

T2M Mamba: Motion Periodicity-Saliency Coupling Approach for Stable Text-Driven Motion Generation

Xingzu Zhan, Chen Xie, Honghang Chen, Yixun Lin, and Xiaochun Mai

Abstract—Text-to-motion generation, which converts motion language descriptions into coherent 3D human motion sequences, has attracted increasing attention in fields, such as avatar animation and humanoid robotic interaction. Though existing models have achieved significant fidelity, they still suffer from two core limitations: (i) They treat motion periodicity and keyframe saliency as independent factors, overlooking their coupling and causing generation drift in long sequences. (ii) They are fragile to semantically equivalent paraphrases, where minor synonym substitutions distort textual embeddings, propagating through the decoder and producing unstable or erroneous motions. In this work, we propose T2M Mamba to address these limitations by (i) proposing Periodicity-Saliency Aware Mamba, which utilizes novel algorithms for keyframe weight estimation via enhanced Density Peaks Clustering and motion periodicity estimation via FFT-accelerated autocorrelation to capture coupled dynamics with minimal computational overhead, and (ii) constructing a Periodic Differential Cross-modal Alignment Module (PDCAM) to enhance robust alignment of textual and motion embeddings. Extensive experiments on HumanML3D and KIT-ML datasets have been conducted, confirming the effectiveness of our approach, achieving an FID of 0.068 and consistent gains on all other metrics.

I. INTRODUCTION

Text-to-motion generation is a fundamental task in various fields including avatar animation, humanoid robots, and gaming, which transforms a natural-language description about motion into a semantically coherent and physically plausible high-dimensional temporal sequence of human poses.

Text-to-motion generation has evolved from early latent-regression networks, through high-fidelity yet costly diffusion transformers, to linear-time Mamba state-space models and finally tokenised motion–language frameworks that exploit large-scale language modelling for long horizon synthesis [1]–[7]. Building on these technological advances, researchers have begun to explore ways of enhancing the quality of generated motion. MMDM [8] and MoMask [9] boost quality by reconstructing randomly masked frames, whereas KMM [10] and KeyMotion [11] mask only algorithmically detected keyframes to focus the model on motion pivots. Researchers have also begun to exploit motion periodicity. For example, DiffusionPhase [12] and FlowMDM [5] introduce

Xingzu Zhan, Chen Xie, Honghang Chen are with Shenzhen University, Shenzhen, China. Xiaochun Mai is with the College of Electronics and Information Engineering, Shenzhen University, Shenzhen, China, with the Shenzhen Key Laboratory of Digital Creative Technology, with the Guangdong Province Engineering Laboratory for Digital Creative Technology, and also with the Guangdong-Hong Kong Joint Laboratory for Big Data Imaging and Communication, Shenzhen, China.

Yixun Lin is with Jinan University, Guangzhou, China.

Corresponding author: Xiaochun Mai (e-mail: maixc@szu.edu.cn).

frequency or phase domain representations to help motion generation. We can see that existing work often overlooks the coupling of keyframes and motion periodicity. However, in fact that keyframe learning and cycle control are not isolated. A single motion sequence often contains multiple motion patterns, which may exhibit periodicity, e.g., running or walking, with keyframes potentially serving as boundaries between these patterns. Besides, recent work SOTA [13] indicates that current models are highly sensitive to semantically equivalent paraphrases, even minor synonym substitutions can trigger large drifts in textual embeddings. Subsequently, the drifts will be progressively amplified during decoding, resulting in divergent or erroneous motion sequences. The main reason for the low stability of textual and motion mapping is that the cross-attention suffers from the time scale mismatch, that is, one word corresponds to multiple motion frames, or multiple words describe an instantaneous action. Hence, motivated by the issues analyzed above, this paper aims to i) *explore the coupling relationship between keyframes and motion periodicity*, then exploit the coupling to *eliminate the mamba's history forgetting* in the long sequence motion generation, and ii) *tackle the time scale mismatch issue* for enhancing the consistency and stability of the text-to-motion mapping.

To this end, we introduce *T2M Mamba*, a novel methodology for human motion generation that comprises two principal modules: *Periodicity-Saliency Aware Mamba* and the *Periodic Differential Cross-modal Alignment Module (PDCAM)*. To address the first issue, T2M Mamba detects physically meaningful turning frames, measures their saliency to obtain keyframe weights. Then T2M Mamba applies an FFT-accelerated autocorrelation function between successive keyframes to estimate the motion period; Periodicity-Saliency Aware Mamba subsequently modifies Mamba's input projection matrix with these keyframe weights, alleviates Mamba's historical forgetting, and injects the estimated period to capture the latent motion rhythm. To confront the second issue, we propose the Periodic Differential Cross-modal Alignment Module (PDCAM). PDCAM integrates keyframes to pinpoint pivotal motion segments. It leverages motion periodicity to exploit repetitive patterns, enabling consistent alignment. Furthermore, it incorporates differential essence to emphasize semantic differences while mitigating noise. As a result, the module reliably extracts key information and semantics. This yields robust cross-modal alignment and enhanced stability across varying time scales. Overall, our contributions can be summarised as follows.

- We propose novel algorithms for keyframe weight estimation and motion periodicity estimation, utilizing enhanced density peaks clustering (DPC) for adaptive keyframe detection within motion segments and FFT-accelerated autocorrelation for identifying dominant periods in motion signals, all while incurring almost no additional computational overhead.
- We introduce Periodicity-Saliency Aware Mamba, which combines keyframe weighting and motion periodicity modeling to alleviate historical forgetting in long sequences, and learn coordinated motion patterns. Furthermore, we propose the Periodic Differential Cross-modal Alignment Module (PDCAM) to bolster robust cross-modal mapping against textual perturbations by integrating keyframe importance and motion phase encoding.
- We conduct extensive experiments on the HumanML3D and KIT-ML datasets to validate the effectiveness of T2M Mamba. We achieve substantial improvements across all metrics, including a notable FID score of 0.068, demonstrating the effectiveness and efficiency of our proposed approach.

