

# TLControl: Trajectory and Language Control for Human Motion Synthesis

Weilin Wan<sup>1,2</sup>

Zhiyang Dou<sup>1,2</sup>

Taku Komura<sup>1</sup>

Wenping Wang<sup>3</sup>

Dinesh Jayaraman<sup>2†</sup>


Lingjie Liu<sup>2†</sup>

<sup>1</sup>The University of Hong Kong


<sup>2</sup>University of Pennsylvania

<sup>3</sup>Texas A&M University

†denotes equal contributions



“the man walks around in a circle counter-clockwise.”



“a man is walking ahead slowly.”

“the person walked forward.”

“a man raised both his hands.”

Figure 1. **TLControl**, a novel method for **Trajectory** and **Language Control** for Human Motion Synthesis. The dotted lines represent the input control trajectories. The corresponding control joints are highlighted in orange. Our method demonstrates versatile multi-joint controls (see Figures 1a to 1c), the ability to handle complex trajectories (see Figure 1d), multi-stage control (see Figure 1e), and the preservation of language semantics while utilizing trajectory controls (see Figure 1f). We highly encourage readers to view our supplementary video to see our results.

## Abstract

Controllable human motion synthesis is essential for applications in AR/VR, gaming, movies, and embodied AI. Existing methods often focus solely on either language or full trajectory control, lacking precision in synthesizing motions aligned with user-specified trajectories, especially for multi-joint control. To address these issues, we present TLControl, a new method for realistic human motion synthesis, incorporating both low-level trajectory and high-level language semantics controls. Specifically, we first train a VQ-VAE to learn a compact latent motion space orga-

nized by body parts. We then propose a Masked Trajectories Transformer to make coarse initial predictions of full trajectories of joints based on the learned latent motion space, with user-specified partial trajectories and text descriptions as conditioning. Finally, we introduce an efficient test-time optimization to refine these coarse predictions for accurate trajectory control. Experiments demonstrate that TLControl outperforms the state-of-the-art in trajectory accuracy and time efficiency, making it practical for interactive and high-quality animation generation. Project Website: <https://tlcontrol.weilinwl.com>

## 1. Introduction

Human motion generation is a fundamental problem in computer graphics, with a wide range of applications in AR/VR, gaming industry and large-scale training data generation for embodied AI. In these applications, it is crucial to support user control at different levels on motion generation. For example, users may want to have high-level semantics control (e.g., text descriptions) for motion synthesis, and, meanwhile, expect flexible and precise spatial control to “drag” any joint on arbitrary trajectories. In this work, we study the problem of realistic and precise human motion generation with both high-level language and low-level trajectory controls.

Recently, great advancements have been made in language-conditioned human motion generation [1, 2, 8, 12, 13, 22, 23, 30, 32, 39, 41, 48, 59]. However, most existing works only support language control [7, 48, 58]. Another line of research explores trajectory-controlled motion generation [5, 38, 40], a.k.a., inverse kinematics, but it lacks high-level language semantics control. Furthermore, inverse kinematics can only infer joint configurations from full (complete) trajectories of end-effectors, while our method generates full motions from partial (incomplete) trajectories. The most related works are PriorMDM[39], GMD [22], and a concurrent work, Omnicontrol [54], that incorporate spatial control signals in a language-conditioned motion generation model, but neither of them can accurately align the motion with specified trajectories, especially when controlling multiple joints simultaneously, as depicted in Figure 4. Moreover, their motion generation process is time-consuming due to inefficient multi-step sampling in the motion diffusion model, restricting its practicality in real-world applications.

To address these issues, we propose TLControl for Trajectory and Language-controlled human body motion generation. We propose to operate in a part-structured latent encoding of human motions, learned with a VQ-VAE exploiting knowledge of the skeletal topology of human body. We then synthesize full-body motions from partial trajectory and language input in two steps: a feed-forward coarse estimation, and a subsequent optimization-based refinement. To efficiently train the feed-forward model, we propose a Masked Trajectory Transformer (MTT) that is trained by synthesizing various possible partial trajectory controls through spatial and temporal masking, enabling highly flexible operation. The optimization step starts at the MTT solution, and then more precisely aligns it with user-specified trajectories through an iterative search procedure.

We conduct extensive experiments to validate the effectiveness of our approach. Compared to existing methods, TLControl exhibits superior performance in both trajectory accuracy and time efficiency. The results further illustrate that our method empowers users to interactively generate

and modify high-quality animations within a brief runtime, affirming its practicality and effectiveness.

In summary, our technical contributions are as follows:

1. We introduce a new method TLControl that generates high-quality and precise human motions with highly flexible language and partial trajectory-based user specifications.
2. TLControl learns a human body morphology-aware structured latent space suitable for flexible motion synthesis.
3. An efficient feed-forward transformer model trained using carefully designed masking strategies initializes iterative optimization for large computational efficiencies in motion synthesis compared to prior work, at comparable or superior performance levels.

## 2. Related Work

There have recently been significant advances in human motion generation, incorporating diverse, multi-modal inputs such as text prompts [1, 8, 12, 13, 21, 23, 29, 30, 32, 41, 48], action labels [10, 31], partial motion sequences [9, 15, 51], control signals [8, 28, 43, 45], music [2, 24, 25], and images [6, 38]. In this section, we conduct a comprehensive review of the existing literature in the field of motion generation. Our reviews covers the studies ranging from unconditional motion generation to conditional approaches, particularly focusing on language-conditioned motion synthesis, known as action-to-motion and text-to-motion. We also explore the integration of trajectory or path controls in these processes, highlighting how these innovations contribute to advancing the field.

**Unconditional Motion Generation** This area of research focuses on the autonomous generation of motion sequences without specific guiding conditions or annotations. Pioneering studies, such as those referenced in [36, 48, 55, 60, 61] have made notable strides in this domain. These works stand out for their unique ability to capture and model the entire motion space, leveraging raw motion data to produce diverse and dynamic motion patterns. These methods exhibit innovation in interpreting and representing complex motion dynamics. However, as our research primarily focuses on conditional motion generation, our discussion of unconditional motion generation only serves as a contextual background to our primary focus.

