

Generating Human Motion from Textual Descriptions with Discrete Representations

Jianrong Zhang^{1,3*}, Yangsong Zhang^{2,3*}, Xiaodong Cun³, Yong Zhang³, Hongwei Zhao¹
 Hongtao Lu², Xi Shen^{3,†}, Shan Ying³

*Equal contribution †Corresponding author

¹Jilin University

²Shanghai Jiao Tong University

³Tencent AI Lab

Abstract

In this work, we investigate a simple and must-known conditional generative framework based on Vector Quantised-Variational AutoEncoder (VQ-VAE) and Generative Pre-trained Transformer (GPT) for human motion generation from textual descriptions. We show that a simple CNN-based VQ-VAE with commonly used training recipes (EMA and Code Reset) allows us to obtain high-quality discrete representations. For GPT, we incorporate a simple corruption strategy during the training to alleviate training-testing discrepancy. Despite its simplicity, our T2M-GPT shows better performance than competitive approaches, including recent diffusion-based approaches. For example, on HumanML3D, which is currently the largest dataset, we achieve comparable performance on the consistency between text and generated motion (R-Precision), but with FID 0.116 largely outperforming MotionDiffuse of 0.630. Additionally, we conduct analyses on HumanML3D and observe that the dataset size is a limitation of our approach. Our work suggests that VQ-VAE still remains a competitive approach for human motion generation. Our implementation is available on the project page: <https://mael-zys.github.io/T2M-GPT/>.

1. Introduction

Generating motion from textual descriptions can be used in numerous applications in the game industry, film-making, and animating robots. For example, a typical way to access new motion in the game industry is to perform motion capture, which is expensive. Therefore automatically generating motion from textual descriptions, which allows producing meaningful motion data, could save time and be more economical.

Motion generation conditioned on natural language is challenging, as motion and text are from different modalities. The model is expected to learn precise mapping from

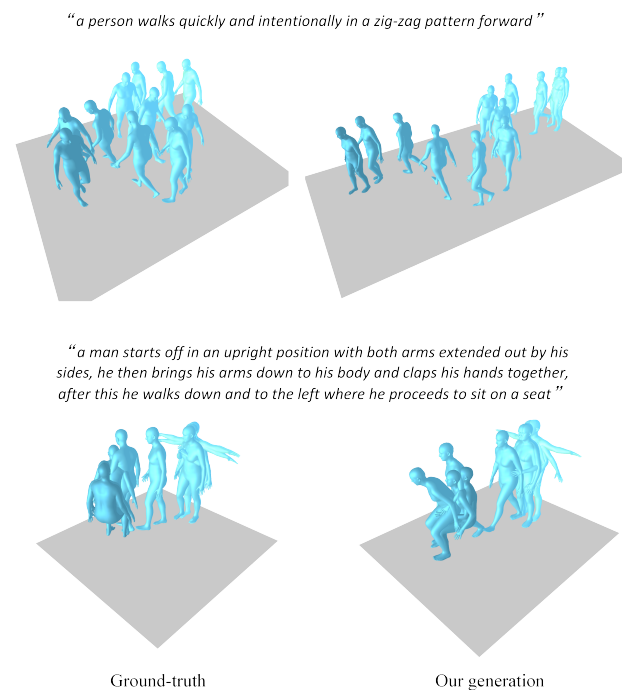


Figure 1. **Visual results on HumanML3D [22]**. Our approach is able to generate precise and high-quality human motion consistent with challenging text descriptions. More visual results are on the project page.

the language space to the motion space. To this end, many works propose to learn a joint embedding for language and motion using auto-encoders [3, 21, 65] and VAEs [52, 53]. MotionClip [65] aligns the motion space to CLIP [55] space. ACTOR [52] and TEMOES [53] propose transformer-based VAEs for action-to-motion and text-to-motion respectively. These works show promising performances with simple descriptions and are limited to producing high-quality motion when textual descriptions become long and complicated. Guo *et al.* [22] and TM2T [23] aim to generate motion

sequences with more challenging textual descriptions. However, both approaches are not straightforward, involve three stages for text-to-motion generation, and sometimes fail to generate high-quality motion consistent with the text (See Figure 4 and more visual results on the project page). Recently, diffusion-based models [32] have shown impressive results on image generation [60], which are then introduced to motion generation by MDM [66] and MotionDiffuse [74] and dominates text-to-motion generation task. However, we find that compared to classic approaches, such as VQ-VAE [68], the performance gain of the diffusion-based approaches [66, 74] might not be that significant. In this work, we are inspired by recent advances from learning the discrete representation for generation [5, 15, 16, 19, 45, 58, 68, 70] and investigate a simple and classic framework based on Vector Quantized Variational Autoencoders (VQ-VAE) [68] and Generative Pre-trained Transformer (GPT) [56, 69] for text-to-motion generation.

Precisely, we propose a two-stage method for motion generation from textual descriptions. In stage 1, we use a standard 1D convolutional network to map motion sequences to discrete code indices. In stage 2, a standard GPT-like model [56, 69] is learned to generate sequences of code indices from pre-trained text embedding. We find that the naive training of VQ-VAE [68] suffers from code collapse. One effective solution is to leverage two standard recipes during the training: *EMA* and *Code Reset*. We provide a full analysis of different quantization strategies. For GPT, the next token prediction brings inconsistency between the training and inference. We observe that simply corrupting sequences during the training alleviates this discrepancy. Moreover, throughout the evolution of image generation, the size of the dataset has played an important role. We further explore the impact of dataset size on the performance of our model. The empirical analysis suggests that the performance of our model can potentially be improved with larger datasets.

Despite its simplicity, our approach can generate high-quality motion sequences that are consistent with challenging text descriptions (Figure 1 and more on the project page). Empirically, we achieve comparable or even better performances than concurrent diffusion-based approaches MDM [66] and HumanDiffuse [74] on two widely used datasets: HumanML3D [22] and KIT-ML [54]. For example, on HumanML3D, which is currently the largest dataset, we achieve comparable performance on the consistency between text and generated motion (R-Precision), but with FID 0.116 largely outperforming MotionDiffuse of 0.630. We conduct comprehensive experiments to explore this area, and hope that these experiments and conclusions will contribute to future developments.

In summary, our contributions include:

- We present a simple yet effective approach for motion generation from textual descriptions. Our ap-

proach achieves state-of-the-art performance on HumanML3D [22] and KIT-ML [54] datasets.