II. RELATED WORK

A. Text-to-Motion Generation

Text-to-motion generation has evolved through VAE-based, diffusion-based, and hybrid methods. VAE-based approaches, such as Action2Motion [14] and TM2T [15], model motion via latent variables to enable probabilistic sampling, promoting diversity and interpretability in continuous or discrete spaces. However, they often produce unrealistic results and suffer from mode collapse due to poor generalization. Diffusion-based methods, like MotionDiffuse [3] and MDM [4], iteratively denoise noisy motion sequences conditioned on text, excelling as motion learners for multi-modal conditioning and high-fidelity outputs, but they are limited in long sequences due to generation drift and high computational demands. Hybrid models, such as MLD [16] and AttT2M [17], combine VAEs for latent compression with diffusion for generation, struggling with fine-grained details and high computational complexity. Based on diffusion framework, our T2M Mamba uniquely couples motion periodicity and keyframe saliency, addressing generation drift in long sequences. Furthermore, it introduces the Periodic Differential Cross-modal Alignment Module (PDCAM) to robustly align textual and motion embeddings, mitigating instability from synonym substitutions.

B. State Space Models

State Space Models (SSMs) describe dynamic systems through intermediate state variables. S4 [18] introduce structured low-rank corrections for stable diagonalization, enabling efficient long-sequence modeling with linear complexity. Based on S4, Mamba [19] incorporates a selective mechanism and hardware-aware algorithms, parameterizing SSMs based on input sequences to achieve content-aware processing while maintaining efficiency. In motion generation, Motion

Mamba [6], TM-Mamba [20] and Infinimotion [21] offer linear-time complexity for long-sequence modeling via selective state updates, but they suffer from historical forgetting in extended sequences, where early inputs lose influence, and overlook coupled dynamics like periodicity and saliency in motion data. By comparison, our Periodicity-Saliency Aware Mamba explicitly considers keyframe importance and injects estimated periodicity into Mamba’s sequences, alleviating historical forgetting more effectively than standard SSMs. By integrating these coupled factors with minimal overhead, it captures latent motion rhythms and coordinated patterns absent in prior SSMs.

III. METHODOLOGY

A. Overview of T2M Mamba

To address the challenges of overlooked coupling between keyframe saliency and motion periodicity, as well as instability in text-to-motion mapping under semantic paraphrases, we propose T2M Mamba, as shown in Fig 1, a framework that integrates Periodicity-Saliency Aware Mamba and Periodic Differential Cross-modal Alignment Module (PDCAM). Given a noisy motion sequence after T diffusion steps, $\{M_T^1, M_T^2, M_T^3, \dots, M_T^N\}$, T2M Mamba conditions on the CLIP-encoded text prompt and the current diffusion timestep T to predict the corresponding clean motion sequence $\{M_0^1, M_0^2, M_0^3, \dots, M_0^N\}$.

B. Motion Periodicity-Saliency Coupling

In text-to-motion generation, capturing keyframe saliency and inherent periodicity is crucial for maintaining physical plausibility and rhythmic coherence. We first introduce methods for keyframe weight estimation and motion periodicity estimation. These estimations serve as foundational cues for the Periodicity-Saliency Aware Mamba and Periodic Differential Cross-modal Alignment Module, as detailed in the following subsections.

a) keyframe weight estimation: To estimate keyframe weights, we partition the motion stream into N equal-length segments and apply an enhanced Density Peaks Clustering (DPC) within each segment to automatically detect keyframes based on local density and separation distance, followed by assigning normalized weights to emphasize their importance. Specifically, we first partition the motion stream into N equal-length temporal segments $\{X_1, \dots, X_N\}$. Within each segment, an improved *Density Peaks Clustering* (DPC) [22] is performed. Given a segment $\{\mathbf{x}_i\}_{i=1}^M$, pairwise similarity is measured by the Euclidean distance $d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|$, the subscript j denotes the index of a sample distinct from i within the current segment, paired with i to compute the Euclidean distance d_{ij} . And the cut-off distance d_c is selected adaptively so that approximately 1–2% of the points lie within this neighbourhood. The local density is defined by

