



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/236953/>

Version: Accepted Version

---

**Proceedings Paper:**

Dang, B., Wu, L., Yang, X. et al. (2026) SegMo: Segment-aligned text to 3D human motion generation. In: 2026 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). 2026 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), 06-10 Mar 2026, Tucson, Arizona, USA. Institute of Electrical and Electronics Engineers (IEEE). ISBN: 9798331555122. ISSN: 2472-6737. EISSN: 2642-9381.

<https://doi.org/10.1109/WACV61042.2026.00671>

---

© 2026 The Authors. Except as otherwise noted, this author-accepted version of a conference paper published in 2026 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) is made available via the University of Sheffield Research Publications and Copyright Policy under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>

**Reuse**

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

# SegMo: Segment-aligned Text to 3D Human Motion Generation

Bowen Dang<sup>1</sup> Lin Wu<sup>2</sup> Xiaohang Yang<sup>3</sup> Zheng Yuan<sup>1</sup> Zhixiang Chen<sup>1\*</sup>

<sup>1</sup>University of Sheffield <sup>2</sup>University of Glasgow <sup>3</sup>Queen Mary University of London

<sup>1</sup>{bdang2, zheng.yuan1, zhixiang.chen}@sheffield.ac.uk <sup>2</sup>l.wu.1@research.gla.ac.uk <sup>3</sup>xiaohang.yang@qmul.ac.uk

## Abstract

Generating 3D human motions from textual descriptions is an important research problem with broad applications in video games, virtual reality, and augmented reality. Recent methods align the textual description with human motion at the sequence level, neglecting the internal semantic structure of modalities. However, both motion descriptions and motion sequences can be naturally decomposed into smaller and semantically coherent segments, which can serve as atomic alignment units to achieve finer-grained correspondence. Motivated by this, we propose SegMo, a novel **Segment-aligned text-conditioned human Motion** generation framework to achieve fine-grained text–motion alignment. Our framework consists of three modules: (1) Text Segment Extraction, which decomposes complex textual descriptions into temporally ordered phrases, each representing a simple atomic action; (2) Motion Segment Extraction, which partitions complete motion sequences into corresponding motion segments; and (3) Fine-grained Text–Motion Alignment, which aligns text and motion segments with contrastive learning. Extensive experiments demonstrate that SegMo improves the strong baseline on two widely used datasets, achieving an improved TOP 1 score of 0.553 on the HumanML3D test set. Moreover, thanks to the learned shared embedding space for text and motion segments, SegMo can also be applied to retrieval-style tasks such as motion grounding and motion-to-text retrieval.

## 1. Introduction

Realistic human motion is important for applications in computer graphics, robotics, and immersive technologies such as video games, film production, virtual reality (VR), and augmented reality (AR). However, large-scale motion data mainly relies on costly motion capture systems with limited coverage, making it hard to represent the full range of human activities. To overcome these limitations, recent




Figure 1. The main idea of our method. We decompose the complex motion description and motion sequence into simpler temporally ordered segments and align them in a shared embedding space to improve the accuracy and realism of generated motions.

research has turned to motion generation [33], where text-conditioned methods are gaining attention due to the intuitive and flexible control offered by natural language [21].

To achieve text-conditioned motion generation, current methods typically rely on the CLIP [20] text encoder to extract features from textual descriptions and use these features as conditions for the generation model. These methods typically pair an entire motion sequence with a single textual description, performing alignment only at the *sequence level*. However, such alignment is inherently coarse because it *neglects the internal structure of both modalities*. Consequently, models may produce generation errors such as missing actions, repeated actions, or incorrect action orders, which limit the accuracy and fidelity of synthesized motions.

To address this problem, we introduce a smaller alignment unit that enables more fine-grained text–motion alignment, thereby enhancing the accuracy and realism of generated motions. In practice, both motion descriptions and motion sequences can be decomposed into smaller, temporally and semantically coherent segments. In motion descriptions, *a sentence is composed of phrases each describing an atomic motion*; in motion sequences, *a sequence can be perceived as distinct actions connected by smooth transitions*. Cognitive studies, such as Event Segmentation Theory [28], further support this view by showing that hu-

\*Corresponding author

Humans naturally segment continuous streams into meaningful events to facilitate perception and understanding. As illustrated in Figure 1, a novel segment-level alignment constraint is imposed on the motion generation model to enable fine-grained text–motion correspondence and improve generation quality.

To achieve segment-level alignment, we first need to get the text and motion segments. On the text side, we leverage large language models (LLMs) to decompose complex textual descriptions into temporally ordered and semantically meaningful phrases. On the motion side, we explore several segmentation strategies, including uniform segmentation, Change-Point-Detection (CPD)-based segmentation, and clustering-based segmentation [14]. Building on these components, we further introduce a fine-grained text–motion alignment module which performs contrastive learning within each sample to achieve more precise cross-modal correspondence. To the best of our knowledge, this is the first work to incorporate segment-level alignment into text-conditioned human motion generation. Extensive experiments demonstrate that our method generates more accurate and natural human motions than existing approaches. Moreover, the contrastive learning framework enables our method to generalize to retrieval-style tasks, such as motion grounding and motion-to-text retrieval.

Our main contributions can be summarized as follows:

- We introduce a novel framework that decomposes both motion descriptions and motion sequences into atomic alignment units and performs fine-grained segment-level alignment via contrastive learning, enabling more accurate and natural motion generation.
- By learning a shared embedding for text and motion segments, our method naturally supports fine-grained retrieval-style tasks, including motion grounding and motion-to-text retrieval.
- Extensive experiments show that our method consistently improves upon strong baselines and achieves results comparable to the state-of-the-art without relying on the language-motion pretraining model.