- We show that GPT-like models incorporating discrete representations still remain a very competitive approach for motion generation.
- We provide a detailed analysis of the impact of quantization strategies and dataset size. We show that a larger dataset might still offer a promising prospect to the community.

Our implementation is available on the project page.

2. Related work

VQ-VAE. Vector Quantized Variational Autoencoders (VQ-VAE), which is a variant of VAE [36], is initially proposed in [68]. VQ-VAE is composed of an AutoEncoder architecture, which aims at learning reconstruction with discrete representations. Recently, VQ-VAE achieves promising performance on generative tasks across different modalities, which includes: image synthesis [19, 70], text-to-image generation [58], speech gesture generation [5], music generation [15, 16] etc. The success of VQ-VAE for generation might be attributed to its decoupling of learning the discrete representation and the prior. A naive training of VQ-VAE suffers from the codebook collapse, *i.e.*, only a number of codes are activated, which importantly limited the performances of the reconstruction as well as generation. To alleviate the problem, a number of techniques can be used during training, including stop-gradient along with some losses [68] to optimize the codebook, exponential moving average (EMA) for codebook update [70], reset inactivated codes during the training (Code Reset [70]), etc.

Human motion synthesis. Research on human motion synthesis has a long history [8]. One of the most active research fields is human motion prediction, which aims at predicting the future motion sequence based on past observed motion. Approaches mainly focus on efficiently and effectively fusing spatial and temporal information to generate deterministic future motion through different models: RNN [12, 20, 50, 51], GAN [9, 30], GCN [49], Attention [48] or even simply MLP [11, 17, 25]. Some approaches aim at generating diverse motion through VAE [4, 26, 73]. In addition to synthesizing motion conditioning on past motion, another related topic is generating motion “in-betweening” that takes both past and future poses and fills motion between them [18, 27, 28, 35, 64]. [51] considers the generation of locomotion sequences from a given trajectory for simple actions such as: walking and running. Motion can also be generated with music to produce 3D dance motion [6, 13, 37, 38, 40, 41]. For unconstrained generations, [72] generates a long sequence altogether by transforming from a sequence of latent




Figure 2. **Overview of our framework for text-driven motion generation.** It includes two modules: Motion VQ-VAE (Figure 2a) and T2M-GPT (Figure 2b). In T2M-GPT, an additional learnable *End* token is inserted to indicate the stop of the generation. During the inference, we first generate code indexes in an auto-regressive fashion and then obtain the motion using the decoder in Motion VQ-VAE.

vectors sampled from a Gaussian process. In graphics literature, many works focus on animator control. Holden *et al.* [34] learn a convolutional autoencoder to reconstruct motion, the learned latent representation can be used to synthesize and edit motion. [33] proposes phase functioned neural network to perform the control task. [63] uses a deep auto-regressive framework to scene interaction behaviors. Starke *et al.* [62] proposes to reconstruct motion through periodic features, the learned periodic embedding improves motion synthesis. Recently, inspired by SinGAN [61] for image synthesis, Li *et al.* [39] propose a generative model approach for motion synthesis from a single sequence.

Text-driven human motion generation. Text-driven human motion generation aims at generating 3D human motion from textual descriptions. Text2Action [2] trains an RNN-based model to generate motion conditioned on a short text. Language2Pose [3] employs a curriculum learning approach to learn a joint embedding space for both text and pose. The decoder can thus take text embedding to generate motion sequences. Ghost *et al.* [21] learn two manifold representations for the upper body and the lower body movements, which shows improved performance compared to Language2Pose [3]. Similarly, MotionCLIP [65] also tends to align text and motion embedding but proposes to utilize CLIP [55] as the text encoder and employ rendered images as extra supervision. It shows the ability to generate out-of-distribution motion and enable latent code editing. However, the generated motion sequences are not in high-quality and are without global translation. ACTOR [52] proposes a transformer-based VAE to generate motion in a non-autoregressive fashion from a pre-defined action class. TEMOS [53] extends the architecture of ACTOR [52] by introducing an additional text encoder and producing diverse motion sequences given text descriptions. TEMOS demonstrates its effect on KIT Motion-Language [54] with

mainly short sentences and suffers from out-of-distribution descriptions [53]. TEACH [7] further extends TEMOS to generate temporal motion compositions from a series of natural language descriptions. Recently, a large-scale dataset HumanML3D is proposed in [22]. Guo *et al.* [22] also propose to incorporate motion length prediction from text to produce motion with reasonable length. TM2T [23] considers text-to-motion and motion-to-text tasks. It also shows additional improvement can be obtained through jointly training both tasks. As concurrent works, diffusion-based models are introduced for text-to-motion generation by MDM [66] and MotionDiffuse [74]. In this work, we show that without any sophisticated designs, the classic VQ-VAE framework could achieve competitive or even better performance with a classical framework and some standard training recipes.

3. Method

Our goal is to generate high-quality motion that is consistent with text descriptions. The overall framework consists of two modules: Motion VQ-VAE and T2M-GPT, which is illustrated in Figure 2. The former learns a mapping between motion data and discrete code sequences, the latter generates code indices conditioned on the text description. With the decoder in Motion VQ-VAE, we are able to recover the motion from the code indices. In Section 3.1, we present the VQ-VAE module. The T2M-GPT is introduced in Section 3.2.

3.1. Motion VQ-VAE

VQ-VAE, proposed in [68], enables the model to learn discrete representations for generative models. Given a motion sequence $X = [x_1, x_2, \dots, x_T]$ with $x_t \in \mathbb{R}^d$, where T is the number of frames and d is the dimension of the motion, we aim to recover the motion sequence through an autoencoder and a learnable codebook containing K codes $C = \{c_k\}_{k=1}^K$ with $c_k \in \mathbb{R}^{d_c}$, where d_c is the dimension of




Figure 3. **Architecture of the motion VQ-VAE.** We use a standard CNN-based architecture with 1D convolution (*Conv1D*), residual block (*ResBlock*) and ReLU activation. ‘ L ’ denotes the number of residual blocks. We use convolution with stride 2 and nearest interpolation for temporal downsampling and upsampling.

codes. The overview of VQ-VAE is presented in Figure 2a. With encoder and decoder of the autoencoder denoted by E and D , the latent feature Z can be computed as $Z = E(X)$ with $Z = [z_1, z_2, \dots, z_{T/l}]$ and $z_i \in \mathbb{R}^{d_c}$, where l represents the temporal downsampling rate of the encoder E . For i -th latent feature z_i , the quantization through C is to find the most similar element in C , which can be properly written as:

$$\hat{z}_i = \arg \min_{c_k \in C} \|z_i - c_k\|_2 \quad (1)$$

Optimization goal. To optimize VQ-VAE, the standard optimization goal [68] \mathcal{L}_{vq} contains three components: a reconstruction loss \mathcal{L}_{re} , the embedding loss \mathcal{L}_{embed} and the commitment loss \mathcal{L}_{commit} .