$$\rho_i = \sum_{j \neq i} e^{-\left(\frac{d_{ij}}{d_c}\right)^2}. \quad (1)$$

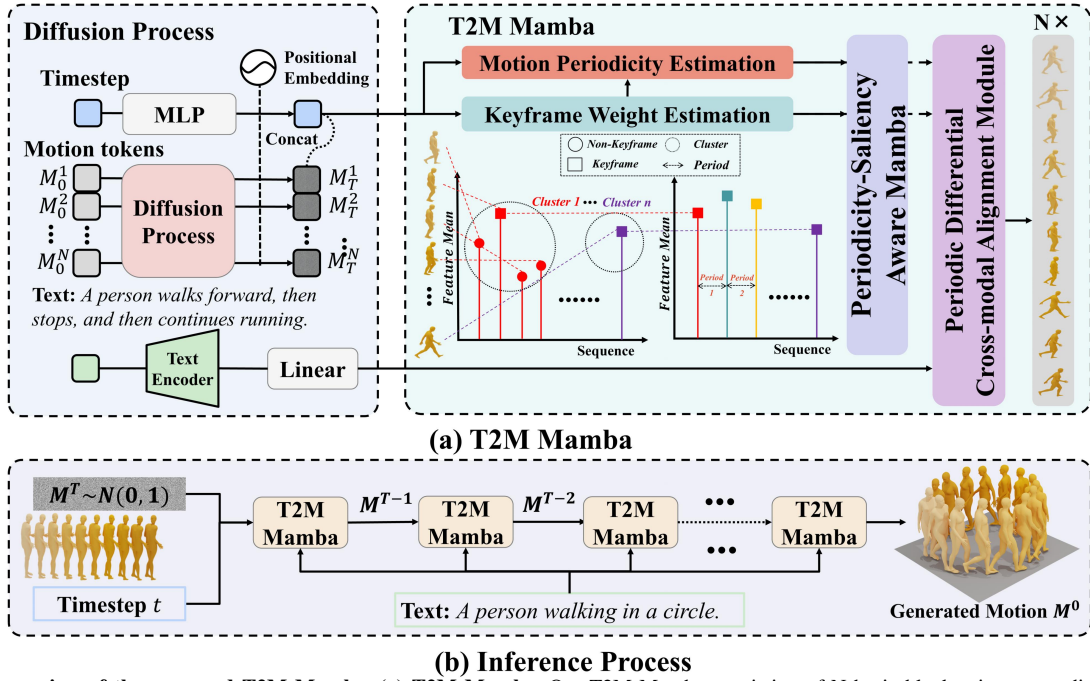


Fig. 1. **The overview of the proposed T2M Mamba.** (a) **T2M Mamba.** Our T2M Mamba consisting of N basic blocks aims to predict clean motion sequence (b) **Inference Process.** Starting from Gaussian noise, the model iteratively denoises to generate a clean motion sequence M^0 semantically aligned with the input text prompt.

After computing the local density ρ_i for each sample, relying solely on density magnitude is insufficient, as a high-density region may contain multiple peaks that could be misidentified as centers without distinction. To address this, Density Peaks Clustering introduces the minimum separation distance δ_i , defined as the smallest Euclidean distance to any point with higher density (or the maximum distance if none exists), i.e.,

$$\delta_i = \begin{cases} \min_{j: \rho_j > \rho_i} d_{ij}, & \text{if } \exists j: \rho_j > \rho_i, \\ \max_j d_{ij}, & \text{otherwise.} \end{cases} \quad (2)$$

The composite peak score is $\gamma_i = \rho_i \delta_i$, ensuring candidate centers are dense (high ρ_i) and well-separated (large δ_i). After computing all γ_i , we sort them descendingly and use the elbow point in the γ curve to infer the number of keyframes automatically. We then construct the keyframe weight matrix $\mathbf{M} \in \mathbb{R}^{L \times 1}$ by concatenating results from the N segments in temporal order, where for each frame i ,

$$\mathbf{M}_i = \begin{cases} \tilde{\gamma}_i, & i \in \mathcal{K} \\ 1, & \text{otherwise,} \end{cases} \quad \tilde{\gamma}_i = \frac{\gamma_i - \gamma_{\min}}{\gamma_{\max} - \gamma_{\min}}, \quad (3)$$

with \mathcal{K} denoting the set of detected keyframes, and γ_{\min} , γ_{\max} the minimum and maximum γ values among detected keyframes in the sequence.

b) Motion Periodicity Estimation: To estimate motion periodicity, we segment the motion sequence into intervals between consecutive keyframes and independently analyze each segment's one-dimensional motion signal using FFT-accelerated autocorrelation to detect dominant periods, determining periodicity based on peak, prominence, and spectral entropy criteria. For each segment s with length L_s , we

extract the signal $\mathbf{x}_s \in \mathbb{R}^{L_s}$ by averaging across feature dimensions and compute its normalized autocorrelation function (ACF) via the Wiener-Khinchin theorem rather than directly evaluating the ACF for efficiency.

$$\mathbf{x}_s = \frac{1}{D} \sum_{d=1}^D \mathbf{F}_d, \quad R(\tau) = \frac{1}{R(0)} \cdot \mathcal{F}^{-1}\{|\mathcal{F}\{\mathbf{x}_s\}|^2\}, \quad (4)$$

where \mathbf{F}_d is the d -th feature dimension, \mathcal{F} and \mathcal{F}^{-1} denote the Fast Fourier Transform and its inverse, τ is the time lag measuring signal shift, and $R(0)$ (the signal's variance) normalizes $R(\tau)$ to $[-1, 1]$ for similarity assessment at various lags.

We identify the dominant period by locating the first significant peak in the autocorrelation function $R(\tau)$ for time lag $\tau > 0$. We employ the peak ratio criterion [23], prominence criterion [24], and spectral entropy criterion [25] to assess periodicity in motion segments. Stringent thresholds across these criteria effectively filter out spurious patterns, preventing erroneous assignment of periodicity to inherently non-periodic motions. A segment is classified as periodic if it satisfies three criteria,

$$\begin{cases} R(\tau_{\max}) > \theta_{\text{peak}} & \text{(peak ratio)} \\ R(\tau_{\max}) - \bar{R} > \theta_{\text{prom}} & \text{(prominence)} \\ H_{\text{spec}} < \theta_{\text{ent}} & \text{(spectral entropy),} \end{cases} \quad (5)$$

where $\tau_{\max} = \arg \max_{\tau > 0} R(\tau)$ denotes the lag position of the maximum peak, $R(\tau_{\max})$ is the autocorrelation value at the peak position, \bar{R} represents the mean value of the autocorrelation function, H_{spec} is the spectral entropy which quantifies the regularity/randomness of the power spectrum. θ_{peak} , θ_{prom} , and θ_{ent} are thresholds for peak ratio, prominence,

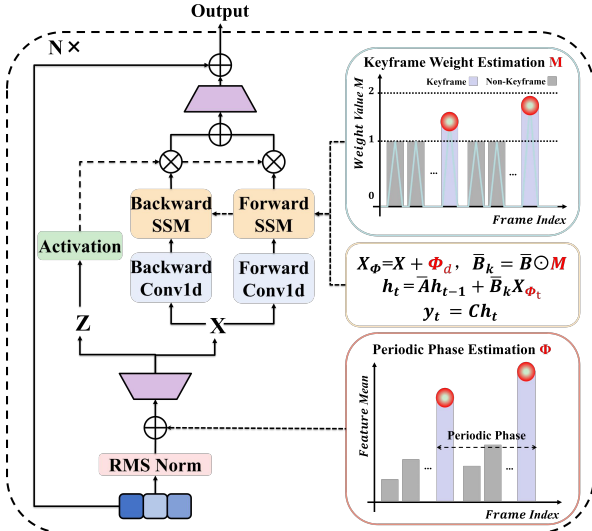


Fig. 2. Illustration of our Periodicity-Saliency Aware Mamba. \odot denotes dot product.