**Conditional Motion Generation** In the realm of conditional, tasks, ACTOR [31] presents a class-agnostic transformer VAE as a baseline. It introduces learnable biases within a transformer VAE to encapsulate action for motion generation. But similar to other works in the field of action-to-motion [10], it is limited to cover changes in the conditional inputs. Recent studies in the task text-to-motion [1, 3, 12, 23, 32, 33, 48, 59] has emerged as a principal driver, reshaping research frontiers with its user-friendly

nature by allowing natural language input. The advancements can be categorized into two methodologies: joint-latent models and diffusion models. Joint-latent models like TEMOS [32] typically employ a motion VAE alongside a text variational encoding. These components are then aligned into a compatible latent space by using divergence loss. This method advances the synthesis of human motion sequences from natural language inputs, representing a significant step forward in the field. However, as observed in the misalignment of Gaussian distributions [47], there are difficulties remained in aligning the structure and distribution of natural language and human motions.

Recent advancements in image synthesis have been marked by the successful application of diffusion generative models [42]. Building on this progress, these models have now been extended to the domain of human motion synthesis. This novel application is being applied in recent works such as those presented in [2, 7, 19, 23, 48, 54]. Notably, MotionDiffuse [59] emerges as the first text-based motion diffusion model, offering fine-grained instructions for individual body parts. MDM [48] introduces a motion diffusion model that operates directly on raw motion data, further expanding the scope of natural language controls in human motion generation. This method represents a significant leap forward in terms of enabling intuitive, natural language inputs to motion synthesis.

**Motion Generation with Spatial Constraints** Despite these previous efforts, the language description itself is still limited to a coarse control over the motion. To tackle the issue, Shafir et al. [39] proposed PriorMDM to generate long-sequence human motion and joint control signals. To further support more flexible fine-grained motion generation, GMD [22] integrates spatial constraints by employing a two-stage diffusion strategy. However, GMD only controls the 2D positions of the pelvis, a limitation that reduces its flexibility in many practical scenarios. Recently, the concurrent work Omnicontrol [54] incorporates flexible spatial control signals over different joints by introducing analytic spatial guidance and realism guidance that ensure the generated motion from the diffusion model can tightly conform to the input control signals. However, Omnicontrol struggles to generate body motion accurately aligned with specified trajectories, as we will illustrate. Furthermore, it exhibits a relatively prolonged motion generation time. In comparison with Omnicontrol, our method excels in producing motion that precisely adheres to the control signal, demonstrating superior accuracy and significantly faster runtime.

### 3. Method

We aim to generate full-body motions to match user specifications, which consist of two parts: a text description  $\mathbf{L}$ , and a partial trajectory  $\mathbf{R}'$ .  $\mathbf{R}'$  has  $T$  frames, corresponding to the target motion length. In each frame  $t$ ,  $\mathbf{R}'$  contains

the 3D positions of some subset  $m_t < n$  of the  $n$  full-body joints, corresponding to the “control joints” which the user chose to specify for that frame. For example, the user might choose to specify only the head joints for all frames  $t$ , or even specify different joints in different frames. The desired final output is the full body motion  $\mathbf{J} \in \mathbb{R}^{T \times M}$  of the character, where  $M$  is the dimension of the single frame pose represented by the joint rotations and positions.

We propose to train such a system, TLControl, using data in which each sample  $(\mathbf{R}, \mathbf{J}, \mathbf{L})$  is a tuple containing a full trajectory specifying some key joints, ground truth full body motions, and language description. To achieve the goal, we first learn a motion embedding by training a VQ-VAE [49] to establish part-based latent spaces, representing human body motion at the body part level (Sec.3.1). We then train a text-conditioned Masked Trajectory Transformer (MTT) to output a coarse completion of partially masked trajectories in this learned part-based embedding (Sec.3.2). This coarse completion is then further refined by a simple yet effective optimization over the learned latent space to minimize the distance between control targets and completed motions, leading to accurate and efficient trajectory control (Sec. 3.3). An overview of our framework is shown in Fig. 2.

#### 3.1. Part-based VQ-VAE

Learning discrete representations has proven effective in motion synthesis tasks in recent works [14, 58]. These methods use VQ-VAE [49] for discretizing a continuous motion space, facilitating the capturing of complex motion patterns in a structured manner and enhancing the generation of coherent and high-fidelity motion sequences. However, previous works [14, 58] encode the full body into the discrete motion space, treating the human body as a holistic entity. Instead, in this paper, we propose to learn a better representation of part-based motion priors, which have demonstrated impressive results for computer character animation [4, 18, 34, 44]. To do so, we propose a part-based VQ-VAE that disaggregates the motion at the human body-part level for learning a well-structured and compact latent space.

Specifically, we first divide all the joints into six joint groups. The first five groups contain the end-effectors and related joints: *Head*, *Left arm*, *Right arm*, *Left leg*, and *Right leg*. The sixth group comprises the *Root* joint only. The input to our VQ-VAE is the ground-truth full body motion  $\mathbf{J} \in \mathbb{R}^{T \times M}$  of the character. By reorganizing the features into these six groups, we can arrange the input  $\mathbf{J}$  to be:  $\mathbf{J} = [\mathbf{J}_{Head}, \mathbf{J}_{Larm}, \mathbf{J}_{Rarm}, \mathbf{J}_{Lleg}, \mathbf{J}_{Rleg}, \mathbf{J}_{Root}]$ .

In our model, we have a unique encoder  $\mathcal{E}_k$  for each joint group  $\mathbf{J}_k$  where  $k \in \{Head, Lhand, Rhand, Lfoot, Rfoot, Root\}$  to learn a codebook  $\mathbf{C}_k$  for each joint group separately. Here,  $\mathbf{C}_k = \{\mathbf{c}_{i,k}\}_{i=1}^{|C|}$  and  $\mathbf{c}_{i,k} \in \mathbb{R}^d$ , where  $|C|$




Figure 2. Overview of TLControl framework: At training stage I, we train the part-based VQ-VAE in 3.1 for reconstructing human motions. In training stage II, the decoder of the part-based VQ-VAE is frozen and we train the masked trajectory transformer (MTT) in 3.2 for predicting code indices from control inputs. Finally, at test time, the MTT receives text description and partial control trajectories to predict an initial VQ-VAE quantized code seed, which is refined by run-time optimization as in 3.3 before decoding with the VQ-VAE into full body motions.

represents the number of discrete codes within the codebook and  $d$  is the dimension of the code. When encoding an input motion  $\mathbf{J}_k$  into its latent representation  $\mathbf{Q}_k = \mathcal{E}_k(\mathbf{J}_k)$ ,  $\mathbf{Q}_k = \{\mathbf{q}_{t,k}\}_{t=1}^{T/s}$  is of a reduced time length, where  $T$  is the total number of frames and  $s$  is a temporal downsampling factor, and  $\mathbf{q}_{t,k} \in \mathbb{R}^d$  is the feature for each time step  $t$ .

