

# Light-T2M: A Lightweight and Fast Model for Text-to-motion Generation

Ling-An Zeng<sup>1</sup>, Guohong Huang<sup>1</sup>, Gaojie Wu<sup>1</sup>, Wei-Shi Zheng<sup>1 2\*</sup>

<sup>1</sup>Sun Yat-sen University

<sup>2</sup> Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, China  
zenglan3@mail2.sysu.edu.cn, huanggh37@mail2.sysu.edu.cn, wugj7@mail2.sysu.edu.cn, wszheng@ieee.org

## Abstract

Despite the significant role text-to-motion (T2M) generation plays across various applications, current methods involve a large number of parameters and suffer from slow inference speeds, leading to high usage costs. To address this, we aim to design a lightweight model to reduce usage costs. First, unlike existing works that focus solely on global information modeling, we recognize the importance of local information modeling in the T2M task by reconsidering the intrinsic properties of human motion, leading us to propose a lightweight Local Information Modeling Module. Second, we introduce Mamba to the T2M task, reducing the number of parameters and GPU memory demands, and we have designed a novel Pseudo-bidirectional Scan to replicate the effects of a bidirectional scan without increasing parameter count. Moreover, we propose a novel Adaptive Textual Information Injector that more effectively integrates textual information into the motion during generation. By integrating the aforementioned designs, we propose a lightweight and fast model named Light-T2M. Compared to the state-of-the-art method, MoMask, our Light-T2M model features just **10%** of the parameters (4.48M vs 44.85M) and achieves a **16%** faster inference time (0.152s vs 0.180s), while surpassing MoMask with an FID of **0.040** (vs. 0.045) on HumanML3D dataset and **0.161** (vs. 0.228) on KIT-ML dataset.

**Code** — <https://github.com/qinghuannn/light-t2m>

## 1 Introduction

Text-to-motion (T2M) generation seeks to produce 3D human motion sequences from textual descriptions, playing a crucial role in diverse fields such as film production, video games, and virtual/augmented reality. The broad applicability of T2M has spurred rapid advancements in various methodologies (Chen et al. 2023; Zhang et al. 2023a, 2024; Guo et al. 2024) in recent years. However, the simplistic and rudimentary design of these models results in high parameter counts and slow inference speeds, which raises the barrier to entry and hampers usability.

Given the imperative to minimize deployment costs and the ubiquitous use of mobile phones, developing lightweight

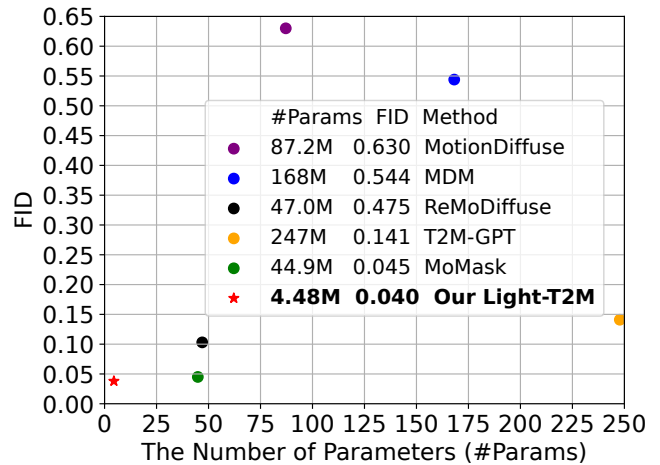


Figure 1: Comparison on FID and the number of parameters. The closer the model is to the origin, the better. Only trainable parameters are calculated.

models is essential for the effective implementation of deep learning technologies (Liu et al. 2024b). Prior research (Tan and Le 2019; Touvron et al. 2021; Belousov 2021) has focused on creating more efficient models for image classification and generation. In the realm of text-to-motion generation, a streamlined T2M model can substantially lower hardware demands during game and film production and improve the AR/VR and gaming experiences on mobile devices. In this work, we aim to design a lightweight T2M model.

To design a lightweight T2M model, we have recognized the overlooked importance of local information modeling in T2M, ignored by recent methods, and use this to guide model design. Intuitively, the essence of generating realistic human motion lies in the smooth transition between adjacent movements, indicating the importance of local information modeling. However, recent existing methods focus only on the global information interaction modeled by Transformer (Vaswani et al. 2017). Thus, we design a lightweight Local Information Modeling Module (LIMM) using the lightweight convolution network. By appropriately replacing some global information modeling layers with our LIMM, we significantly reduce the number of parameters without sacrificing performance.

\*Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Secondly, we introduce Mamba (Gu and Dao 2023; Dao and Gu 2024) into the T2M task to reduce the number of parameters and GPU memory requirements, and design a novel Pseudo-bidirectional Scan to achieve bidirectional scanning without adding extra parameters. By replacing the Transformer heavily relied upon in previous methods with Mamba, we further streamline resource usage. Besides, since the single-direction scan in Mamba, intended for causal sequences, is inadequate for global information modeling, a simple yet effective Pseudo-bidirectional Scan is designed to achieve a bidirectional scan. In detail, we feed the concatenation of the reverse sequence and original sequence into Mamba and, during a single-direction scan, each element in the original sequence can see the information of the element on its left or right. In this way, we achieve the effect of a bidirectional scan without increasing parameter count.

Furthermore, different from existing T2M works that rely on token mixing in self/cross-attention mechanisms to integrate textual information into motion frames, we propose a novel Adaptive Textual Information Injector (ATII) that employs a more effective method for embedding textual information during motion generation. Inspired by the gating mechanism in LSTM (Graves and Graves 2012), our ATII adaptively extracts segment-aware semantics from the text embedding token derived from CLIP (Radford et al. 2021), and incorporates it into corresponding motion segments. Specifically, for each motion segment, the ATII predicts channel-wise weights for the text embedding token and uses these weights to modulate the text token. Subsequently, the modulated text token and the motion segment token are combined in a fusion layer, enabling adaptive injection of textual information.