and spectral entropy, respectively. The spectral entropy is calculated as

$$H_{\text{spec}} = -\frac{1}{\log N} \sum_{k=0}^{N-1} p_k \log p_k, \quad p_k = \frac{|X_k|^2}{\sum_{j=0}^{N-1} |X_j|^2}, \quad (6)$$

where X_k denotes the k -th frequency component of the Fourier transform, p_k is the normalized power spectrum representing the relative power at frequency k , and N is the FFT size. The normalization factor $1/\log N$ ensures $H_{\text{spec}} \in [0, 1]$.

For periodic segments, the detected period is $T = \tau_{\text{max}}$. Otherwise, the period T is set as the segment length, $T = L_s$. Finally, we compute the instantaneous phase for each frame, i.e.,

$$\phi(t) = \frac{2\pi t}{T}, \quad \Phi = [\sin \phi(t), \cos \phi(t)]. \quad (7)$$

By concatenating the obtained $\phi(t)$ and Φ along the time order, we form the radian vector $\phi \in \mathbb{R}^{L \times 1}$, utilized in the Periodic Differential Cross-modal Alignment Module (subsection III-D), and the phase encoding matrix $\Phi \in \mathbb{R}^{L \times 2}$, employed in the Periodicity-Saliency Aware Mamba module (subsection III-C).

C. Periodicity-Saliency Aware Mamba

Due to the state-space recurrence in Mamba [19], which continuously compresses historical information into the hidden state $h(t)$, physically critical frames become severely attenuated, and the model fails to capture latent periodicity in cyclic motions such as walking, running, and dancing, resulting in beat drift and pose jitter. To mitigate these limitations, we propose Periodicity-Saliency Aware Mamba, as shown in Fig 2 which leverages the detected keyframe weights and periodicity estimates from the previous section. Specifically, it constructs a saliency matrix from the keyframe weights to replace the input projection \bar{B} of Mamba, explicitly amplifying the influence of important frames; meanwhile, it injects sine-cosine phase encodings derived from the

estimated periods at every timestep. This dual strategy simultaneously reinforces keyframe supervision and enforces global rhythmic coherence, thereby addressing deficiencies caused by information compression. We detail the Periodicity-Saliency Aware Mamba formulation below.

Given a motion sequence $X \in \mathbb{R}^{L \times D}$, a keyframe weight matrix $M \in \mathbb{R}^{L \times 1}$, and a phase encoding matrix $\Phi \in \mathbb{R}^{L \times 2}$, we first linearly project the phase encoding matrix to $\Phi_d \in \mathbb{R}^{L \times D}$ via $\Phi_d = \Phi W_\phi$, where $W_\phi \in \mathbb{R}^{2 \times D}$ is a learnable projection matrix. The input sequence X is then added to Φ_d to obtain an explicitly rhythm-enhanced input X_Φ . To enable Mamba to *explicitly increase attention on keyframes*, we multiply the obtained keyframe weight matrix F with the input projection matrix $\bar{B} \in \mathbb{R}^{N \times 1}$ in Mamba's state space equation, yielding \bar{B}_k . This allows more keyframe information to be stored in the historical state h_t . The overall formulation of Periodicity-Saliency Aware Mamba is thus given by

$$\begin{cases} X_\Phi = X + \Phi_d, & \bar{B}_k = \bar{B} \odot M, \\ h_t = \bar{A}h_{t-1} + \bar{B}_k X_{\Phi_t}, \\ y_t = Ch_t. \end{cases} \quad (8)$$

Where \odot denotes element-wise multiplication. Note that the keyframe detection and periodicity detection modules are computationally efficient in the training, our Periodicity-Saliency Aware Mamba achieves explicit attention to keyframes and learning of periodic rhythms with almost no additional computational time overhead.

Periodicity-Saliency Aware Mamba maintains computational complexity comparable to standard Mamba and preserves its linear time complexity while addressing limitations in long-sequence motion generation. It integrates explicit keyframe weighting and motion phase encoding to mitigate historical forgetting from information compression and to capture rhythmic patterns in periodic actions.

D. Periodic Differential Cross-modal Alignment Module

In text-to-motion generation, cross-attention often suffers from time-scale mismatch, where one word maps to multiple motion frames or multiple words describe an instantaneous action, leading to unstable cross-modal alignment. To address this issue, we improve the differential attention from the Differential Transformer [26], by using keyframes, motion periodicity, and differential essence to extract key information and semantics, thereby achieving robust alignment and enhancing cross-modal stability.