Subsequently, we quantize  $\mathbf{Q}_k$  using  $\mathbf{C}_k$ . For each time step feature  $\mathbf{q}_{t,k}$ , quantization is achieved by locating the closest code  $c_{i,k}$  in the codebook by

$$\hat{\mathbf{q}}_{t,k} = \arg \min_{c_{i,k} \in \mathbf{C}_k} \|\mathbf{c}_{i,k} - \mathbf{q}_{t,k}\|_2, \quad (1)$$

where the quantized feature  $\hat{\mathbf{q}}_{t,k}$  is reassembled in their original time order to form  $\hat{\mathbf{Q}}_k$ .

The decoder  $\mathcal{D}$  of our part-based VQ-VAE is designed to consider the features of all groups comprehensively to reconstruct the full body motion. Hence the input to  $\mathcal{D}$  is the concatenation  $\hat{\mathbf{Q}} = \bigoplus_k \hat{\mathbf{Q}}_k$  of quantized features from all groups. This enables the decoder to reconstruct the full body motion as  $\hat{\mathbf{J}} = \mathcal{D}(\hat{\mathbf{Q}})$ .

For training our part-based VQ-VAE, we adopt standard loss terms including quantization, commitment and reconstruction losses:

$$\mathcal{L} = \sum_k (\beta \|\text{sg}[\hat{\mathbf{Q}}_k] - \mathbf{Q}_k\|_2 + \|\hat{\mathbf{Q}}_k - \text{sg}[\mathbf{Q}_k]\|_2) + \|\mathbf{J} - \hat{\mathbf{J}}\|_2 \quad (2)$$

where  $\beta$  represents a weight hyper-parameter, while  $\text{sg}$  denotes the stop-gradient operator. Also, suggested by previous works in this field [58], we apply code reset and the

exponential moving average [46] to prevent codebook collapse during training.

### 3.2. Masked Trajectory Transformer

Having trained this part-based discrete embedding space, we now train a feed-forward transformer network to produce an initial coarse guess  $\hat{\mathbf{Q}}_0$  of the full-body motion embedded in this space, given the specification  $(\mathbf{R}', \mathbf{L})$ .

Specifically, our Masked Trajectory Transformer (MTT) takes CLIP embeddings of the text prompts, as in prior work [48, 58]. For trajectories, we adopt the same joint partition in 3.1 and treat the trajectories of the five end-effectors and the root joint as the ground truth “full” trajectories  $\mathbf{R}$  for training. Recall that in our application settings, we would like to enable users to provide rough partial sketches of the desired trajectories to control motion generation. To simulate such specifications when training the MTT, we apply two masking strategies to  $\mathbf{R}$  to obtain the masked control trajectories  $\mathbf{R}'$ :

▷ **Continuous Trajectory Masking.** Rather than treating each waypoint’s masking as an independent event, we consider the influence of adjacent points, which simulates the common scenario where a user’s drawn trajectory might have breaks or be sketched in segments. Our method involves first determining a proportion of the trajectory points that need to be masked. Then, we randomly select segments of varying lengths to mask, ensuring that the total number of masked points matches the predetermined proportion.



▷ **Joint-Level Masking.** Recognizing that users may not always provide trajectories for all six joints, we include joint-level masking. This approach simulates scenarios where only a subset of joint trajectories are specified. Specifically, we randomly select a number of joints and mask all corresponding trajectory points during training.

These masking strategies enable the network to learn the overall relationship among human motion, joint trajectories, and textual information. This approach strengthens the model’s capability to handle real user specifications, particularly when detailed and complete joint trajectories are unavailable.

We train the MTT encoder  $\mathcal{T}$  to encode user-provided observations into a compressed latent space  $\mathbf{Z} = \mathcal{T}(\mathbf{R}', L)$ . Then, we employ a decoders on  $\mathbf{Z}$ , which focuses on inferring a list of code indices  $\hat{I}$  in the codebooks from  $\mathbf{Z}$ , where the codes of these indices should form the quantized latent codes  $\hat{\mathbf{Q}}_0$  that can be decoded to corresponding motion sequences through our decoder. The loss function for training the MTT is:

$$\mathcal{L} = \mathbb{E}_{\hat{I} \sim P(I)} [-\log P(\hat{I} | R', L)] + \|\mathbf{D}(\hat{\mathbf{Q}}_0) - \mathbf{J}\|_2 \quad (3)$$

### 3.3. Test-Time Optimization

At test time, the MTT takes a partially specified control trajectory  $\mathbf{R}'$  to produce an initial coarse feed-forward prediction  $\hat{\mathbf{Q}}_0$  in the quantized latent code space as above. We then use test-time optimization over the learned latent space to refine this initial coarse prediction. Specifically, we solve:

$$\hat{\mathbf{Q}} = \arg \min_{\mathbf{Q}} \sum_j \|\mathcal{P}_j(\mathcal{D}(\mathbf{Q})) - \mathbf{R}'_j\|_2 \quad (4)$$

where  $j$  indexes the provided control joints, and  $\mathcal{P}_j$  is a projection function that converts the human pose representation into the global joint positions of joint  $j$  which are in the same frame as the waypoints in  $\mathbf{R}'$ . By applying  $\mathcal{D}(\hat{\mathbf{Q}})$ , we are able to achieve the final body motion controlled by the given trajectories. The optimization is performed using search initialized with the coarse MTT output prediction  $\hat{\mathbf{Q}}_0$ .

Since our body-part level motion embedding effectively captures the motion distribution in a compact space, this simple yet effective optimization framework produces high-quality human motion with the desired controllability while maintaining semantic coherence with the textual description. In contrast to the closest prior work OmniControl [54] where limited control specifications are coupled with the network during training, this framework allows for flexible controls, such as simultaneously controlling multiple joints at different time steps; See Sec. 4.1. Moreover, it is much more computationally efficient compared with SOTA methods, which we investigate in Sec. 4.2.

## 4. Experiment

We conduct extensive experiments of our method and compare it with state-of-the-art methods in terms of motion quality, controllability, accuracy, and efficiency. We further investigate the key design choices by conducting ablation studies.

**Datasets** Our study utilizes two prominent datasets for tasks involving text-based motion generation: KIT Motion Language (KIT-ML) [35], and HumanML3D [11]. The KIT-ML dataset encompasses 3,911 unique human motion sequences and 6,278 individual text annotations, with a frame rate of 12.5 FPS for the motion sequences. On the other hand, HumanML3D offers a more extensive collection of human motions, featuring 14,616 unique motion capture data paired with 44,970 textual descriptions. For uniformity, all motion sequences in both KIT-ML and HumanML3D are padded to a length of 196 frames.