By integrating the aforementioned designs, we have developed Light-T2M, a lightweight and fast model tailored for text-to-motion generation. As shown in Figure 1, Light-T2M employs only 4.48M parameters—approximately **10%** of those used by the state-of-the-art method, MoMask (Guo et al. 2024). Despite its smaller size, it excels, achieving an FID of 0.040 (vs. 0.045 of MoMask) on the HumanML3D dataset and 0.161 (vs. 0.228) on the KIT-ML dataset. Owing to its efficient design, Light-T2M also delivers a **16%** faster inference time (0.152s vs. 0.180s). Thanks to its lightweight design and fast inference speed, our Light-T2M reduces the cost of use and paves the road for the widespread use of text-to-motion generation. In summary, our main contributions can be summarized as follows:

- We introduce Light-T2M, a lightweight model for text-to-motion generation, aiming to reduce the cost of use and pave the road for the widespread use of text-to-motion generation.
- We propose a novel Adaptive Textual Information Injector to inject the textual information into the motion during generation, which is an efficient and effective way to control the generated motion through text.
- We realize the importance of local information modeling and introduce Mamba into the text-to-motion task, both reducing the number of model parameters and speeding up inference time, along with a novel Pseudo-

bidirectional Scan to achieve the effect of the bidirectional scan without increasing parameters.

## 2 Related Work

**Text-to-motion Generation.** Early works such as (Ghosh et al. 2021; Plappert, Mandery, and Asfour 2018; Yamada, Matsunaga, and Ogata 2018) utilize encoder-decoder RNN models to achieve text-to-motion generation. Subsequent studies (Guo et al. 2022a; Petrovich, Black, and Varol 2022; Zhong et al. 2023) adopt a variational autoencoder (VAE) (Kingma and Welling 2014) to encode motion into a latent distribution and transform text into the target latent distribution. Following the success of diffusion models, most recent works (Kim, Kim, and Choi 2023; Zhang et al. 2024; Tevet et al. 2023; Kong et al. 2023; Shafir et al. 2024; Xie et al. 2024; Han et al. 2024) resort to diffusion models and build the network based on Transformer to achieve text-driven motion generation. Some researchers (Chen et al. 2023; Liu et al. 2023) design diffusion models to predict motion in latent spaces, while others (Zhang et al. 2023a; Jiang et al. 2024) develop autoregressive models to sequentially generate motion frames. A few studies (Wang et al. 2023; Zhang et al. 2023b) propose specially designed modules instead of directly using Transformers. Recently, Guo et al. (Guo et al. 2024) treated the target motion as masked tokens and developed a Masked Transformer to progressively predict all masked tokens. However, due to simplistic and rudimentary architectural designs, these methods require a large number of parameters. In contrast, we aim to design a lightweight model based on our well-designed modules and proposed Adaptive Textual Information Injector.

**Lightweight Network.** Reducing computational costs and enhancing storage capabilities have elevated the importance of lightweight network design in deep learning (Liu et al. 2024a). Projects like the MobileNet series (Howard et al. 2017; Mehta and Rastegari 2022; Chen et al. 2022; Sandler et al. 2018) demonstrate manual construction of lightweight models informed by prior knowledge. MobileNet, for instance, features a streamlined architecture with depth-wise separable convolutions for creating efficient neural networks (Howard et al. 2017). Techniques like neural architecture search aim to automate model development while minimizing human input (Wu et al. 2023; Hu et al. 2024), and neural network compression offers additional routes to model efficiency (Xiao, Wang, and Rajasekaran 2019; Gholami et al. 2021; Yan et al. 2024). Yet, the trade-offs between performance and parameters in these methods often lead to suboptimal results. In this work, we tailor our model to meet the unique demands of the T2M task without sacrificing performance.

**State Space Models.** The Structured State-Space Sequence (S4) model is attractive for its linear scaling with sequence length, prompting further exploration. Recently, Mamba (Gu and Dao 2023; Dao and Gu 2024), leveraging a data-dependent SSM layer, has outperformed Transformers in large-scale data, demonstrating its potential for sequence modeling. Following Mamba’s success, numerous studies have investigated its application in specific fields. For in-