## 2. Related Work

**Text-Guided Human Motion Generation:** Existing human motion generation methods generally follow two paradigms: discrete models and continuous models [3].

Discrete models treat motion sequences as discrete tokens. Typically, they adopt a two-stage generation process: first, a quantization model such as the Vector Quantized Variational Autoencoder (VQ-VAE) [24] discretizes continuous motion into a sequence of tokens; then, a generation model generates motion token sequences conditioned on the input text. Inspired by the success of GPT [19], T2M-GPT [30] uses causal attention to generate motion autoregressively. However, autoregressive models are uni-

directional and cannot capture bidirectional dependencies in motion data; they also require increased time for inference. Inspired by the masked modeling strategy in BERT-style models [1, 4], some works adopt a transformer decoder structure and achieve promising results. For example, MMM [17] applies masked modeling to support motion editing tasks such as in-betweening, upper-body modification, and long-sequence generation. MoMask [9] introduces a motion residual VQ-VAE [29] to reduce quantization error. Accordingly, a mask transformer and a residual transformer are proposed to generate base and residual tokens, respectively. Building on MoMask, we propose to impose a segment-level alignment constraint to achieve more accurate and natural motion generation.

In contrast to discrete models, which operate in the discrete motion space, continuous models directly model motion in continuous space using diffusion models. These models generate motion by gradually denoising an initial random signal into a realistic sequence through a learned conditional denoising process. MotionDiffuse [32] is the first diffusion-based model for text-guided human motion generation. MDM [22] proposes predicting the sample directly instead of noise, enabling the use of geometric losses on motion positions and velocities. MLD [2] introduces a latent diffusion model that operates in a latent space learned via a Variational Autoencoder (VAE) [12]. MARDM [16] revisits diffusion-based motion generation by addressing redundant representations and introducing a masked autoregressive framework with improved performance.

**Text-Motion Alignment:** Text–motion alignment is a critical aspect of evaluating text-guided human motion generation. To ensure that the generated motion accurately reflects the textual description, most existing methods directly use the CLIP [20] text encoder to extract text features as conditioning signals. Some approaches exploit more advanced text representations for fine-grained control. GraphMotion [11] constructs a hierarchical semantic graph from textual descriptions and applies coarse-to-fine diffusion based on the graph. CoMo [10] leverages LLMs to generate body-part-specific descriptions, enhancing the model’s controllability.

Inspired by contrastive learning in text–image modeling, several methods adopt similar ideas to improve text–motion alignment. MotionCLIP [23] aligns motions with the CLIP latent space to obtain semantically rich embeddings that support high-quality text-to-motion generation, editing, and recognition. MotionPatches [27] applies contrastive learning to train a motion-language model, enabling both text-to-motion and motion-to-text retrieval. CAR [5] similarly decomposes textual descriptions into multiple phrases, but only for data augmentation in retrieval tasks. In contrast, we incorporate segment-level alignment into the motion gener-

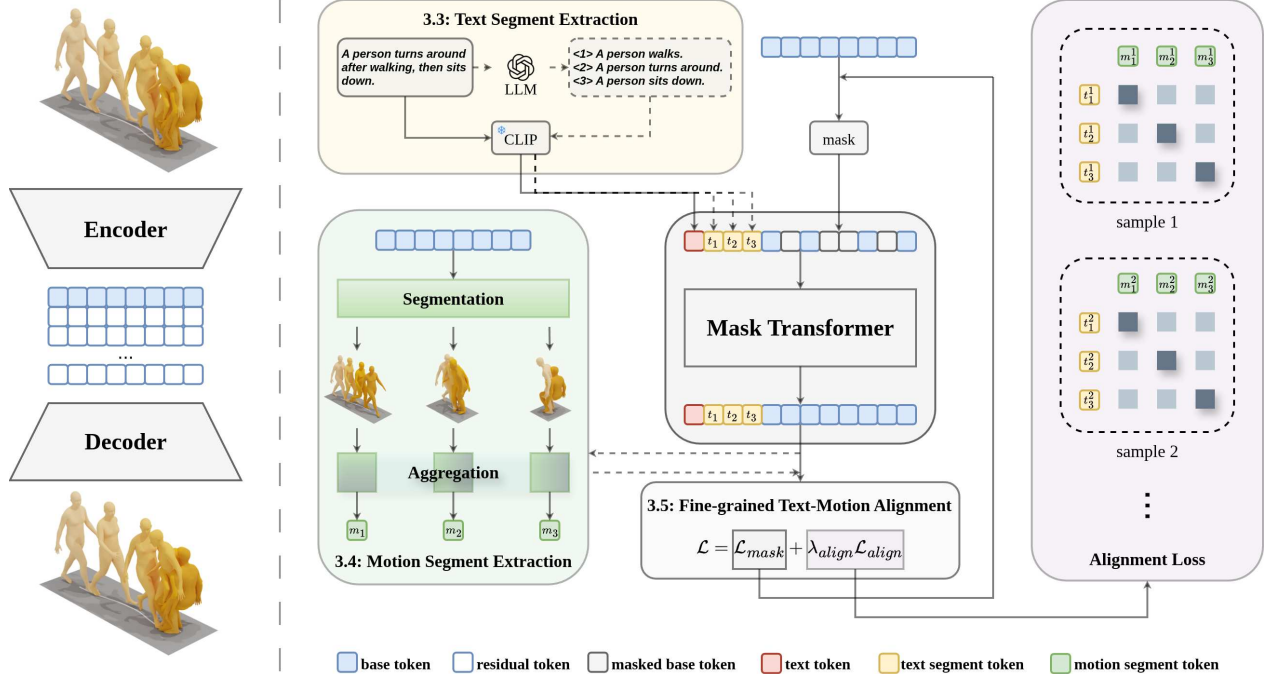


Figure 2. The overview of our method. **Left:** The Residual VQ-VAE encodes a continuous motion sequence into discrete motion tokens, including base tokens and residual tokens, which will be generated by the mask transformer and residual transformer, respectively. **Right:** The Mask Transformer predicts the masked base tokens conditioned on the textual description. To achieve segment-level fine-grained alignment, we introduce a *Text Segment Extraction* module and a *Motion Segment Extraction* module, which extract text and motion segments respectively, and align them through the *Fine-grained Text-Motion Alignment* module.

ation task, leading to more accurate and natural motions.