$$\mathcal{L}_{vq} = \mathcal{L}_{re} + \underbrace{\|Z - sg[\hat{Z}]\|_2}_{\mathcal{L}_{embed}} + \beta \underbrace{\|sg[Z] - \hat{Z}\|_2}_{\mathcal{L}_{commit}} \quad (2)$$

where β is a hyper-parameter for the commitment loss and sg is the stop-gradient operator. For the reconstruction, we find that L1 smooth loss \mathcal{L}_1^{smooth} performs best and an additional regularization on the velocity enhances the generation quality. Let X_{re} be the reconstructed motion of X , i.e., $X_{re} = D(\hat{Z})$, $V(X)$ be the velocity of X where $V = [v_1, v_2, \dots, v_{T-1}]$ with $v_i = x_{i+1} - x_i$. Therefore, the objective of the reconstruction is as follows:

$$\mathcal{L}_{re} = \mathcal{L}_1^{smooth}(X, X_{re}) + \alpha \mathcal{L}_1^{smooth}(V(X), V(X_{re})) \quad (3)$$

where α is a hyper-parameter to balance the two losses. We provide an ablation study on α as well as different reconstruction losses (\mathcal{L}_1 , \mathcal{L}_1^{smooth} and \mathcal{L}_2) in Section 2 of the supplementary material.

Quantization strategy. A naive training of VQ-VAE suffers from codebook collapse [59, 68]. Two training recipes [59] are commonly used to improve the codebook utilization: exponential moving average (*EMA*) and codebook reset (*Code Reset*). *EMA* makes the codebook C evolve smoothly: $C^t \leftarrow \lambda C^{t-1} + (1 - \lambda)C^t$, where C^t is the codebook at iteration t and λ is the exponential moving constant. *Code Reset* finds inactive codes during the training and re-assigns them according to input data. We provide an ablation study on the quantization strategy in Section 4.3.

Architecture. We use a simple convolutional architecture composed of 1D convolution, residual block [29], and ReLU. Our VQ-VAE architecture is illustrated in Figure 3. The architecture is inspired by [19, 41]. We use convolution with stride 2 and nearest interpolation for temporal downsampling and upsampling respectively. The downsampling rate is thus $l = 2^L$, where L denotes the number of residual blocks. We provide an ablation study on the architecture in Section 4.3. The detail of the architecture is provided in Section 6 of the supplementary material.

3.2. T2M-GPT

With a learned motion VQ-VAE, a motion sequence $X = [x_1, x_2, \dots, x_T]$ can be mapped to a sequence of indices $S = [s_1, s_2, \dots, s_{T/l}, End]$ with $s_i \in [1, 2, \dots, s_{T/l}]$, which are indices from the learned codebook. Note that a special *End* token is added to indicate the stop of the motion, which is different from [22] that leverages an extra module to predict motion length. By projecting S back to their corresponding codebook entries, we obtain $\hat{Z} = [\hat{z}_1, \hat{z}_2, \dots, \hat{z}_{T/l}]$ with $\hat{z}_i = c_{s_i}$, which can be decoded to a motion X_{re} through the decoder D . Therefore, text-to-motion generation can be formulated as an autoregressive next-index prediction: given previous $i - 1$ indices, i.e., $S_{<i}$, and text condition c , we aim to predict the distribution of possible next indices $p(S_i | c, S_{<i})$, which can be addressed with transformer [69]. The overview of our transformer is shown in Figure 2b.

Optimization goal. Denoting the likelihood of the full sequence as $p(S|c) = \prod_{i=1}^{|S|} p(S_i | c, S_{<i})$, we directly maximize the log-likelihood of the data distribution:

$$\mathcal{L}_{trans} = \mathbb{E}_{S \sim p(S)} [-\log p(S|c)] \quad (4)$$

We leverage CLIP [55] to extract text embedding c , which has shown its effectiveness in relevant tasks [14, 57, 65].

Causal Self-attention. We apply the causal self-attention [56] in T2M-GPT. Precisely, the output of the causal self-attention is calculated as follows:

$$\text{Attention} = \text{Softmax} \left(\frac{QK^T \times \text{mask}}{\sqrt{d_k}} \right) \quad (5)$$

Methods	R-Precision \uparrow			FID \downarrow	MM-Dist \downarrow	Diversity \uparrow	MModality \uparrow
	Top-1	Top-2	Top-3				
Real motion	0.511 \pm .003	0.703 \pm .003	0.797 \pm .002	0.002 \pm .000	2.974 \pm .008	9.503 \pm .065	-
Our VQ-VAE (Recons.)	0.501 \pm .002	0.692 \pm .002	0.785 \pm .002	0.070 \pm .001	3.072 \pm .009	9.593 \pm .079	-
Seq2Seq [42]	0.180 \pm .002	0.300 \pm .002	0.396 \pm .002	11.75 \pm .035	5.529 \pm .007	6.223 \pm .061	-
Language2Pose [3]	0.246 \pm .002	0.387 \pm .002	0.486 \pm .002	11.02 \pm .046	5.296 \pm .008	7.676 \pm .058	-
Text2Gesture [10]	0.165 \pm .001	0.267 \pm .002	0.345 \pm .002	5.012 \pm .030	6.030 \pm .008	6.409 \pm .071	-
Hier [21]	0.301 \pm .002	0.425 \pm .002	0.552 \pm .004	6.532 \pm .024	5.012 \pm .018	8.332 \pm .042	-
MoCoGAN [67]	0.037 \pm .000	0.072 \pm .001	0.106 \pm .001	94.41 \pm .021	9.643 \pm .006	0.462 \pm .008	0.019 \pm .000
Dance2Music [37]	0.033 \pm .000	0.065 \pm .001	0.097 \pm .001	66.98 \pm .016	8.116 \pm .006	0.725 \pm .011	0.043 \pm .001
TEMOS [§] [53]	0.424 \pm .002	0.612 \pm .002	0.722 \pm .002	3.734 \pm .028	3.703 \pm .008	8.973 \pm .071	0.368 \pm .018
TM2T [23]	0.424 \pm .003	0.618 \pm .003	0.729 \pm .002	1.501 \pm .017	3.467 \pm .011	8.589 \pm .076	2.424 \pm .093
Guo <i>et al.</i> [22]	0.455 \pm .003	0.636 \pm .003	0.736 \pm .002	1.087 \pm .021	3.347 \pm .008	9.175 \pm .083	2.219 \pm .074
MLD [§] [71]	0.481 \pm .003	0.673 \pm .003	0.772 \pm .002	0.473 \pm .013	3.196 \pm .010	9.724 \pm .082	2.413 \pm .079
MDM [66] [§]	-	-	0.611 \pm .007	0.544 \pm .044	5.566 \pm .027	9.559 \pm .086	2.799 \pm .072
MotionDiffuse [74] [§]	0.491 \pm .001	0.681 \pm .001	0.782 \pm .001	0.630 \pm .001	3.113 \pm .001	9.410 \pm .049	1.553 \pm .042
Our GPT ($\tau = 0$)	0.417 \pm .003	0.589 \pm .002	0.685 \pm .003	0.140 \pm .006	3.730 \pm .009	9.844 \pm .095	3.285 \pm .070
Our GPT ($\tau = 0.5$)	0.491 \pm .003	0.680 \pm .003	0.775 \pm .002	0.116 \pm .004	3.118 \pm .011	9.761 \pm .081	1.856 \pm .011
Our GPT ($\tau \in \mathcal{U}[0, 1]$)	0.492 \pm .003	0.679 \pm .002	0.775 \pm .002	0.141 \pm .005	3.121 \pm .009	9.722 \pm .082	1.831 \pm .048