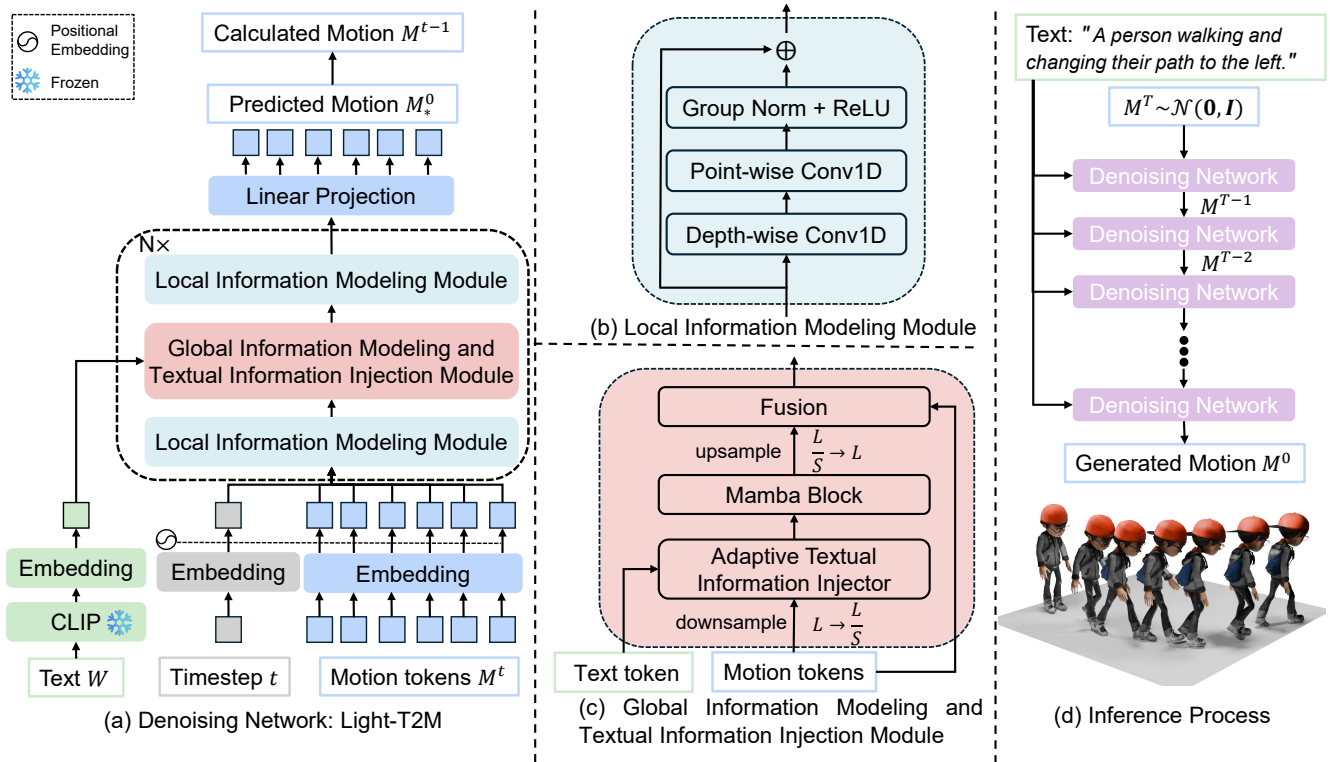


Figure 2: Overview of our Light-T2M. (a) Our Light-T2M consisting of  $N$  basic blocks aims to predict  $M_*^0$ , and then  $M^{t-1}$  can be calculated via Eq. 2. (b) The structure of our lightweight Local Information Modeling Module. (c) The motion is downsampled to obtain segments containing local semantic information. Next, a novel Adaptive Textual Information Injector and a Mamba Block are adopted to adaptively inject semantics into each segment and model the global information, respectively. The upsampled motion and the original motion are fused by a fusion layer. (d) The overview of the inference process.

stance, Vision Mamba (Zhu et al. 2024) and VMamba (Liu et al. 2024c) have developed unique scanning techniques for 2D image processing. In our work, we aim to extend Mamba’s advantages to the T2M task by designing a model that is both lightweight and fast.

### 3 Methodology

#### 3.1 Problem Formulation

Given a text  $W = \{w_i\}$  describing motion, our goal is to generate a corresponding 3D human pose sequence  $M = \{m_i\}$  of length  $L$ . According to existing works (Guo et al. 2022a, 2024), each 3D human pose is represented by a  $D_m$ -dimensional vector, which includes the root angular velocity along the Y-axis, root linear velocities on the XZ-plane, root height, and local joint positions, velocities, and rotations relative to the root space. For more detailed information on human pose representation, please see T2M (Guo et al. 2022a).

#### 3.2 Overview of Light-T2M

Figure 2 illustrates the architecture of our Light-T2M, which consists of  $N$  blocks. Each block contains two Local Information Modeling Modules and one Global Information Modeling and Textual Information Injection Module. The former modules model local information within each motion frame’s local window, ensuring the consistency of adjacent

frames—a critical aspect of generating realistic human motion. The latter module manages global information across all frames and injects textual information, ensuring coherence throughout the sequence and alignment with the given text. Additionally, a fixed CLIP (Radford et al. 2021) model and three embedding layers extract text features and harmonize various features within the same space.

**Training.** As shown in Figure 2, our Light-T2M is a diffusion-based model. Following (Ho, Jain, and Abbeel 2020), we define a Markov chain of  $T$  diffusion steps to slowly add random noise  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  to motion. Given a motion  $M$ , and the diffusion time-step  $t \in [1, T]$ , our denoising network (Light-T2M)  $\phi_\theta(\cdot)$  is trained via the following training objective:

$$\mathcal{L} = \mathbb{E}_{M^0, t, \epsilon} \|M^0 - \phi_\theta(M^t, t, c)\|^2, \quad (1)$$

$$M^t = \sqrt{\alpha_t} M^0 + \sqrt{1 - \alpha_t} \epsilon, \quad (2)$$

where  $M^0 = M$ ,  $c$  is the condition text, and a predefined  $\alpha_t$  aims to control the variance. While predicting  $M^0$  and predicting noise  $\epsilon$  are mathematically equivalent, we find that predicting  $M^0$  yields better performance in the T2M task. To both learn conditioned or unconditioned generation during training, we randomly set  $c = \emptyset$  with a probability  $\tau$ .

**Inference.** During generation, we reverse the diffusion process to construct desired data samples from the noise  $M^T \sim$




Figure 3: In our pseudo-bidirectional scan, each element in the original sequence can obtain the information from elements originally on its right, achieving the effect of bidirectional scanning without increasing parameters.

$\mathcal{N}(\mathbf{0}, \mathbf{I})$ . For each time step  $t$  and noise motion  $M^t$ , we first use the trained Light-T2M to predict the conditional sample  $\hat{M}_*^0$  and unconditional sample  $\bar{M}_*^0$  by setting  $c$  as the given text  $W$  and  $\emptyset$ . Using  $M^t$ ,  $\hat{M}_*^0$ , and  $\bar{M}_*^0$ , we can calculate the conditional noise  $\epsilon_c$  and unconditional noise  $\epsilon_u$  via Equation 2. Finally, we adopt *classifier-free guidance* method (Ho and Salimans 2022) for conditional generation:

$$\hat{\epsilon} = (1 + s) \cdot \epsilon_c - s \cdot \epsilon_u, \quad (3)$$

where the guidance scale  $s$  controls the strength of the condition. The obtained noise  $\hat{\epsilon}$  is then used to calculate  $M^{t-1}$ . More details are shown in the Appendix.

### 3.3 Learning to Model Local Information

Recognizing that modeling the interaction of location information across adjacent frames can enhance the smoothness of generated motion and be implemented efficiently, we have integrated location information modeling into the T2M task. We believe that strategically replacing some global information modeling layers with location information modeling layers can significantly reduce the parameter count without sacrificing performance.

Inspired by the success of CNN in local information modeling, we design our Local Information Modeling Modules (LIMM) based on lightweight 1D convolution layers. As shown in Figure 2.b, to reduce the number of parameters and FLOPs, our LIMM mainly consists of a 1D point-wise convolution  $f^p(\cdot)$  and a 1D depth-wise convolution  $f^d(\cdot)$ . Formally, our LIMM can be defined as:

$$\hat{X} = X + \text{ReLU}(\text{Norm}(f^p(f^d(X)))), \quad (4)$$

where  $X \in \mathbb{R}^{D \times L}$  and  $\text{Norm}(\cdot)$  are the input motion and the Group Normalization (Wu and He 2020).

