2.521
MLD	0.481 $\pm$ .003	0.673 $\pm$ .003	0.772 $\pm$ .002	0.473 $\pm$ .013	3.196 $\pm$ .010	9.724 $\pm$ .082	2.413 $\pm$ .079
MDM	0.320 $\pm$ .005	0.498 $\pm$ .004	0.611 $\pm$ .007	0.544 $\pm$ .044	5.566 $\pm$ .027	9.559 $\pm$ .086	2.799 $\pm$ .072
MotionDiffuse	0.491 $\pm$ .001	0.681 $\pm$ .001	0.782 $\pm$ .001	0.630 $\pm$ .001	3.113 $\pm$ .001	9.410 $\pm$ .049	1.553 $\pm$ .042
ReMoDiffuse	<b>0.510</b> $\pm$ .005	<b>0.698</b> $\pm$ .006	<b>0.795</b> $\pm$ .004	0.103 $\pm$ .004	<b>2.974</b> $\pm$ .016	9.018 $\pm$ .075	1.795 $\pm$ .043
<b>MWNet(ours)</b>	<u>0.502</u> $\pm$ .002	<u>0.692</u> $\pm$ .004	<u>0.788</u> $\pm$ .006	<b>0.053</b> $\pm$ .007	<u>3.037</u> $\pm$ .003	9.585 $\pm$ .082	0.8104 $\pm$ .023

Table 1: Quantitative results on the HumanML3D test set. The symbol  $\rightarrow$  denotes that the results are more favorable when the metric closely approximates the distribution of real motions (i.e., the metrics of authentic movements). The methodologies are categorized based on their reliance on DDPM. A demarcation line is utilized to distinguish the approaches, with those situated below the line being DDPM-based. The **bold fonts** indicates the best results, while the underline fonts denotes the next to best. Due to the inherent randomness of the metrics, most methods were evaluated 20 times to calculate the mean and variance (as superscripts), while TM2D and MoFusion didn't provide the variance.

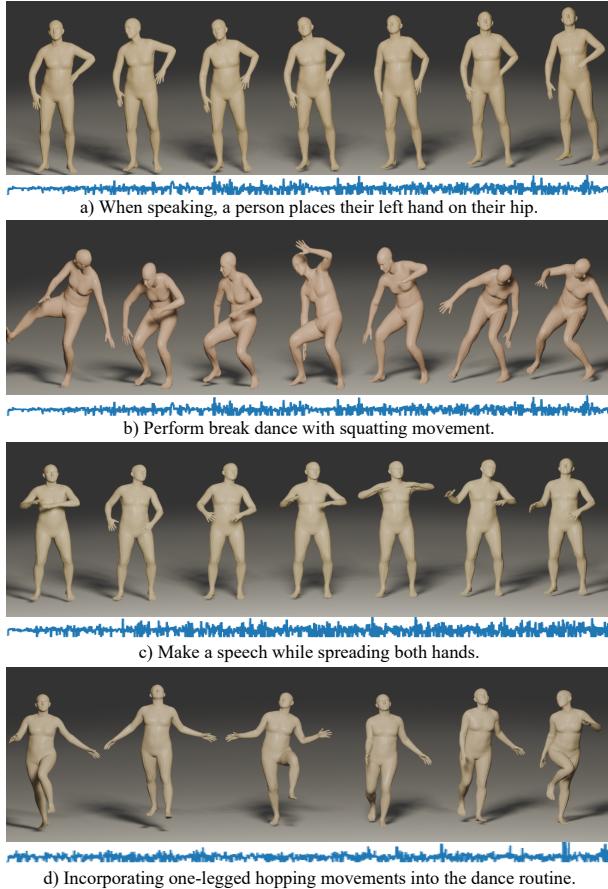


Figure 4: Text-sound multi-condition motion synthesis with MCM (MWNet as the main branch). Each sample is obtained using a segment of text and a segment of audio as inputs.

Data	beat align	text match	motion quality
AIST++	95.0%	95.5%	93.0%
Wild	69.0%	74.0%	62.0%

Table 2: We request users to compare the quality of the multi-condition HMS from three aspects: beat alignment, dance motion quality, and semantic match. Each cell in the table represents the win rate of MCM in each comparative aspect.

#### 4.6 Text-sound Multi-condition Generation

The outputs of text-sound multi-condition generation are illustrated in Figure 4. Under the multi-condition HMS scenario, we generate corresponding dance, or co-speech movements, under the joint control of textual and auditory conditions. For instance, we adjust specific gestures during a speech and particular dance moves in a dance through text, achieving interesting control effects.