**Metrics** Following the experiments setting in [54], we evaluate the fidelity of our result using Fréchet Inception Distance (FID), R-Precision and Diversity. **FID** measures the quality of motion created by generative techniques. It calculates the disparity between actual and produced distributions within the feature space of a pre-trained model as being used in [11]. **R-Precision** measures a motion sequence against 32 text narratives, one of which is accurate and the remaining 31 are random descriptions. It involves computing the Euclidean distances between the embeddings of the motion and the textual narratives. The accuracy of retrieving the correct text from the motion is then reported based on the top-3 matches. **Diversity** is evaluated by randomly pairing all generated sequences from the test texts. The average cumulative difference within each pair is then calculated to determine diversity.

We also include an analysis of several metrics to evaluate control accuracy. Following [7, 54], we adopt metrics including Trajectory error, Location error, and Average error. These metrics evaluate the 3D control accuracy of joint positions in keyframes when generating motions with trajectory controls. **Trajectory Error** quantifies the proportion of trajectories that is unsuccessful, characterized by any keyframe’s joint location deviating beyond a set threshold. **Location Error** calculates the percentage of keyframe locations not attained within a specified proximity limit. **Average Error** is computed as the mean Euclidean distance between the joint positions in the generated motion and the given control trajectories at each keyframe motion step.

**Implementation Details** In the part-based VQ-VAE, the settings of the encoder and the decoder module are based on T2M-GPT [58] with the downsampling  $s$  factor set to 4. For

Method	Control Joint	FID↓	R-precision↑ (Top-3)	Diversity→	Traj. Err.↓ (50 cm, %)	Loc. Err.↓ (50 cm, %)	Avg. Err. (cm)↓
Real	-	0.002	0.797	9.503	0.00	0.00	0.00
MDM	Pelvis	0.698	0.602	9.197	40.22	30.76	59.59
PriorMDM		0.475	0.583	9.156	34.57	21.32	44.17
GMD		0.576	0.665	9.206	9.31	3.21	14.39
OmniControl		<b>0.218</b>	0.687	9.422	3.87	0.96	3.38
<b>Ours</b>		<b>0.271</b>	<b>0.779</b>	<b>9.569</b>	<b>0.00</b>	<b>0.00</b>	<b>1.08</b>
OmniControl	Head	0.335	0.696	<b>9.480</b>	4.22	0.79	3.49
<b>Ours</b>		<b>0.279</b>	<b>0.778</b>	9.606	<b>0.00</b>	<b>0.00</b>	<b>1.10</b>
OmniControl	Left Hand	0.304	0.680	<b>9.436</b>	8.01	1.34	5.29
<b>Ours</b>		<b>0.135</b>	<b>0.789</b>	9.757	<b>0.00</b>	<b>0.00</b>	<b>1.08</b>
OmniControl	Right Hand	0.299	0.692	<b>9.519</b>	8.13	1.27	5.19
<b>Ours</b>		<b>0.137</b>	<b>0.787</b>	9.734	<b>0.00</b>	<b>0.00</b>	<b>1.09</b>
OmniControl	Left Foot	<b>0.280</b>	0.696	<b>9.553</b>	5.94	0.94	3.14
<b>Ours</b>		0.368	<b>0.768</b>	9.774	<b>0.00</b>	<b>0.00</b>	<b>1.14</b>
OmniControl	Right Foot	<b>0.319</b>	0.701	<b>9.481</b>	6.66	1.20	3.34
<b>Ours</b>		0.361	<b>0.775</b>	9.778	<b>0.00</b>	<b>0.00</b>	<b>1.16</b>
OmniControl	All Joints above	2.614	0.606	8.594	75.59	12.30	23.67
<b>Ours</b>		<b>0.032</b>	<b>0.794</b>	<b>9.750</b>	<b>0.00</b>	<b>0.00</b>	<b>1.57</b>

Table 1. Quantitative results of comparison with state-of-the-art methods on Humanml3D test set. The best scores are highlighted in red.

Method	Control Joint	FID ↓	R-precision ↑ (Top-3)	Diversity →	Traj. Err. ↓ (50 cm, %)	Loc. Err. ↓ (50 cm, %)	Avg. Err. (cm) ↓
Real	-	0.031	0.779	11.08	0.000	0.000	0.000
PriorMDM	Pelvis	0.851	0.397	10.518	33.10	14.00	23.05
GMD		1.565	0.382	9.664	54.43	30.03	40.70
OmiControl		0.702	0.397	<b>10.927</b>	11.05	3.37	7.59
<b>Ours</b>		<b>0.432</b>	<b>0.757</b>	10.723	<b>0.28</b>	<b>0.11</b>	<b>2.76</b>
OmiControl		Average	0.788	0.379	<b>10.841</b>	14.33	3.68
<b>Ours</b>		<b>0.487</b>	<b>0.751</b>	10.716	<b>0.52</b>	<b>0.15</b>	<b>2.98</b>

Table 2. Quantitative results of comparison with state-of-the-art methods on KIT test set. The best scores are highlighted in red.

the masked trajectory transformer, we use a frozen CLIP-ViT-B/32 [37] model for pre-processing the text prompt  $L$ , and then a 4-layer transformer processes the text information and control trajectories  $R'$  into the latent  $Z$ , followed by a 3-layer transformer decodes  $Z$  to the codes indices. The embedding dimension of these two stage transformers are 512 and 256 respectively. The networks are trained using AdamW optimizer with learning rate decays from  $1e-4$  to  $1e-7$  with a batch size of 64. For the run-time optimization step, we use Limited-Memory BFGS [26] as the optimization algorithm with learning rate set to 0.1. The accuracy criteria is set to  $1E-6$  and we provide an ablation study of the criteria setting in 4.3.4.

In this experiments section, we use a machine with RTX4090 GPU for running the tests except in 4.2 for a fair comparison. The testing batch size is 32 following previous works using HumanML3D [11] and KIT-ML [35] dataset.

#### 4.1. Controllable Human Motion Generation

We compare our method in controllable human motion generation with the current SOTA methods. When comparing with MDM [48], PriorMDM [39], and GMD [22], we focus on controlling the pelvis only for fair comparisons. We further compare with Omnicontrol [54] with more settings of different combinations of control joints. The results are shown in 1 and 2. For the single joint controlling task, our method achieves better scores in R-precision while

having comparable results in other fidelity metrics. However, our approach consistently surpasses existing methodologies in the metrics measuring control accuracy. On the Humanml3D test set, the results generated by our approach steadily remained within a 50 cm range of the intended control signals, showcasing remarkable accuracy and stability. We also demonstrate enhanced control accuracy on KIT dataset for using the pelvis only and the average of controlling the joints individually, where we continues to outperform other methods in accuracy.