### 3.4 Learning to Model Global Information

Beyond modeling local information, modeling global information and injecting textual information are crucial for text-driven motion generation. The former ensures rationality and semantic consistency across the overall motion, while the latter aligns the motion semantically with the provided text. Given their interrelated nature, we have consolidated global information modeling and textual information injection into a single module, namely the Global Information Modeling and Textual Information Injection Module.




Figure 4: Illustration of our Adaptive Textual Information Injector.  $\odot$  and  $\oplus$  denote dot product and concatenation, respectively.

As shown in Figure 2.c, we first downsample the input motion sequence in the time dimension by a ratio of  $S$  via a lightweight point-wise convolution layer. We then inject textual information into each segment using our proposed Adaptive Textual Information Injector (ATII). Next, we utilize Mamba to model the global information using our proposed pseudo-bidirectional scan. Subsequently, we upsample the segments back to their original length through repeated padding. Finally, we fuse the original motion sequence  $X$  with the upsampled motion sequence  $\hat{X}$  using a fully connected layer  $h(\cdot)$ :

$$\hat{X} = h(X + \bar{X}). \quad (5)$$

**Benefits of Downsample.** We note that the downsample operation has two benefits. Firstly, the semantic information in a single frame is typically weak, making it challenging to extract corresponding semantics from the given text. Furthermore, modeling global information at the frame level is inefficient. Moreover, downsampling reduces the length of the motion sequence from  $L$  to  $\frac{L}{S}$ , significantly decreasing the computational load of subsequent operations.

**Pseudo-bidirectional Scan.** Since the original single-direction scan used in Mamba limits information exchange, we propose a simple yet effective method to achieve a bidirectional scan without adding extra parameters. As shown in Figure 3, we reverse the input motion sequence and feed both the reversed  $X^r$  and the original  $X^o$  sequences into Mamba. Then, during the sequential scan, each element in the original sequence  $X^o$  can access information from adjacent elements on either side. After scanning, only the original sequence  $X^o$  is retained. With this approach, we achieve the effect of a bidirectional scan, as in (Zhu et al. 2024), and a cross-scan, as in (Liu et al. 2024c), without increasing the number of parameters. We conduct experiments that demonstrate the superior efficacy of our pseudo-bidirectional scan.

### 3.5 Injecting Textual Information into Motion

Unlike most existing T2M works that rely on self/cross-attention in Transformers for conditional generation, we propose a novel Adaptive Textual Information Injector (ATII) to inject textual information into the motion during generation more effectively.

Methods	#Params	AIT↓	FID↓	R-Precision↑			MM. Dist.↓	MM. ↑
				Top1	Top2	Top3		
<i>On the HumanML3D dataset.</i>								
TM2T (Guo et al.)	41.05M	0.680s	1.501 $\pm$ .017	0.424 $\pm$ .003	0.618 $\pm$ .003	0.729 $\pm$ .002	3.467 $\pm$ .011	2.424 $\pm$ .093
T2M (Guo et al.)	31.16M	<b>0.038s</b>	1.087 $\pm$ .021	0.455 $\pm$ .003	0.636 $\pm$ .003	0.736 $\pm$ .002	3.347 $\pm$ .008	2.219 $\pm$ .074
MDM (Tevet et al.)	17.88M	14.32s	0.544 $\pm$ .044	-	-	0.611 $\pm$ .007	5.566 $\pm$ .027	<b>2.799</b> $\pm$ .072
MLD (Chen et al.)	26.38M	0.250s	0.473 $\pm$ .013	0.481 $\pm$ .003	0.673 $\pm$ .003	0.772 $\pm$ .002	3.196 $\pm$ .010	2.413 $\pm$ .079
MotionDiffuse (Zhang et al.)	87.15M	17.36s	0.630 $\pm$ .001	0.491 $\pm$ .001	0.681 $\pm$ .001	0.782 $\pm$ .001	3.113 $\pm$ .001	1.553 $\pm$ .042
T2M-GPT (Zhang et al.)	247.80M	0.434s	0.141 $\pm$ .005	0.492 $\pm$ .003	0.679 $\pm$ .002	0.775 $\pm$ .002	3.121 $\pm$ .009	1.831 $\pm$ .048
ReMoDiffuse (Zhang et al.)	46.97M	0.475s	0.103 $\pm$ .004	0.510 $\pm$ .005	0.698 $\pm$ .006	0.795 $\pm$ .004	2.974 $\pm$ .016	1.795 $\pm$ .043
MoMask (Guo et al.)	44.85M	0.180s	0.045 $\pm$ .002	<b>0.521</b> $\pm$ .002	<b>0.713</b> $\pm$ .002	<b>0.807</b> $\pm$ .002	<b>2.958</b> $\pm$ .008	1.241 $\pm$ .040
<b>Our Light-T2M</b>	<b>4.48M</b>	<u>0.151s</u>	<b>0.040</b> $\pm$ .002	<u>0.511</u> $\pm$ .003	<u>0.699</u> $\pm$ .002	<u>0.795</u> $\pm$ .002	3.002 $\pm$ .008	1.670 $\pm$ .061
<i>On the KIT-ML dataset.</i>								
TM2T (Guo et al.)	-	-	3.599 $\pm$ .153	0.280 $\pm$ .005	0.463 $\pm$ .006	0.587 $\pm$ .005	4.591 $\pm$ .026	<b>3.292</b> $\pm$ .081
T2M (Guo et al.)	-	-	3.022 $\pm$ .107	0.361 $\pm$ .005	0.559 $\pm$ .007	0.681 $\pm$ .007	3.488 $\pm$ .028	2.052 $\pm$ .107
MDM (Tevet et al.)	-	-	0.497 $\pm$ .021	-	-	0.396 $\pm$ .004	9.191 $\pm$ .022	1.907 $\pm$ .214
MLD (Chen et al.)	-	-	0.404 $\pm$ .027	0.390 $\pm$ .008	0.609 $\pm$ .008	0.734 $\pm$ .007	3.204 $\pm$ .027	2.192 $\pm$ .071
MotionDiffuse (Zhang et al.)	-	-	1.954 $\pm$ .062	0.417 $\pm$ .004	0.621 $\pm$ .004	0.739 $\pm$ .004	2.958 $\pm$ .005	0.730 $\pm$ .013
T2M-GPT (Zhang et al.)	-	-	0.514 $\pm$ .029	0.416 $\pm$ .006	0.627 $\pm$ .006	0.745 $\pm$ .006	3.007 $\pm$ .023	1.570 $\pm$ .039
ReMoDiffuse (Zhang et al.)	-	-	<b>0.155</b> $\pm$ .006	0.427 $\pm$ .014	0.641 $\pm$ .004	0.765 $\pm$ .055	2.814 $\pm$ .012	1.239 $\pm$ .028
MoMask (Guo et al.)	-	-	0.204 $\pm$ .011	0.433 $\pm$ .007	0.656 $\pm$ .005	0.781 $\pm$ .005	2.779 $\pm$ .022	1.131 $\pm$ .043
<b>Our Light-T2M</b>	-	-	<u>0.161</u> $\pm$ .009	<b>0.444</b> $\pm$ .006	<b>0.670</b> $\pm$ .007	<b>0.794</b> $\pm$ .005	<b>2.746</b> $\pm$ .016	1.005 $\pm$ .036

