Motion Mamba: Efficient and Long Sequence Motion Generation

Zeyu Zhang^{12*†}, Akide Liu^{1*}, Ian Reid³, Richard Hartley², Bohan Zhuang¹, and Hao Tang^{4⊠}

 ¹ Monash University
 ² The Australian National University
 ³ Mohamed bin Zayed University of Artificial Intelligence
 ⁴ National Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University

https://steve-zeyu-zhang.github.io/MotionMamba



Fig. 1: Motion Mamba has achieved significantly superior performance on long squence modeling and motion generation efficiency compared with other well-designed state-of-the-art methods such as MLD [6], MotionDiffuse [54], and MDM [49].

Abstract. Human motion generation stands as a significant pursuit in generative computer vision, while achieving long-sequence and efficient motion generation remains challenging. Recent advancements in state space models (SSMs), notably Mamba, have showcased considerable promise in long sequence modeling with an efficient hardware-aware design, which appears to be a significant direction upon building motion generation model. Nevertheless, adapting SSMs to motion generation faces hurdles since the lack of a specialized design architecture to model motion sequence. To address these challenges, we propose **Motion Mamba**, a simple yet efficient approach that presents the pioneering motion generation model utilized SSMs. Specifically, we design a **Hierarchical Temporal Mamba** (**HTM**) block to process temporal data by ensemble varying numbers of isolated SSM modules across a symmetric U-Net architecture aimed at preserving motion consistency between frames. We also design a **Bidirectional Spatial Mamba** (**BSM**) block

^{*}Equal contribution.

[†]Work done while being a research assistant at Monash University.

[⊠]Corresponding author: bjdxtanghao@gmail.com

to bidirectionally process latent poses, to enhance accurate motion generation within a temporal frame. Our proposed method achieves up to 50% FID improvement and up to 4 times faster on the HumanML3D and KIT-ML datasets compared to the previous best diffusion-based method, which demonstrates strong capabilities of high-quality long sequence motion modeling and real-time human motion generation.

Keywords: Human Motion Generation · Selective State Space Models · Latent Diffusion Models

1 Introduction

Human motion generation involves creating high-quality 3D human motion based on given conditions, such as a text prompt, which generates Euler angles for each joint in a human skeleton. To emulate human motion effectively, virtual characters must respond to the conditional context, exhibit natural movement, and perform motion accurately. Recent motion generation models are categorized into four main approaches: autoencoder-based [1,10,17,36,48,56], utilizing transformers for latent space compression and motion synthesis; GAN-based [4,20,30], using discriminators to enhance the realism of generated motions; autoregressive models [24], treating motion sequences as languages with specialized codebooks; and diffusion-based [6,49,54], employing denoising steps for motion generation. Challenges vary across methods, with autoencoder models struggling to generate accurate motions from detailed descriptions due to textual information compression, GAN-based models facing training difficulties, especially in conditional tasks, and diffusion-based models relying on complex transformer-based architectures results in inefficient of motion prediction.

Although diffusion-based models excel at generating motion with robust performance and often exhibit superior diversity, they encounter two limitations. 1) Convolutional or transformer-based diffusion methods exhibit limitations in generating long-range motion sequences. Previous transformer-based methodologies [6,49,54] have focused on modeling long-range dependencies and acquiring comprehensive global context information. Despite these advances, they are frequently associated with a substantial increase in computational requirements. Furthermore, transformer architectures are not intrinsically designed for temporal sequential modeling, which poses an inherent limitation. 2) The efficiency of inference in transformer-based diffusion methods is constrained. Although prior research has attempted to leverage the Variational Autoencoder for denoising operations in the latent space [6], the inference speed remains adversely affected by the attention mechanism's quadratic scaling, leading to inefficient motion generation. Consequently, exploring a new architectural paradigm that accommodates long-range dependencies and maintains a linear computational complexity is crucial for sustaining motion generation tasks.

Recent advances have sparked renewed interest in state space models (SSMs) [14, 15], a field that originated from the foundational classic state space model

² Zeyu and Akide et al.

3

[25]. Modern versions of SSMs stand out due to their ability to effectively capture long-range dependencies, a capability greatly improved by the introduction of parallel training techniques. This evolution has led to various methodologies based on SSM, notably the linear state space layers (LSSL) [15], the structured state-space sequence model (S4) [14], the diagonal state space (DSS) [19], and S4D [13]. These methods have been carefully designed to handle sequential data across various tasks and modalities, paying special attention to modeling longrange dependencies. Their efficacy in managing long sequences is attributed to the implementation of convolutional computations [15] and near-linear computational strategies, such as mamba [12], marking a significant stride in sequentially oriented tasks, including large language model decoding and motion sequence generation.

Adapting selective state space modules for motion generation tasks presents notable challenges, primarily due to the lack of specialized design in SSMs for capturing the sensitive motion details required for temporal representation and the complexities involved in aggregating latent space. In response to these challenges, we have meticulously developed a motion generation architecture, specifically tailored to address the intricacies of long-term sequence generation, while optimizing for computational efficiency with near-linear-time complexity. This innovation is embodied in the Motion Mamba model, a simple vet potent approach to motion generation. The Motion Mamba framework pioneers a diffusion-based generative system, incorporating two key components oriented toward SSM as shown in Figure. 2: (1) a Hierarchical Temporal Mamba (HTM) block: This component is ingeniously crafted to arrange motion frames in sequential order, using hierarchically adjusted scanning. It is adept at identifying temporal dependencies at various depths, thereby facilitating a thorough comprehension of the dynamics inherent in motion sequences. (2) a Bidirectional Spatial Mamba (BSM) block: This block is designed to unravel the structured latent skeleton by evaluating data from both forward and reverse directions. Its primary goal is to safeguard the continuity of information flow, significantly bolstering the model's capacity for precise motion generation through the retention of dense informational exchange.

The Motion Mamba introduces a new approach to motion generation that strikes an exceptional trade-off between accuracy and efficiency, shown in Fig. 1. Our experimental results underscore the significant improvements brought about by Motion Mamba, showcasing a remarkable improvement in the Fréchet Inception Distance (FID), with a reduction of up to 50% from the prior state-of-the-art metric of 0.473 to an impressive 0.281 on the HumanML3D dataset [17]. Furthermore, we emphasize our framework's unparalleled inference speed, which is four times faster than previous methods, achieving an average inference time of only 0.058 seconds per sequence compared to the 0.217 seconds required by the MLD [6] method per sequence. These outcomes unequivocally establish Motion Mamba's state-of-the-art performance, concurrently ensuring fast inference speeds for conditional human motion generation tasks.

