

Fg-T2M: Fine-Grained Text-Driven Human Motion Generation via Diffusion Model

Yin Wang¹, Zhiying Leng^{1,2}, Frederick W. B. Li³, Shun-Cheng Wu², Xiaohui Liang^{1,4,*}

¹State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing, China

²Technical University of Munich, Germany

³Department of Computer Science, University of Durham, U.K.

⁴Zhongguancun Laboratory, Beijing, China.

{wang_yin, zhiyingleng, liang_xiaohui}@buaa.edu.cn, frederick.li@durham.ac.uk, shuncheng.wu@tum.de

Abstract

Text-driven human motion generation in computer vision is both significant and challenging. However, current methods are limited to producing either deterministic or imprecise motion sequences, failing to effectively control the temporal and spatial relationships required to conform to a given text description. In this work, we propose a fine-grained method for generating high-quality, conditional human motion sequences supporting precise text description. Our approach consists of two key components: 1) a linguistics-structure assisted module that constructs accurate and complete language feature to fully utilize text information; and 2) a context-aware progressive reasoning module that learns neighborhood and overall semantic linguistics features from shallow and deep graph neural networks to achieve a multi-step inference. Experiments show that our approach outperforms text-driven motion generation methods on HumanML3D and KIT test sets and generates better visually confirmed motion to the text conditions.

1. Introduction

Human motion generation is a crucial task in computer vision with various applications in animation production, gaming, robot control, and movie script visualization. Obtaining human motion sequences through traditional software is a labor-intensive and tedious process, while motion capture is complex and expensive. Recently, with the advancements in deep learning and computer vision, learning-based human motion generation has emerged as a solution to this problem, leading to the development of associated generation methods based on multimodal data. The input multimodal data include music [12, 15, 25, 27], motion cat-




Figure 1: Our approach generates human motion sequences that grasp fine-grained details.

egories [4, 9, 20], text [2, 3, 5, 7, 8, 16, 21, 23, 28, 33], among others. Text-driven human motion generation has been a popular research topic, because of its convenience and human-friendliness. In particular, natural language comprises nouns, verbs, adverbs, etc. The mutual connections among different words in a sentence establish its semantics. Verbs define the action’s category, while adverbs control the fineness of the action. The interaction between words in syntax plays a vital role in determining the structure and meaning of a sentence. Failure to fully incorporate these text features may result in inadequate text modeling, causing the generated motion sequence to deviate from the intended meaning of the original text.

Existing methods can be divided into two branches, including 1) cross-modal alignment of motion and text [2, 3, 5, 7, 8, 16, 21]; 2) conditional diffusion models [28, 33]. In the first methods, text sequences and motion sequences are mapped onto separate feature spaces and forcibly aligned, leading to a loss of original information from both domains. In the second methods, the diffusion model incorporates

*Corresponding author.

text information as a conditioning factor to learn the probability mapping of human motions. However, the model interacts with only one text feature at each time step of the inference process, lacking a progressive approach. Moreover, text modeling only involves simple manipulation, which ignores the importance of certain fine-grained words and leads to incomplete semantic understanding, making it challenging to learn focus points at each step. Overall, existing methods only use text information to a limited extent, which in turn affects the accuracy of motion generation based on the corresponding text content, especially for the motion in which texts contain fine-grained words. For instance, comprehending the sentence “A man is walking forward while waving his right hand” can be a difficult task, and expecting the model to grasp the fine-grained meaning of the terms “while” and “right hand” is even more demanding.

To tackle the aforementioned issues, we propose a fine-grained text-driven method for generating human motion sequences that precisely align with text prompts in Figure 1. Typically, people initially read a sentence to gain an overall semantic understanding before focusing on the fine-grained details of individual words. To replicate this process, our method includes a linguistics-structure assisted module and a context-aware progressive reasoning module to fully utilize text information. Firstly, we utilize linguistics structure to facilitate information exchange between each text word. We use dependency parsing [19] to analyze the relationships among words in each sentence and construct a dependency tree, allowing each node to effectively communicate based on its dependent nodes and relationships. Then, the dependency tree nodes are passed to multi-layer graph neural networks to learn information aggregation. The multi-layer graph neural networks allow shallow network to learn neighborhood features as it can comprehend nearby details, and allow deep network to grasp overall semantic features because it is capable of aggregating information from entire nodes. Additionally, our GAT captures rich inter-word relationships while preserving the text linguistics structure by designing adaptive weights for each part-of-speech and dependency relation due to their distinctive role in the Text to Motion (T2M) task.

Secondly, achieving the purpose of fine-grained interaction, context-aware progressive reasoning module performs a multi-step inference process with the progressive fusion of global and local information between text and motion, which is unprecedented in the T2M task. This involves utilizing hierarchical semantic features to simulate the way humans comprehend sentences. We adopt the diffusion model framework and stack the hierarchical semantic features obtained from deep to shallow networks at each step to capture high-order relationships at different semantic levels. We evaluate our method on HumanML3D dataset [7] and KIT dataset [22]. Experiments show that our approach out-

performs the state-of-the-art methods and generates better visual motion. Our main contributions include:

- To the best of our knowledge, we are the first to bring NLP methods into T2M task. Utilizing the structured understanding of the natural language prompts to help T2M models achieve better reasoning skills, which brings new ideas for the text-to-X community from a textual perspective.
- We propose the Linguistics-Structure Assisted Module (LSAM), which utilizes a dependency parsing tree and graph networks to facilitate effective information exchange and data aggregation. It can obtain both neighborhood and overall semantic linguistic features.
- We propose a Context-Aware Progressive Reasoning Module (CAPR) that implements a multi-step progressive inference strategy within the diffusion model framework, mimicking the human reading process by moving from global to local relationships.
- Experimental results demonstrate that our proposed method outperforms previous methods, and achieves competitive performance on the HumanML3D and KIT datasets.