In the task that utilize all six joints for controlling, Omni-control struggles with multi-joint controls. In contrast, our method can handle multi-joint control effectively, producing highly realistic outcomes with such abundant control signals. Visual comparison is shown in Fig 4.

Methods	Ours	OmniControl	MDM	GMD
Time (s/frame) ↓	<b>0.016</b>	0.619	0.200	0.561
FPS ↑	<b>61.7</b>	1.6	5	1.8

Table 3. Runtime Efficiency of different methods.

#### 4.2. Run-time Performance

We further compared our method with the SOTA methods in terms of run-time performance. We run our method on a machine with an RTX2080Ti GPU and record the average process time for a sample using one control joint. As shown




Figure 3. Qualitative results of our method. Figure 3a and Figure 3d demonstrate that our method enables separate controls using language and joint-level trajectories. Figure 3b and Figure 3c showcase the capability of our method to manage multi-joint control simultaneously. Please refer to our supplementary for more qualitative results




Figure 4. Qualitative comparison results. We compare our results with Omnicontrol[54]. Our results are shown in Yellow, while those of Omnicontrol are depicted in Green. Please refer to our supplementary for more comparison results

in Tab. 3, due to our structured and compact latent space and efficient optimization framework, our method consistently outperforms all existing methods with a significant margin.

### 4.3. Ablation Study

#### 4.3.1 Motion Latent Space

We investigate the learned part-based VQ-VAE for motion embedding. In Fig. 5, we randomly choose 10 test cases from Humanml3D test set, and process them with MMT for achieving the corresponding features. We present the re-




Figure 5. t-SNE visualization of learned motion latent space. Each color stands for a textual description of the motion. We randomly select ten sentences for visualization.

sult of using t-SNE [50] for visualizing these features, offering insights into the structure of the acquired motion latent space. Our part-based VQ-VAE contributes a compact motion latent space while preserving the semantics described by language inputs, as evidenced in Fig. 5.

Next, we visualize the learned body part-level motion prior. As shown in Fig. 6, our framework for motion embedding learns a well-structured latent space. This results in generating motions that conform to the language description while capturing a diverse distribution of motions under each text condition.

#### 4.3.2 Trajectory Incompleteness

We investigate the performance of our method under different levels of trajectory incompleteness, for which we simulate various incompleteness scenarios by adjusting the ratio of random masking. As revealed in Fig. 7, our model’s performance degrades gracefully even as the control trajectories become more incomplete. As a reference, note that even with full joint trajectories, the performance of previous methods [54] in terms of FID and R-precision are 2.614 and 0.606.

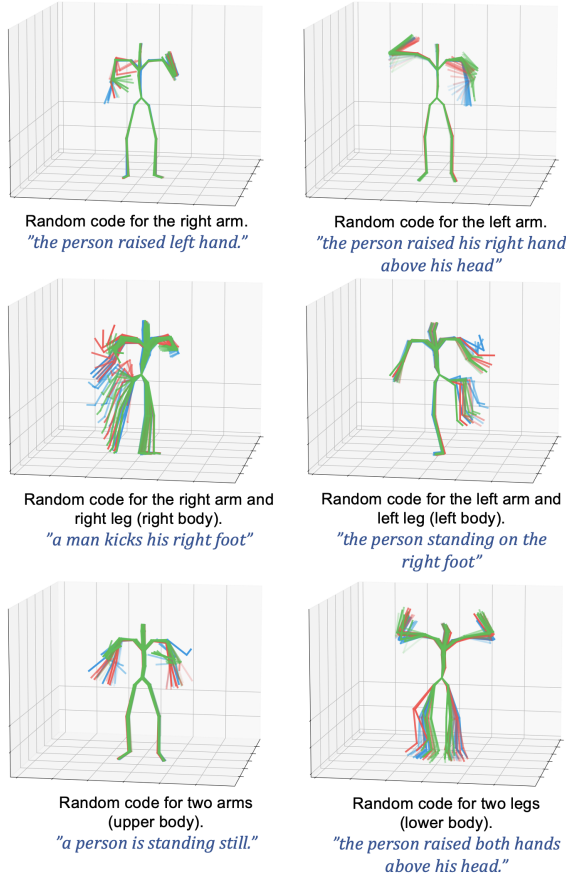


Figure 6. Motion variations controlled by language (text in blue) and randomly sampled body part level latent codes. Three sample motions in red, blue, and green are shown.




Figure 7. Influence of different trajectory incompleteness. We simulate the incompleteness by applying random masking. The left vertical axis represents the FID metric, while the right vertical axis indicates the R-precision metric.

Optimize Acc Setting	1E-4	1E-5	1E-6	1E-7	1E-8
Runtime (per batch, s)	3.48	11.89	28.86	61.43	127.90
Joint Name	Avg. Err (cm)				
Pelvis	4.92	2.03	1.20	0.71	0.45
Head	5.87	2.77	1.46	0.77	0.48
Left Hand	8.41	4.14	1.83	0.89	0.56
Right Hand	8.45	4.08	1.80	0.88	0.56
Left Foot	6.32	3.05	1.57	0.92	0.61
Right Foot	6.37	3.04	1.55	0.91	0.60
Avg.	6.72	3.19	1.57	0.85	0.54

Table 4. Average Error for Different Joints v.s. Accuracy criteria during optimization when controlling all the joints.




Figure 8. Running time statistics of our optimization when applying different controlled joints. “Upper Body” includes the joints of the hands and the head joint. “Lower Body” includes the joints of two feet and the joint of the pelvis.

### 4.3.3 Runtime of different joint control strategy

We present the runtime associated with different joint control strategies. In Figure 8, we record the average runtime for processing a batch under the condition of using different combinations of joint trajectories. Due to the complexity of hand motions, optimization that exclusively using both hands trajectories is the most time-consuming.

### 4.3.4 Optimization Scheme

Finally, we present the average error in controlling all six joints of our method under varying accuracy criteria during optimization. These errors for each joint are recorded in Table 4, utilizing the Humanml3D dataset test set. As indicated in Table 4, higher optimization accuracy settings result in lower accuracy errors but require more optimization time. In this paper, we consistently set the optimization accuracy to 1E-6.