Our contributions to the field of motion generation can be summarized as:

- 4 Zeyu and Akide et al.
- 1. We introduce a simple yet effective framework, named *Motion Mamba*, which is a pioneering method integrates a selective scanning mechanism into motion generation tasks.
- 2. Motion Mamba is comprised of two modules: Hierarchical Temporal Mamba (HTM) and Bidirectional Spatial Mamba (BSM), which are designed for temporal and spatial modeling, respectively. HTM blocks are tasked with processing temporal motion data, aiming to enhance motion consistency across frames. BSM blocks are engineered to bidirectionally capture the channel-wise flow of hidden information within the latent pose representations.
- Motion Mamba framework demonstrated exceptional performance on text-tomotion generation task, through experimental validation on the HumanML3D [17] and KIT-ML [38] datasets. Our methodology achieved state-of-the-art generation quality and significantly improved long-squence modeling, meanwhile optimizing inference speed.

2 Related Works

Human Motion Generation. Generating human motion is a significant application of computer vision, essential for various applications like 3D modelling and robot manipulation. Recently, the predominant method of achieving human motion generation, known as the *Text-to-Motion* task, involves learning a common latent space for both language and motion.

DVGAN [30] creates the GAN [11] discriminator by densely validating at each time-scale and perturbing the discriminator input for translation invariance, enabling motion generation and completion. ERD-QV [20] enhances latent representations through two additive modifiers: a time-to-arrival embedding applied universally and an additive scheduled target noise vector used during extended transitions. It further improves transition quality by incorporating a GAN framework with two discriminators operating at different timescales. HP-GAN [4], trained with a modified version of the improved WGAN-GP [16], utilizes a custom loss function designed for human motion prediction. It learns a probability density function of future human poses conditioned on previous poses.

Autoencoders [27, 43] are notable generative models known for their ability to represent data robustly by compressing high-dimensional data into a latent space, which is widely adopted in human motion generation. JL2P [1] uses RNN-based autoencoders [23] to learn a combined representation of language and pose. It restricts the direct mapping from text to motion to a one-to-one relationship. MotionCLIP [48] uses Transformer-based Autoencoders [51] to reconstruct motion while ensuring alignment with the corresponding text label in the CLIP [40] space. This alignment effectively integrates the semantic knowledge from CLIP into the human motion manifold. TEMOS [36] and T2M [17] combine a Transformer-based VAE [26] with a text encoder to generate distribution parameters that work within the VAE latent space. AttT2M [56] and TM2D [10] incorporate a body-part spatio-temporal encoder into VQ-VAE [50] for enhanced learning of a discrete latent space with increased expressiveness.

5

Diffusion models [7, 22, 41, 47] have recently surpassed GANs and VAEs in generating 2D images. Developing a motion generation model based on diffusion models is obviously an attractive direction. MotionDiffuse [54] introduces the inaugural framework for text-driven motion generation based on diffusion models. It showcases several desirable properties, including probabilistic mapping, realistic synthesis, and multi-level manipulation. MDM [49] utilizes a classifierfree Transformer-based diffusion model for the human motion domain to predict sample rather than noise in each diffusion step. MLD [6] performs a diffusion process in latent motion space, rather than using a diffusion model to establish connections between raw motion sequences and conditional inputs.

State Space Models. Recently, state space sequence models (SSMs) [14, 15], drawing inspiration from classical state-space models [25], have emerged as a promising architecture for sequence modeling. Mamba [12] introduces a selective SSM architecture, integrating time-varying parameters into the SSM framework, and proposes a hardware-aware algorithm to facilitate highly efficient training and inference processes. Some research works leverage SSM in computer vision to process 2D data. The 2D SSM [3] introduces an SSM block at the beginning of each transformer block [8, 51]. This approach aims to achieve efficient parameterization, accelerated computation, and a suitable normalization scheme. SGConvNeXt [29] presents a structured global convolution method inspired by ConvNeXt [32], incorporating multi-scale sub-kernels to achieve both parameterization efficiency and effective long sequence modeling. ConvSSM [46] integrates the tensor modeling principles of ConvLSTM [45] with SSMs, elucidating the utilization of parallel scans in convolutional recurrences. This approach enables subquadratic parallelization and rapid autoregressive generation. Vim [57] introduces a bidirectional SSM block [52] for efficient and versatile visual representation learning, achieving performance comparable to established ViT [8] methods. VMamba [31] introduces a Cross-Scan Module (CSM) designed to traverse the spatial domain and transform any non-causal visual image into ordered patch sequences. This approach achieves linear complexity while preserving global receptive fields. There have also been attempts to utilize SSMs to handle higher-dimensional data. Mamba-ND [28] explores various combinations of SSM and different scan directions within the SSM block to adapt Mamba [12] to higher-dimensional tasks. Recent efforts have sought to replace the traditional transformer-based U-Net within the diffusion denoiser with the SSM block, with the aim of enhancing image generation efficiency. DiffuSSM [53] adeptly manages higher resolutions without relying on global compression, thus maintaining detailed image representation throughout the diffusion process.

3 The Proposed Method

In this section, we delineate the architecture and operational principles of the *Motion Mamba* framework, designed for generating human motion over long ranges efficiently from textual descriptions. Initially, we discuss the foundational concepts underpinning our approach, including the Mamba Model [12] and the

latent diffusion model [6]. Following this, we detail our uniquely crafted architecture that leverages the Mamba model to enhance motion generation efficiency. This architecture comprises two principal components: the Hierarchical Temporal Mamba (HTM) block, which addresses temporal aspects, and the Bidirectional Spatial Mamba (BSM) block, focusing on spatial dynamics.

3.1 Preliminaries

Selective Structured State Space Sequence Model. SSMs particularly through the contributions of structured state space sequence models (S4) and Mamba, have demonstrated exceptional proficiency in handling long sequences. These models operationalize the mapping of a 1-D function or sequence, $x(t) \in \mathbb{R} \mapsto y(t) \in \mathbb{R}$, through a hidden state $h(t) \in \mathbb{R}^N$, employing $\mathbf{A} \in \mathbb{R}^{N \times N}$ as the evolution parameters, $\mathbf{B} \in \mathbb{R}^{N \times 1}$ and $\mathbf{C} \in \mathbb{R}^{1 \times N}$ as the projection parameters, respectively.