2. Related Work

2.1. Motion generation model

Generative models play a crucial role in motion synthesis by generating high-quality human motion. Generative Adversarial Nets (GAN) [6] use two sub-models: a generator model that produces new samples, and a discriminator model that attempts to classify samples as either real or fake. These two models compete against each other during training. However, the interpretability of GAN is poor because the learned data distribution lacks an explicit expression, resembling a black box mapping function.

Auto-Encoding Variational Bayes (VAE) [13] is a widely used generative model in motion synthesis. Its primary objective is to generate new samples from the learned distribution of objects by learning latent attributes from the probability distribution of the latent variable space, thereby constructing new examples. Despite its usefulness, the quality of samples generated by VAE can be improved.

Recently proposed diffusion models [10, 18, 26] have shown immense potential in modeling and present an exciting opportunity to expand into text-driven motion generation. These models utilize the stochastic diffusion process modeled in thermodynamics, which gradually adds noise to the samples of the data distribution. The deep learning model then learns the reverse process of denoising the samples gradually. Diffusion models have the advantage over previous models as they do not make any assumptions about

the target distribution, leading to a more diverse generation and better suitability for our task. Therefore, we propose a novel fine-grained human motion generation method that employs the Denoising Diffusion Probability Model [10].

2.2. Text-driven human motion generation

The task of text-driven human motion generation involves generating 3D human motion sequences that conform to a textual description. Several previous works have tackled this task. Initially, Text2Action [1] proposed short-text conditioned motion generation based on an RNN model. Subsequently, Ahuja *et al.* [2] and Ghosh *et al.* [5] focused on creating a joint representation of text and motion by projecting both features into a shared latent space. However, these methods involve a one-to-one mapping between text and motion, implying that given the same text, they can only produce fixed motion sequence results.

To increase generated result diversity, TEMOS [21] introduced a VAE architecture that finds a joint latent space for motion and text under Gaussian distribution constraints. Guo *et al.* [7] used a temporal VAE to autoregressively generate motion sequences based on text features. However, these methods have a significant drawback of mapping text and motion sequences to separate feature spaces and forcibly align them, leading to a loss of information in both domains. Recently, diffusion models have shown great potential in image generation and have inspired the development of diffusion models for human motion generation. Tevet *et al.* [28] and Zhang *et al.* [33] encoded text descriptions using pre-trained models and estimated Gaussian noise or the original motion sequence at each reverse diffusion step. However, their text modeling is often crude and does not fully leverage linguistic structure for sentence semantics. Also, they do not have a progressive process that allows the diffusion model to focus on different content at different time stamps, and they only interact with fixed text features during the inference process.

3. Preliminaries

The diffusion process is a Markov process consisting of a forward process and a reverse process. The forward process starts with the real data X_0 at step 0 and proceeds in a Markovian manner by adding Gaussian noise at each step. Over t steps, X_0 is transformed into X_t , which is close to the Gaussian distribution $N(0, I)$. As a result, the original motion sequence is converted into a complete noise distribution, which can be expressed as:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t I) \quad (1)$$

where β_t is a hyper-parameter that controls the diffusion rate. The entire diffusion forward process is formulated as:

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}) \quad (2)$$

where T denotes total steps in diffusion. The diffusion reverse process samples from the Gaussian distribution X_t as the initial input and attempts to gradually remove the noise on a reverse Markov-chain, which can be defined as follows:

$$p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) \quad (3)$$

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t)) \quad (4)$$

where μ_θ is the estimated item by model, t is the timestep indicating where the denoising process has conducted.

4. Method

To enhance the modeling of fine-grained human motion generation, we present our proposed method in Figure 2. Previous methods [1, 2, 5] have adopted a coarse approach to text modeling, leading to an underutilization of text information. Moreover, they treat all words equally without considering their importance and uniqueness. In contrast, we leverage linguistic structures in sentences to further enhance text encoding and differentiate between overall semantic information and detailed features using a multi-step progressive reasoning strategy.

Given a text prompt, $W = \{w_1, w_2, \dots, w_n\}$, $W \in \mathbb{R}^{N \times L}$ where N represents the number of words and L is the dimension of word vector. Our goal is to generate a human motion sequence, denoted as $M = \{m_1, m_2, \dots, m_t\}$, where $M \in \mathbb{R}^{T \times D}$. Here, T refers to the sequence length and D is the motion representation dimension. To achieve this, we introduce Fg-T2M, a method that generates motion sequences that align well with the corresponding textual content. In the following, we provide an overview of our approach in section 4.1, followed by the introduction of the Linguistics-Structure Assisted Module (LSAM) in section 4.2. Lastly, we present the Context-Aware Progressive Reasoning (CAPR) module in section 4.3.