Table 1: Quantitative evaluation on the HumanML3D and KIT-ML test set. Following previous works, we replicated the experiment 20 times to calculate the average results, presented with a 95% confidence interval (denoted by  $\pm$ ). The best result is bolded and the second is underlined. Average Inference time (AIT) is calculated from the average across 100 samples using the same RTX 3090Ti GPU.

As shown in Figure 4, the concatenation of the text embedding token  $y$  and a segment token  $X_i$  is fed into a fully connected layer  $g^c(\cdot)$  followed by a Sigmoid function to obtain channel-wise weights for the text embedding token, inspired by the gating mechanism in LSTM (Graves and Graves 2012), and then the channel-wise weights are used to reweight the text token:

$$\hat{y} = \text{Sigmoid}(g^c(X_i, y)) \odot y, \quad (6)$$

where  $\odot$  denotes the point-wise production. Thus, the obtained feature  $\hat{y}$  contains the segment-aware semantics. After that, another fully connected layer  $g^f(\cdot)$  is adopted to fuse obtained feature  $\hat{y}$  and the segment token  $X_i$ :

$$\hat{X}_i = g^f(X_i, \hat{y}). \quad (7)$$

In this way, we successfully extract adaptive semantics from the text token for each motion segment and then inject such information into the motion segment.

## 4 Experiments

### 4.1 Datasets and Evaluation Metrics

**Datasets.** We conduct experiments on two most common public text-motion datasets, i.e., the HumanML3D dataset (Guo et al. 2022a) and the KIT-ML dataset (Plappert, Mandery, and Asfour 2016). The **HumanML3D** dataset, constructed based on the HumanAct12 (Guo et al. 2020) and AMASS (Mahmood et al. 2019), contains 14,616 motion

sequences and 44,970 text descriptions. This dataset contains diverse activities including exercising, sports, and acrobatics. The **KIT-ML** dataset is smaller and contains 3,911 motion sequences and 6,278 text descriptions. For both datasets, the preprocessing procedure and the train-test-validation split remain consistent with (Guo et al. 2022a).

**Evaluation Metrics.** We strictly adhere to the evaluation pipeline and adopt the same evaluation metrics used in (Guo et al. 2022a). Each motion-text pair is fed into a pretrained network to extract feature embeddings and then measured by the following metrics: (1) *Frechet Inception Distance (FID)* measures the overall motion quality by calculating the similarity between features of generated and ground truth motions. (2) *R-Precision* and *multimodal distance (MM. Dist.)* assess the semantic alignment between the input text and the generated motions. (3) *Multimodality (MM)* aims to assess the diversity of motions generated from the same text. Following recommendations from MoMask (Guo et al. 2024), FID and R-Precision serve as the primary performance metrics, and the diversity metric is not included.

### 4.2 Implementation Details

The max diffusion step  $T$  is 1000 and the linearly varying variances  $\beta_t$  range from  $10^{-4}$  to  $10^{-2}$ . During inference, we adopt UniPC (Zhao et al. 2023) with 10 time steps for the fast sampling. The number of blocks  $N$ , the hidden dim  $D$ , and the downsampling factor  $S$  are 4, 256, and 8, respec-