The discretized system can then be expressed as follows, incorporating a step size Δ :

$$h_t = \mathbf{A}h_{t-1} + \mathbf{B}x_t,$$

$$y_t = \mathbf{C}h_t.$$
(1)

This adaptation facilitates the computation of output through global convolution, leveraging a structured convolutional kernel $\overline{\mathbf{K}}$, which encompasses the entire length M of the input sequence \mathbf{x} :

$$\overline{\mathbf{K}} = (\mathbf{C}\overline{\mathbf{B}}, \mathbf{C}\overline{\mathbf{A}}\overline{\mathbf{B}}, \dots, \mathbf{C}\overline{\mathbf{A}}^{M-1}\overline{\mathbf{B}}),$$

$$\mathbf{y} = \mathbf{x} * \overline{\mathbf{K}}.$$
 (2)

Selective models like Mamba introduce time-varying parameters, deviating from the linear time invariance (LTI) assumption and complicating parallel computation. However, hardware-aware optimizations, such as associative scans, have been developed to address these computational challenges, highlighting the ongoing evolution and application of SSMs in modeling complex temporal dynamics. Latent Motion Diffusion Model. Diffusion probabilistic models offer a significant advancement in motion generation by gradually reducing noise from a Gaussian distribution to a target data distribution p(x) through a T-length learned Markov process [7,22,41,44,47,49,54], giving $\{\boldsymbol{x_t}\}_{t=1}^T$. In the motion generation, we define our trainable diffusion models with a denoiser $\epsilon_{\theta}(x_t, t)$ which anneal the random noise to motion sequence $\{\hat{x}_t^{1:N}\}_{t=1}^T$ iteratively. To address the inefficiencies of applying diffusion models directly to raw motion sequences, we employ a low-dimensional motion latent space for the diffusion process. Given an input condition c, such as a descriptive sentence $w^{1:N} = \{w^i\}_{i=1}^N$, an action label a from a predefined set \mathcal{A} , or an empty condition $c = \emptyset$, and the motion representation that combines 3D joint rotations, positions, velocities, and foot contact as proposed in [17]. The frozen CLIP [40] text encoder τ_{θ}^{w} has been employed to obtain projected text embedding $\tau_{\theta}^{w}(w^{1:N}) \in \mathbb{R}^{1 \times d}$, thereby conditional denoiser comprised in term of $\epsilon_{\theta}(z_t, t, \tau_{\theta}(c))$. The latent diffusion model $\epsilon_{\theta}(x_t, t)$ aimed



Fig. 2: This figure illustrates the architecture of the proposed Motion Mamba model. Each of encoder and decoder blocks consists of a Hierarchical Temporal Mamba block (HTM) and a Bidirectional Spatial Mamba (BSM) block, which possess hierarchical scan and bidirectional scan within SSM layers respectively. This symmetric distribution of scans ensure a balanced and coherence framework across the encoder-decoder architecture.

to generate the human motion sequence in terms of $\hat{x}^{1:L} = \{\hat{x}^i\}_{i=1}^L$, where L denotes the sequence length or number of frames [34,35,37,55]. After that we reused the motion Variational AutoEncoder (VAE) $\mathcal{V} = \{\mathcal{E}, \mathcal{D}\}$ proposed in MLD [6] to manipulate the motion sequence in latent space $z = \mathcal{E}(x^{1:L})$, and decompress the intermediate representation to motion sequence by $\hat{x}^{1:L} = \mathcal{D}(z) = \mathcal{D}\mathcal{E}(x^{1:L})$ [26, 42, 51]. Finally, our latent diffusion model is trained with an objective focusing on minimization of MSE between true and predicted noise in the latent space, facilitating efficient and high-quality motion generation [2, 22].

3.2 Motion Mamba

The architecture of the proposed *Motion Mamba* framework is illustrated in Figure. 2. At its core, Motion Mamba utilizes a denoising U-Net architecture, which is distinguished for its effectiveness in modeling the continuous, temporal sequences of motion frames. This effectiveness is attributed to the inherent long-sequence modeling capacity of the Mamba model. The denoiser, denoted by ϵ_{θ} , comprises N blocks including encoder $E_{1..N}$ and decoder $D_{1..N}$. Additionally, the architecture is enhanced with a transformer-based attention mixer block M, designed to augment the model's ability to capture complex temporal dynamics.

$$\epsilon_{\theta}(x) \equiv \{E_{1\dots N}, M, D_{1\dots N}\}.$$
(3)

The encoder blocks are represented as $E_{1..N}$, arranged sequentially, and the decoder blocks as $D_{1..N}$, configured in reverse order to facilitate effective bottomup and top-down information flow. Given that selective operations have signif-

Algorithm 1 Hierarchical Temporal Mamba (HTM) Block.

Require: compressed latent representations z : (T, B, C)**Ensure:** transformed representations z_{HTM} : (T, B, E)1: /* linear projection layer */ 2: $x, z : (T, B, E) \leftarrow \text{Linear}(z)$ 3: /* set of scans and memory matrices */ 4: $K = \{S_{2N_n-1}, S_{2N_{n-1}-1}, \dots, S_1\}$ 5: Memory matrices: $\{A_1, \ldots, A_k\}$ 6: for each scan S_i in K do $x'_o: (T, B, E) \leftarrow \text{Conv1D}(x)$ 7: $B_o, C_o, \Delta_o \leftarrow \operatorname{Linear}(x'_o)$ 8: Transform \overline{A}_o and \overline{B}_o using Δ_o 9: 10: $O_i \leftarrow \text{SSM}_{A_i,x}(x'_o)$ 11: end for 12: /* aggregation of outputs */ 13: z_{HTM} : $(T, B, E) \leftarrow \text{Linear}(\text{Aggregate}(\{O_1, \dots, O_k\}))$ 14: Return: $z_{\rm HTM}$

icantly lower computational complexity compared to attention-based methods, we have increased the number of scans to achieve higher quality generations. Concurrently, it is imperative to maintain a balance between the model's parameters and its efficiency. Thereby, a novel aspect of our model is the introduction of a hierarchical scan strategy, characterized by a sequence of scan numbers as,

$$K = \{S_{2N-1}, S_{2(N-1)-1}, \dots, S_1\}.$$
(4)

This sequence specifies the number of scans allocated to each layer, in descending order of complexity. For instance, the uppermost encoder layer, E_1 , and the lowermost decoder layer, D_N , are allocated S_{2N-1} scans, indicating the highest scanning complexity. Conversely, the lowest encoder layer, E_N , and the uppermost decoder layer, D_1 , are assigned S_1 scans, reflecting the lowest level of scanning complexity.

$$E_i(S) = \begin{cases} S_{2N-1} & \text{for } i = 1\\ S_{2(N-i)-1} & \text{for } i = 2, \dots, N-1\\ S_1 & \text{for } i = N \end{cases}$$
(5)

$$D_{j}(S) = \begin{cases} S_{2N-1} & \text{for } j = N \\ S_{2(N-j)-1} & \text{for } j = N-1, \dots, 2 \\ S_{1} & \text{for } j = 1 \end{cases}$$
(6)