4.1. Overview

Figure 2 shows our pipeline for generating motion sequences. We randomly sample \mathbf{x}_T from distribution $N(0, I)$, input \mathbf{x}_T , current step T , and text control condition c to obtain \mathbf{x}_{T-1} , iterating T rounds until we get \mathbf{x}_0 . In the denoising process, the condition c is fed into our LSAM text encoder. By leveraging dependency relationships between words, we use graph neural networks to aggregate data and extract hierarchical semantic features. The motion decoder

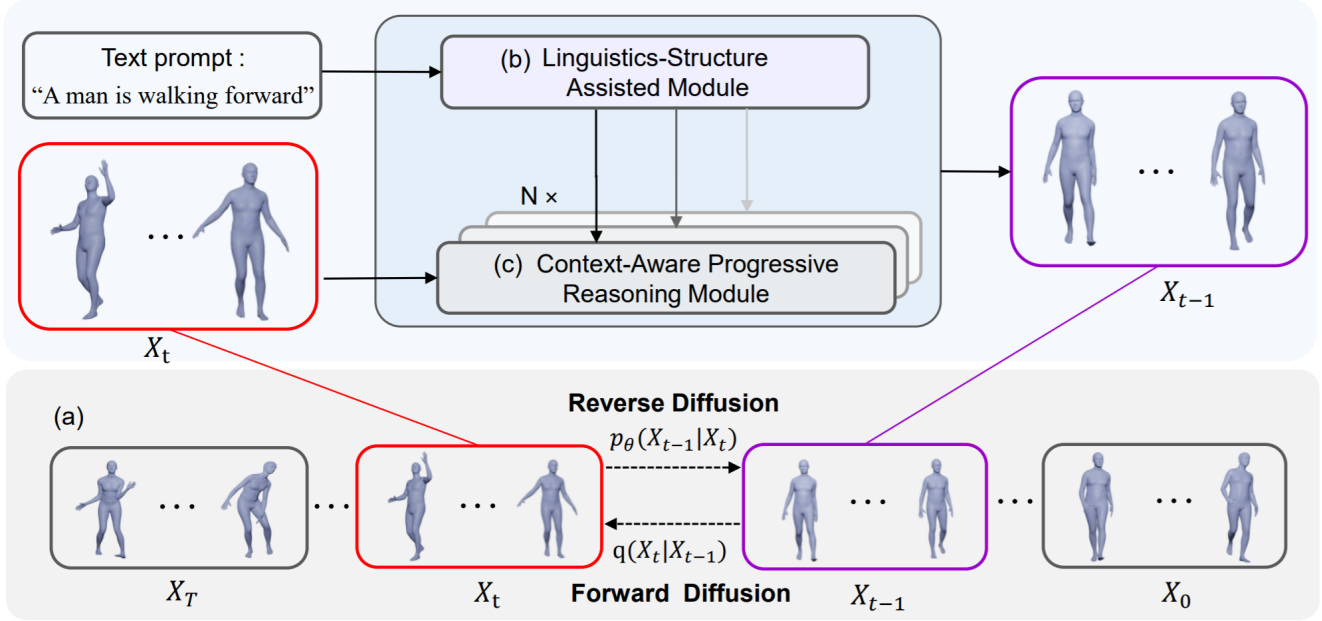


Figure 2: **Overall Pipeline of our proposed method Fg-T2M:** (a) the reverse denoising process of the diffusion model from X_T to X_0 ; (b) the text encoder with proposed linguistics-structure assisted module with dependency parsing in section 4.2; (c) the motion decoder with introduced context-aware progressive reasoning module in section 4.3.

is stacked with CAPR. Multiple layers of these modules enhance the quality of the generated motion sequence.

To train the motion diffusion model, we optimize the objective to predict the original data, representing as follows:

$$\mathcal{L} = \mathbb{E}_{t \in [1, T], \mathbf{x}_0 \sim q(\mathbf{x}_0), \epsilon \sim \mathcal{N}(0, I)} [\| \mathbf{x}_0 - \hat{\mathbf{x}}_0(\mathbf{x}_t, t, c) \|_2^2]. \quad (5)$$

The regular L2 loss can improve performance for all geometric losses mentioned above.

4.2. Linguistics-Structure Assisted Module

The current text modeling method is limited in its sensitivity to fine-grained words, which hinders the effectiveness of subsequent generations. To address this, we enhance the context modeling of sentences by identifying the phrase structure and syntactic relationships between phrases using dependency parsing, as shown in Figure 3. Linguistic structure, which comprises phrase structure and syntax relationships, helps us better understand the essence of text and the differences and similarities between sentences. Dependency parsing [19] precisely analyzes vocabulary and syntax to identify the dependency relationships between words in a sentence. In this approach, each word is treated as a node, and the dependency relationships between words are represented by edges that indicate syntactic connections. The resulting node and edge features representing the dependency relationships are input into a graph network to

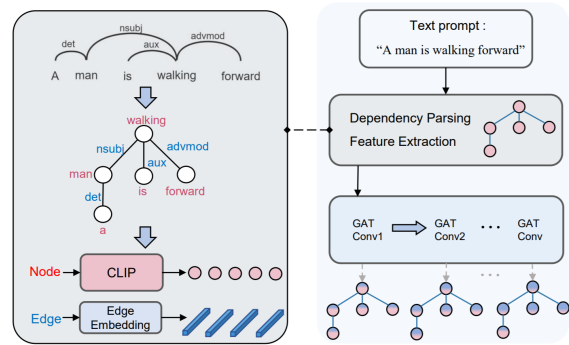


Figure 3: **Architecture of LSAM.** The left panel illustrates the process of obtaining node and edge features for one sentence in dependency analysis feature extraction, while the right panel demonstrates the workflows that utilize multiple layers of Graph Attention Network to extract feature information with varying degrees.

obtain multi-level semantics. In summary, given a sentence, the dependency parsing analyses to build a text tree, nodes for words, and links between nodes for their linguistic relationships. The text tree structure as the graph initialization for graph topology and edge features provides GAT a better prior. The GAT aims to capture inter-word relationships while maintaining the linguistic structure of the text.