Table 1. Comparison with the state-of-the-art methods on HumanML3D [22] test set. We compute standard metrics following Guo *et al.* [22]. For each metric, we repeat the evaluation 20 times and report the average with 95% confidence interval. Red and Blue indicate the best and the second best result. [§] reports results using ground-truth motion length.

where $Q \in \mathbb{R}^{T \times d_k}$ and $K \in \mathbb{R}^{T \times d_k}$ are query and key respectively, while $mask$ is the causal mask with $mask_{i,j} = -\infty \times \mathbf{1}(i > j) + \mathbf{1}(i \leq j)$, where $\mathbf{1}(\cdot)$ is the indicator function. This causal mask ensures that future information is not allowed to attend the calculation of current tokens. For inference, we start from the text embedding and generate indices in an autoregressive fashion, the generation process will be stopped if the model predicts the *End* token. Note that we are able to generate diverse motions by sampling from the predicted distributions given by the transformer.

Corrupted sequences for the training-testing discrepancy. There is a discrepancy between training and testing. For training, $i - 1$ correct indices are used to predict the next index. While for inference, there is no guarantee that indices serving as conditions are correct. To address this problem, we adopt a simple data augmentation strategy: we replace $\tau \times 100\%$ ground-truth code indices with random ones during training. τ can be a hyper-parameter or randomly sampled from $\tau \in \mathcal{U}[0, 1]$. We provide an ablation study on this strategy in Section 3 of the supplementary material.

4. Experiment

In this section, we present our experimental results. In Section 4.1, we introduce standard datasets as well as evaluation metrics. We compare our results to competitive approaches in Section 4.2. Finally, we provide analysis and

discussion in Section 4.3.

4.1. Datasets and evaluation metric

We conduct experiments on two standard datasets for text-driven motion generations: HumanML3D [22] and KIT Motion-Language (KIT-ML) [54]. Both datasets are commonly used in the community. We follow the evaluation protocol proposed in [22].

KIT Motion-Language (KIT-ML). KIT-ML [54] contains 3,911 human motion sequences and 6,278 textual annotations. The total vocabulary size, that is the number of unique words disregarding capitalization and punctuation, is 1,623. Motion sequences are selected from KIT [47] and CMU [1] datasets but downsampled into 12.5 frame-per-second (FPS). Each motion sequence is described by from 1 to 4 sentences. The average length of descriptions is approximately 8. Following [22, 23], the dataset is split into training, validation, and test sets with proportions of 80%, 5%, and 15%, respectively. We select the model that achieves the best FID on the validation set and report its performance on the test set.

HumanML3D. HumanML3D [22] is currently the largest 3D human motion dataset with textual descriptions. The dataset contains 14,616 human motions and 44,970 text descriptions. The entire textual descriptions are composed of 5,371 distinct words. The motion sequences are originally from AMASS [46] and HumanAct12 [24] but with specific pre-processing: motion is scaled to 20 FPS; those that are longer than 10 seconds are randomly cropped to 10-second

Methods	R-Precision \uparrow			FID \downarrow	MM-Dist \downarrow	Diversity \uparrow	MModality \uparrow
	Top-1	Top-2	Top-3				
Real motion	0.424 \pm .005	0.649 \pm .006	0.779 \pm .006	0.031 \pm .004	2.788 \pm .012	11.08 \pm .097	-
Our VQ-VAE (Recons.)	0.399 \pm .005	0.614 \pm .005	0.740 \pm .006	0.472 \pm .011	2.986 \pm .027	10.994 \pm .120	-
Seq2Seq [42]	0.103 \pm .003	0.178 \pm .005	0.241 \pm .006	24.86 \pm .348	7.960 \pm .031	6.744 \pm .106	-
Language2Pose [3]	0.221 \pm .005	0.373 \pm .004	0.483 \pm .005	6.545 \pm .072	5.147 \pm .030	9.073 \pm .100	-
Text2Gesture [10]	0.156 \pm .004	0.255 \pm .004	0.338 \pm .005	12.12 \pm .183	6.964 \pm .029	9.334 \pm .079	-
Hier [21]	0.255 \pm .006	0.432 \pm .007	0.531 \pm .007	5.203 \pm .107	4.986 \pm .027	9.563 \pm .072	-
MoCoGAN [67]	0.022 \pm .002	0.042 \pm .003	0.063 \pm .003	82.69 \pm .242	10.47 \pm .012	3.091 \pm .043	0.250 \pm .009
Dance2Music [37]	0.031 \pm .002	0.058 \pm .002	0.086 \pm .003	115.4 \pm .240	10.40 \pm .016	0.241 \pm .004	0.062 \pm .002
TEMOS § [53, 71]	0.353 \pm .002	0.561 \pm .002	0.687 \pm .002	3.717 \pm .028	3.417 \pm .008	10.84 \pm .071	0.532 \pm .018
TM2T [23]	0.280 \pm .006	0.463 \pm .007	0.587 \pm .005	3.599 \pm .051	4.591 \pm .019	9.473 \pm .100	3.292 \pm .034
Guo <i>et al.</i> [22]	0.361 \pm .006	0.559 \pm .007	0.681 \pm .007	3.022 \pm .107	3.488 \pm .028	10.72 \pm .145	2.052 \pm .107
MLD § [71]	0.390 \pm .008	0.609 \pm .008	0.734 \pm .007	0.404 \pm .027	3.204 \pm .027	10.80 \pm .117	2.192 \pm .071
MDM [66] §	-	-	0.396 \pm .004	0.497 \pm .021	9.191 \pm .022	10.847 \pm .109	1.907 \pm .214
MotionDiffuse [74] §	0.417 \pm .004	0.621 \pm .004	0.739 \pm .004	1.954 \pm .062	2.958 \pm .005	11.10 \pm .143	0.730 \pm .013
Our GPT ($\tau = 0$)	0.392 \pm .007	0.600 \pm .007	0.716 \pm .006	0.737 \pm .049	3.237 \pm .027	11.198 \pm .086	2.309 \pm .055
Our GPT ($\tau = 0.5$)	0.402 \pm .006	0.619 \pm .005	0.737 \pm .006	0.717 \pm .041	3.053 \pm .026	10.862 \pm .094	1.912 \pm .036
Our GPT ($\tau \in \mathcal{U}[0, 1]$)	0.416 \pm .006	0.627 \pm .006	0.745 \pm .006	0.514 \pm .029	3.007 \pm .023	10.921 \pm .108	1.570 \pm .039