### 3. Method

#### 3.1. Overview

Our task is to generate 3D human motions from textual descriptions. Given a sentence describing a person performing certain actions, the goal is to produce the corresponding motion sequence that accurately reflects the input while remaining natural and realistic.

An overview of our method is shown in Figure 2. We adopt MoMask [9] as our baseline, which consists of two components: (1) a residual VQ-VAE (RVQ-VAE) [24, 29] that encodes the continuous motion sequence into discrete motion tokens, including tokens and residual tokens, and (2) a mask transformer and a residual transformer that generate the base and residual tokens, respectively. During generation, these tokens are fed into the RVQ-VAE decoder to reconstruct the motion sequence.

To enable fine-grained alignment, we define text segments and motion segments as the basic alignment units. A text segment is a standalone phrase that forms part of the full input text, while a motion segment denotes the corresponding sub-sequence of the motion. For example, in Figure 2, the sentence “A person turns around after walking,

then sits down.” can be decomposed into three temporally ordered text segments: [“A person walks.”, “A person turns around.”, “A person sits down.”], each corresponding to a sub-sequence of the motion.

Our main modification focuses on the mask transformer, where we introduce a *Text Segment Extraction* module (Section 3.3) and a *Motion Segment Extraction* module (Section 3.4), which extract the text and motion segments, respectively. These segments are then aligned through a *Fine-grained Text-Motion Alignment* module (Section 3.5), enabling segment-aligned motion generation.

#### 3.2. Model Structure

**Residual VQ-VAE:** Given an input continuous motion sequence  $M \in \mathbb{R}^{N \times D}$ , where  $N$  is the number of frames and  $D$  is the motion dimension, we first project it into the latent space using a 1D convolutional encoder:

$$V = \text{Encoder}(M) \in \mathbb{R}^{n \times d}, \quad (1)$$

where  $n/N$  is the downsampling ratio and  $d$  is the latent dimension. Each latent vector  $V_i$  is then quantized using a residual quantization technique:

$$\tilde{V}_i = \sum_{j=0}^k \mathbf{z}_i^j, \quad (2)$$

where  $\mathbf{z}_i^0$  comes from the base codebook,  $\mathbf{z}_i^j$  for  $j > 0$  quantizes the residual  $V_i - \sum_{l=0}^{j-1} \mathbf{z}_i^l$  coming from the residual codebook, and  $k$  is the layer number. Finally, the quantized latent sequence is projected back to the motion space using a 1D convolutional decoder:

$$\hat{M} = \text{Decoder}(\tilde{V}) \in \mathbb{R}^{N \times D}. \quad (3)$$

By using residual VQ-VAE to encode the motion, we can represent the motion sequence using base token sequence  $\mathbf{x}^0$  and residual token sequences  $\mathbf{x}^{1:k}$ .

**Mask Transformer:** The mask transformer is used to generate the base tokens, where our alignment module is applied. The input to the mask transformer is the text token  $T$ , text segment tokens  $\mathbf{t}_{1:A}$ , and masked base tokens  $\hat{\mathbf{x}}^0$ . The output is the predicted base tokens  $\mathbf{x}^0$ :

$$\mathbf{x}^0 = \text{MaskTrans}(T, \mathbf{t}_{1:A}, \hat{\mathbf{x}}^0). \quad (4)$$

$A$  denotes the maximum segment number. During generation, starting from a fully masked empty sequence, we iteratively predict the tokens under a cosine scheduler [1, 9].

**Residual Transformer:** The residual transformer is used to generate the residual tokens and has a similar structure to the mask transformer. The input of the residual transformer is the text token  $T$ , text segment tokens  $\mathbf{t}_{1:A}$ , first  $i-1$  layer tokens  $\mathbf{x}^{0:i-1}$  and target layer index  $i$ . The output is the  $i$ -th layer tokens:

$$\mathbf{x}^i = \text{ResTrans}(T, \mathbf{t}_{1:A}, \mathbf{x}^{0:i-1}, i). \quad (5)$$

During training, a random target layer  $i$  is selected, and all the tokens in the preceding  $i-1$  layers are embedded and summed up as the token embedding input. During generation, after generating the base tokens, the residual tokens are produced layer by layer.

### 3.3. Text Segment Extraction

The text segment extraction module leverages the power of LLMs to decompose the raw textual description into temporally ordered text segments. When designing the prompting strategy, we adhere to the following principles to ensure that the decomposed text segments not only preserve the original information but also provide temporal structure information, which facilitates segment-level alignment between text and motion.

- **Temporal ordering:** Segments should follow the order of actions, rather than the order of appearance in the raw text input. For example, “*A person runs after walking.*” is decomposed into [“*A person walks.*”, “*A person runs.*”].

- **Preserving details:** Segments should retain important modifiers such as “*slowly*”, “*two times*”, or “*as if + description*”. For instance, “*A person runs quickly, then jumps two times.*” is decomposed into [“*A person runs quickly.*”, “*A person jumps two times.*”].
- **Simultaneous actions:** If multiple actions occur simultaneously, they are treated as a single segment. For example, “*A person is standing and waving the hands.*” remains as one segment.