## 5. Conclusion

We introduce TLControl, a method for controllable human motion generation using a combination of joint trajectory and language inputs. TLControl leverages a disentangled latent space for diverse human motion, allowing for effective optimization and producing high-fidelity motion aligned with both language descriptions and specified trajectories. The proposed Masked Trajectories Transformer captures the correspondence between trajectories and textual information, leading to a well-structured compact latent space. We further introduce an efficient optimization framework for joint-level motion editing while maintaining motion semantics encoded by the language, thus enabling interactive generation and modification of high-quality motion within a short run-time. Extensive experiments validate the effectiveness of our framework, highlighting its capability to enable users to quickly generate and modify high-fidelity human motion interactively.



## References

- [1] Chaitanya Ahuja and Louis-Philippe Morency. Language2pose: Natural language grounded pose forecasting. In *2019 International Conference on 3D Vision (3DV)*, pages 719–728. IEEE, 2019. 2
- [2] Simon Alexanderson, Rajmund Nagy, Jonas Beskow, and Gustav Eje Henter. Listen, denoise, action! audio-driven motion synthesis with diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–20, 2023. 2, 3
- [3] Nikos Athanasiou, Mathis Petrovich, Michael J Black, and Gül Varol. Teach: Temporal action composition for 3d humans. In *2022 International Conference on 3D Vision (3DV)*, pages 414–423. IEEE, 2022. 2
- [4] Jinseok Bae, Jungdam Won, Donggeun Lim, Cheol-Hui Min, and Young Min Kim. Pmp: Learning to physically interact with environments using part-wise motion priors. *arXiv preprint arXiv:2305.03249*, 2023. 3
- [5] Raphael Bensedou, Shir Gur, Nitsan Blau, and Lior Wolf. Neural inverse kinematic. In *International Conference on Machine Learning*, pages 1787–1797. PMLR, 2022. 2
- [6] Xin Chen, Zhuo Su, Lingbo Yang, Pei Cheng, Lan Xu, Bin Fu, and Gang Yu. Learning variational motion prior for video-based motion capture. *arXiv preprint arXiv:2210.15134*, 2022. 2
- [7] 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, 3, 5
- [8] Zhiyang Dou, Xuelin Chen, Qingnan Fan, Taku Komura, and Wenping Wang. C-ase: Learning conditional adversarial skill embeddings for physics-based characters. *arXiv preprint arXiv:2309.11351*, 2023. 2, 3
- [9] Yinglin Duan, Tianyang Shi, Zhengxia Zou, Yanan Lin, Zhehui Qian, Bohan Zhang, and Yi Yuan. Single-shot motion completion with transformer. *arXiv preprint arXiv:2103.00776*, 2021. 2
- [10] 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. 2
- [11] 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 (CVPR)*, pages 5152–5161, 2022. 5, 6, 1, 2
- [12] 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 (CVPR)*, pages 5152–5161, 2022. 2
- [13] 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 *ECCV*, 2022. 2
- [14] 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. 3
- [15] Félix G Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. *ACM Transactions on Graphics (TOG)*, 39(4):60–1, 2020. 2
- [16] Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael J Black. Stochastic scene-aware motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11374–11384, 2021. 3
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 2
- [18] Deok-Kyeong Jang, Soomin Park, and Sung-Hee Lee. Motion puzzle: Arbitrary motion style transfer by body part. *ACM Transactions on Graphics (TOG)*, 41(3):1–16, 2022. 3
- [19] Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. Motion2pt: Human motion as a foreign language. *arXiv preprint arXiv:2306.14795*, 2023. 3
- [20] Yifeng Jiang, Jungdam Won, Yuting Ye, and C Karen Liu. Drop: Dynamics responses from human motion prior and projective dynamics. *arXiv preprint arXiv:2309.13742*, 2023. 3
- [21] Jordan Juravsky, Yunrong Guo, Sanja Fidler, and Xue Bin Peng. Padl: Language-directed physics-based character control. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. 2
- [22] Korrawe Karunratanakul, Konpat Preechakul, Supasorn Suwajanakorn, and Siyu Tang. Guided motion diffusion for controllable human motion synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2151–2162, 2023. 2, 3, 6, 1
- [23] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. Flame: Free-form language-based motion synthesis & editing. *arXiv preprint arXiv:2209.00349*, 2022. 2, 3
- [24] Buyu Li, Yongchi Zhao, Shi Zhelun, and Lu Sheng. Danceformer: Music conditioned 3d dance generation with parametric motion transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1272–1279, 2022. 2
- [25] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13401–13412, 2021. 2
- [26] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989. 6, 3
- [27] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)*, 37(4):1–14, 2018. 3
- [28] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Trans. Graph.*, 40(4), 2021. 2