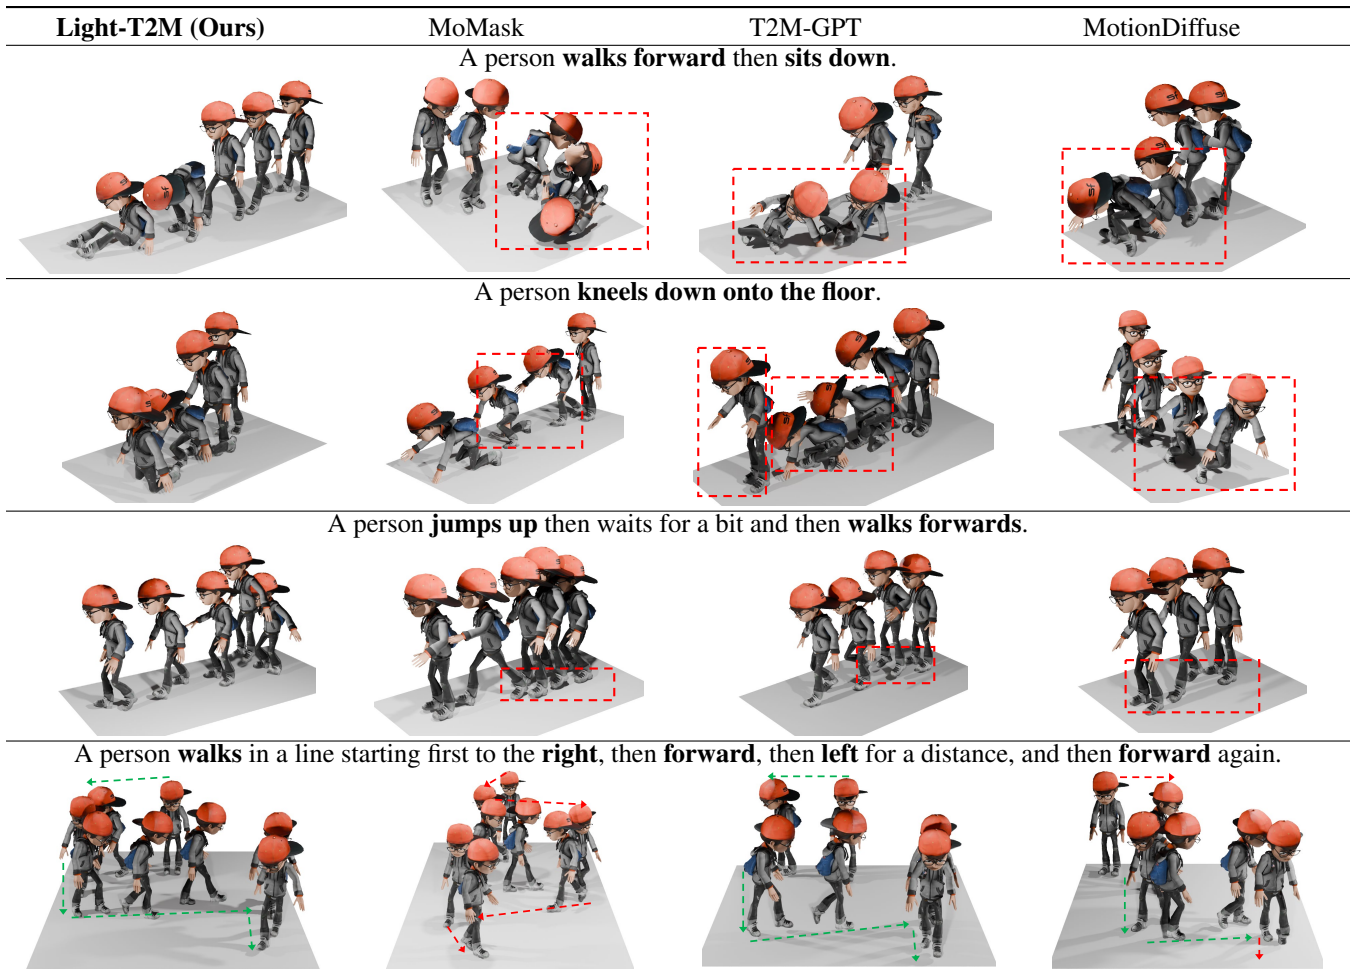


Figure 5: Qualitative comparisons on the HumanML3D dataset. The areas highlighted in red indicate where the generated content does not correspond to the given text or where there are issues such as limb distortion. We also use dashed lines to display the character’s movement path, with green and red indicating whether it corresponds to the given text, respectively.

tively. The guidance scale  $s$  and the text dropout ratio  $\tau$  are set to 4 and 0.2, respectively. Our Light-T2M is optimized by Adamw (Loshchilov and Hutter 2019) with a learning rate of  $2e-4$ , a cosine annealing schedule, and a batch size of 256 on 2 RTX 3090Ti GPUs. Light-T2M is trained with 3000/5000 epochs on the HumanML3D/KIT-ML datasets. More details are shown in the Appendix.

### 4.3 Comparison with State-of-the-arts

Since the model parameters are nearly identical across both datasets and some works do not provide pretrained models on the KIT-ML dataset, we report only the number of trainable parameters and the inference time for the HumanML3D dataset. Note that for models utilizing VAE or VQVAE, the parameters of these architectures are included in the count. Due to limited space, we present the performance of state-of-the-art methods and a part of the works in the table.

**Quantitative Comparison.** The quantitative results are shown in Table 1. Compared to MoMask, our Light-T2M only uses 10% of trainable parameters (4.48M vs 44.85M)

and achieves a 16% faster inference time (0.152s vs 0.180s), while outperforming MoMask in FID on both datasets. In terms of R-Precision and multimodal distance, Light-T2M performs better on the KIT-ML dataset but falls short on the HumanML3D dataset. Regarding the suboptimal FID on the KIT-ML dataset, ReMoDiffuse utilizes complex data retrieval from a large database to enhance motion quality, as noted in MoMask. These quantitative results confirm that Light-T2M not only reduces usage costs but also improves the feasibility of deployment on mobile devices.

**Qualitative Comparison.** Figure 5 shows a qualitative comparison with previous works. MotionDiffuse fails to generate poses for actions such as “sit down” and “knees down”, and it does not accurately interpret the concept of direction. T2M-GPT, while generating semantically correct motions, often produces poses with noticeable limb distortions. The motions generated by MoMask do not align well with the given text and also exhibit limb distortions. In contrast, our model consistently generates correct and reasonable motion poses that accurately reflect the given text and understand

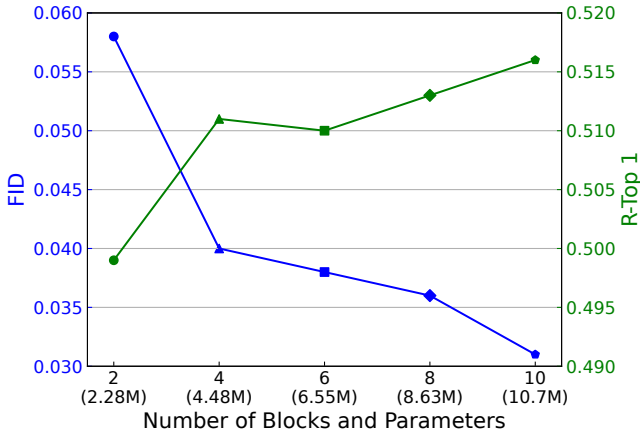


Figure 6: Impact of The Number of Parameters.

directional concepts. Overall, the motion quality produced by our model exceeds that of the existing works.

#### 4.4 Ablation Studies

All ablation studies are conducted on the HumanML3D dataset and confidence intervals are omitted for simplicity.

**Impact of the Number of Parameters.** As shown in Figure 6, we evaluate the model’s performance using different numbers of basic blocks shown in Figure 2. As the number of modules increases, so does the model’s performance. However, while the parameter count rises by 139% when the number of modules increases from 4 to 10, the FID improvement is only 22.3%. Consequently, the model with 4 modules is selected as the final configuration. Surprisingly, our Light-T2M with just 2 basic blocks, comprising 2.28M parameters, achieves an FID of 0.086 and an R-Top 1 of 0.478, demonstrating the method’s effectiveness and efficiency.