Following these rules, the input textual description is parsed into multiple text segments using an LLM, where each segment corresponds to a short motion sequence. We tested different medium-size LLMs and found that Qwen 3:8B [26] produces reliable and consistent results. The results using different LLMs and the detailed prompt are provided in the supplementary material (Section A).

In practice, the processing step is conducted offline. After obtaining the text segments, we use the CLIP [20] text encoder to extract features for both the raw textual description and its segments, resulting in one text token  $T$  and up to  $A$  text segment tokens  $\mathbf{t}_{1:A}$ . These features are then fed into both the mask transformer and residual transformer as the conditioning signals.

### 3.4. Motion Segment Extraction

After obtaining the text segments of the textual description, the next step is to extract the corresponding motion segments, which will serve as the atomic alignment units.

The segmentation module partitions the motion into the same number of segments as the text. We adopt a uniform segmentation strategy as our default choice, where motion tokens are evenly distributed across the text segments, ensuring balanced segment lengths that facilitate stable text–motion alignment and consistent model training. Given the predicted base tokens  $\mathbf{x}^0$  from the mask transformer, along with the start index  $s_i$  and end index  $e_i$  of the  $i$ -th motion segment from the segmentation module, the motion segment aggregation module outputs the corresponding motion segment  $\mathbf{m}_i$ :

$$\begin{aligned} \mathbf{m}_i &= \text{Agg}(\mathbf{x}^0, s_i, e_i) \\ &= \text{MLP}(\text{Concat}(\text{mean}(\mathbf{x}_{s_i:e_i}^0), \text{max}(\mathbf{x}_{s_i:e_i}^0))). \end{aligned} \quad (6)$$

For comparison, we also explored other feasible segmentation methods, including Change-Point-Detection (CPD)-based and Clustering-based segmentation approaches. Details are provided in Section 4.3. While these methods can be applied, uniform segmentation consistently produces the most stable alignment and training performance, resulting in superior generation quality compared to alternative methods.

Method	Text Encoder	R-Precision $\uparrow$			FID $\downarrow$	MM-Dist $\downarrow$	Diversity $\rightarrow$
		Top 1	Top 2	Top 3			
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
TM2T [8]	GRU	0.424 $\pm$ .003	0.618 $\pm$ .003	0.729 $\pm$ .002	1.501 $\pm$ .017	3.467 $\pm$ .011	8.589 $\pm$ .076
T2M [7]	GRU	0.455 $\pm$ .003	0.636 $\pm$ .003	0.736 $\pm$ .002	1.087 $\pm$ .021	3.347 $\pm$ .008	9.175 $\pm$ .083
MotionDiffuse [32]	Transformer	0.491 $\pm$ .001	0.681 $\pm$ .001	0.782 $\pm$ .001	0.630 $\pm$ .001	3.113 $\pm$ .001	9.410 $\pm$ .049
MDM [22]	CLIP	-	-	0.611 $\pm$ .007	0.544 $\pm$ .044	5.566 $\pm$ .027	9.559 $\pm$ .086
MLD [2]	CLIP	0.481 $\pm$ .003	0.673 $\pm$ .003	0.772 $\pm$ .002	0.473 $\pm$ .013	3.196 $\pm$ .010	9.724 $\pm$ .082
T2M-GPT [30]	CLIP	0.491 $\pm$ .003	0.680 $\pm$ .003	0.775 $\pm$ .002	0.116 $\pm$ .004	3.118 $\pm$ .011	9.761 $\pm$ .081
ReMoDiffuse [31]	CLIP	0.510 $\pm$ .005	0.698 $\pm$ .006	0.795 $\pm$ .004	0.103 $\pm$ .004	2.974 $\pm$ .016	9.018 $\pm$ .075
MMM [17]	CLIP	0.515 $\pm$ .002	0.708 $\pm$ .002	0.804 $\pm$ .002	0.089 $\pm$ .005	2.926 $\pm$ .007	9.577 $\pm$ .050
CoMo [10]	CLIP	0.502 $\pm$ .002	0.692 $\pm$ .007	0.790 $\pm$ .002	0.262 $\pm$ .004	3.032 $\pm$ .015	9.936 $\pm$ .066
LAMP [13]	LAMP	0.557 $\pm$ .003	0.751 $\pm$ .002	0.843 $\pm$ .001	0.032 $\pm$ .002	2.759 $\pm$ .007	9.571 $\pm$ .069
MoMask (Baseline) [9]	CLIP	0.521 $\pm$ .002	0.713 $\pm$ .002	0.807 $\pm$ .002	0.045 $\pm$ .002	2.958 $\pm$ .008	-
SegMo (Ours)	CLIP	0.553 $\pm$ 0.003	0.748 $\pm$ 0.003	0.838 $\pm$ 0.002	0.042 $\pm$ 0.002	2.782 $\pm$ 0.008	9.632 $\pm$ 0.079

Table 1. Comparison of text-conditional human motion generation on the HumanML3D test set. For each metric, we repeat the evaluation 20 times and report the average with a 95% confidence interval. Red and Blue indicate the best and the second-best results.

### 3.5. Fine-Grained Text-Motion Alignment

After obtaining the text and motion segments, the next step is to align them in a shared embedding space. We design an alignment module to establish fine-grained correspondence between text and motion. Here, we introduce the loss function used to train the mask transformer under this segment-level alignment constraint.

**Mask Loss:** We denote the text condition as  $c$  and randomly mask a portion of the base tokens to obtain  $\mathbf{x}^0$ . The mask loss minimizes the negative log-likelihood of the target predictions for the masked base tokens:

$$\mathcal{L}_{mask} = - \sum_{i \in \mathcal{M}} \log p_{\theta}(\mathbf{x}_i^0 | c, \hat{\mathbf{x}}^0). \quad (7)$$

Here,  $\mathcal{M}$  denotes the index set of the masked base tokens.