This hierarchical scanning approach ensures that processing capabilities are evenly distributed throughout the encoder-decoder architecture., facilitating a detailed and nuanced analysis of temporal sequences. Within this structured framework, each denoiser is equipped with a specialized Hierarchical Temporal Mamba (HTM) block, which serves to augment the model's ability to process Algorithm 2 Bidirectional Spatial Mamba (BSM) Block. **Require:** compressed latent representations z : (T, B, C)**Ensure:** transformed representations z_{BSM} : (C, B, E)1: /* dimension rearrangement */ 2: $z': (C, B, T) \leftarrow \text{Rearrange}(z)$ 3: /* linear projection after normalization */ 4: $z' : (C, B, T) \leftarrow \operatorname{Norm}(z)$ 5: $x, z : (C, B, E) \leftarrow \text{Linear}(z')$ 6: for o in {forward, backward} do $x'_o: (C, B, E) \leftarrow \text{Conv1D}(x)$ 7: $B_o, C_o, \Delta_o \leftarrow \operatorname{Linear}(x'_o)$ 8: Transform \overline{A}_o and \overline{B}_o using Δ_o 9: $y_o \leftarrow \text{SSM}(\overline{A}_o, \overline{B}_o, C_o)$ 10: 11: end for 12: /* gating and summing outputs */13: $z_{\text{BSM}} : (C, B, T) \leftarrow \text{GateAndSum}(y_{forward}, y_{backward}, z)$ 14: $z'_{\text{BSM}} : (T, B, C) \leftarrow \text{Rearrange}(z)$ 15: Return: z'_{BSM}

temporal information effectively. Additionally, the proposed Motion Mamba incorporates an attention-based mixer block denoted as M, strategically integrated to enhance conditional fusion.

Hierarchical Temporal Mamba (HTM) block processes compressed latent representations, denoted as z, with the dimensions (T, B, C), of which procedure shown in Algorithm 1. Here, T signifies the temporal dimension, as specified in the Variational AutoEncoder (VAE) framework. Initially, the input z is subjected to a linear projection layer, producing transformed representations xand z with dimension E. Our analysis revealed an increased density of motion within the lower-level feature spaces. Consequently, we developed a hierarchical scanning methodology that is executed at various depths of the network. This approach not only accommodates the diverse motion densities, but also significantly reduces computational overhead. This step utilizes a hierarchically structured set of scans, $K = \{S_{2N_n-1}, S_{2N_{n-1}-1}, \dots, S_1\}$, in conjunction with a corresponding series of memory matrices $\{A_1, \ldots, A_k\}$. Each sub-SSM scan first applies a 1-D convolution to x, resulting in x'_o . x'_o is then linearly projected to derive B_o, C_o , and Δ_o . These projections B_o, C_o use Δ_o to effect transformations in \overline{A}_o and \overline{B}_o , respectively. After executing a sequence of SSM scans $\{SSM_{A_1,x}, SSM_{A_2,x}, \ldots, SSM_{A_k,x}\}$, a set of outputs $\{O_1, \ldots, O_k\}$ is compiled. This collection is subsequently aggregated via a linear projection to obtain the final output of the HTM block.

Bidirectional Spatial Mamba (BSM) block focuses on enhancing latent representation learning through a novel approach of dimension rearrangement and bidirectional scanning, of which the process is shown in Algorithm 2. Initially, it alters the original input dimensions from (T, B, C) to (C, B, T), effectively swap-

10 Zeyu and Akide et al.

ping the temporal and channel dimensions. After this rearrangement, the input, now denoted z', undergoes a linear projection after normalization, resulting in dimensions x and z of size E. The process involves bidirectional scanning of the latent channel dimension, where \mathbf{x} is subjected to a 1-D convolution, yielding \mathbf{x}'_o for both forward and backward directions. Each \mathbf{x}'_o is then linearly projected to obtain B_o , C_o , and Δ_o , which are utilized to transform \overline{A}_o and \overline{B}_o , respectively. The final output token sequence, \mathbf{z}_1 , is computed by gating and summing the forward y_{forward} and backward y_{backward} output with \mathbf{z} . This component is engineered to decode the structured latent skeleton by analyzing data from both forward and reverse viewpoints. Its main objective is to ensure the seamless continuity of information flow, thereby substantially enhancing the model's ability to generate accurate motion. This is achieved through the maintenance of a dense informational exchange, which is critical for the model's performance.

4 Experiments

4.1 Datasets

We evaluate our proposed Motion Mamba on two prominent *Text-to-Motion* synthesis benchmarks as follows:

HumanML3D. The HumanML3D [17] dataset aggregates 14,616 motions sourced from the AMASS [33] and HumanAct12 [18] datasets, with each motion accompanied by three textual descriptions, culminating in 44,970 scripts. This dataset spans a wide range of actions such as exercising, dancing, and acrobatics, presenting a rich motion-language corpus.

KIT-ML. The KIT-ML dataset [38] is comprised of 3,911 motions paired with 6,278 textual descriptions, serving as a compact yet effective benchmark for evaluation. For both datasets, the pose representation adopted is derived from T2M [17], ensuring consistency in motion representation across evaluations.

4.2 Evaluation Metrics

We adapt the standard evaluation metrics on following aspects throughout our experiments, including: **Generation Quality.** We implement a Fréchet inception distance (FID) [21] to quantify the realism and diversity of motion generated by models. Moreover, we use multi-modal distance (MM Dist) to measure the distance between motions and texts and assess motion-text alignment. **Diversity.** We use the diversity metric to measure motion diversity, which calculates variance in features extracted from the motions. Additionally, we employ multi-modality (MModality) to assess diversity within generated motions sharing the same text description.

4.3 Comparative Studies

We evaluate our method against the state-of-the-art methods on the HumanML3D [17] and KIT-ML [38] datasets. We train our Motion Mamba with HTM arrangement strategy MM ($\{S_{2N_n-1}, \ldots, S_1\}$), BSM bidirectional block strategy on the

Table 1: Comparison of text-conditional motion synthesis on HumanML3D [17]. These metrics are evaluated by the motion encoder from [17]. Empty MModality indicates the non-diverse generation methods. We employ real motion as a reference and sort all methods by descending FIDs. The right arrow \rightarrow means that the closer to the real motion, the better. **Bold** and <u>underline</u> indicate the best and second best result.