Unlike the Differential Transformer, which builds two softmax attention maps on the same query-key pair, PDCAM constructs two linear-attention maps from phase-rotated query halves and value, then subtracts them. PDCAM yielding a lightweight differential operator that accentuates discriminative cross-modal cues, acting as a differential amplifier. Additionally, PDCAM uses keyframe weight matrix M from Eq. 3 and radian matrix ϕ from Eq. 7 to enhance the generating effects of motion sequences. Specifically, given a motion sequence $X \in \mathbb{R}^{L_m \times D}$ and a textual sequence

Methods	R Precision \uparrow			FID \downarrow	MM Dist \downarrow	Diversity \rightarrow	MModality \uparrow
	Top 1	Top 2	Top 3				
<i>On the HumanML3D</i>	$0.511 \pm .003$	$0.703 \pm .003$	$0.797 \pm .002$	$0.002 \pm .000$	$2.974 \pm .008$	$9.503 \pm .065$	-
T2M [2]	$0.457 \pm .002$	$0.639 \pm .003$	$0.740 \pm .003$	$1.067 \pm .002$	$3.340 \pm .008$	$9.188 \pm .002$	$2.090 \pm .083$
MDM [4]	$0.320 \pm .005$	$0.498 \pm .004$	$0.611 \pm .007$	$0.544 \pm .044$	$5.566 \pm .027$	$9.559 \pm .086$	2.799 $\pm .072$
MotionDiffuse [3]	$0.491 \pm .001$	$0.681 \pm .001$	$0.782 \pm .001$	$0.630 \pm .001$	$3.113 \pm .001$	$9.410 \pm .049$	$1.553 \pm .042$
MLD [16]	$0.481 \pm .003$	$0.673 \pm .003$	$0.772 \pm .002$	$0.473 \pm .013$	$3.196 \pm .010$	$9.724 \pm .082$	$2.413 \pm .079$
Motion Mamba [6]	$0.502 \pm .003$	$0.693 \pm .002$	$0.792 \pm .002$	$0.281 \pm .009$	$3.060 \pm .058$	$9.871 \pm .084$	$2.294 \pm .058$
T2M-GPT [7]	$0.492 \pm .003$	$0.679 \pm .002$	$0.775 \pm .002$	$0.141 \pm .005$	$3.121 \pm .009$	$9.722 \pm .082$	$1.831 \pm .048$
AttT2M [17]	$0.499 \pm .003$	$0.690 \pm .002$	$0.786 \pm .002$	$0.112 \pm .006$	$3.038 \pm .007$	$9.700 \pm .090$	$2.452 \pm .051$
MoMask [9]	0.521 $\pm .002$	0.713 $\pm .002$	0.807 $\pm .002$	0.045 $\pm .002$	2.958 $\pm .008$	-	$1.241 \pm .040$
T2M Mamba (Ours)	<u>0.506</u> $\pm .002$	<u>0.696</u> $\pm .002$	<u>0.793</u> $\pm .002$	<u>0.068</u> $\pm .004$	<u>3.034</u> $\pm .007$	9.497 $\pm .010$	$2.310 \pm .065$
<i>On the KIT-ML</i>	$0.424 \pm .005$	$0.649 \pm .006$	$0.779 \pm .006$	$0.031 \pm .004$	$2.788 \pm .012$	$11.08 \pm .097$	-
T2M [2]	$0.370 \pm .005$	$0.569 \pm .007$	$0.693 \pm .007$	$2.770 \pm .109$	$3.401 \pm .008$	$10.91 \pm .119$	$1.482 \pm .065$
MDM [4]	$0.164 \pm .004$	$0.291 \pm .004$	$0.396 \pm .004$	$0.497 \pm .021$	$9.191 \pm .022$	$10.85 \pm .109$	$1.907 \pm .214$
MotionDiffuse [3]	$0.417 \pm .004$	$0.621 \pm .004$	$0.739 \pm .004$	$1.954 \pm .062$	$2.958 \pm .005$	11.10 $\pm .143$	$0.730 \pm .013$
MLD [16]	$0.390 \pm .008$	$0.609 \pm .008$	$0.734 \pm .007$	$0.404 \pm .027$	$3.204 \pm .027$	$10.80 \pm .117$	<u>2.192</u> $\pm .071$
Motion Mamba [6]	$0.419 \pm .006$	$0.645 \pm .005$	$0.765 \pm .006$	$0.307 \pm .041$	$3.021 \pm .025$	<u>11.02</u> $\pm .098$	$1.678 \pm .064$
T2M-GPT [7]	$0.416 \pm .006$	$0.627 \pm .006$	$0.745 \pm .006$	$0.514 \pm .029$	$3.007 \pm .023$	$10.92 \pm .108$	$1.570 \pm .039$
AttT2M [17]	$0.413 \pm .006$	$0.632 \pm .006$	$0.751 \pm .006$	$0.870 \pm .039$	$3.039 \pm .021$	$10.96 \pm .123$	2.281 $\pm .047$
MoMask [9]	<u>0.433</u> $\pm .007$	<u>0.656</u> $\pm .005$	<u>0.781</u> $\pm .005$	<u>0.204</u> $\pm .011$	<u>2.779</u> $\pm .022$	-	$1.131 \pm .043$
T2M Mamba (Ours)	0.437 $\pm .006$	0.659 $\pm .006$	0.784 $\pm .006$	0.195 $\pm .021$	2.756 $\pm .018$	$10.92 \pm .096$	$1.512 \pm .088$

TABLE I

COMPARISON OF RESULTS ON THE HUMANML3D AND KIT-ML TEST SET. THE RIGHT ARROW (\rightarrow) INDICATES THAT VALUES CLOSER TO REAL MOTIONS CORRESPOND TO BETTER PERFORMANCE. EACH EVALUATION WAS RUN 20 TIMES, AND “ \pm ” DENOTES THE 95% CONFIDENCE INTERVAL. BOLD TEXT HIGHLIGHTS THE BEST SCORES, WHILE UNDERLINED ENTRIES MARK THE SECOND-BEST.