Table 2. **Comparison with the state-of-the-art methods on KIT-ML [54] test set.** We compute standard metrics following Guo *et al.* [22]. For each metric, we repeat the evaluation 20 times and report the average with 95% confidence interval. Red and Blue indicate the best and the second best result. § reports results using ground-truth motion length.

ones; they are then re-targeted to a default human skeletal template and properly rotated to face Z+ direction initially. Each motion is paired with at least 3 precise textual descriptions. The average length of descriptions is approximately 12. According to [22], the dataset is split into training, validation, and test sets with proportions of 80%, 5%, and 15%, respectively. We select the best FID model on the validation set and report its performance on the test set.

Implementation details. For Motion VQ-VAE, the codebook size is set to 512×512 . The downsampling rate l is 4. We provide an ablation on the number of codes in Section 4 of the supplementary material. For both HumanML3D [22] and KIT-ML [54] datasets, the motion sequences are cropped to $T = 64$ for training. We use AdamW [44] optimizer with $[\beta_1, \beta_2] = [0.9, 0.99]$, batch size of 256, and exponential moving constant $\lambda = 0.99$. We train the first 200K iterations with a learning rate of $2e-4$, and 100K with a learning rate of $1e-5$. β and α in \mathcal{L}_{vq} and \mathcal{L}_{re} are set to 1 and 0.5, respectively. Following [22], the dataset KIT-ML and HumanML3D are extracted into motion features with dimensions 251 and 263 respectively, which correspond to local joints position, velocity, and rotations in root space as well as global translation and rotations. These features are computed from 21 and 22 joints of SMPL [43]. More details about the motion representations are provided in Section 5.1 of the supplementary material.

For the T2M-GPT, we employ 18 transformer [69] layers

with a dimension of 1,024 and 16 heads. The ablation for different scales of the transformer is provided in Section 1 of the supplementary material. Following Guo *et al.* [22], the maximum length of Motion is 196 on both datasets, and the minimum lengths are 40 and 24 for HumanML3D [22] and KIT-ML [54] respectively. The maximum length of the code index sequence is $T' = 50$. We train an extra *End* token as a signal to stop index generation. The transformer is optimized using AdamW [44] with $[\beta_1, \beta_2] = [0.5, 0.99]$ and batch size 128. The initialized learning rate is set to $1e-4$ for 150K iterations and decayed to $5e-6$ for another 150K iterations. Training Motion VQ-VAE and T2M-GPT take about 14 hours and 78 hours respectively on a single Tesla V100-32G GPU.

Evaluation metric. Following [22], global representations of motion and text descriptions are first extracted with the pre-trained network in [22], and then measured by the following five metrics:

- *R-Precision.* Given one motion sequence and 32 text descriptions (1 ground-truth and 31 randomly selected mismatched descriptions), we rank the Euclidean distances between the motion and text embeddings. Top-1, Top-2, and Top-3 accuracy of motion-to-text retrieval are reported.
- *Frechet Inception Distance (FID).* We calculate the distribution distance between the generated and real