In order to conduct a fair comparison with existing methods capable of integrating textual and auditory conditions. We sampled 10 music segments from the AIST++ dataset and 5 segments from outside the AIST++ dataset, and designed distinct dance motion descriptions for each piece of music. We employed TM2D and MCM to generate dance movements under multi-conditional constraints, specifically, to produce dance sequences that conform to textual descriptions and align with musical rhythms. Subsequently, we invited 20 users to conduct a comparative assessment of the generated outcomes. The subjects were asked to make comparisons from three aspects: beat alignment, text match, and motion quality. Importantly, the entire evaluation process was meticulously designed to be anonymous, thereby ensuring that participants were unable to discern the origins of the data samples, which in turn guaranteed an unbiased assessment of the different models. As shown in Table 2, the results of the

Methods	R Precision $\uparrow$			FID $\downarrow$	MultiModal Dist $\downarrow$	Diversity $\rightarrow$	MultiModality $\uparrow$
	Top 1	Top 2	Top 3				
T/CA/F	0.389	0.547	0.668	1.081	3.772	<b>9.393</b>	<b>3.033</b>
T/CA/F/T/F	0.423	0.600	0.708	0.825	3.577	9.018	2.788
T/F/T/CA/F	0.431	0.617	0.720	0.586	3.472	9.157	2.152
T/CA/F/CS/F	<b>0.455</b>	<b>0.642</b>	<b>0.744</b>	0.751	3.399	8.933	1.707
CS/F/T/CA/F (MWNet)	<b>0.455</b>	0.640	0.742	<b>0.377</b>	<b>3.349</b>	9.312	1.481

Table 3: The results on the HumanML3D dataset after training Transformer Decoder modules with different structures for 500 epochs.

user study reveal a distinct preference among users for our outcomes on the AIST++ dataset. Moreover, our results also demonstrate a relative advantage in external datasets. This is attributed to the fact that in MCM, both text and audio conditions concurrently influence the entire motion sequence. In contrast, TM2D’s post-fusion approach results in each motion token being driven effectively by only a single modality. Consequently, the actions generated from the text do not align with the rhythm of the music. Simultaneously, our superiority in semantic relevance is largely attributable to the significantly stronger semantic correlation capabilities of our main branch compared to TM2D, as evident from the evaluation results on the HumanML3D dataset.

## 4.7 Ablation Study

### Dual-branch vs Single-branch

In this section, we further explore the superiority of the Dual-branch structure in transferring from text-condition scenarios to sound-condition scenarios. The dual branch structure trains the main branch on the text2motion task, and the control branch on the sound condition; whereas for the single branch structure, we directly optimize the parameters of the main branch during the control stage, and the audio condition is added directly to the motion latent vector as input. In the lower half of Table 4, we compare the performance differences on the AIST++ dataset between dual-branch and single-branch structures equipped with different main branches. The results show that all dual-branch outcomes surpass their single-branch counterparts. Additionally, the single-branch structure no longer possesses the semantic

association obtained from text-to-motion pretraining.

### Multi-wise Self-attention Module Design

We designed experiments to verify the necessity and superiority of multi-wise self-attention module. Our Multi-wise self-attention module consists of a time-wise self-attention (denoted as TA), a channel-wise self-attention (CS), a cross-attention (CA), and two FFNs (F). By modifying or reordering these modules, we compared the performance metrics of different configurations on the HumanML3D benchmark. In this experiment, each trial was limited to training for only 500 epochs to conserve experimental time. We experimented with arranging or adjusting the attention modules for different categories. As shown in Table 3, comparing T/CA/F/T/F and T/CA/F/CS/F, it can be observed that substituting CS for T results in significantly improved motion quality, as CS compensates for the spatial information deficiencies of T. It was found that using both T and CS in the module simultaneously helps significantly in enhancing semantic relevance and motion quality. Additionally, we observed that placing CS before T and CA can improve motion quality with a slight loss in semantic relevance. We also found that the downside of CS is that it leads to reduced MultiModality, which we aim to explore and address in future work. Ultimately, we chose the option that offered the highest motion quality as the module design of MWNet.