**Alignment Loss:** As illustrated in Figure 2, the alignment module ensures that each text segment is paired with its corresponding motion segment within a sample, avoiding semantic interference across different samples. This differs from CLIP-style implementations [20], which compute contrastive loss across all samples in a batch. The alignment loss is defined as follows:

$$\begin{aligned} \mathcal{L}_{t2m} &= - \frac{1}{B \cdot A} \sum_{i=1}^B \sum_{j=1}^A \log \frac{\exp(\text{sim}(\mathbf{t}_j^i, \mathbf{m}_j^i) / \tau)}{\sum_{k=1}^A \exp(\text{sim}(\mathbf{t}_j^i, \mathbf{m}_k^i) / \tau)}, \\ \mathcal{L}_{m2t} &= - \frac{1}{B \cdot A} \sum_{i=1}^B \sum_{j=1}^A \log \frac{\exp(\text{sim}(\mathbf{t}_j^i, \mathbf{m}_j^i) / \tau)}{\sum_{k=1}^A \exp(\text{sim}(\mathbf{t}_k^i, \mathbf{m}_j^i) / \tau)}, \\ \mathcal{L}_{align} &= \frac{1}{2} (\mathcal{L}_{t2m} + \mathcal{L}_{m2t}). \end{aligned} \quad (8)$$

Here,  $B$  denotes the batch size,  $A$  denotes the maximum number of segments per sample,  $\text{sim}(\mathbf{t}_j^i, \mathbf{m}_k^i)$  denotes the cosine similarity between the  $j$ -th text segment and the  $k$ -th motion segment in the  $i$ -th sample, and  $\tau$  denotes the temperature parameter. The alignment loss enforces the text/motion segment to be closer to the paired motion/text segment, while being far away from other motion/text segments.

**Overall Objective:** The total loss for training the mask Transformer combines the mask loss and alignment loss:

$$\mathcal{L} = \mathcal{L}_{mask} + \lambda_{align} \mathcal{L}_{align}. \quad (9)$$

Here,  $\lambda_{align}$  denotes the weight for the alignment loss.

## 4. Experiments

### 4.1. Experiment Settings

**Datasets:** We conduct experiments on two widely used datasets: HumanML3D [7] and KIT Motion-Language (KIT-ML) [18]. HumanML3D combines motion sequences from the HumanAct12 [6] and AMASS [15] datasets, pairing each sequence with three distinct textual descriptions, resulting in 14.6k motion sequences and 44.9k textual annotations. KIT-ML is a smaller dataset containing 3.9k motion sequences and 6.3k textual annotations. The segment number distributions for both datasets are provided in the supplementary material (Section E). For both datasets, we adopt the pose representation from T2M [7].

**Evaluation Metrics:** We follow the standard evaluation protocols from T2M [7] to assess the motion generation model. The metrics are categorized into three groups: (1)

Method	Text Encoder	R-Precision $\uparrow$			FID $\downarrow$	MM-Dist $\downarrow$	Diversity $\rightarrow$
		Top 1	Top 2	Top 3			
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
TM2T [8]	GRU	0.280 $\pm$ .005	0.463 $\pm$ .006	0.587 $\pm$ .005	3.599 $\pm$ .153	4.591 $\pm$ .026	9.473 $\pm$ .117
T2M [7]	GRU	0.361 $\pm$ .005	0.559 $\pm$ .007	0.681 $\pm$ .007	3.022 $\pm$ .107	3.488 $\pm$ .028	10.72 $\pm$ .145
MotionDiffuse [32]	Transformer	0.417 $\pm$ .004	0.621 $\pm$ .004	0.739 $\pm$ .004	1.934 $\pm$ .064	2.958 $\pm$ .005	11.10 $\pm$ .143
MDM [22]	CLIP	-	-	0.396 $\pm$ .004	0.497 $\pm$ .021	9.191 $\pm$ .022	10.847 $\pm$ .109
MLD [2]	CLIP	0.390 $\pm$ .008	0.609 $\pm$ .008	0.734 $\pm$ .007	0.404 $\pm$ .027	3.204 $\pm$ .027	10.80 $\pm$ .117
T2M-GPT [30]	CLIP	0.402 $\pm$ .006	0.619 $\pm$ .005	0.737 $\pm$ .006	0.717 $\pm$ .041	3.053 $\pm$ .026	10.86 $\pm$ .094
ReMoDiffuse [31]	CLIP	0.427 $\pm$ .014	0.641 $\pm$ .004	0.765 $\pm$ .055	0.155 $\pm$ .006	2.814 $\pm$ .012	10.80 $\pm$ .105
MMM [17]	CLIP	0.404 $\pm$ .005	0.621 $\pm$ .005	0.744 $\pm$ .004	0.316 $\pm$ .028	2.977 $\pm$ .019	10.910 $\pm$ .101
CoMo [10]	CLIP	0.422 $\pm$ .009	0.638 $\pm$ .007	0.765 $\pm$ .011	0.332 $\pm$ .045	2.873 $\pm$ .021	10.95 $\pm$ .196
LAMP [13]	LAMP	0.479 $\pm$ .006	0.691 $\pm$ .005	0.826 $\pm$ .005	0.141 $\pm$ .013	2.704 $\pm$ .018	10.929 $\pm$ .101
MoMask (Baseline) [9]	CLIP	0.433 $\pm$ .007	0.656 $\pm$ .005	0.781 $\pm$ .005	0.204 $\pm$ .011	2.779 $\pm$ .022	-
SegMo (Ours)	CLIP	0.443 $\pm$ 0.005	0.663 $\pm$ 0.006	0.784 $\pm$ 0.005	0.163 $\pm$ 0.007	2.733 $\pm$ 0.012	10.797 $\pm$ 0.113

Table 2. Comparison of text-conditional human motion generation on the KIT-ML test set. For each metric, we repeat the evaluation 20 times and report the average with a 95% confidence interval. Red and Blue indicate the best and the second-best results.