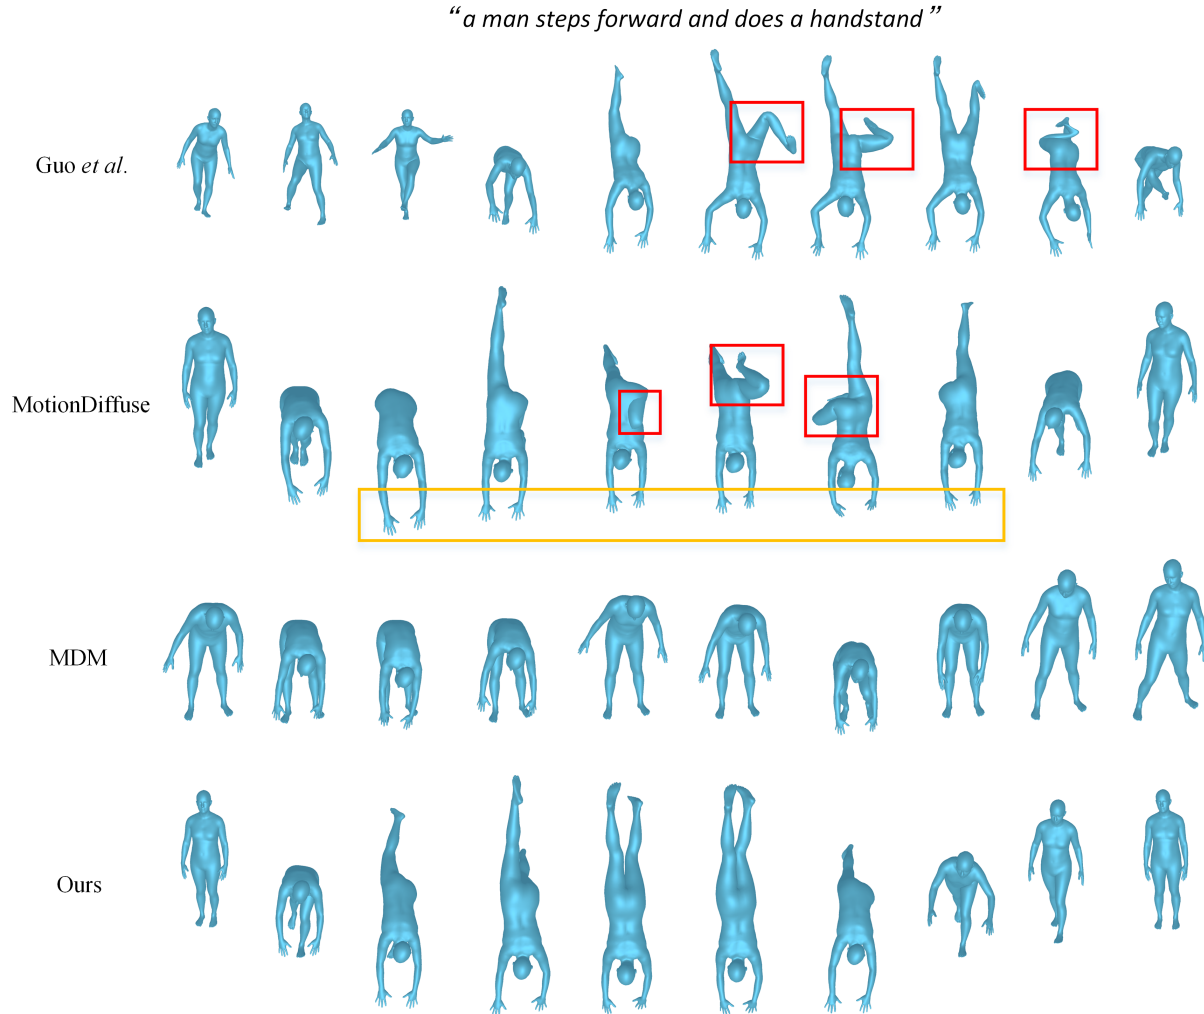


Figure 4. **Visual results on HumanML3D [22] dataset.** We compare our generation with Guo *et al.* [22], MotionDiffuse [74], and MDM [66]. Distorted motions (red) and sliding (yellow) are highlighted. More visual results can be found on the project page project page.

motion using FID [31] on the extracted motion features.

- **Multimodal Distance (MM-Dist).** The average Euclidean distances between each text feature and the generated motion feature from this text.
- **Diversity.** From a set of motions, we randomly sample 300 pairs of motion. We extract motion features and compute the average Euclidean distances of the pairs to measure motion diversity in the set.
- **Multimodality (MModality).** For one text description, we generate 20 motion sequences forming 10 pairs of motion. We extract motion features and compute the average Euclidean distances of the pairs. We finally report the average over all the text descriptions.

Note that more details about the evaluation metrics are provided in Section 5 of the supplementary material.

4.2. Comparison to state-of-the-art approaches

Quantitative results. We show the comparison results in Table 1 and Table 2 on HumanML3D [22] test set and KIT-ML [54] test set. On both datasets, our reconstruction with VQ-VAE reaches close performances to real motion, which suggests high-quality discrete representations learned by our VQ-VAE. For the generation, our approach achieves comparable performance on text-motion consistency (R-Precision and MM-Dist) compared to the state-of-the-art method MotionDiffuse [74], while significantly outperforms MotionDiffuse with FID metric. KIT-ML [54] and HumanML3D [22] are in different scales, which demonstrates the robustness of the proposed approach. Manually corrupting sequences during the training of GPT brings consistent improvement ($\tau = 0.5$ v.s. $\tau = 0$). A more detailed analysis is provided in Section 3 of the supplementary material. Unlike Guo *et*

Quantizer		Reconstruction		Generation	
Code Reset	EMA	FID ↓	Top-1 ↑	FID ↓	Top-1 ↑
	✓	0.492±.004	0.436±.003	42.797±.156	0.048±.001
✓	✓	0.097±.001	0.499±.002	0.176±.008	0.490±.002
✓	✓	0.102±.001	0.494±.003	0.248±.009	0.461±.002
	✓	0.070±.001	0.501±.002	0.116±.004	0.491±.003

Table 3. **Analysis of VQ-VAE quantizers on HumanML3D [22] test set.** For all the quantizers, we set $\tau = 0.5$ and use the same architectures (VQ-VAE and GPT) described in Section 4.1. We report FID and Top-1 for both reconstruction and generation. For each metric, we repeat the evaluation 20 times and report the average with 95% confidence interval.

al. [22] involving an extra module to predict motion length, we implicitly learn the motion length through an additional *End* token, which is more straightforward and shown to be more effective. Note that MDM [66] and MotionDiffuse [74] evaluate their models with the ground-truth motion length, which is not practical for real applications.

Qualitative comparison. Figure 4 shows visual results on HumanML3D [22]. We compare our generations with the current state-of-the-art models: Guo *et al.* [22], MDM [66] and MotionDiffuse [74]. From the example in Figure 4, one can figure out that our model generates human motion with better quality than the others, and we highlight in red for unrealistic motion generated by Guo *et al.* [22] and MotionDiffuse [74]. Moreover, the generated motion of MDM [66] is not related to the text description. Note that more visual results and the failure case are provided on the project page.

4.3. Discussion

Quantization strategies. We first investigate the impact of different quantization strategies presented in Section 3.1. The results are illustrated in Table 3 for both reconstruction and generation. We notice that naive VQ-VAE training is not able to reconstruct nor generate high-quality motion. However, training with *EMA* or *Code Reset* can importantly boost the performances for both reconstruction and generation.

Impact of dataset size. We further analyze the impact of dataset size. To understand whether the largest dataset HumanML3D [22] contains enough data for motion generation, we train our motion VQ-VAE and T2M-GPT on different subsets of the training data, which consists of 10%, 20%, 50%, 80% and 100% of the training data respectively. The trained models are evaluated on the entire test set. The results are illustrated in Figure 5. We evaluate reconstruction for our motion VQ-VAE and generation for our T2M-GPT using four metrics: FID, MM-Dist, Top-1, and Top-3 accuracies. Several insights can be figured out: *i)* metric for motion quality (FID) and metric for motion-text consistency (MM-Dist, Top-1, and Top-3) should be considered at the same time.