**Analysis of Model Design.** Table 2 shows the performance when replacing basic blocks in our Light-T2M with other basic blocks. Notably, when using alternate basic blocks, the concatenation of frame tokens with the text token is processed through the global information layer for textual information injection. From this data, we draw several conclusions: (1) The comparisons between “T”, and “LTL” highlight the significance of local information modeling. (2) The comparisons among “LTL”, “LML”, and “LM\*L” underscore the efficiency of the original Mamba and the value of the bidirectional scan. (3) The contrasts between “LTL”, “LML”, and “LGL” demonstrate that our Light-T2M achieves considerable performance enhancements with only a modest increase in parameters. These findings validate our claims made in the introduction. Please refer to the Appendix for more details about the model design.

**Comparison of Different Scan Methods.** We evaluate the performance when replacing our pseudo-bidirectional scan with different scan methods. As shown in the upper part of Table 3, the significant performance gap between the single-directional scan and the bidirectional scan underscores the necessity of bidirectional information modeling. Furthermore, compared to the bidirectional scan in (Zhu et al.

Block Design	#Params	FID↓	R-Top1↑	R-Top3↑
T	3.49M	0.214	0.437	0.727
LTL	4.03M	0.160	0.469	0.763
LML	<b>2.63M</b>	0.165	0.456	0.736
LM*L	4.38M	<u>0.114</u>	<u>0.501</u>	<u>0.783</u>
LGL (Ours)	4.48M	<b>0.040</b>	<b>0.511</b>	<b>0.795</b>

Table 2: Analysis of Model Design. We evaluate the performance when replacing basic blocks in Light-T2M with other basic blocks. “T” denotes one Transformer encoder layer. “L” denotes our Local Information Modeling Module. “G” denotes our Global Information Modeling and Injection Textual Information Module. “M” and “M\*” denote an original Mamba block and a BiMamba block in (Zhu et al. 2024).

	#Params	FID↓	R-Top1↑	R-Top3↑
SDS	<b>4.48M</b>	0.058	0.495	0.781
BDS	4.66M	0.042	0.509	0.793
PBDS (Ours)	<b>4.48M</b>	<b>0.040</b>	<b>0.511</b>	<b>0.795</b>
w/o ATII	<b>3.43M</b>	0.102	0.489	0.778
w/o gating	3.95M	0.074	0.497	0.781
w/ ATII (Ours)	4.48M	<b>0.040</b>	<b>0.511</b>	<b>0.795</b>

Table 3: Evaluation of Different Scans in Mamba and Analysis on Adaptive Textual Information Injector (ATII). SDS denotes the original single-directional scan in Mamba. BDS denotes the bidirectional scan in (Zhu et al. 2024). PBDS denotes our pseudo-bidirectional scan. “w/o gating” denotes that we directly feed the concatenation of segment token and text token into the fusion layer.

2024), our pseudo-bidirectional scan achieves superior performance without increasing the number of parameters.

**Analysis of Adaptive Textual Information Injector (ATII).** The lower part of Table 3 shows the experimental analysis of our ATII. Note that when removing our ATII, the concatenation of frame tokens with the text token is fed into the global information layer for textual information injection. Besides, in experiments involving the removal of our ATII or the gating mechanism, the downsampling and up-sampling processes are retained. After removing our ATII, the FID and R-Top 1 dropped significantly. When removing the gating mechanism, the performance will also decrease. These results validate the effectiveness of our ATII.

## 5 Conclusion

In this work, we introduce Light-T2M, a lightweight and fast model for text-to-motion generation that underscores the importance and effectiveness of local information modeling. Furthermore, Light-T2M incorporates the Mamba and pseudo-bidirectional scan for global information modeling, both of which significantly reduce parameters. Light-T2M features an adaptive textual information injector, an efficient way to control the generated motion through text. Through these designs, Light-T2M achieves superior performance while using only a small number of parameters.

## Acknowledgments

This work was supported partially by NSFC(92470202, U21A20471) and Guangdong NSF Project (No. 2023B1515040025).

## References

- Belousov, S. 2021. MobileStyleGAN: A lightweight convolutional neural network for high-fidelity image synthesis. *arXiv preprint arXiv:2104.04767*.
- Chen, X.; Jiang, B.; Liu, W.; Huang, Z.; Fu, B.; Chen, T.; 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*, 18000–18010.
- Chen, Y.; Dai, X.; Chen, D.; Liu, M.; Dong, X.; Yuan, L.; and Liu, Z. 2022. MobileFormer: Bridging MobileNet and Transformer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18–24, 2022*, 5260–5269. IEEE.
- Dao, T.; and Gu, A. 2024. Transformers are SSMS: Generalized Models and Efficient Algorithms Through Structured State Space Duality. In *International Conference on Machine Learning (ICML)*.
- Gholami, A.; Kim, S.; Dong, Z.; Yao, Z.; Mahoney, M. W.; and Keutzer, K. 2021. A Survey of Quantization Methods for Efficient Neural Network Inference. *CoRR*, abs/2103.13630.
- Ghosh, A.; Cheema, N.; Oguz, C.; Theobalt, C.; and Slusallek, P. 2021. Synthesis of Compositional Animations from Textual Descriptions. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021*, 1376–1386. IEEE.
- Graves, A.; and Graves, A. 2012. Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, 37–45.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-Time Sequence Modeling with Selective State Spaces. *arXiv preprint arXiv:2312.00752*.
- Guo, C.; Mu, Y.; Javed, M. G.; Wang, S.; and Cheng, L. 2024. Momask: Generative masked modeling of 3d human motions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1900–1910.
- Guo, C.; Zou, S.; Zuo, X.; Wang, S.; Ji, W.; Li, X.; and Cheng, L. 2022a. 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.; and Cheng, L. 2022b. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *European Conference on Computer Vision*, 580–597. Springer.
- 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 Chen, C. W.; Cucchiara, R.; Hua, X.; Qi, G.; Ricci, E.; Zhang, Z.; and Zimmermann, R., eds., *MM '20: The 28th ACM International Conference on Multimedia, Virtual Event / Seattle, WA, USA, October 12–16, 2020*, 2021–2029. ACM.
- Han, B.; Peng, H.; Dong, M.; Ren, Y.; Shen, Y.; and Xu, C. 2024. AMD: Autoregressive Motion Diffusion. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20–27, 2024, Vancouver, Canada, 2022–2030*. AAAI Press.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 6840–6851. Curran Associates, Inc.
- Ho, J.; and Salimans, T. 2022. Classifier-Free Diffusion Guidance. *CoRR*, abs/2207.12598.
- Howard, A. G.; Zhu, M.; Chen, B.; Kalenichenko, D.; Wang, W.; Weyand, T.; Andreetto, M.; and Adam, H. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.
- Hu, L.; Wang, Z.; Li, H.; Wu, P.; Mao, J.; and Zeng, N. 2024. DARTS: Light-weight differentiable architecture search with robustness enhancement strategy. *Knowl. Based Syst.*, 288: 111466.
- Jiang, B.; Chen, X.; Liu, W.; Yu, J.; Yu, G.; and Chen, T. 2024. Motiongpt: Human motion as a foreign language. *Advances in Neural Information Processing Systems*, 36.
- Kim, J.; Kim, J.; and Choi, S. 2023. Flame: Free-form language-based motion synthesis & editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 8255–8263.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In Bengio, Y.; and LeCun, Y., eds., *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings*.
- Kong, H.; Gong, K.; Lian, D.; Mi, M. B.; and Wang, X. 2023. Priority-centric human motion generation in discrete latent space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14806–14816.
- Liu, H.; Galindo, M.; Xie, H.; Wong, L.; Shuai, H.; Li, Y.; and Cheng, W. 2024a. Lightweight Deep Learning for Resource-Constrained Environments: A Survey. *ACM Comput. Surv.*, 56(10): 267.
- Liu, H.-I.; Galindo, M.; Xie, H.; Wong, L.-K.; Shuai, H.-H.; Li, Y.-H.; and Cheng, W.-H. 2024b. Lightweight Deep Learning for Resource-Constrained Environments: A Survey. *ACM Computing Surveys*.
- Liu, X.; Li, Y.; Zeng, A.; Zhou, Z.; You, Y.; and Lu, C. 2023. Bridging the Gap between Human Motion and Action Semantics via Kinematic Phrases. *CoRR*, abs/2310.04189.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024c. VMamba: Visual State Space Model. *CoRR*, abs/2401.10166.

- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Mahmood, N.; Ghorbani, N.; Troje, N. F.; Pons-Moll, G.; and Black, M. J. 2019. AMASS: Archive of Motion Capture As Surface Shapes. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, 5441–5450. IEEE.
- Mehta, S.; and Rastegari, M. 2022. MobileViT: Lightweight, General-purpose, and Mobile-friendly Vision Transformer. In *International Conference on Learning Representations*.
- Petrovich, M.; Black, M. J.; and Varol, G. 2022. TEMOS: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision*, 480–497. Springer.
- Plappert, M.; Mandery, C.; and Asfour, T. 2016. The KIT Motion-Language Dataset. *Big Data*, 4(4): 236–252.
- Plappert, M.; Mandery, C.; and Asfour, T. 2018. Learning a bidirectional mapping between human whole-body motion and natural language using deep recurrent neural networks. *Robotics and Autonomous Systems*, 109: 13–26.
- 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.
- Sandler, M.; Howard, A. G.; Zhu, M.; Zhmoginov, A.; and Chen, L. 2018. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 4510–4520. Computer Vision Foundation / IEEE Computer Society.
- Shafir, Y.; Tevet, G.; Kapon, R.; and Bermano, A. H. 2024. Human motion diffusion as a generative prior. In *International Conference on Learning Representations (ICLR)*.
- Tan, M.; and Le, Q. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, 6105–6114. PMLR.
- Tevet, G.; Raab, S.; Gordon, B.; Shafir, Y.; Cohen-or, D.; and Bermano, A. H. 2023. mdm. In *The Eleventh International Conference on Learning Representations*.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jegou, H. 2021. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, volume 139, 10347–10357.
- 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.
- Wang, Y.; Leng, Z.; Li, F. W.; Wu, S.-C.; and Liang, X. 2023. Fg-t2m: Fine-grained text-driven human motion generation via diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22035–22044.
- Wu, G.; Zheng, W.; Lu, Y.; and Tian, Q. 2023. PSLT: A Light-Weight Vision Transformer With Ladder Self-Attention and Progressive Shift. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(9): 11120–11135.
- Wu, Y.; and He, K. 2020. Group Normalization. *Int. J. Comput. Vis.*, 128(3): 742–755.
- Xiao, X.; Wang, Z.; and Rajasekaran, S. 2019. AutoPrune: Automatic Network Pruning by Regularizing Auxiliary Parameters. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 13681–13691.
- Xie, Y.; Jampani, V.; Zhong, L.; Sun, D.; and Jiang, H. 2024. OmniControl: Control Any Joint at Any Time for Human Motion Generation. In *The Twelfth International Conference on Learning Representations*.
- Yamada, T.; Matsunaga, H.; and Ogata, T. 2018. Paired recurrent autoencoders for bidirectional translation between robot actions and linguistic descriptions. *IEEE Robotics and Automation Letters*, 3(4): 3441–3448.
- Yan, Y.; Zhou, Z.; Wang, Z.; Gao, J.; and Yang, X. 2024. Dialoguerf: Towards realistic avatar face-to-face conversation video generation. *Visual Intelligence*, 2(1): 24.
- Zhang, J.; Zhang, Y.; Cun, X.; Zhang, Y.; Zhao, H.; Lu, H.; Shen, X.; and Shan, Y. 2023a. Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14730–14740.
- Zhang, M.; Cai, Z.; Pan, L.; Hong, F.; Guo, X.; Yang, L.; and Liu, Z. 2024. Motiodiffuse: Text-driven human motion generation with diffusion model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(6): 4115–4128.
- Zhang, M.; Guo, X.; Pan, L.; Cai, Z.; Hong, F.; Li, H.; Yang, L.; and Liu, Z. 2023b. Remodiffuse: Retrieval-augmented motion diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 364–373.
- Zhao, W.; Bai, L.; Rao, Y.; Zhou, J.; and Lu, J. 2023. UniPC: A Unified Predictor-Corrector Framework for Fast Sampling of Diffusion Models. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Zhong, C.; Hu, L.; Zhang, Z.; and Xia, S. 2023. Att2m: Text-driven human motion generation with multi-perspective attention mechanism. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 509–519.
- Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision Mamba: Efficient Visual Representation Learning with Bidirectional State Space Model. *CoRR*, abs/2401.09417.