R-Precision and Multimodal Distance (MM-Dist), which measure how well the generated motions align with the input textual descriptions. R-Precision is reported with Top-1, Top-2, and Top-3 accuracy. (2) Fréchet Inception Distance (FID), which evaluates the realism of the generated motions by computing the distribution distance between the generated and ground-truth motion features. (3) Diversity, which is calculated by averaging the Euclidean distances of 300 randomly sampled motion pairs. The closer this value is to that of real motions, the better.

**Implementation Details:** Our model is implemented in PyTorch and evaluated on an NVIDIA A100 GPU. We follow the same training setup as MoMask [9] for the residual VQ-VAE, the mask transformer, and the residual transformer. We calculate the average segment number of both datasets and find that setting the maximum number of segments to 5 can cover most samples. The temperature for computing the alignment loss is fixed at 0.1, and the loss weight is set to 1.0 on the HumanML3D dataset and 0.1 on the KIT-ML dataset.

## 4.2. Evaluation

**Quantitative Comparison:** We evaluate our method against a range of existing approaches, including discrete models [7–10, 13, 17, 30] and continuous models [2, 22, 31, 32]. The results on the HumanML3D dataset are presented in Table 1. Compared to MoMask [9], a strong baseline, our approach achieves significant improvements, particularly on text-motion alignment metrics such as R-Precision and MM-Dist, while also maintaining or improving performance on other metrics. Furthermore, our method achieves performance comparable to LAMP [13], the current state-of-the-art (SOTA) approach, without relying on a

language-motion pretraining model. The results on the KIT-ML dataset are shown in Table 2. Although most sequences in KIT-ML contain only a single action, our alignment module improves all metrics over the baseline method, while achieving a comparable FID to the SOTA method.

**Qualitative Comparison:** Figure 3 presents a qualitative comparison of our approach with T2M-GPT [30] and MoMask. In Figure 3 (a), both T2M-GPT and MoMask start from a standing pose and fail to generate the “stand up” action. In Figure 3 (b), existing methods fail to accurately capture or distinguish critical descriptors such as “left” and “right”, resulting in mismatched motions. In Figure 3 (c), the compared methods do not “continue walking forward”, leading to incorrect motion segments. Overall, our method effectively captures segment-level details, generating motions that are both natural and consistent with the textual descriptions.

## 4.3. Discussion

Table 3 presents an ablation study evaluating the impact of different design choices by replacing individual modules.

**Motion Segmentation Module:** Human motion is continuous and often contains long transition phases, making action boundaries inherently ambiguous. This motivates exploring different segmentation strategies, including uniform, Change-Point-Detection (CPD)-based, and clustering-based approaches. Further details are reported in the supplementary material (Section B). As reported in Table 3 and the supplementary material (Table 2), the mean segmentation error appears to have little effect on performance, since transition frames can plausibly belong to either neighboring action. In contrast, the variance of seg-




Figure 3. Qualitative comparison of T2M-GPT, MoMask, and Ours on the HumanML3D test set. A fixed number of keyframes is shown for each motion sequence. Please refer to the supplementary video for additional comparison results.

mentation errors is more critical: *high variance leads to inconsistent boundaries and noisy alignments, destabilizing optimization, whereas low variance produces more consistent pairs, enabling stable training and high-quality generation.* Among all methods, uniform segmentation naturally achieves the lowest error variance, yielding the most consistent improvements in generation quality, as shown in Table 3.

To further evaluate how segmentation affects text–motion alignment, we introduce the Intra-Segment Consistency (ISC) score, defined as the cosine similarity between paired text and motion segments in the learned embedding space. ISC reflects both segmentation quality and the model’s alignment ability, depending jointly on the trained model and the segmentation method. Figure 4 highlights two key observations: (1) all models achieve the highest ISC when evaluated with uniform segmentation, whereas clustering-based segmentation introduces high instability and reduces alignment quality; (2) when evaluated with the same segmentation method used during training, models trained with uniform segmentation achieve the highest ISC, consistent with the generation metrics.

**Text-Motion Alignment Module:** For the text-motion alignment module, we also evaluated different strategies: (1) Batch-Level Alignment, where text and motion seg-

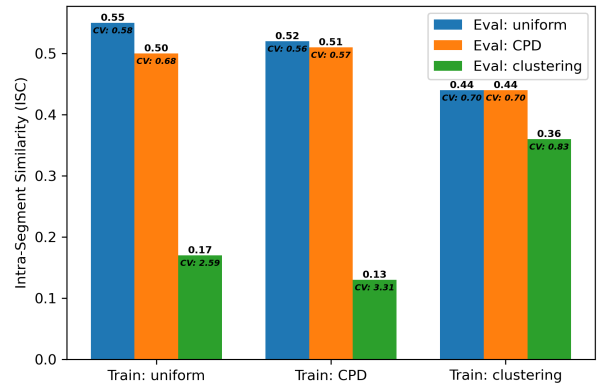


Figure 4. The Intra-Segment Consistency (ISC) of all models evaluated under different segmentation methods on the HumanML3D test set. The Coefficient of Variation (CV), defined as  $\text{std}(ISC)/\text{mean}(ISC)$ , is reported to assess stability. “Train” denotes the segmentation method used for training, while “Eval” denotes the segmentation method used for evaluation.