To extract word associations, we parse the dependency

of phrase and obtain the hierarchical syntactic relationships using Spacy¹ for dependency parsing, as shown in the left panel of Figure 3. Spacy is a natural language processing software library for text processing, including lexical analysis, syntactic analysis, and more. In dependency parsing, each word is treated as a node, and the edges represent the labels of the dependency relationships between words, allowing us to construct a tree of language structure for the given text. We extract features using a graph attention network (GAT) [31]. For the input nodes of the graph network, we use word features obtained from the CLIP [24] model. For the adjacency matrix, we set a value of 1 or 0 to represent the presence or absence of a dependency relationship between nodes. Since the number of dependency relationships is fixed, we use one-hot encoding to obtain a one-hot label for each relationship, which is then fed into an embedding layer. This can be described as:

$$\mathbf{v}_e = \beta_i(\varphi_e(\text{F}_{\text{onehot}}(R_i))) \quad (6)$$

where $\mathbf{v}_e \in \mathbb{R}^{D_e}$, D_e is the dimension of edge feature, φ_e denotes an embedding layer, β_i is adaptive weight parameters to be learned for every R_i , while R_i indicates the edge relation between two nodes. Since each word may be adjacent to multiple nodes, and different nodes contribute differently to semantics, it is necessary to distinguish them during the text feature extraction process. We achieve this by applying GAT [31], which extracts multi-level features from the obtained node information, adjacency matrix, and edge information. GAT [31] considers differences between nodes during information aggregation and handles irregularities between different nodes effectively, which is described as:

$$x'_i = \alpha_{i,i}\Theta x_i + \sum_{j \in N(i)} \alpha_{i,j}\Theta x_j \quad (7)$$

$$\alpha_{i,j} = \frac{\exp(\text{F}(\omega^T[\Theta x_i || \Theta x_j || \Theta_e e_{i,j}]))}{\sum_{k \in N(i) \cup \{i\}} \exp(\text{F}(\omega^T[\Theta x_i || \Theta x_k || \Theta_e e_{i,k}]))} \quad (8)$$

We have $x, \alpha \in \mathbb{R}^{N \times L}$, where $\alpha_{i,j}$ represents attention coefficients, F denotes LeakyReLU, e is edge features, ω and Θ represents the weight parameters to be learned, and $||$ is the concatenation operation. GAT [31] can model high-order dependency relationships by stacking multiple graph attention layers to capture global and local graph topology effectively. For example, stacking three layers of GAT [31] results in x_i^1 , x_i^2 , and x_i^3 for each node x_i , as follows:

$$x_i^{l+1} = \alpha_{i,i}^l \Theta x_i^l + \sum_{j \in N(i)} \alpha_{i,j}^l \Theta x_j^l \quad (9)$$

The shallow network x_i^1 learns detailed features of the neighborhood as it can only gather information from directly adjacent nodes in the first step of node aggregation.

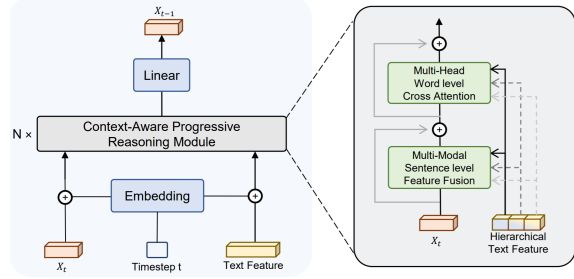


Figure 4: **Architecture of CAPR.** The left panel illustrates inferring X_{t-1} from X_t , t , and text features. The right panel provides the detailed view of one context-aware progressive reasoning module, which comprises two parts: sentence-level feature fusion and word-level cross-attention.

This process captures more local and fine-grained information. By comparison, the deep network x_i^3 can learn overall semantic features by aggregating information from distant nodes after multiple aggregation steps. This enables it to capture more global and holistic information.

4.3. Context-Aware Progressive Reasoning Module

Guided by the Linguistics-Structure Assisted Module (LSAM) in the previous section, we stack Context-Aware Progressive Reasoning Modules to perform multi-step progressive reasoning in a structured manner. Unlike previous methods that learn fixed features [28, 33], our stacked modules grasp features from global to local. Each block receives distinct contextual information, with higher-level blocks utilizing deeper context features from the LSAM, and lower-level blocks utilizing shallower features. This results in the model perceiving hierarchical information during the inference process, which greatly benefits its ability to hierarchically comprehend the meaning of text content and sense its fine-grained words.

The CAPR comprises two parts: Multi-Modal Sentence-level Feature-Fusion and Multi-Head Word-level Cross-Attention. Figure 4 shows that we start with the motion feature $x_t \in \mathbb{R}^{T \times D}$, text feature $W \in \mathbb{R}^{N \times L}$, and timestep t . To better capture the unique characteristics of each time step in the diffusion model, we first perform sinusoidal time embedding through linear layers to obtain timestep embedding emb_t . We then add the motion and text features to emb_t to incorporate different time information at each timestep.