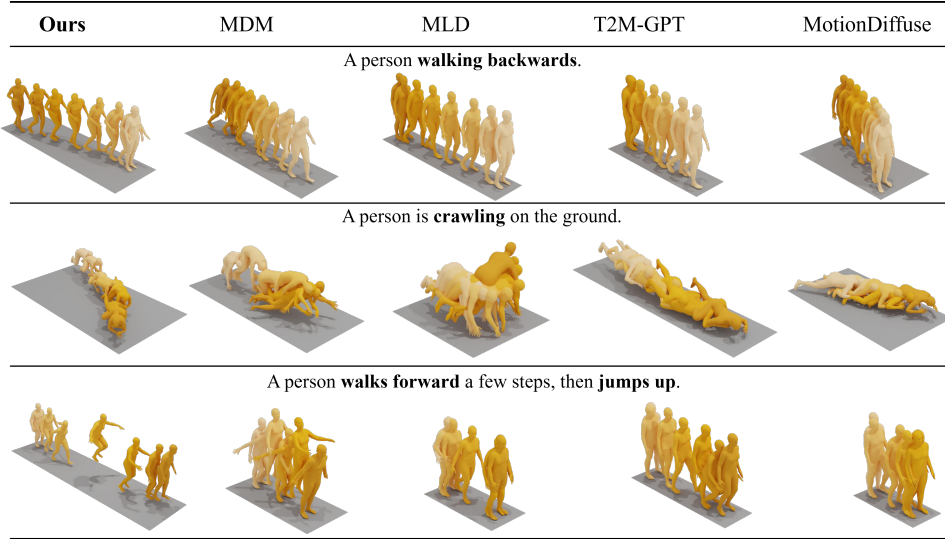


Fig. 4. Qualitative comparison of prominent state-of-the-art methods.

computational efficiency. Following the recommendations of MoMask [9], we adopt FID and R-Precision as the primary evaluation metrics in our ablation study, and additionally report the important metric Multi-Modal Distance. All experiments follow the same evaluation protocol as the main results, with metrics averaged over 20 runs.

a) Ablation on Key Modules: To evaluate the effectiveness of our module designs, we perform ablations on the two core components, that is, Periodicity-Saliency Aware Mamba, which integrates keyframe weights M and motion periodicity phase Φ to enhance sequence modeling, and Periodic Dif-

ferential Cross-modal Alignment Module (PDCAM), which reinjects M and Φ to achieve robust cross-modal alignment.

As shown in Table II, ablating key components in the PDCAM path leads to performance degradation. Removing M increases FID from 0.068 to 0.088 (a 29.4% rise), with MM Dist rising by 0.088. Removing ϕ elevates FID to 0.112 and slightly reduces R-Top1 to 0.497. These results indicate that keyframe learning and periodicity control are interdependent rather than isolated. Joint removal yields an FID of 0.108. Replacing PDCAM with standard cross-attention decreases R-Top3 from 0.793 to 0.755 and increases MM Dist from 3.034

Visual Comparison of T2M Mamba Under Subtle Semantic Variations

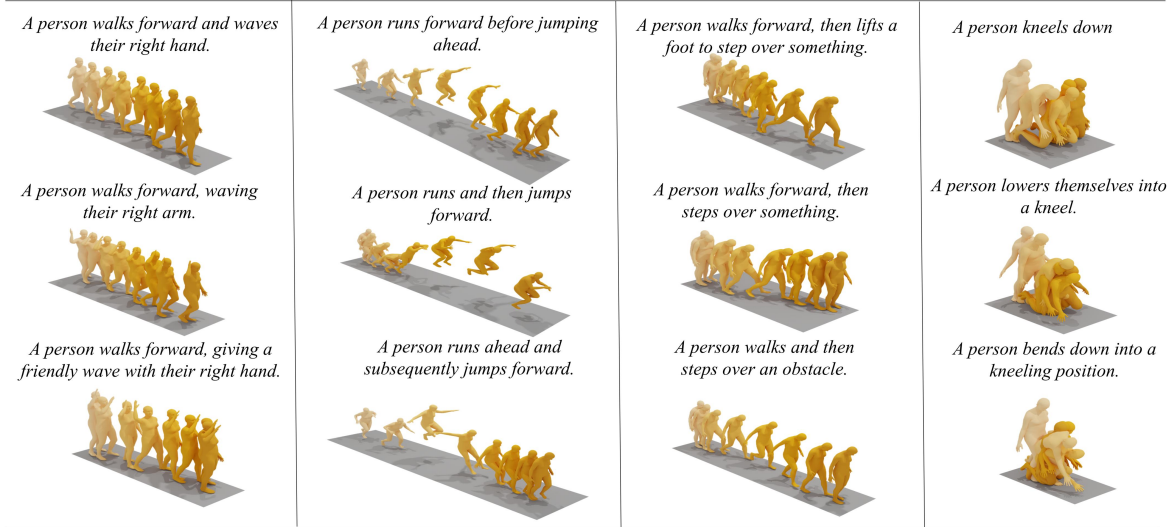


Fig. 5. Visual comparison of T2M Mamba under subtle semantic variations

	R-Top1 \uparrow	R-Top3 \uparrow	FID \downarrow	MM Dist \downarrow
w/o M	0.494	0.784	0.088	3.122
w/o ϕ	0.497	0.782	0.112	3.119
w/o M & ϕ	0.496	0.785	0.108	3.112
CroAtten	0.456	0.755	0.131	3.240
PDCAM(Ours)	0.506	0.793	0.068	3.034
w/o M	0.494	0.781	0.091	3.125
w/o ϕ	0.499	0.787	0.079	3.092
w/o M & ϕ	0.500	0.787	0.125	3.086
w/ M & ϕ (Ours)	0.506	0.793	0.068	3.034

TABLE II

ABLATION STUDY ON T2M MAMBA’S CORE MODULES ON HUMANML3D. UPPER: PDCAM VARIANTS; LOWER: PERIODICITY-SALIENCEY AWARE MAMBA ABLATIONS (W/O M: NO KEYFRAME WEIGHTING; W/O ϕ : NO MOTION PERIODICITY PHASE; W/O M & ϕ : NEITHER; CROATTEN: PDCAM REPLACED BY STANDARD CROSS-ATTENTION). BOLD VALUES HIGHLIGHT BEST PERFORMANCE BY FULL MODELS.

to 3.240, demonstrating that PDCAM substantially enhances cross-modal mapping.