Method	R Precision \uparrow			FID	MM Dist.	Diversity→	MModalitv↑
	Top 1	Top 2	Top 3	ΠDΨ	11111 121504	Diversity /	
Real	$0.511^{\pm.003}$	$0.703^{\pm.003}$	$0.797^{\pm.002}$	$0.002^{\pm.000}$	$2.974^{\pm.008}$	$9.503^{\pm.065}$	-
Seq2Seq [39]	$0.180^{\pm.002}$	$0.300^{\pm.002}$	$0.396^{\pm.002}$	$11.75^{\pm.035}$	$5.529^{\pm.007}$	$6.223^{\pm.061}$	-
LJ2P [1]	$0.246^{\pm.001}$	$0.387^{\pm.002}$	$0.486^{\pm.002}$	$11.02^{\pm.046}$	$5.296^{\pm.008}$	$7.676^{\pm.058}$	-
T2G [5]	$0.165^{\pm.001}$	$0.267^{\pm.002}$	$0.345^{\pm.002}$	$7.664^{\pm.030}$	$6.030^{\pm.008}$	$6.409^{\pm.071}$	-
Hier [9]	$0.301^{\pm.002}$	$0.425^{\pm.002}$	$0.552^{\pm.004}$	$6.532^{\pm.024}$	$5.012^{\pm.018}$	$8.332^{\pm.042}$	-
TEMOS [37]	$0.424^{\pm.002}$	$0.612^{\pm.002}$	$0.722^{\pm.002}$	$3.734^{\pm.028}$	$3.703^{\pm.008}$	$8.973^{\pm.071}$	$0.368^{\pm.018}$
T2M [17]	$0.457^{\pm.002}$	$0.639^{\pm.003}$	$0.740^{\pm.003}$	$1.067^{\pm.002}$	$3.340^{\pm.008}$	$9.188^{\pm.002}$	$2.090^{\pm .083}$
MDM [49]	$0.320^{\pm.005}$	$0.498^{\pm.004}$	$0.611^{\pm.007}$	$0.544^{\pm.044}$	$5.566^{\pm.027}$	$9.559^{\pm.086}$	$2.799^{\pm.072}$
MotionDiffuse [54]	$0.491^{\pm.001}$	$0.681^{\pm.001}$	$0.782^{\pm.001}$	$0.630^{\pm.001}$	$3.113^{\pm.001}$	$9.410^{\pm.049}$	$1.553^{\pm.042}$
MLD [6]	$0.481^{\pm.003}$	$0.673^{\pm.003}$	$0.772^{\pm.002}$	$\underline{0.473}^{\pm.013}$	$\overline{3.196}^{\pm.010}$	$9.724^{\pm.082}$	$2.413^{\pm .079}$
Motion Mamba (Ours)	$0.502^{\pm.003}$	$0.693^{\pm.002}$	$0.792^{\pm.002}$	$0.281^{\pm.009}$	$3.060^{\pm.058}$	$9.871^{\pm.084}$	$2.294^{\pm.058}$

Table 2: We involve KIT-ML [38] dataset and evaluate the SOTA methods on the text-to-motion task.

Method	Ι	R Precision \uparrow			MM Dist	Diversity→	MModalitv↑
	Top 1	Top 2	Top 3	1124	11111 21004	Direibity ,	initio durity
Real	$0.424^{\pm.005}$	$0.649^{\pm.006}$	$0.779^{\pm.006}$	$0.031^{\pm.004}$	$2.788^{\pm.012}$	$11.08^{\pm.097}$	-
Seq2Seq [39]	$0.103^{\pm.003}$	$0.178^{\pm.005}$	$0.241^{\pm.006}$	$24.86^{\pm.348}$	$7.960^{\pm.031}$	$6.744^{\pm.106}$	-
T2G [5]	$0.156^{\pm.004}$	$0.255^{\pm.004}$	$0.338^{\pm.005}$	$12.12^{\pm.183}$	$6.964^{\pm.029}$	$9.334^{\pm.079}$	-
LJ2P [1]	$0.221^{\pm.005}$	$0.373^{\pm.004}$	$0.483^{\pm.005}$	$6.545^{\pm.072}$	$5.147^{\pm.030}$	$9.073^{\pm.100}$	-
Hier [9]	$0.255^{\pm.006}$	$0.432^{\pm.007}$	$0.531^{\pm.007}$	$5.203^{\pm.107}$	$4.986^{\pm.027}$	$9.563^{\pm.072}$	$2.090^{\pm .083}$
TEMOS [37]	$0.353^{\pm.006}$	$0.561^{\pm.007}$	$0.687^{\pm.005}$	$3.717^{\pm.051}$	$3.417^{\pm.019}$	$10.84^{\pm.100}$	$0.532^{\pm.034}$
T2M [17]	$0.370^{\pm.005}$	$0.569^{\pm.007}$	$0.693^{\pm.007}$	$2.770^{\pm.109}$	$3.401^{\pm.008}$	$10.91^{\pm.119}$	$1.482^{\pm.065}$
MDM [49]	$0.164^{\pm.004}$	$0.291^{\pm.004}$	$0.396^{\pm.004}$	$0.497^{\pm.021}$	$9.191^{\pm.022}$	$10.85^{\pm.109}$	$1.907^{\pm.214}$
MotionDiffuse [54]	$0.417^{\pm.004}$	$0.621^{\pm.004}$	$0.739^{\pm.004}$	$1.954^{\pm.062}$	$2.958^{\pm.005}$	$11.10^{\pm.143}$	$0.730^{\pm.013}$
MLD [6]	$0.390^{\pm.008}$	$0.609^{\pm.008}$	$0.734^{\pm.007}$	$0.404^{\pm .027}$	$3.204^{\pm.027}$	$10.80^{\pm.117}$	$2.192^{\pm.071}$
Motion Mamba (Ou	urs) $0.419^{\pm.006}$	$0.645^{\pm.005}$	$0.765^{\pm.006}$	$\textbf{0.307}^{\pm.041}$	$3.021^{\pm.025}$	$11.02^{\pm.098}$	$1.678^{\pm.064}$

latent dimension = 2 with 11 layers. We evaluate our model and previous works with suggested metrics in HumanML3D [17] and calculate 95% confidence interval by repeat evaluation 20 times. The results for the HumanML3D dataset are presented in Table 1. Our model outperforms other methods significantly across various evaluation metrics, including FID, R precision, multi-modal distance, and diversity. For instance, our Motion Mamba outperforms previous best diffusion based motion generation MLD by 40.5% in terms of FID, and up to 10% improvement on R Precision, we aslo obtained best MModality by 3.060. The results for the KIT-ML dataset are presented in Table 2. We have also outperformed other well-established methods in FID and multi-modal distance.