Multi-Modal Sentence-Level Feature-Fusion module fuses sentence-level text features and motions to obtain multi-modal features, as shown in Figure 5. For the i -th Context-Aware Progressive Reasoning Module block, the text W_i is transformed to a sentence-level feature S using convolution, which is derived from the i -th LSAM layer:

$$S_i = \text{Conv}_{1d}(W_i) \quad (10)$$

¹Spacy: <https://spacy.io/>




Figure 5: Illustration of **Multi-Modal Sentence-Level Feature-Fusion Module**.

where $S_i \in \mathbb{R}^{1 \times D_s}$, $W_i \in \mathbb{R}^{N \times D_w}$, D_s and D_w denotes the dimension of sentence feature and words feature, and $Conv_{1d}$ is the 1d convolution. We then perform a matrix multiplication to obtain an attention list A :

$$A = X_t S_i^T \quad (11)$$

where $X_t \in \mathbb{R}^{T \times D}$ and $A \in \mathbb{R}^{T \times 1}$. Here, X_t represents the motion feature at timestep t , and A calculates the feature relevance between the sentence and each frame on the feature map. The resulting cross-modal feature, X'_t , highlights the sentence-relevant frame feature channels:

$$X'_t = X_t + \lambda(X_t \odot (\sigma(A))) \quad (12)$$

where λ is a hyper-parameter, \odot is element-wise multiplication and σ is a sigmoid activation function. To enhance modeling the correlation between different fused information, a self-attention mechanism [30] is added to strengthen the connection between multiple frames. Multi-head self-attention is conducted on the fusion feature X'_t as follows:

$$Q = W_q X'_t, \quad K = W_k X'_t, \quad \text{and} \quad V = W_v X'_t \quad (13)$$

where W_q , W_k and W_v are trainable weights to generate Q , K and V , respectively. Obtain the attention scores using the formula below, where \otimes is the matrix multiplication.

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{Q \otimes K^T}{\sqrt{d}}\right)V \quad (14)$$

Multi-Head Word-Level Cross-Attention module learns the cross-interaction between motion sequences and multi-level contextual text features. It uses the formulas mentioned in the self-attention module for calculation, but modifies the motion features of K and V to text features. Therefore, on the stacked Context-Aware Progressive Reasoning modules B_1, B_2, \dots, B_i , Block B_i utilizes different word-level semantic features by:

$$\hat{X}_t = X'_t + \text{Attention}(X'_t, W_i, W_i) \quad (15)$$

where X'_t is motion features and W_i is the text features. X'_t forms the query vector Q , W_i forms the key vector K and value vector V . They pass through the process in Equation 14. Finally, several MLP layers further processes the above features to output the predicted target $X_t \in \mathbb{R}^{T \times D}$.

Methods	R-TOPI ↑	R-TOPI2 ↑	R-TOPI3 ↑	FID↓	MM Dist↓	Diversity↑	MModality↑
Real	0.511±.003	0.703±.003	0.797±.002	0.002±.000	2.974±.008	9.503±.065	-
Seq2Seq [16]	0.180±.002	0.300±.002	0.396±.002	11.75±.035	5.529±.007	6.223±.061	-
L2P [2]	0.246±.002	0.387±.002	0.486±.002	11.02±.046	5.296±.008	7.676±.058	-
T2G[3]	0.165±.001	0.267±.002	0.345±.002	7.664±.030	6.030±.008	6.409±.071	-
Hier [5]	0.301±.002	0.425±.002	0.552±.004	6.532±.024	5.012±.018	8.332±.042	-
MoCoGAN [29]	0.037±.000	0.072±.001	0.106±.001	94.41±.021	9.643±.006	0.462±.008	0.019±.000
Dance2Music [14]	0.033±.000	0.065±.001	0.097±.001	66.98±.016	8.116±.006	0.725±.011	0.043±.001
TEMOS [21]	0.424±.002	0.612±.002	0.722±.002	3.734±.028	3.703±.008	8.973±.071	0.368±.018
Temporal VAE [7]	0.455±.003	0.636±.003	0.740±.003	1.067±.002	3.340±.008	9.188±.002	2.090±.083
TM2T [8]	0.424±.003	0.618±.003	0.729±.002	1.501±.017	3.467±.011	8.589±.076	2.424±.093
MotionDiffuse [33]	0.491±.001	0.681±.001	0.782±.001	0.630±.001	3.113±.001	9.410±.049	1.553±.042
MDM [28]	0.320±.005	0.498±.004	0.611±.007	0.544±.044	5.566±.027	9.559±.086	2.799±.072
T2M-GPT [32]	0.491±.003	0.680±.003	0.775±.002	0.116±.004	3.118±.011	9.761±.081	1.856±.011
Fg-T2M	0.492±.002	0.683±.003	0.783±.002	0.243±.019	3.109±.007	9.278±.072	1.614±.049

Table 1: Comparison of text-conditional motion synthesis on HumanML3D [7] dataset. **Red** and **Blue** indicate the best and the second best result, respectively.

Methods	R-TOPI ↑	R-TOPI2 ↑	R-TOPI3 ↑	FID↓	MM Dist↓	Diversity↑	MModality↑
Real	0.424±.005	0.649±.006	0.779±.006	0.031±.004	2.788±.012	11.08±.097	-
Seq2Seq[16]	0.103±.003	0.178±.005	0.241±.006	24.86±.348	7.960±.031	6.744±.106	-
T2G[3]	0.156±.004	0.255±.004	0.338±.005	12.12±.183	6.964±.029	9.334±.079	-
L2P [2]	0.221±.005	0.373±.004	0.483±.005	6.545±.072	5.147±.030	9.073±.100	-
Hier [5]	0.255±.006	0.432±.007	0.531±.007	5.203±.107	4.986±.027	9.563±.072	-
MoCoGAN [29]	0.022±.002	0.042±.003	0.063±.003	82.69±.242	10.47±.012	3.091±.043	0.250±.009
Dance2Music [14]	0.031±.002	0.058±.002	0.086±.003	115.4±.240	10.40±.016	0.241±.004	0.062±.002
TEMOS [21]	0.353±.006	0.561±.007	0.687±.005	3.717±.051	3.417±.019	10.84±.100	0.532±.034
Temporal VAE [7]	0.361±.006	0.559±.007	0.693±.007	2.770±.109	3.401±.008	10.91±.119	1.482±.065
TM2T [8]	0.280±.005	0.463±.006	0.587±.005	3.599±.153	4.591±.026	9.473±.117	3.292±.081
MotionDiffuse [33]	0.417±.004	0.621±.004	0.739±.004	1.954±.062	2.958±.005	11.10±.143	0.730±.013
MDM [28]	0.164±.004	0.291±.004	0.396±.004	0.497±.021	9.19±.022	10.847±.109	1.907±.214
T2M-GPT [32]	0.416±.006	0.627±.006	0.745±.006	0.514±.029	3.007±.023	10.92±.108	1.570±.039
Fg-T2M	0.418±.005	0.626±.004	0.745±.004	0.571±.047	3.114±.015	10.93±.083	1.019±.029