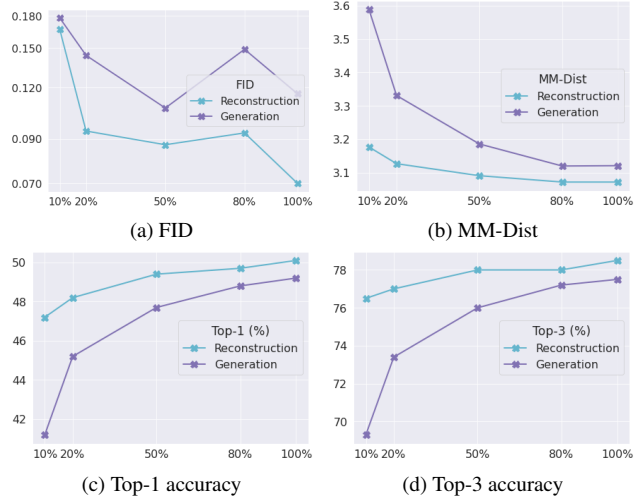


Figure 5. **Impact of dataset size on HumanML3D [22].** We train our motion VQ-VAE (*Reconstruction*) and T2M-GPT (*Generation*) on the subsets of HumanML3D [22] composed of 10%, 20%, 50%, 80%, and 100% training set respectively. All the models are evaluated on the entire test set. We report FID, MM-Dist, Top-1, and Top-3 accuracy for all the models. Results suggest that our model might benefit from more training data.

With only 10% data, the motion might be of good quality, however, the model is not able to generate a correct motion that corresponds to the text description; *ii)* the performances become better with more training data. This trend suggests that additional training data could bring non-negligible improvement to both reconstruction and generation.

5. Conclusion

In this work, we investigated a classic framework based on VQ-VAE and GPT to synthesize human motion from textual descriptions. Our method achieved comparable or even better performances than concurrent diffusion-based approaches, suggesting that this classic framework remains a very competitive approach for motion generation. We explored in detail the effect of various quantization strategies on motion reconstruction and generation. Moreover, we provided an analysis of the dataset size. Our finding suggests that a larger dataset could still bring additional improvement to our approach.

Acknowledgement We thank Mathis Petrovich, Yuming Du, Yingyi Chen, Dexiong Chen, and Xuelin Chen for inspiring discussions and valuable feedback. This paper is supported by NSF of China (No. 62176155) and Jilin Province (20200201037JC), *etc.* More funding information is provided in the supplementary material.

References

- [1] Cmu graphics lab motion capture database. <http://mocap.cs.cmu.edu/>. Accessed: 2022-11-11. **5**
- [2] Hyemin Ahn, Timothy Ha, Yunho Choi, Hwiyeon Yoo, and Songhwai Oh. Text2action: Generative adversarial synthesis from language to action. In *International Conference on Robotics and Automation (ICRA)*, 2018. **3**
- [3] Chaitanya Ahuja and Louis-Philippe Morency. Language2pose: Natural language grounded pose forecasting. In *International Conference on 3D Vision (3DV)*, 2019. **1, 3, 5, 6**
- [4] Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Lars Petersson, and Stephen Gould. A stochastic conditioning scheme for diverse human motion prediction. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. **2**
- [5] Tenglong Ao, Qingzhe Gao, Yuke Lou, Baoquan Chen, and Libin Liu. Rhythmic gesticulator: Rhythm-aware co-speech gesture synthesis with hierarchical neural embeddings. In *SIGGRAPH Asia*, 2022. **2**
- [6] Andreas Aristidou, Anastasios Yiannakidis, Kfir Aberman, Daniel Cohen-Or, Ariel Shamir, and Yiorgos Chrysanthou. Rhythm is a dancer: Music-driven motion synthesis with global structure. *arXiv*, 2021. **2**
- [7] Nikos Athanasiou, Mathis Petrovich, Michael J. Black, and Gül Varol. TEACH: Temporal Action Compositions for 3D Humans. In *International Conference on 3D Vision (3DV)*, 2022. **3**
- [8] Norman I Badler, Cary B Phillips, and Bonnie Lynn Webber. *Simulating humans: computer graphics animation and control*. Oxford University Press, 1993. **2**
- [9] Emad Barsoum, John Kender, and Zicheng Liu. Hp-gan: Probabilistic 3d human motion prediction via gan. In *Proceedings of the Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018. **2**
- [10] Uttaran Bhattacharya, Nicholas Rewkowski, Abhishek Banerjee, Pooja Guhan, Aniket Bera, and Dinesh Manocha. Text2gestures: A transformer-based network for generating emotive body gestures for virtual agents. In *Virtual Reality and 3D User Interfaces (VR)*, 2021. **5, 6**
- [11] Arij Bouazizi, Adrian Holzbock, Ulrich Kressel, Klaus Dietmayer, and Vasileios Belagiannis. Motionmixer: Mlp-based 3d human body pose forecasting. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2022. **2**
- [12] Judith Butepage, Michael J Black, Danica Kragic, and Hedvig Kjellström. Deep representation learning for human motion prediction and classification. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. **2**
- [13] Kang Chen, Zhipeng Tan, Jin Lei, Song-Hai Zhang, Yuan-Chen Guo, Weidong Zhang, and Shi-Min Hu. Choreomaster: choreography-oriented music-driven dance synthesis. *ACM Transactions on Graphics (TOG)*, 2021. **2**
- [14] Katherine Crowson, Stella Biderman, Daniel Kornis, Dashiell Stander, Eric Hallahan, Louis Castricato, and Edward Raff. Vqgan-clip: Open domain image generation and editing with natural language guidance. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. **4**
- [15] Prafulla Dhariwal, Heewoo Jun, Christine Payne, Jong Wook Kim, Alec Radford, and Ilya Sutskever. Jukebox: A generative model for music. *arXiv*, 2020. **2**
- [16] Sander Dieleman, Aaron van den Oord, and Karen Simonyan. The challenge of realistic music generation: modelling raw audio at scale. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. **2**
- [17] Yuming Du, Robin Kips, Albert Pumarola, Sebastian Starke, Ali Thabet, and Arsiom Sanakoyeu. Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. **2**
- [18] Yinglin Duan, Tianyang Shi, Zhengxia Zou, Yenan Lin, Zhehui Qian, Bohan Zhang, and Yi Yuan. Single-shot motion completion with transformer. *arXiv*, 2021. **2**
- [19] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. **2, 4**
- [20] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2015. **2**
- [21] Anindita Ghosh, Noshaba Cheema, Cennet Oguz, Christian Theobalt, and Philipp Slusallek. Synthesis of compositional animations from textual descriptions. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. **1, 3, 5, 6**
- [22] 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 Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. **1, 2, 3, 4, 5, 6, 7, 8**
- [23] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. **1, 3, 5, 6**
- [24] 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 ACM International Conference on Multimedia (ACMMM)*, 2020. **5**
- [25] Wen Guo, Yuming Du, Xi Shen, Vincent Lepetit, Xavier Alameda-Pineda, and Francesc Moreno-Noguer. Back to mlp: A simple baseline for human motion prediction. In *Proceedings of the Winter Conference on Applications of Computer Vision (WACV)*, 2022. **2**
- [26] Ikhsanul Habibie, Daniel Holden, Jonathan Schwarz, Joe Yearsley, and Taku Komura. A recurrent variational autoencoder for human motion synthesis. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2017. **2**
- [27] Félix G Harvey and Christopher Pal. Recurrent transition networks for character locomotion. In *SIGGRAPH Asia 2018 Technical Briefs*, 2018. **2**