4.4 Ablation Studies

We concluded the ablation studies including long sequence evaluation, hierarchical design with HTM, bidirectional design in the BSM, number of latent dimensions, and number of layers of our proposed motion mamba in Table 4.

12 Zeyu and Akide et al.

Table 3: In order to evaluate the models' capability in long sequence motion generation, we compared our method with an existing approach on the recently introduced HumanML3D-LS dataset. This dataset comprises motion sequences longer than 190 frames from the original evaluation set. Our model demonstrates superior performance compared to other methods.

Method	R Precision \uparrow			FID.L	MM Dist.	Diversity→	MModalitv↑
	Top 1	Top 2	Top 3	+			
Real	$0.437^{\pm.003}$	$0.622^{\pm.004}$	$0.721^{\pm.004}$	$0.004^{\pm.000}$	$3.343^{\pm.015}$	$8.423^{\pm.090}$	-
MDM [49] MotionDiffuse [54]	$0.368^{\pm.005}$ $0.367^{\pm.004}$	$0.553^{\pm.006}$ $0.521^{\pm.004}$	$0.672^{\pm.005}$ $0.623^{\pm.004}$	$\frac{0.802^{\pm.044}}{2.460^{\pm.062}}$	$3.860^{\pm.025}$ $3.789^{\pm.005}$	$\frac{8.817^{\pm.068}}{8.707^{\pm.143}}$	$1.602^{\pm.013}$
MLD [6] Motion Mamba (Ours)	$0.403^{\pm .000}$ 0.417 ^{±.003}	$0.584^{\pm.000}$ 0.606 ^{±.003}	$0.690^{\pm .000}$	0.952 ^{±.020}	$3.580^{\pm.015}$	$9.050^{\pm.000}$ $9.021^{\pm.070}$	$2.711^{\pm .001}$ $2.373^{\pm .084}$

Long Sequence Motion Generation. The HumanML3D [17] dataset exhibits a longtailed and right-skewed distribution with a significant proportion of long-sequence human motions, as shown in Figure 3. We suggest previous study overlooked the challenges in the long-sequence generation problem. Thus, we introduce a new dataset variant, *HumanML3D-LS*, comprising motion sequences longer than 190 frames extracted from the original test set. This addition allows us to showcase our capability in generating long-sequence motions. Subsequently, we eval-



Fig. 3: The figure shows a long tail distribution of the HumanML3D [17], which has a significant proportion of long-sequence human motions.

uate the performance of our method on HumanML3D-LS and compare it with other diffusion-based motion generation approaches. The comparative results are presented in Table 3. Motion Mamba by leverage the benefits on long-range dependency modeling make it well suitable for long sequence motion generation.

Hierarchical Design with HTM. In our ablation studies, we observed a slight improvement upon reversing the scan order from a lower to a higher level, specifically transitioning from MM $\{S_1, \ldots, S_N\}$ to MM $\{S_N, \ldots, S_1\}$. This enhancement suggests a correlation with the increase in temporal motion density within the lower-level feature spaces. Furthermore, to achieve the optimal result, we introduce the hierarchical design to arrange the scanning frequency, resulting in the sequence MM $\{S_{2N_n-1}, \ldots, S_1\}$. This expansion in the number of scans led to a performance increase. We attribute this enhancement to the observation that individual selective scan operations significantly reduce the parameter count, especially when compared to the parameter-intensive constructs of self-attention and feedforward network blocks prevalent in transformer architectures.

Bidirectional Design in BSM. We developed three distinct variations of latent scanning mechanisms, differentiated by their scanning directions. In the context of motion generation tasks, we posit that the flow of hidden information within the structured latent skeleton holds significance, an aspect previously underexplored. Our ablation study reveals that a *single scan* across the latent skeleton holds.

Table 4: Evaluation of text-based motion synthesis on HumanML3D [17]: we use metrics in Table 1 and provides real reference, we evaluate the various HTM and BSM design choices, the dimension of the latent input, the different number of layer of Motion Mamba model.

Models	R Precision Top 3↑	FID↓	MM Dist.↓	${\rm Diversity} {\rightarrow}$	MModality↑
Real	$0.797^{\pm.002}$	$0.002^{\pm.000}$	$2.974^{\pm.008}$	$9.503^{\pm.065}$	-
$ \begin{array}{c} \operatorname{MM} \left(\{S_1, \dots, S_N\} \right) \\ \operatorname{MM} \left(\{S_N, \dots, S_1\} \right) \\ \operatorname{MM} \left(\{S_1, \dots, S_{2N_{n-1}}\} \right) \\ \operatorname{MM} \left(\{S_{2N_n-1}, \dots, S_1\} \right) \end{array} $	$\begin{array}{c} 0.673^{\pm.003} \\ 0.738^{\pm.002} \\ 0.698^{\pm.002} \\ 0.792^{\pm.002} \end{array}$	$\begin{array}{c} 1.278^{\pm.012}\\ 0.962^{\pm.011}\\ 0.856^{\pm.008}\\ 0.281^{\pm.009} \end{array}$	$\begin{array}{c} 3.802^{\pm.041} \\ 3.433^{\pm.003} \\ 3.624^{\pm.037} \\ 3.060^{\pm.058} \end{array}$	$\begin{array}{c} 8.678^{\pm.096} \\ 9.180^{\pm.071} \\ 9.229^{\pm.067} \\ 9.871^{\pm.084} \end{array}$	$\begin{array}{c} 3.127^{\pm.024} \\ 2.723^{\pm.033} \\ 2.826^{\pm.017} \\ 2.294^{\pm.058} \end{array}$
MM (SingleScan) MM (BiScan, layer) MM (B iScan, block)	$\begin{array}{c} 0.736^{\pm.003} \\ 0.735^{\pm.004} \\ 0.792^{\pm.002} \end{array}$	$\begin{array}{c} 1.063^{\pm.010} \\ 0.789^{\pm.007} \\ 0.281^{\pm.009} \end{array}$	$\begin{array}{c} 3.443^{\pm.026} \\ 3.408^{\pm.034} \\ 3.060^{\pm.058} \end{array}$	$\begin{array}{c} 9.180^{\pm.067} \\ 9.374^{\pm.059} \\ 9.871^{\pm.084} \end{array}$	$2.676^{\pm.041} \\ 2.591^{\pm.046} \\ 2.294^{\pm.058}$
MM (Dim, 1) MM (Dim, 2) MM (Dim, 5) MM (Dim, 7) MM (Dim, 10)	$\begin{array}{c} 0.706^{\pm.003}\\ 0.792^{\pm.002}\\ 0.741^{\pm.008}\\ 0.738^{\pm.004}\\ 0.715^{\pm.003}\end{array}$	$\begin{array}{c} 0.652^{\pm.011} \\ 0.281^{\pm.009} \\ 0.728^{\pm.009} \\ 0.599^{\pm.007} \\ 0.628^{\pm.008} \end{array}$	$\begin{array}{c} 3.541^{\pm.072}\\ 3.060^{\pm.058}\\ 3.307^{\pm.027}\\ 3.359^{\pm.068}\\ 3.548^{\pm.043} \end{array}$	$\begin{array}{c} 9.141^{\pm.082}\\ 9.871^{\pm.084}\\ 9.427^{\pm.099}\\ 9.166^{\pm.075}\\ 9.200^{\pm.075}\end{array}$	$\begin{array}{c} 2.612^{\pm.055} \\ 2.294^{\pm.058} \\ 2.314^{\pm.062} \\ 2.488^{\pm.037} \\ 2.884^{\pm.096} \end{array}$
MM (9 layers) MM (11 layers) MM (27 layers) MM (37 layers)	$\begin{array}{c} 0.755^{\pm.002} \\ 0.792^{\pm.002} \\ 0.750^{\pm.003} \\ 0.754^{\pm.005} \end{array}$	$\begin{array}{c} 1.080^{\pm.012}\\ 0.281^{\pm.009}\\ 0.975^{\pm.008}\\ 0.809^{\pm.010}\end{array}$	$3.309^{\pm.057} \\ 3.060^{\pm.058} \\ 3.336^{\pm.096} \\ 3.338^{\pm.061}$	$\begin{array}{c} 9.721^{\pm.081} \\ 9.871^{\pm.084} \\ 9.249^{\pm.071} \\ 9.355^{\pm.062} \end{array}$	$2.974^{\pm.039} \\ 2.294^{\pm.058} \\ 2.821^{\pm.063} \\ 2.741^{\pm.077}$