Table 2: Comparison of text-conditional motion synthesis on KIT [22] dataset. **Red** and **Blue** indicate the best and the second best result, respectively.

5. Experiments

Our Fg-T2M framework is evaluated on text-driven motion generation task in this section. We first describe the dataset used and the evaluation metrics in section 5.1. In section 5.2, we provide implementation details. We compare our framework with the current state-of-the-art methods in section 5.3. Finally, we present qualitative results and visualization for comparison in section 5.4.

5.1. Datasets and Evaluation Metrics

Several datasets exist for conditional motion generation, such as proposed in [7, 9, 11, 22]. However, datasets such as [9] and [11] are based on action categories and do not provide complete text sentences as conditioning inputs, making them unsuitable for our method. Instead, we use text-driven datasets, specifically, **HumanML3D** dataset [7] and **KIT Motion-Language** dataset [22], for our experiments.

The HumanML3D dataset [7] is a combination of the HumanAct12 [9] and AMASS [17] datasets, comprising 14,616 motions and 44,970 text descriptions across various human activities, such as daily activities, sports, acrobatics, etc., with a total duration of about 28.59 hours. The KIT Motion-Language dataset [22] consists of 3,911 motion sequences and 6,353 natural language descriptions, with a total duration of around 10.33 hours.

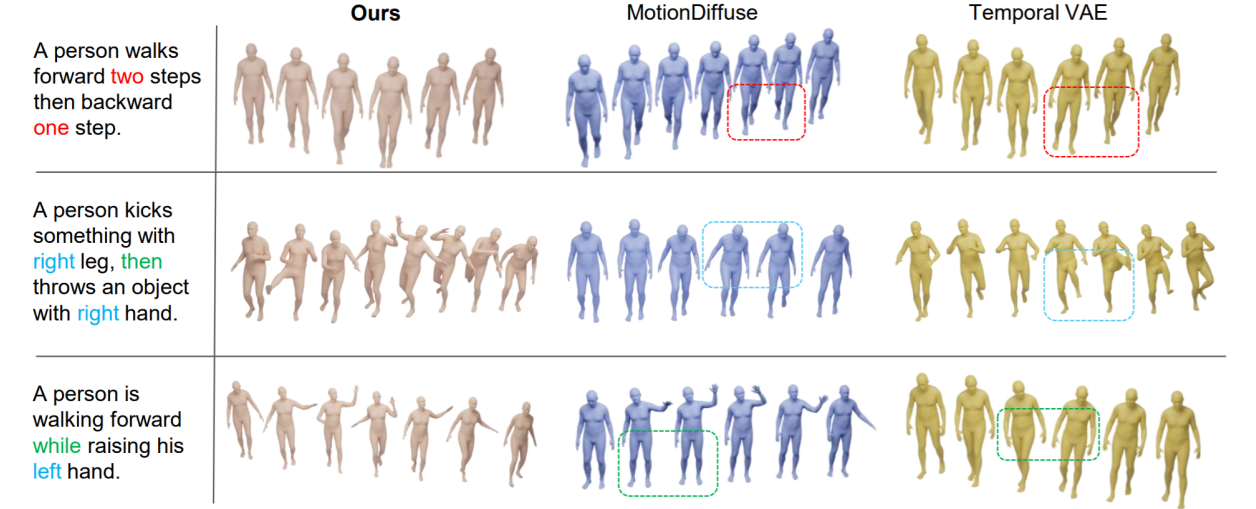


Figure 6: **Qualitative results:** Our method is compared with two state-of-the-art methods: MotionDiffuse [33] and Temporal VAE [7]. Motion frames are ordered from left to right. Those not matching with text prompt are marked with a box.

Evaluation Metrics are followed [7]. (1) *R-precision*. For each inferred text-motion pair, 31 mismatched descriptions are randomly selected from the test set. The average top-k precision is obtained by calculating and ranking the Euclidean distance between the motion and each of the 32 descriptions. (2) *Frechet Inception Distance (FID)*. FID measures the similarity between the feature distributions extracted from the generated motions and ground truth motions. (3) *Multi-Modal Distance*. The multimodal distance is computed between the text feature and the relevant generated motion feature, concerning the given description. (4) *Diversity*. Diversity evaluates the dissimilarities among all generated motions across all descriptions by computing the mean pairwise Euclidean distance between randomly partitioned groups of motions. (5) *Multimodality*. For a given text description, 32 motion sequences are generated randomly, and multimodality quantifies the dissimilarities among these generated motion sequences. We primarily value R-precision and FID as pivotal performance metrics, which serve as important measures for evaluating the overall quality of generated motions.