## 5 Conclusion

We propose MCM, a novel framework for the multi-conditioned motion synthesis method that spans multiple scenarios. Utilizing MCM, methods based on DDPM for text2motion tasks can concurrently adapt to multi-conditions without necessitating any structural modifications. Additionally, we introduce a Transformer-based architecture MWNet that incorporates multi-wise self-attention, enhancing the modeling of spatial information and inter-joint correlations. We quantitatively evaluate our approach across tasks based on various modal conditions. In the text-to-motion task, our results surpass all existing methods in terms of motion quality and are comparable in semantic correlation capabilities. In the music-to-dance task, we achieved the best beat alignment and  $FID_k$  metrics, along with comparable dance diversity and quality. In the multi-condition HMS domain, our approach accomplishes the concurrent generation of dance movements based on both text and sound, as well as co-speech motion.

Methods	$FID_k \downarrow$	$FID_g \downarrow$	$Div_k \rightarrow$	$Div_g \rightarrow$	$BAS \uparrow$
Ground Truth	17.10	10.60	8.19	7.45	0.2374
DanceNet	69.18	25.49	2.86	2.85	0.143
DanceRevolution	73.42	25.92	2.86	3.52	0.195
FACT	35.35	22.11	5.94	6.18	0.2209
Bailando	28.16	<u>9.62</u>	<b>7.83</b>	6.34	0.233
UDE	<u>17.25</u>	<b>8.69</b>	7.78	5.81	0.231
TM2D	19.01	20.09	9.45	<u>6.36</u>	0.204
MoFusion	50.31	-	9.09	-	0.253
EDGE	-	23.04	-	-	0.270
MotionDiffuse + finetune	45.22	13.21	7.27	5.39	<u>0.273</u>
MotionDiffuse + MCM	44.27	13.43	9.24	5.29	0.266
MDM + finetune	47.39	22.07	8.94	5.16	0.258
MDM + MCM	39.21	19.55	5.81	6.34	0.265
MWNet + finetune	34.73	20.25	5.87	<b>6.79</b>	0.249
MWNet + MCM	<b>15.57</b>	25.85	6.50	5.74	<b>0.275</b>

Table 4: Results on AIST++ validation and test set. The method labeled “finetune” employs the single-branch structure introduced in Section 4.7, replacing the dual-branch structure of MCM.

## Acknowledgments

This work was supported in part by the Zhejiang Province (China) Research Project 2023C01045 and National Natural Scientific Foundation of China 61972346. Corresponding author: Weidong Geng. † Equal contribution.

## References

[Azadi *et al.*, 2023] Samaneh Azadi, Akbar Shah, Thomas Hayes, Devi Parikh, and Sonal Gupta. Make-an-animation: Large-scale text-conditional 3d human motion generation. In *International Conference on Computer Vision*, 2023.

[Ba *et al.*, 2016] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

[Chen *et al.*, 2023] 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.

[Chung *et al.*, 2014] Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Workshop on Deep Learning, December 2014*, 2014.

[Dabral *et al.*, 2023] Rishabh Dabral, Muhammad Hamza Mughal, Vladislav Golyanik, and Christian Theobalt. MoFusion: A framework for denoising-diffusion-based motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9760–9770, 2023.

[Dhariwal *et al.*, 2020] Prafulla Dhariwal, Heewoo Jun, Christine Payne, Jong Wook Kim, Alec Radford, and Ilya Sutskever. Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*, 2020.

[Ding *et al.*, 2022] Mingyu Ding, Bin Xiao, Noel Codella, Ping Luo, Jingdong Wang, and Lu Yuan. DaViT: Dual attention vision transformers. In *European Conference on Computer Vision*, pages 74–92. Springer, 2022.

[Gong *et al.*, 2023a] Kehong Gong, Dongze Lian, Heng Chang, Chuan Guo, Zihang Jiang, Xinxin Zuo, Michael Bi Mi, and Xinchao Wang. TM2D: Bimodality Driven 3D Dance Generation via Music-Text Integration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9942–9952, 2023.

[Gong *et al.*, 2023b] Kehong Gong, Dongze Lian, Heng Chang, Chuan Guo, Zihang Jiang, Xinxin Zuo, Michael Bi Mi, and Xinchao Wang. TM2D: Bimodality driven 3d dance generation via music-text integration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9942–9952, 2023.

[Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[Gopinath and Won, 2020] Deepak Gopinath and Jungdam Won. Fairmotion-tools to load, process and visualize motion capture data. <https://github.com/facebookresearch/fairmotion>, 2020.

[Guo *et al.*, 2022a] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5152–5161, 2022.