- [28] Félix G Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. *ACM Transactions on Graphics (TOG)*, 2020. 2
- [29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4
- [30] Alejandro Hernandez, Jurgen Gall, and Francesc Moreno-Noguer. Human motion prediction via spatio-temporal inpainting. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [31] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 7
- [32] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [33] Daniel Holden, Taku Komura, and Jun Saito. Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)*, 2017. 3
- [34] Daniel Holden, Jun Saito, and Taku Komura. A deep learning framework for character motion synthesis and editing. *ACM Transactions on Graphics (TOG)*, 2016. 3
- [35] Manuel Kaufmann, Emre Aksan, Jie Song, Fabrizio Pece, Remo Ziegler, and Otmar Hilliges. Convolutional autoencoders for human motion infilling. In *International Conference on 3D Vision (3DV)*, 2020. 2
- [36] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations (ICLR)*, 2014. 2
- [37] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. Dancing to music. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2, 5, 6
- [38] Jiaman Li, Yihang Yin, Hang Chu, Yi Zhou, Tingwu Wang, Sanja Fidler, and Hao Li. Learning to generate diverse dance motions with transformer. *arXiv*, 2020. 2
- [39] Peizhuo Li, Kfir Aberman, Zihan Zhang, Rana Hanocka, and Olga Sorkine-Hornung. Ganimator: Neural motion synthesis from a single sequence. *ACM Transactions on Graphics (TOG)*, 2022. 3
- [40] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [41] Siyao Li, Weijiang Yu, Tianpei Gu, Chunze Lin, Quan Wang, Chen Qian, Chen Change Loy, and Ziwei Liu. Bailando: 3d dance generation by actor-critic gpt with choreographic memory. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 4
- [42] Angela S Lin, Lemeng Wu, Rodolfo Corona, Kevin Tai, Qixing Huang, and Raymond J Mooney. Generating animated videos of human activities from natural language descriptions. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 5, 6
- [43] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 2015. 6
- [44] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019. 6
- [45] Thomas Lucas, Fabien Baradel, Philippe Weinzaepfel, and Grégory Rogez. Posegpt: Quantization-based 3d human motion generation and forecasting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [46] 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 International Conference on Computer Vision (ICCV)*, 2019. 5
- [47] Christian Mandery, Ömer Terlemez, Martin Do, Nikolaus Vahrenkamp, and Tamim Asfour. The kit whole-body human motion database. In *International Conference on Robotics and Automation (ICRA)*, 2015. 5
- [48] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. History repeats itself: Human motion prediction via motion attention. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [49] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 2
- [50] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [51] Dario Pavllo, David Grangier, and Michael Auli. Quaternet: A quaternion-based recurrent model for human motion. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2018. 2
- [52] Mathis Petrovich, Michael J. Black, and Gül Varol. Action-conditioned 3D human motion synthesis with transformer VAE. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 1, 3
- [53] Mathis Petrovich, Michael J. Black, and Gül Varol. TEMOS: Generating diverse human motions from textual descriptions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 1, 3, 5, 6
- [54] Matthias Plappert, Christian Mandery, and Tamim Asfour. The kit motion-language dataset. *Big data*, 2016. 2, 3, 5, 6, 7
- [55] 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 (ICML)*, 2021. 1, 3, 4
- [56] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 2, 4

- [57] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv*, 2022. 4
- [58] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning (ICML)*, 2021. 2
- [59] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 4
- [60] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [61] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. Singan: Learning a generative model from a single natural image. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [62] Sebastian Starke, Ian Mason, and Taku Komura. Deepphase: periodic autoencoders for learning motion phase manifolds. *ACM Transactions on Graphics (TOG)*, 2022. 3
- [63] Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. Neural state machine for character-scene interactions. *ACM Transactions on Graphics (TOG)*, 2019. 3
- [64] Xiangjun Tang, He Wang, Bo Hu, Xu Gong, Ruifan Yi, Qilong Kou, and Xiaogang Jin. Real-time controllable motion transition for characters. *ACM Transactions on Graphics (TOG)*, 2022. 2
- [65] Guy Tevet, Brian Gordon, Amir Hertz, Amit H Bermano, and Daniel Cohen-Or. Motionclip: Exposing human motion generation to clip space. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 1, 3, 4
- [66] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Amit H Bermano, and Daniel Cohen-Or. Human motion diffusion model. *arXiv*, 2022. 2, 3, 5, 6, 7, 8
- [67] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 5, 6
- [68] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 2, 3, 4
- [69] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 2, 4, 6
- [70] Will Williams, Sam Ringer, Tom Ash, David MacLeod, Jamie Dougherty, and John Hughes. Hierarchical quantized autoencoders. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [71] Chen Xin, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, Jingyi Yu, and Gang Yu. Executing your commands via motion diffusion in latent space. *arXiv*, 2022. 5, 6
- [72] Sijie Yan, Zhizhong Li, Yuanjun Xiong, Huahan Yan, and Dahua Lin. Convolutional sequence generation for skeleton-based action synthesis. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 2
- [73] Xinchen Yan, Akash Rastogi, Ruben Villegas, Kalyan Sunkavalli, Eli Shechtman, Sunil Hadap, Ersin Yumer, and Honglak Lee. Mt-vae: Learning motion transformations to generate multimodal human dynamics. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [74] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiandiffuse: Text-driven human motion generation with diffusion model. *arXiv*, 2022. 2, 3, 5, 6, 7, 8