5.2. Implementation Details

The diffusion model uses 1000 diffusion steps and a linearly varying variance β_t ranging from 0.0001 to 0.02. LSAM employs a 3-layer GAT network with a corresponding CAPR layer also set to 3. Hyper-parameter λ in CAPR is 0.1. Training is performed with the Adam optimizer using a fixed learning rate of $5e-5$, a batch size of 128, and NVIDIA GeForce RTX 3090 hardware. The KIT dataset is trained for approximately 40K iterations, while the HumanML3D dataset is trained for about 80K iterations.

5.3. Comparison with State-of-the-arts

We compared our method with several state-of-the-art models, including Lin *et al.* [16], Language2Pose [2], Ghosh *et al.* [5], MoCoGAN [29], Dance2Music [14], TEMOS [21], TM2T [8], Text2Gesture [3], Guo *et al.* [7], MotionDiffuse [33], and MDM [28]. Quantitative comparisons of our method with these models on the HumanML3D [7] and KIT [22] datasets are shown in Tables 1 and 2, respectively.

Our method achieves competitive performance between text and motion features, as measured by MM Dist, to state-of-the-art methods, while exhibiting significantly higher scores in R-precision and FID. This demonstrates the ability of our method to generate high-quality motions that align with the text prompts. On the other hand, other approaches showcase remarkable competitiveness in diversity and multimodality. However, these aspects should be grounded in accuracy (R-precision) and precision (FID, MMDist) to strengthen their persuasiveness. Otherwise, the diversity or multimodality would be rendered meaningless if the generated motion fails to align with the desired outcome. Therefore, based on our experiments, our method has achieved advanced experimental results and demonstrates robustness in terms of model performance on the two datasets.

Meanwhile, we design two experiments to evaluate fine-grained control. The first one is conducted on a Harder-HumanML3D dataset. We compress the HumanML3D test set of 4382 data into a Harder-HumanML3D set with 2582 data, by searching for sentences that contain more fine-grained words, like “left”, “right”, and so on. The comparison results are shown in Table 3. Our method exhibits significantly higher scores, indicating a better ability to capture




Figure 7: The result of user study.

Methods	R-Precision \uparrow			FID \downarrow
	Top-1	Top-2	Top-3	
MotionDiffuse [33]	0.439 \pm .006	0.614 \pm .004	0.725 \pm .005	0.732 \pm .021
Fg-T2M	0.460 \pm .005	0.660 \pm .004	0.763 \pm .004	0.421 \pm .013

Table 3: Experiment results on Harder-HumanML3D split from the HumanML3D [7] test set.

fine-grained details. The second one is a user study in which we collect user preferences with T2M-GPT [32] and MotionDiffuse [33]. The statistics of the user study are shown in Figure 7. Compared with others, our method achieves superior performance in R-Precision, MM Dist, and yields competitive results in FID, which generates motions with comparable quality. Furthermore, for both questions, especially in the fine-grained aspect, ours is preferred over others and even competitive to the ground truth motions.

We present ablation results in Table 4 to further understand the role of LSAM and CAPR in our method. The results demonstrate that models without the CAPR and LSAM modules exhibit performance degradation. And we also show the comparison of some qualitative examples on the ablation study of this two modules. Additionally, we experiment with the impact of different GAT semantic layers of text information. Since a common dependency tree of motion-text prompt often has three or four depths, a three-layer GAT is sufficient to capture overall nodes information. Hence, a shallow or deep GAT layer hinders the global or local semantic comprehension, resulting in diminished results. We further evaluate the proposed two parts in CAPR. The results show that without CAPR-2 part remarkably reduces the results, which also reflects the importance of fine-grained words for T2M tasks. Finally, the hyper-parameter λ in CAPR controls the degree of text information fusion. A larger λ dilutes the inherent properties of the motion features, leading to a decrease in precision.

5.4. Qualitative Results

Visual results on the HumanML3D [7] dataset are presented in Figure 6, where our method is compared with the state-of-the-art models of MotionDiffuse [33] and Temporal VAE [7]. As can be seen from the examples in the figure, our method generates human motions that more accurately reflect the text prompts. In contrast, the methods of MotionDiffuse [33] and Temporal VAE [7] often result in unrealistic movements. Specifically, these methods tend to

Methods	R-TOP3 \uparrow	FID \downarrow
Fg-T2M	0.745 \pm .004	0.571 \pm .047
Fg-T2M (w/o LSAM)	0.722 \pm .005	1.077 \pm .101
Fg-T2M (w/o CAPR)	0.727 \pm .007	0.943 \pm .089
Fg-T2M (one layer)	0.729 \pm .006	0.951 \pm .092
Fg-T2M (two layers)	0.738 \pm .003	0.692 \pm .086
Fg-T2M (four layers)	0.740 \pm .004	0.636 \pm .102
Fg-T2M(w/o CAPR-1)	0.739 \pm .006	0.652 \pm .035
Fg-T2M(w/o CAPR-2)	0.732 \pm .004	0.764 \pm .039
Fg-T2M ($\lambda = 0.2$)	0.738 \pm .008	0.686 \pm .062
Fg-T2M ($\lambda = 0.3$)	0.732 \pm .006	0.830 \pm .074
Fg-T2M ($\lambda = 0.5$)	0.729 \pm .011	1.086 \pm .104

Table 4: Ablation analysis on KIT [22]. “layers” indicates the number of layers of GAT in LSAM. “(w/o) CAPR-1” means without multi-modal sentence-level feature-fusion. “(w/o) CAPR-2” means without multi-head word-level cross-attention. “ λ ” is the hyper-parameter of Multi-Modal Sentence-Level Feature-Fusion module in CAPR.

only understand one motion within a long text prompt and are not sensitive to specific numerical values, such as ‘one’ or ‘two’. As depicted in example one, MotionDiffuse [33] only performed a backward motion without considering the required number of steps. On the other hand, Temporal VAE [7] successfully achieved the desired effect of walking forward and then backward, but its backward motion step count did not meet the expected requirement. Moreover, They also make errors in spatial orientation, such as ‘left’ or ‘right’, and lack a thorough understanding of temporal issues, such as ‘while’ and ‘then’, related to movement. Overall, our proposed method outperforms these models, especially in terms of the issues mentioned above. More diverse samples are presented in the supplementary material.