tent dimension yields minimal improvement. Subsequently, we investigated both layer-based and block-based bidirectional scans. Our findings indicate that the block-based bidirectional scan achieves optimal performance. This suggests that spatial information flows are encoded within the latent spaces and that bidirectional scanning facilitates the exchange of this information, thereby enhancing the efficacy of motion generation tasks.

Architecture Design for Motion Mamba. The proposed Motion Mamba which is grounded in a standardized motion latent diffusion system. We delved into the interplay between dimensional aspects and the module's capacity (measured by the number of layers) to ascertain their impact on system performance. Experimental results demonstrate that the Motion Mamba achieves superior performance at a latent dimension of 2, diverging from prior works where the optimal dimension was identified as 1. We attribute this discrepancy to our HTM, which necessitates multiple scans correlating with the sequence length, thus implicating dimensionality as a pivotal factor. A marginal increase in dimensionality enabled us to attain peak performance, simultaneously enhancing efficiency compared to models with a dimensionality of 10. Furthermore, we conducted experiments to determine the optimal layer count for Motion Mamba, inspired by the design of its selective scanning mechanism. Notably, a single Mamba layer achieves a parameter reduction of approximately 75% compared to a conventional transformer encoder block. By increasing the number of layers, we aim to uncover the relationship between model capacity and its performance. Our findings reveal that, through the integration of our specially designed HTM and BSM (Bidirectional Scanning Module) blocks, the Motion Mamba reaches its optimal performance with 11 layers. This represents a slight increase over the MLD [6] baseline. How-

14 Zeyu and Akide et al.

ever, due to the reduced parameter count in each layer, Motion Mamba exhibits significantly greater efficiency than previous methodologies.

4.5 Inference Time

Inference time remains a significant challenge for diffusion-based methods. To address this, we enhance the inference speed by incorporating the efficient Mamba block within a lightweight architecture. Compared to the previous strong baseline, such as the MLD model cited in [6], which reports an average inference time of 0.217 seconds, our Motion Mamba model achieves a notable reduction in computational overhead, as shown in Figure 4. Specifically, it requires four times less computational effort, thereby facilitating faster and real-time inference speeds.



Fig. 4: The figure shows the average inference time per sentence (AIT) vs FID, our proposed motion mamba obtained 0.058s AIT and 0.281 FID overall outperform previous methods. We evaluate all methods on a single V100 GPU.

5 Discussion and Conclusion

In this study, we introduced Motion Mamba, a novel framework designed for efficient and extended sequence motion generation. Our approach represents the inaugural integration of the Mamba model within the domain of motion generation, featuring significant advancements including the implementation of Hierarchical Temporal Mamba (HTM) blocks. These blocks are specifically engineered to enhance temporal alignment through hierarchically organized selective scanning. Furthermore, Bidirectional Spatial Mamba (BSM) blocks have been developed to amplify the exchange of information flow within latent spaces, thereby augmenting the model's ability to bidirectionally capture skeleton-level density features with greater precision. Compared to previous diffusion-based motion generation methodologies that predominantly utilize transformer blocks, our Motion Mamba framework achieves SOTA performance, evidencing an improvement of up to 50% in FID scores and a quadrupled improvement in inference speed. Through comprehensive experimentation across a variety of human motion generation tasks, the effectiveness and efficiency of our proposed Motion Mamba model have been robustly demonstrated, marking a significant leap forward in the field of human motion generation.

Acknowledgements

This work is partially supported by the Fundamental Research Funds for the Central Universities, Peking University.

References

- Ahuja, C., Morency, L.P.: Language2pose: Natural language grounded pose forecasting. In: 2019 International Conference on 3D Vision (3DV). pp. 719–728. IEEE (2019)
- Bao, F., Li, C., Cao, Y., Zhu, J.: All are worth words: a vit backbone for score-based diffusion models. arXiv preprint arXiv:2209.12152 (2022)
- Baron, E., Zimerman, I., Wolf, L.: 2-d ssm: A general spatial layer for visual transformers. arXiv preprint arXiv:2306.06635 (2023)
- Barsoum, E., Kender, J., Liu, Z.: Hp-gan: Probabilistic 3d human motion prediction via gan. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 1418–1427 (2018)
- Bhattacharya, U., Rewkowski, N., Banerjee, A., Guhan, P., Bera, A., Manocha, D.: Text2gestures: A transformer-based network for generating emotive body gestures for virtual agents. In: 2021 IEEE virtual reality and 3D user interfaces (VR). pp. 1–10. IEEE (2021)
- Chen, X., Jiang, B., Liu, W., Huang, Z., Fu, B., Chen, T., Yu, G.: Executing your commands via motion diffusion in latent space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18000–18010 (2023)
- Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in Neural Information Processing Systems 34, 8780–8794 (2021)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020)
- Ghosh, A., Cheema, N., Oguz, C., Theobalt, C., Slusallek, P.: Synthesis of compositional animations from textual descriptions. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1396–1406 (2021)
- Gong, K., Lian, D., Chang, H., Guo, C., Jiang, Z., Zuo, X., Mi, M.B., Wang, X.: Tm2d: Bimodality driven 3d dance generation via music-text integration. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9942–9952 (2023)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems 27 (2014)
- 12. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752 (2023)
- Gu, A., Goel, K., Gupta, A., Ré, C.: On the parameterization and initialization of diagonal state space models. Advances in Neural Information Processing Systems 35, 35971–35983 (2022)
- Gu, A., Goel, K., Ré, C.: Efficiently modeling long sequences with structured state spaces. arXiv preprint arXiv:2111.00396 (2021)
- Gu, A., Johnson, I., Goel, K., Saab, K., Dao, T., Rudra, A., Ré, C.: Combining recurrent, convolutional, and continuous-time models with linear state space layers. Advances in neural information processing systems 34, 572–585 (2021)
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of wasserstein gans. Advances in neural information processing systems **30** (2017)
- 17. Guo, C., Zou, S., Zuo, X., Wang, S., Ji, W., Li, X., Cheng, L.: Generating diverse and natural 3d human motions from text. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5152–5161 (2022)