In the Periodicity-Saliency Aware Mamba Module, lacking M raises FID to 0.091, underscoring keyframe weighting’s role in mitigating historical forgetting. Lacking ϕ increases FID to 0.079, highlighting the importance of periodicity phase for capturing latent motion rhythms. Removing both results in an FID of 0.125 (an 83.8% degradation). When both M and ϕ are retained, the two paths achieve the lowest FID of 0.068 and the highest R-Top3 of 0.793.

b) Ablation on the Layer Number of T2M Mamba:

To assess the effect of model depth, we vary the number of layers in the Mamba backbone from 4 to 10 while keeping other hyperparameters fixed. Results are presented in Table III. Performance improves with increasing depth up to 6 layers, achieving the best scores across most metrics. Beyond 6 layers, metrics degrade slightly, with FID rising to 0.091 at 7 layers and R-Top3 dropping to 0.783. This suggests that 6 layers strike an optimal balance between capacity and overfitting, as deeper models may introduce redundancy

Layer Number N	R-Top1 \uparrow	R-Top3 \uparrow	FID \downarrow	MM Dist \downarrow
4	0.498	0.791	0.089	3.074
5	0.497	0.789	0.095	3.081
6	0.506	0.793	0.068	3.034
7	0.497	0.783	0.091	3.101
8	0.497	0.779	0.088	3.126
10	0.494	0.782	0.137	3.113

TABLE III

THE EFFECT OF LAYER NUMBER ON T2M MAMBA PERFORMANCE ON HUMANML3D. BOLD VALUES HIGHLIGHT THE BEST PERFORMANCE ACHIEVED BY OUR FULL MODELS.

without proportional gains in sequence modeling.

c) Analysis of Keyframe Weight Estimation and Motion

Periodicity Estimation: To assess the impact of keyframe weight estimation and motion periodicity estimation, we ablate by removing all related components, converting Periodicity-Saliency Aware Mamba to standard bidirectional Mamba without keyframe weighting or periodicity injection, and PDCAM to plain differential attention without keyframe saliency and phase modulation. As shown in Table IV, this raises FID from 0.068 to 0.139 (104.4% increase) and drops R-Top3 from 0.793 to 0.775. These modules incur negligible overhead (AITS increases from 0.40 to 0.41, with unchanged model parameters), yet significantly enhance motion quality and stability via explicit keyframe attention and periodic rhythm learning. periodic rhythms.

d) Paraphrase Robustness Enabled by PDCAM:

As shown in Figure 5, we prompt an LLM to produce four pairs of action descriptions, each pair containing three sentences with minor semantic variations, and feed them to T2M Mamba to synthesize motions. The results show our model accurately follows the descriptions, and such small changes do not cause divergent or erroneous motions. This indicates that T2M Mamba overcomes the paraphrase sensitivity of prior methods, where even minor synonym substitutions can induce embedding drift and decoding errors.

	#Params	AITs ↓	R-Top3↑	FID↓
w/o ϕ	24.6M	0.41	0.782	0.098
w/o M	24.6M	0.41	0.785	0.134
w/o M & ϕ	24.6M	0.40	0.775	0.139
w/ M & ϕ (Ours)	24.6M	0.41	0.793	0.068

TABLE IV

ABLATION STUDY ON KEYFRAME WEIGHTING (M) AND MOTION PERIODICITY PHASE (ϕ) IN T2M MAMBA. AITS IS THE INFERENCE TIME AND BOLD VALUES HIGHLIGHT THE BEST PERFORMANCE.

D. Implementation Details

All experiments are conducted on a single NVIDIA RTX 4090 GPU. Training uses a linear beta schedule with 1000 diffusion steps, 140,000 iterations, AdamW optimizer [29] (LR 2×10^{-4} , weight decay 10^{-2} , gradient clip norm 1), and batch size 128; LR decays by 0.9 every 5,000 steps. Classifier-free guidance masks conditioning with 0.1 probability. Inference employs UniPC [30] with 10 timesteps. Average training time is 20.5 hours.

V. CONCLUSION

In this work, we introduced T2M Mamba, a novel framework that couples motion periodicity and keyframe saliency via Periodicity-Saliency Aware Mamba and PDCAM to enable stable text-driven motion generation. By mitigating historical forgetting, capturing latent rhythms, and ensuring robust cross-modal alignment, our approach achieves significant performance on HumanML3D and KIT-ML datasets, with an FID of 0.068 and consistent gains across metrics. Future efforts will explore extensions to real-time applications.

REFERENCES

- [1] X. Zhan, C. Xie, H. Chen, H. Sun, and X. Mai, "Multi-granular body modeling with redundancy-free spatiotemporal fusion for text-driven motion generation," *arXiv preprint arXiv:2503.06897*, 2025.
- [2] C. Guo, S. Zou, X. Zuo, S. Wang, W. Ji, X. Li, and L. Cheng, "Generating diverse and natural 3D human motions from text," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2022, pp. 5152–5161.
- [3] M. Zhang, Z. Cai, L. Pan, F. Hong, X. Guo, L. Yang, and Z. Liu, "Motiondiffuse: Text-driven human motion generation with diffusion model," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, no. 6, pp. 4115–4128, 2024.
- [4] G. Tevet, S. Raab, B. Gordon, Y. Shafir, D. Cohen-Or, and A. H. Bermanto, "Human motion diffusion model," *arXiv preprint arXiv:2209.14916*, 2022. [Online]. Available: <https://arxiv.org/abs/2209.14916>
- [5] G. Barquero, S. Escalera, and C. Palmero, "Seamless human motion composition with blended positional encodings," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2024, pp. 457–469.
- [6] Z. Zhang, A. Liu, I. Reid, R. Hartley, B. Zhuang, and H. Tang, "Motion mamba: Efficient and long sequence motion generation," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*. Springer, 2024, pp. 265–282.
- [7] J. Zhang, Y. Zhang, X. Cun, S. Huang, Y. Zhang, H. Zhao, H. Lu, and X. Shen, "T2m-GPT: Generating human motion from textual descriptions with discrete representations," *arXiv preprint arXiv:2301.06052*, 2023. [Online]. Available: <https://arxiv.org/abs/2301.06052>
- [8] X. Chen, "Text-driven human motion generation with motion masked diffusion model," *arXiv preprint arXiv:2409.19686*, 2024. [Online]. Available: <https://arxiv.org/abs/2409.19686>
- [9] C. Guo, Y. Mu, M. G. Javed, S. Wang, and L. Cheng, "Momask: Generative masked modeling of 3D human motions," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2024, pp. 1900–1910.