ments from different samples within a batch are concatenated, and the alignment loss is computed within the batch; (2) Global Alignment, where the alignment loss is calculated between the entire textual description and the complete motion sequence within a batch; (3) Token Align-

Method	R-Precision $\uparrow$			FID $\downarrow$	MM-Dist $\downarrow$	Diversity $\rightarrow$
	Top 1	Top 2	Top 3			
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
SegMo (CPD-based Seg)	0.548 $\pm$ 0.002	0.744 $\pm$ 0.002	0.834 $\pm$ 0.002	0.046 $\pm$ 0.002	2.804 $\pm$ 0.008	9.602 $\pm$ 0.070
SegMo (clustering-based Seg)	0.548 $\pm$ 0.002	0.741 $\pm$ 0.002	0.833 $\pm$ 0.002	0.051 $\pm$ 0.002	2.811 $\pm$ 0.008	9.585 $\pm$ 0.071
SegMo (No Align)	0.538 $\pm$ 0.003	0.732 $\pm$ 0.003	0.825 $\pm$ 0.002	0.065 $\pm$ 0.003	2.863 $\pm$ 0.009	9.636 $\pm$ 0.065
SegMo (Batch-Level Align)	0.546 $\pm$ 0.002	0.741 $\pm$ 0.002	0.831 $\pm$ 0.002	0.056 $\pm$ 0.003	2.817 $\pm$ 0.008	9.506 $\pm$ 0.081
SegMo (Global Align)	0.539 $\pm$ 0.002	0.732 $\pm$ 0.002	0.826 $\pm$ 0.003	0.072 $\pm$ 0.003	2.861 $\pm$ 0.007	9.632 $\pm$ 0.090
SegMo (Token Align)	0.536 $\pm$ 0.002	0.730 $\pm$ 0.002	0.824 $\pm$ 0.002	0.056 $\pm$ 0.002	2.870 $\pm$ 0.007	9.654 $\pm$ 0.066
SegMo	0.553 $\pm$ 0.003	0.748 $\pm$ 0.003	0.838 $\pm$ 0.002	0.042 $\pm$ 0.002	2.782 $\pm$ 0.008	9.632 $\pm$ 0.079

Table 3. Ablation results of replacing the segmentation and the alignment module on the HumanML3D test set. For each metric, we repeat the evaluation 20 times and report the average with a 95% confidence interval. Red and Blue indicate the best and the second-best results.

ment, where the alignment loss is computed directly between motion tokens and their corresponding text segments within a sample. Further details are provided in the supplementary material (Section D). In summary, our proposed alignment module provides a more balanced design: it preserves fine-grained text–motion correspondences while maintaining sufficient semantic context, leading to more robust alignment and better overall performance.

#### 4.4. Applications

Our alignment module enforces each text segment and its corresponding motion segment to be close in a shared embedding space via contrastive learning. This property naturally enables our model to extend beyond generation tasks to retrieval-style applications. In particular, we explore two tasks: motion grounding and motion-to-text retrieval.

**Motion Grounding:** Motion grounding aims to localize the corresponding motion segment for a query text within a motion sequence [25]. Specifically, we input both the query text and the motion sequence into our model to map them into the shared embedding space. Then, we compute the cosine similarity between the query text segment  $t_q$  and all candidate motion segments  $\mathcal{M}$ . The candidate motion segments can be obtained using a sliding window. The motion segment with the highest similarity is regarded as the localized result:  $\mathbf{m}^* = \arg \max_{\mathbf{m} \in \mathcal{M}} \text{sim}(t_q, \mathbf{m})$ . Figure 5 illustrates an example of motion grounding, showing how our model effectively matches a query text with its corresponding motion segment, while the baseline fails to establish the correct correspondence.

**Motion-to-Text Retrieval:** In the reverse direction, our model can also perform motion-to-text retrieval, where the goal is to find the most matching text description for a given motion segment. Given a query motion segment  $\mathbf{m}_q$ , we retrieve the text segment with the highest similarity from the candidate set  $\mathcal{T}$ :  $\mathbf{t}^* = \arg \max_{\mathbf{t} \in \mathcal{T}} \text{sim}(\mathbf{t}, \mathbf{m}_q)$ .

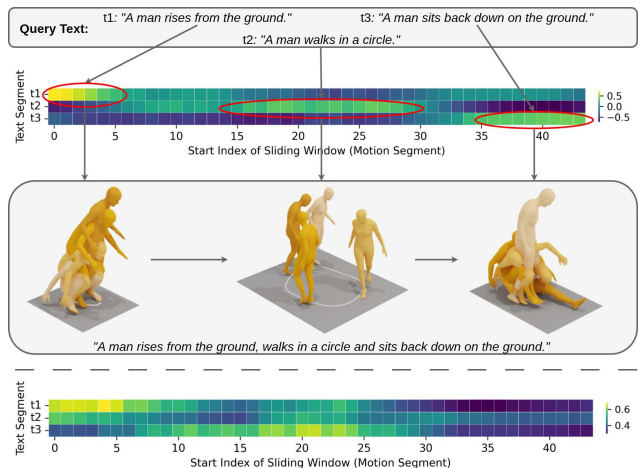


Figure 5. Example of motion grounding. **Top:** Results using the similarity map generated by our method. **Bottom:** Similarity map generated by MoMask [9]. In each map, the x-axis denotes the start index of the sliding window and the y-axis denotes the text segment. The motion length is 49, and the window size is 5.

## 5. Conclusion

In this work, we introduce SegMo, a segment-aligned text-conditional human motion generation method. Our main idea is to minimize the alignment unit by decomposing complex motion descriptions and motion sequences into temporally ordered segments and aligning them in a shared embedding space. Evaluation on standard benchmarks shows that SegMo achieves fine-grained text–motion alignment, generating more accurate and natural human motions.