[Guo *et al.*, 2022b] 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.

[Han *et al.*, 2023] Bo Han, Yi Ren, Hao Peng, Teng Zhang, Zeyu Ling, Xiang Yin, and Feilin Han. EnchantDance: Unveiling the Potential of Music-Driven Dance Movement. *arXiv preprint arXiv:2312.15946*, 2023.

[Ho *et al.*, 2020] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[Hudson and Zitnick, 2021] Drew A Hudson and Larry Zitnick. Generative adversarial transformers. In *International conference on machine learning*, pages 4487–4499. PMLR, 2021.

[Jiang *et al.*, 2023] 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*, 2023.

[Kingma and Welling, 2013] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. 2013.

[Li *et al.*, 2021] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. AI choreographer: Music conditioned 3d dance generation with AIST++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13401–13412, 2021.

[Liu *et al.*, 2022] Haiyang Liu, Zihao Zhu, Naoya Iwamoto, Yichen Peng, Zhengqing Li, You Zhou, Elif Bozkurt, and Bo Zheng. BEAT: A large-scale semantic and emotional multi-modal dataset for conversational gestures synthesis. In *European Conference on Computer Vision*, pages 612–630. Springer, 2022.

[Lu *et al.*, 2023] 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.

[Nichol *et al.*, 2021] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *International Conference on Machine Learning*, 2021.

[Onuma *et al.*, 2008] Kensuke Onuma, Christos Faloutsos, and Jessica K Hodgins. FMDistance: A Fast and Effective Distance Function for Motion Capture Data. In *Eurographics (Short Papers)*, pages 83–86, 2008.

[Perez *et al.*, 2018] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual reasoning with a general conditioning layer, 2018.

[Petrovich *et al.*, 2022] Mathis Petrovich, Michael J Black, and Gülcin Varol. TEMOS: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision*, pages 480–497. Springer, 2022.

[Radford *et al.*, 2019] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.

[Radford *et al.*, 2021] 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.

[Raffel *et al.*, 2020] 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.

[Rombach *et al.*, 2022] 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.

[Siyao *et al.*, 2022] Li Siyao, Weijiang Yu, Tianpei Gu, Chunze Lin, Quan Wang, Chen Qian, Chen Change Loy, and Ziwei Liu. Bailando: 3D dance generation by actor-critic GPT with choreographic memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11050–11059, 2022.

[Sun *et al.*, 2020] Guofei Sun, Yongkang Wong, Zhiyong Cheng, Mohan S Kankanhalli, Weidong Geng, and Xiangdong Li. DeepDance: music-to-dance motion choreography with adversarial learning. *IEEE Transactions on Multimedia*, 23:497–509, 2020.

[Tevet *et al.*, 2022] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. In *International Conference on Learning Representations*, 2022.

[Tseng *et al.*, 2023] Jonathan Tseng, Rodrigo Castellon, and Karen Liu. EDGE: Editable dance generation from music. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 448–458, 2023.

[Van Den Oord *et al.*, 2017] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.

[Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[Yoon *et al.*, 2020] Youngwoo Yoon, Bok Cha, Joo-Haeng Lee, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geeyuk Lee. Speech gesture generation from the trimodal context of text, audio, and speaker identity. *ACM Transactions on Graphics*, 39(6):1–16, 2020.

[Zhang *et al.*, 2023a] 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. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.

[Zhang *et al.*, 2023b] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023.

[Zhang *et al.*, 2023c] Mingyuan Zhang, Xinying Guo, Liang Pan, Zhongang Cai, Fangzhou Hong, Huirong Li, Lei Yang, and Ziwei Liu. ReMoDiffuse: Retrieval-Augmented Motion Diffusion Model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 364–373, 2023.

[Zhang *et al.*, 2024] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. MotionDiffuse: Text-Driven Human Motion Generation With Diffusion Model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–15, 2024.

[Zhou and Wang, 2023a] Zixiang Zhou and Baoyuan Wang. UDE: A Unified Driving Engine for Human Motion Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5632–5641, 2023.

[Zhou and Wang, 2023b] Zixiang Zhou and Baoyuan Wang. UDE: A unified driving engine for human motion generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5632–5641, 2023.

[Zhuang *et al.*, 2022] Wenlin Zhuang, Congyi Wang, Jinxiang Chai, Yangang Wang, Ming Shao, and Siyu Xia. Music2Dance: Dancenet for music-driven dance generation. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 18(2):1–21, 2022.