- 16 Zeyu and Akide et al.
- Guo, C., Zuo, X., Wang, S., Zou, S., Sun, Q., Deng, A., Gong, M., Cheng, L.: Action2motion: Conditioned generation of 3d human motions. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 2021–2029 (2020)
- Gupta, A., Gu, A., Berant, J.: Diagonal state spaces are as effective as structured state spaces. Advances in Neural Information Processing Systems 35, 22982–22994 (2022)
- Harvey, F.G., Yurick, M., Nowrouzezahrai, D., Pal, C.: Robust motion inbetweening. ACM Transactions on Graphics (TOG) 39(4), 60–1 (2020)
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems 33, 6840–6851 (2020)
- Hopfield, J.J.: Neural networks and physical systems with emergent collective computational abilities. Proceedings of the national academy of sciences 79(8), 2554– 2558 (1982)
- Jiang, B., Chen, X., Liu, W., Yu, J., Yu, G., Chen, T.: Motiongpt: Human motion as a foreign language. Advances in Neural Information Processing Systems 36 (2024)
- Kalman, R.E.: A New Approach to Linear Filtering and Prediction Problems. Journal of Basic Engineering 82(1), 35–45 (1960)
- 26. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. stat 1050, 1 (2014)
- Kramer, M.A.: Nonlinear principal component analysis using autoassociative neural networks. AIChE journal 37(2), 233–243 (1991)
- Li, S., Singh, H., Grover, A.: Mamba-nd: Selective state space modeling for multidimensional data. arXiv preprint arXiv:2402.05892 (2024)
- Li, Y., Cai, T., Zhang, Y., Chen, D., Dey, D.: What makes convolutional models great on long sequence modeling? In: The Eleventh International Conference on Learning Representations (2022)
- Lin, X., Amer, M.R.: Human motion modeling using dvgans. arXiv preprint arXiv:1804.10652 (2018)
- Liu, Y., Tian, Y., Zhao, Y., Yu, H., Xie, L., Wang, Y., Ye, Q., Liu, Y.: Vmamba: Visual state space model. arXiv preprint arXiv:2401.10166 (2024)
- Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11976–11986 (2022)
- Mahmood, N., Ghorbani, N., Troje, N.F., Pons-Moll, G., Black, M.J.: Amass: Archive of motion capture as surface shapes. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5442–5451 (2019)
- Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2019)
- Petrovich, M., Black, M.J., Varol, G.: Action-conditioned 3D human motion synthesis with transformer VAE. In: International Conference on Computer Vision (ICCV) (2021)
- Petrovich, M., Black, M.J., Varol, G.: Temos: Generating diverse human motions from textual descriptions. In: European Conference on Computer Vision. pp. 480– 497. Springer (2022)

- Petrovich, M., Black, M.J., Varol, G.: TEMOS: Generating diverse human motions from textual descriptions. In: European Conference on Computer Vision (ECCV) (2022)
- Plappert, M., Mandery, C., Asfour, T.: The kit motion-language dataset. Big data 4(4), 236–252 (2016)
- Plappert, M., Mandery, C., Asfour, T.: 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)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
- 43. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning internal representations by error propagation. Parallel Distributed Processing pp. 318–362 (1986)
- Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image superresolution via iterative refinement. IEEE Transactions on Pattern Analysis and Machine Intelligence (2022)
- 45. Shi, X., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.K., Woo, W.c.: Convolutional lstm network: A machine learning approach for precipitation nowcasting. Advances in neural information processing systems 28 (2015)
- Smith, J., De Mello, S., Kautz, J., Linderman, S., Byeon, W.: Convolutional state space models for long-range spatiotemporal modeling. Advances in Neural Information Processing Systems 36 (2024)
- 47. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: International conference on machine learning. pp. 2256–2265. PMLR (2015)
- Tevet, G., Gordon, B., Hertz, A., Bermano, A.H., Cohen-Or, D.: Motionclip: Exposing human motion generation to clip space. In: European Conference on Computer Vision. pp. 358–374. Springer (2022)
- Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-or, D., Bermano, A.H.: Human motion diffusion model. In: The Eleventh International Conference on Learning Representations (2022)
- Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. Advances in neural information processing systems 30 (2017)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
- 52. Wang, J., Yan, J.N., Gu, A., Rush, A.M.: Pretraining without attention. arXiv preprint arXiv:2212.10544 (2022)
- Yan, J.N., Gu, J., Rush, A.M.: Diffusion models without attention. arXiv preprint arXiv:2311.18257 (2023)
- Zhang, M., Cai, Z., Pan, L., Hong, F., Guo, X., Yang, L., Liu, Z.: Motiondiffuse: Text-driven human motion generation with diffusion model. IEEE Transactions on Pattern Analysis and Machine Intelligence (2024)

- 18 Zeyu and Akide et al.
- Zhang, Y., Black, M.J., Tang, S.: We are more than our joints: Predicting how 3d bodies move. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3372–3382 (2021)
- Zhong, C., Hu, L., Zhang, Z., Xia, S.: Attt2m: Text-driven human motion generation with multi-perspective attention mechanism. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 509–519 (2023)
- 57. Zhu, L., Liao, B., Zhang, Q., Wang, X., Liu, W., Wang, X.: Vision mamba: Efficient visual representation learning with bidirectional state space model. arXiv preprint arXiv:2401.09417 (2024)