- [29] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics (TOG)*, 40(4):1–20, 2021. 2, 3
- [30] Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. Ase: Large-scale reusable adversarial skill embeddings for physically simulated characters. *ACM Transactions On Graphics (TOG)*, 41(4):1–17, 2022. 2, 3
- [31] Mathis Petrovich, Michael J. Black, and Gül Varol. Action-conditioned 3D human motion synthesis with transformer VAE. In *International Conference on Computer Vision (ICCV)*, 2021. 2
- [32] Mathis Petrovich, Michael J. Black, and Gül Varol. TEMOS: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision (ECCV)*, 2022. 2, 3
- [33] Mathis Petrovich, Michael J Black, and Gül Varol. Tmr: Text-to-motion retrieval using contrastive 3d human motion synthesis. *arXiv preprint arXiv:2305.00976*, 2023. 2
- [34] Huaijin Pi, Sida Peng, Minghui Yang, Xiaowei Zhou, and Hujun Bao. Hierarchical generation of human-object interactions with diffusion probabilistic models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15061–15073, 2023. 3
- [35] Matthias Plappert, Christian Mandery, and Tamim Asfour. The kit motion-language dataset. *Big data*, 4(4):236–252, 2016. 5, 6, 2
- [36] Sigal Raab, Inbal Leibovitch, Peizhuo Li, Kfir Aberman, Olga Sorkine-Hornung, and Daniel Cohen-Or. Modi: Unconditional motion synthesis from diverse data. *arXiv preprint arXiv:2206.08010*, 2022. 2
- [37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021. 6, 2
- [38] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J Guibas. Humor: 3d human motion model for robust pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11488–11499, 2021. 2
- [39] Yonatan Shafir, Guy Tevet, Roy Kapon, and Amit H. Bermano. Human motion diffusion as a generative prior, 2023. 2, 3, 6
- [40] Mingyi Shi, Sebastian Starke, Yuting Ye, Taku Komura, and Jungdam Won. Phasemp: Robust 3d pose estimation via phase-conditioned human motion prior. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14725–14737, 2023. 2
- [41] Yi Shi, Jingbo Wang, Xuekun Jiang, and Bo Dai. Controllable motion diffusion model. *arXiv preprint arXiv:2306.00416*, 2023. 2
- [42] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 3
- [43] Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. Neural state machine for character-scene interactions. *ACM Trans. Graph.*, 38(6):209–1, 2019. 2, 3
- [44] Sebastian Starke, Yiwei Zhao, Fabio Zinno, and Taku Komura. Neural animation layering for synthesizing martial arts movements. *ACM Transactions on Graphics (TOG)*, 40(4):1–16, 2021. 3
- [45] Sebastian Starke, Ian Mason, and Taku Komura. Deepphase: periodic autoencoders for learning motion phase manifolds. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. 2
- [46] James W Taylor. Exponential smoothing with a damped multiplicative trend. *International journal of Forecasting*, 19(4): 715–725, 2003. 4
- [47] Guy Tevet, Brian Gordon, Amir Hertz, Amit H Bermano, and Daniel Cohen-Or. Motionclip: Exposing human motion generation to clip space. *arXiv preprint arXiv:2203.08063*, 2022. 3
- [48] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 2, 3, 4, 6
- [49] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 3
- [50] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 7
- [51] Weilin Wan, Lei Yang, Lingjie Liu, Zhuoying Zhang, Ruixing Jia, Yi-King Choi, Jia Pan, Christian Theobalt, Taku Komura, and Wenping Wang. Learn to predict how humans manipulate large-sized objects from interactive motions. *IEEE Robotics and Automation Letters*, 7(2):4702–4709, 2022. 2
- [52] Jingbo Wang, Sijie Yan, Bo Dai, and Dahua Lin. Scene-aware generative network for human motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12206–12215, 2021. 3
- [53] Jingbo Wang, Ye Yuan, Zhengyi Luo, Kevin Xie, Dahua Lin, Umar Iqbal, Sanja Fidler, and Sameh Khamis. Learning human dynamics in autonomous driving scenarios. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20796–20806, 2023. 3
- [54] Yiming Xie, Varun Jampani, Lei Zhong, Deqing Sun, and Huaizu Jiang. Omnicontrol: Control any joint at any time for human motion generation. *arXiv preprint arXiv:2310.08580*, 2023. 2, 3, 5, 6, 7, 1
- [55] Sijie Yan, Zhizhong Li, Yuanjun Xiong, Huahan Yan, and Dahua Lin. Convolutional sequence generation for skeleton-based action synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4394–4402, 2019. 2
- [56] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. *arXiv preprint arXiv:2212.02500*, 2022. 3
- [57] He Zhang, Sebastian Starke, Taku Komura, and Jun Saito. Mode-adaptive neural networks for quadruped motion control. *ACM Transactions on Graphics (TOG)*, 37(4):1–11, 2018. 3

- [58] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Shaoli Huang, Yong Zhang, Hongwei Zhao, Hongtao Lu, and Xi Shen. T2m-gpt: Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [2](#), [3](#), [4](#), [5](#), [1](#)
- [59] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. [2](#), [3](#)
- [60] Yan Zhang, Michael J Black, and Siyu Tang. Perpetual motion: Generating unbounded human motion. *arXiv preprint arXiv:2007.13886*, 2020. [2](#)
- [61] Rui Zhao, Hui Su, and Qiang Ji. Bayesian adversarial human motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6225–6234, 2020. [2](#)

# TLControl: Trajectory and Language Control for Human Motion Synthesis

## Supplementary Material

This supplementary material covers: the ablation study of our part-based VQ-VQE embedding (Sec. 1); an experiments of our multi-joint control results (Sec. 2); more qualitative comparisons with the SOTA methods (Sec. 3); more implementation details (Sec. 4); and a discussion of limitations and future work (Sec. 5).

We also include a video showing more visual results in our supplementary. We highly encourage our readers to view the supplementary video for our qualitative comparison and results.

### 1. Ablation Study on Part-based VQ-VAE

To validate the effectiveness of the proposed part-based VQ-VAE, we compared our network structure with a counterpart VQ-VAE without body-part level disentangling, which we referred to as **Unsplit VQ-VAE**. To guarantee a fair comparison, both models were configured with an identical feature dimension that represents the full-body motion and is utilized for the runtime optimization. As summarized in Tab. 1, our proposed part-based VQ-VAE demonstrated superior performance in generating controllable motion, showing its efficacy in preserving language semantics alongside facilitating flexible trajectory control. This is attributed to our disentanglement design of our VQ-VAE that allows an efficient representation of motion dynamics at the body-part level. Moreover, our part-based VQ-VAE exhibits the capability to encode motions into a more compact latent space. This compactness not only enhances the model’s accuracy in handling complex motion but also significantly reduces processing time during the runtime optimization, making it more suitable for real-time applications where rapid motion generation and adaptation are crucial.

Figure 1 illustrates the comparative analysis of the average processing times for motion synthesis using our part-based embedding versus the unsplit embedding under various control strategies. On average, our embedding design demonstrates a significant reduction in processing time, specifically a 27.3% decrease compared to the unsplit VQ-VAE embedding. This efficiency gain is primarily attributed to the optimized structure of our part-based embedding, which streamlines the synthesis process by effectively handling the complexities of motion data.

### 2. Multi-Joint Control Results

In this section, we present the performance metrics of our method when applied to various groups of joints. This test is performed in the test set of HumanML3D dataset [11], and the results are shown in Tab. 2. Our approach demonstrates




Figure 1. Per batch running time statistics of our embedding comparing to the unsplit embedding. “Upper Body” includes the joints of the hands and the head joint. “Lower Body” includes the joints of two feet and the joint of the pelvis.

notable efficiency and accuracy in managing multiple joints simultaneously. As we integrate additional control trajectories, the motion representation becomes increasingly detailed. This enhancement allows our method to generate more realistic results, showcasing its adaptability and precision in complex joint operations.

### 3. More Qualitative Comparisons on Controllable Motion Synthesis

We present more detailed qualitative results in this section, comparing our method with the SOTA methods in terms of controllable human motion synthesis: GMD [22] and Omnicontrol [54].

Given that GMD only concentrates on the trajectory control of the root joint, we demonstrate cases where only the root path is assigned, as depicted in Figure 2. Omnicontrol allows control over multiple joints simultaneously. To illustrate this, we compare tasks that involve controlling different joints concurrently with Omnicontrol, with the results displayed in Figure 3. Our method demonstrates superior accuracy in adhering to the control trajectories and achieves this with a notably reduced runtime. Our approach exhibits enhanced adaptability and precision in complex motion patterns, outperforming the compared methods in dynamic motion scenarios where precise joint control is critical.