**Limitation:** The limitation of our method lies in the motion segmentation module. While uniform segmentation is stable and outperforms other approaches, it is not precise enough for fine-grained alignment. In the future, we will improve segmentation accuracy and leverage datasets with detailed segmentation annotations to enhance the model’s alignment precision.

## Acknowledgement

This work made use of the facilities of the N8 Centre of Excellence in Computationally Intensive Research (N8 CIR) provided and funded by the N8 research partnership and EPSRC (Grant No. EP/T022167/1). The Centre is coordinated by the Universities of Durham, Manchester and York.

## References

- [1] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11315–11325, 2022. 2, 4
- [2] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18000–18010, 2023. 2, 5, 6
- [3] Jungbin Cho, Junwan Kim, Jisoo Kim, Minseo Kim, Mingyu Kang, Sungeun Hong, Tae-Hyun Oh, and Youngjae Yu. Discord: Discrete tokens to continuous motion via rectified flow decoding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14602–14612, 2025. 2
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186, 2019. 2
- [5] Kent Fujiwara, Mikihiro Tanaka, and Qing Yu. Chronologically accurate retrieval for temporal grounding of motion-language models. In *European Conference on Computer Vision*, pages 323–339. Springer, 2024. 2
- [6] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2021–2029, 2020. 5
- [7] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5152–5161, 2022. 5, 6
- [8] Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *European Conference on Computer Vision*, pages 580–597. Springer, 2022. 5, 6
- [9] Chuan Guo, Yuxuan Mu, Muhammad Gohar Javed, Sen Wang, and Li Cheng. Momask: Generative masked modeling of 3d human motions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1900–1910, 2024. 2, 3, 4, 5, 6, 8
- [10] Yiming Huang, Weilin Wan, Yue Yang, Chris Callison-Burch, Mark Yatskar, and Lingjie Liu. Como: Controllable motion generation through language guided pose code editing. In *European Conference on Computer Vision*, pages 180–196. Springer, 2024. 2, 5, 6
- [11] Peng Jin, Yang Wu, Yanbo Fan, Zhongqian Sun, Wei Yang, and Li Yuan. Act as you wish: Fine-grained control of motion diffusion model with hierarchical semantic graphs. *Advances in Neural Information Processing Systems*, 36: 15497–15518, 2023. 2
- [12] Diederik P Kingma, Max Welling, et al. Auto-encoding variational bayes, 2013. 2
- [13] Zhe Li, Weihao Yuan, Lingteng Qiu, Shenhao Zhu, Xiaodong Gu, Weichao Shen, Yuan Dong, Zilong Dong, Laurence Tianruo Yang, et al. Lamp: Language-motion pretraining for motion generation, retrieval, and captioning. In *The Thirteenth International Conference on Learning Representations*. 5, 6
- [14] Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982. 2
- [15] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. 5
- [16] Zichong Meng, Yiming Xie, Xiaogang Peng, Zeyu Han, and Huaizu Jiang. Rethinking diffusion for text-driven human motion generation: Redundant representations, evaluation, and masked autoregression. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 27859–27871, 2025. 2
- [17] Ekkasit Pinyoanuntapong, Pu Wang, Minwoo Lee, and Chen Chen. Mmm: Generative masked motion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1546–1555, 2024. 2, 5, 6
- [18] Matthias Plappert, Christian Mandery, and Tamim Asfour. The kit motion-language dataset. *Big data*, 4(4):236–252, 2016. 5
- [19] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 2
- [20] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 1, 2, 4, 5
- [21] Ali Rida Sahili, Najett Neji, and Hedi Tabia. Text-driven motion generation: Overview, challenges and directions. *arXiv preprint arXiv:2505.09379*, 2025. 1
- [22] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*. 2, 5, 6
- [23] Guy Tevet, Brian Gordon, Amir Hertz, Amit H Bermano, and Daniel Cohen-Or. Motionclip: Exposing human motion generation to clip space. In *European Conference on Computer Vision*, pages 358–374. Springer, 2022. 2

- [24] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [25] Xinghan Wang, Zixi Kang, and Yadong Mu. Text-controlled motion mamba: Text-instructed temporal grounding of human motion. *IEEE Transactions on Image Processing*, 34: 7079–7092, 2025. 8
- [26] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. 4
- [27] Qing Yu, Mikihiro Tanaka, and Kent Fujiwara. Exploring vision transformers for 3d human motion-language models with motion patches. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 937–946, 2024. 2
- [28] Jeffrey M Zacks and Khena M Swallow. Event segmentation. *Current directions in psychological science*, 16(2):80–84, 2007. 1
- [29] Neil Zeghidour, Alejandro Luebs, Ahmed Omran, Jan Skoglund, and Marco Tagliasacchi. Soundstream: An end-to-end neural audio codec. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:495–507, 2021. 2, 3
- [30] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Yong Zhang, Hongwei Zhao, Hongtao Lu, Xi Shen, and Ying Shan. Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14730–14740, 2023. 2, 5, 6
- [31] Mingyuan Zhang, Xinying Guo, Liang Pan, Zhongang Cai, Fangzhou Hong, Huirong Li, Lei Yang, and Ziwei Liu. Remodiffuse: Retrieval-augmented motion diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 364–373, 2023. 5, 6
- [32] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *IEEE transactions on pattern analysis and machine intelligence*, 46(6):4115–4128, 2024. 2, 5, 6
- [33] Wentao Zhu, Xiaoxuan Ma, Dongwoo Ro, Hai Ci, Jinlu Zhang, Jiabin Shi, Feng Gao, Qi Tian, and Yizhou Wang. Human motion generation: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4):2430–2449, 2023. 1