6. Conclusion

We present a novel method for text-driven human motion generation using the diffusion model, which offers several advantages over existing techniques. Specifically, our method leverages two key modules - a linguistics-structure assisted module and a context-aware progressive reasoning module - to effectively model fine-grained words in the text. The former module extracts dependency parsing relationships in the text, while the latter performs hierarchical effective information feature fusion based on graph neural networks. Our quantitative and qualitative results demonstrate that our method outperforms existing techniques in text-driven motion generation tasks.

Acknowledgements. This work was supported by the National Natural Science Foundation of China (Project Number: 62272019).

References

- [1] Hyemin Ahn, Timothy Ha, Yunho Choi, Hwiyeon Yoo, and Songhwai Oh. Text2action: Generative adversarial synthesis from language to action. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5915–5920. IEEE, 2018. [3](#)
- [2] Chaitanya Ahuja and Louis-Philippe Morency. Language2pose: Natural language grounded pose forecasting. In *2019 International Conference on 3D Vision (3DV)*, pages 719–728. IEEE, 2019. [1](#), [3](#), [6](#), [7](#)
- [3] Uttaran Bhattacharya, Nicholas Rewkowski, Abhishek Banerjee, Pooja Guhan, Aniket Bera, and Dinesh Manocha. Text2gestures: A transformer-based network for generating emotive body gestures for virtual agents. In *2021 IEEE virtual reality and 3D user interfaces (VR)*, pages 1–10. IEEE, 2021. [1](#), [6](#), [7](#)
- [4] Pablo Cervantes, Yusuke Sekikawa, Ikuro Sato, and Koichi Shinoda. Implicit neural representations for variable length human motion generation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVII*, pages 356–372. Springer, 2022. [1](#)
- [5] Anindita Ghosh, Noshaba Cheema, Cennet Oguz, Christian Theobalt, and Philipp Slusallek. Synthesis of compositional animations from textual descriptions. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1396–1406, 2021. [1](#), [3](#), [6](#), [7](#)
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [2](#)
- [7] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5152–5161, 2022. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [8] 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 *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 580–597. Springer, 2022. [1](#), [6](#), [7](#)
- [9] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2021–2029, 2020. [1](#), [6](#)
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. [2](#), [3](#)
- [11] Yanli Ji, Feixiang Xu, Yang Yang, Fumin Shen, Heng Tao Shen, and Wei-Shi Zheng. A large-scale rgb-d database for arbitrary-view human action recognition. In *Proceedings of the 26th ACM international Conference on Multimedia*, pages 1510–1518, 2018. [6](#)
- [12] Hsuan-Kai Kao and Li Su. Temporally guided music-to-body-movement generation. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 147–155, 2020. [1](#)
- [13] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [2](#)
- [14] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. Dancing to music. *Advances in neural information processing systems*, 32, 2019. [6](#), [7](#)
- [15] 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. [1](#)
- [16] Angela S. Lin, Lemeng Wu, and Qixing Huang Raymond J. Mooney Rodolfo Corona, Kevin Tai. Generating animated videos of human activities from natural language descriptions. In *Proceedings of the Visually Grounded Interaction and Language Workshop at NeurIPS 2018*, December 2018. [1](#), [6](#), [7](#)
- [17] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. [6](#)
- [18] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. [2](#)
- [19] Joakim Nivre. Algorithms for deterministic incremental dependency parsing. *Computational Linguistics*, 34(4):513–553, 2008. [2](#), [4](#)
- [20] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10985–10995, 2021. [1](#)
- [21] Mathis Petrovich, Michael J Black, and Gül Varol. 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*, pages 480–497. Springer, 2022. [1](#), [3](#), [6](#), [7](#)
- [22] Matthias Plappert, Christian Mandery, and Tamim Asfour. The kit motion-language dataset. *Big data*, 4(4):236–252, 2016. [2](#), [6](#), [7](#), [8](#)
- [23] Matthias Plappert, Christian Mandery, and Tamim Asfour. Learning a bidirectional mapping between human whole-body motion and natural language using deep recurrent neural networks. *Robotics and Autonomous Systems*, 109:13–26, 2018. [1](#)
- [24] 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. [5](#)
- [25] Xuanchi Ren, Haoran Li, Zijian Huang, and Qifeng Chen. Self-supervised dance video synthesis conditioned on music. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 46–54, 2020. [1](#)

- [26] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2
- [27] Sebastian Starke, Ian Mason, and Taku Komura. Deepphase: Periodic autoencoders for learning motion phase manifolds. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. 1
- [28] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 1, 3, 5, 6, 7
- [29] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. 6, 7
- [30] 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. 6
- [31] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017. 5
- [32] 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. *arXiv preprint arXiv:2301.06052*, 2023. 6, 8
- [33] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. 1, 3, 5, 6, 7, 8