- [10] Z. Zhang, H. Gao, A. Liu, Q. Chen, F. Chen, Y. Wang, D. Li, R. Zhao, Z. Li, Z. Zhou, *et al.*, "Kmm: Key frame mask mamba for extended motion generation," *arXiv preprint arXiv:2411.06481*, 2024. [Online]. Available: <https://arxiv.org/abs/2411.06481>
- [11] Z. Geng, C. Han, Z. Hayder, J. Liu, M. Shah, and A. Mian, "Text-guided 3D human motion generation with keyframe-based parallel skip transformer," *arXiv preprint arXiv:2405.15439*, 2024. [Online]. Available: <https://arxiv.org/abs/2405.15439>
- [12] W. Wan, Y. Huang, S. Wu, T. Komura, W. Wang, D. Jayaraman, and L. Liu, "DiffusionPhase: Motion diffusion in frequency domain," *arXiv preprint arXiv:2312.04036*, 2023. [Online]. Available: <https://arxiv.org/abs/2312.04036>
- [13] W. Chen, H. Xiao, E. Zhang, L. Hu, L. Wang, M. Liu, and C. Chen, "Sato: Stable text-to-motion framework," in *Proc. ACM Int. Conf. Multimedia (MM)*, Oct. 2024, pp. 6989–6997. [Online]. Available: <http://dx.doi.org/10.1145/3664647.3681034>
- [14] C. Guo, X. Zuo, S. Wang, S. Zou, Q. Sun, A. Deng, M. Gong, and L. Cheng, "Action2motion: Conditioned generation of 3D human motions," in *Proc. ACM Int. Conf. Multimedia (MM)*, Oct. 2020, pp. 2021–2029. [Online]. Available: <http://dx.doi.org/10.1145/3394171.3413635>
- [15] C. Guo, X. Zuo, S. Wang, and L. Cheng, "Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3D human motions and texts," *arXiv preprint arXiv:2207.01696*, 2022. [Online]. Available: <https://arxiv.org/abs/2207.01696>
- [16] X. Chen, B. Jiang, W. Liu, Z. Huang, B. Fu, T. Chen, and G. Yu, "Executing your commands via motion diffusion in latent space," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2023, pp. 18 000–18 010.
- [17] C. Zhong, L. Hu, Z. Zhang, and S. Xia, "AttrT2M: Text-driven human motion generation with multi-perspective attention mechanism," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 2023, pp. 509–519.
- [18] A. Gu, K. Goel, and C. Ré, "Efficiently modeling long sequences with structured state spaces," *arXiv preprint arXiv:2111.00396*, 2022. [Online]. Available: <https://arxiv.org/abs/2111.00396>
- [19] A. Gu and T. Dao, "Mamba: Linear-time sequence modeling with selective state spaces," *arXiv preprint arXiv:2312.00752*, 2023. [Online]. Available: <https://arxiv.org/abs/2312.00752>
- [20] X. Wang, Z. Kang, and Y. Mu, "Text-controlled motion mamba: Text-instructed temporal grounding of human motion," *arXiv preprint arXiv:2404.11375*, 2024. [Online]. Available: <https://arxiv.org/abs/2404.11375>
- [21] Z. Zhang, A. Liu, Q. Chen, F. Chen, I. Reid, R. Hartley, B. Zhuang, and H. Tang, "Infinimotion: Mamba boosts memory in transformer for arbitrary long motion generation," *arXiv preprint arXiv:2407.10061*, 2024. [Online]. Available: <https://arxiv.org/abs/2407.10061>
- [22] H. Tang, L. Ding, S. Wu, B. Ren, N. Sebe, and P. Rota, "Deep unsupervised key frame extraction for efficient video classification," *arXiv preprint arXiv:2211.06742*, 2022. [Online]. Available: <https://arxiv.org/abs/2211.06742>
- [23] Q. Wen, L. Yang, and L. Sun, "Robust dominant periodicity detection for time series with missing data," *arXiv preprint arXiv:2303.03553*, 2023. [Online]. Available: <https://arxiv.org/abs/2303.03553>
- [24] P. Ravbar, N. Zhang, and J. H. Simpson, "Behavioral evidence for nested central pattern generator control of Drosophila grooming," *eLife*, vol. 10, p. e71508, 2021.
- [25] X. Huang, H. L. Shang, and D. Pitt, "Permutation entropy and its variants for measuring temporal dependence," *Aust. N. Z. J. Stat.*, vol. 64, pp. 442–477, 2022.
- [26] T. Ye, L. Dong, Y. Xia, Y. Sun, Y. Zhu, G. Huang, and F. Wei, "Differential transformer," *arXiv preprint arXiv:2410.05258*, 2025. [Online]. Available: <https://arxiv.org/abs/2410.05258>
- [27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.
- [28] M. Plappert, C. Mandery, and T. Asfour, "The KIT motion-language dataset," *Big Data*, vol. 4, no. 4, pp. 236–252, 2016.
- [29] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017. [Online]. Available: <https://arxiv.org/abs/1711.05101>
- [30] W. Zhao, L. Bai, Y. Rao, J. Zhou, and J. Lu, "UniPC: A unified predictor-corrector framework for fast sampling of diffusion models," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 36, 2023, pp. 49 842–49 869.