### 4. Implementation Details

In this section, we provide a more detailed explanation of the implementation of our method. Following [58], our encoder and the decoder consist of standard 1D convolutional layers, residual blocks [17], and ReLU activation functions.



Method	Control Joint	FID↓	R-precision↑ (Top-3)	Diversity→	Avg. Err. (cm)↓	Runtime Per Batch (s)↓
Real	-	0.002	0.797	9.503	0.00	-
Unsplit VQ-VAE	Pelvis	0.334	0.767	9.681	1.52	22.81
<b>Ours</b>		<b>0.271</b>	<b>0.779</b>	<b>9.569</b>	<b>1.08</b>	<b>17.67</b>
Unsplit VQ-VAE	Head	0.325	0.768	9.687	1.61	26.61
<b>Ours</b>		<b>0.279</b>	<b>0.778</b>	<b>9.606</b>	<b>1.10</b>	<b>19.80</b>
Unsplit VQ-VAE	Left Hand	0.156	0.786	9.776	1.57	30.22
<b>Ours</b>		<b>0.135</b>	<b>0.789</b>	<b>9.757</b>	<b>1.08</b>	<b>25.35</b>
Unsplit VQ-VAE	Right Hand	0.145	0.783	9.793	1.60	29.89
<b>Ours</b>		<b>0.137</b>	<b>0.787</b>	<b>9.734</b>	<b>1.09</b>	<b>25.17</b>
Unsplit VQ-VAE	Left Foot	0.433	0.753	9.839	1.59	25.29
<b>Ours</b>		<b>0.368</b>	<b>0.768</b>	<b>9.774</b>	<b>1.14</b>	<b>20.90</b>
Unsplit VQ-VAE	Right Foot	0.412	0.756	9.802	1.64	25.34
<b>Ours</b>		<b>0.361</b>	<b>0.775</b>	<b>9.778</b>	<b>1.16</b>	<b>21.15</b>
Unsplit VQ-VAE	All Joints above	0.066	0.790	9.780	2.51	38.84
<b>Ours</b>		<b>0.032</b>	<b>0.794</b>	<b>9.750</b>	<b>1.57</b>	<b>28.86</b>

Table 1. Quantitative results of comparison with different embedding design on Humanml3D test set. The best results are highlighted in red.




Figure 2. Qualitative comparison results with GMD [22] in the task of controlling root path. Our results are shown in Yellow, while the results of GMD are depicted in Green. Please refer to our supplementary video for more details of the comparison.

The settings of each encoder and the decoder of our part-based VQ-VAE are based on T2M-GPT [58] using standard

Control Joint	FID ↓	R-precision ↑ (Top-3)	Diversity →	Avg. Err. (cm) ↓
Both Hands	0.108	0.789	9.747	1.25
Upper Body	0.088	0.791	9.740	1.35
Both Feet	0.320	0.775	9.748	1.29
Lower Body	0.249	0.777	9.746	1.38

Table 2. Quantitative results of multi-joint control. “Upper Body” includes the joints of hands and the head joint. “Lower Body” includes the joints of two feet and the joint of pelvis.

1D convolutional layers, residual blocks [17], and ReLU activation functions.

For each joint encoder, we use a codebook of 126 code vectors and each code of dimension 126. Our input data first passes through an 1D convolutional layer, followed by a ReLU activation function. By setting the temporal downsampling rate to 4, the feature is processed by 2 sequences, where each sequence is a compound module that includes a 1D convolutional layer with a stride of 2, which downsamples temporal dimension of the data by a scale of 2, then followed by a residual block. The encoded feature is also normalized by the norm of itself before being quantized. By combining the quantized codes from all 6 joint encoders, we achieve a full-body latent code of size  $126 \times 6 = 756$ . The decoder has the same structure as the encoder, where the convolution with stride is replaced with nearest interpolation for temporal upsampling, and the full-body pose output dimension is 263 for Humanml3D [11] dataset and 251 for KIT [35] dataset.

In implementing the masked trajectory transformer, we employ a frozen CLIP-ViT-B/32 [37] model for the initial processing of the text prompt, which translates it into a 512-dimensional vector representative of language features. The joint trajectories are grouped with every 4 closest waypoints being bundled into a single temporal token. This bundling is reflective of our choice to use a downsampling rate of 4 within the part-based VQ-VAE framework. Each way-




Figure 3. Qualitative comparison results with Omnicontrol[54]. Our results are shown in **Yellow**, while the results of Omnicontrol are depicted in **Green**. Please refer to our supplementary video for more details of the comparison.

point group is then mapped onto a 512-dimensional feature space. A standard 4-layer transformer encoder is used to integrate the language and trajectory features into a cohesive latent space, where tokens comprise both the language feature vector and the feature representations of the trajectory at each temporal token. Following this, a 3-layer transformer encoder is designed to transform the latent space into a sequence of logits, which are used to determine the code indices. We then apply a softmax function followed by an argmax operation to select the most probable code indices, which serve as the initial full-body state for the VQ-decoder. For the masking strategies, each iteration within a batch has a 50% probability of using continuous trajectory masking or joint-level masking. The proportion of continuous trajectory masking is gradually increased, starting at 0% and steadily advancing to 75%.

During the run-time optimization stage, the Limited-Memory BFGS [26] method is utilized as the optimization technique. The learning rate is configured at 0.1, and we set the precision target at  $1E-6$ . We limit the process to a maximum of 1000 iterations, maintain an update history size of 200, and employ the ‘strong\_wolfe’ condition for the line search function.

## 5. Limitations and Future Work

Although TLControl achieves controllable motion synthesis that is faithfully aligned with language semantics and control trajectories with high efficiency, there are still some limitations. In particular, the absence of a physics-based simulation in the current version of TLControl may produce human motions that deviate from established physics rules.

Thus, one promising direction is to incorporate physics-based simulation [8, 27, 29, 30] into the generation process, as evidenced by [20, 56].

Furthermore, our approach facilitates detailed motion synthesis, enabling its application in the generation of motions responsive to scenes or objects. This can be achieved by defining control trajectories that align with the movements of specific objects or scenes. However, more effort can be made to realize scene-aware motion synthesis that accurately responds to dynamic changes, e.g., simulating walking on terrains with diverse elevations [16, 43, 52, 53, 57], which can be an intriguing avenue for future research.