

# OMG: Towards Open-vocabulary Motion Generation via Mixture of Controllers

Han Liang<sup>1</sup> Jiacheng Bao<sup>1</sup> Ruichi Zhang<sup>1</sup> Sihan Ren<sup>1</sup> Yuecheng Xu<sup>1</sup>  
 Sibeiyang<sup>1</sup> Xin Chen<sup>2</sup> Jingyi Yu<sup>1</sup> Lan Xu<sup>1</sup>  
<sup>1</sup>ShanghaiTech University <sup>2</sup>Tencent PCG

## Abstract

We have recently seen tremendous progress in realistic text-to-motion generation. Yet, the existing methods often fail or produce implausible motions with unseen text inputs, which limits the applications. In this paper, we present *OMG*, a novel framework, which enables compelling motion generation from zero-shot open-vocabulary text prompts. Our key idea is to carefully tailor the pretrain-then-finetune paradigm into the text-to-motion generation. At the pre-training stage, our model improves the generation ability by learning the rich out-of-domain inherent motion traits. To this end, we scale up a large unconditional diffusion model up to 1B parameters, so as to utilize the massive unlabeled motion data up to over 20M motion instances. At the subsequent fine-tuning stage, we introduce motion ControlNet, which incorporates text prompts as conditioning information, through a trainable copy of the pre-trained model and the proposed novel Mixture-of-Controllers (MoC) block. MoC block adaptively recognizes various ranges of the sub-motions with a cross-attention mechanism and processes them separately with the text-token-specific experts. Such a design effectively aligns the CLIP token embeddings of text prompts to various ranges of compact and expressive motion features. Extensive experiments demonstrate that our *OMG* achieves significant improvements over the state-of-the-art methods on zero-shot text-to-motion generation. Project page: <https://tr3e.github.io/omg-page>.

## 1. Introduction

Generating vivid motions of human characters is a long-standing task in computer vision, with numerous applications in movies, games, robotics, and VR/AR. The text-to-motion setting aims to democratize motion generation for novices and has recently received substantive attention.

Benefited from the existing text-annotated motion datasets [19, 47], recent advances generate diverse motions from text prompts, equipped with various generative models like VAEs [56, 76], autoregressive models [33, 87], and diffusion models [12, 13, 77, 89]. However, these meth-



Figure 1. Our Open-vocabulary Motion Generation (OMG) approach is capable of generating high-quality motions in response to unseen text prompts.

ods heavily rely on the paired text-motion data with limited text annotations, and hence fall short of generalizing to unseen open-vocabulary text inputs. Imagine generating “Gollum breakdance footwork”, it requires out-of-domain generation ability with rich human knowledge to understand the motion traits in various words like characters (Gollum, ninja, etc.) or skills (breakdance, spinkick, etc.). To tackle the open-vocabulary texts, recent works [8, 28, 42] utilize the zero-shot text-image alignment ability of the large-scale pre-trained models, e.g., CLIP [57]. They adopt text-pose alignments and hence remove the reliance on pair-wise text and continuous motions. Yet, the discrete poses are insufficient to represent the rich motion traits, leading to unrealistic motion generation. On the other hand, recent advances [9, 21, 63, 88] enable impressive and controllable text-driven image and even video generation, even from open-vocabulary text inputs. Revisiting their huge success, both the pretrain-then-finetune paradigm as well as scaling up the model have played an important role. Yet, adopting

these strategies for text-to-motion generation is challenging. First, there exists a huge imbalance between the quantities of unpaired motion data and paired text-motion data. Second, for open-vocabulary texts, various tokens correspond to various motion traits, constituting a complex many-to-many problem.

To tackle the challenges, we propose *OMG* – a novel scheme to generate motions of human characters from open-vocabulary texts (see Fig. 1). Our key idea is to carefully tailor the pretrain-then-finetune paradigm into text-to-motion generation. For pre-training, we adopt a minimalistic design to scale up a large unconditional diffusion model, so as to utilize the massive unlabeled motion data. For fine-tuning, we first freeze the pre-trained large model. We then adopt the trainable copy and treat the text prompts as extra conditions through a novel design, named mixture-of-controllers, to learn to predict the conditional residuals. It adaptively fuses the motion traits corresponding to various tokens from the input text prompts, so as to handle the ambiguity between the text and motion modalities.

Specifically, at the pre-training stage, we adopt the diffusion transformer [52] as the backbone with over-parameterized motion representation. We then scale up the models with parameters ranging from 88M to 1B, leveraging over 20M unlabeled motion instances from a diverse range of 13 publicly available datasets. During training, we also adopt a sliding random window strategy to crop various motion sequences, to improve the generation ability for arbitrary lengths of motion. To this end, the pre-trained model learns the rich inherent motion features with a large motion manifold to ensure the realism of the generated motions. For the subsequent stage, to incorporate the text prompts as conditioning information, we adopt a fine-tuning scheme called motion ControlNet, analogous to the famous ControlNet [88] for the text-image generation task. It includes a trainable copy of the large-scale pre-trained model, which serves as a strong backbone to retain the motion features, as well as a novel block called Mixture-of-Controllers (MoC). Such MoC block can effectively inject residual information into the pre-trained model, based on the motion features and the CLIP-based text embeddings. Our key design in the MoC block is the cross-attention mechanism between text and motion, as well as the text-token-specific Mixture-of-Experts [16, 67], which are to be detailed in later sections. Such a design effectively aligns the token embeddings of text prompts to various ranges of compact and expressive motion features, which are warmed up from a powerful pre-trained model. As a result, our *OMG* approach achieves compelling generation from open-vocabulary texts, significantly outperforming prior arts. In particular, only fine-tuned on the HumanML3D dataset [19], it can generate vivid motions of various human characters with diverse motion skills from the Mixamo dataset [1], as shown Fig. 1.

To summarize, our main contributions include:

- We propose a text-to-motion generation approach with a pretrain-then-finetune paradigm to scale up both data and model, achieving state-of-the-art zero-shot performance.
- We experimentally demonstrate that pre-training on large-scale unlabeled motion data improves the generation results from diverse and open-vocabulary texts.
- We propose a fine-tuning scheme for text conditioning, utilizing a mixture of controllers to effectively improve the alignment between text and motion.

## 2. Related Work

**Conditional Motion Synthesis.** Being able to generate realistic and contextually relevant motion sequences based on various types of conditions, conditional motion synthesis has received increasing attention in the field of motion generation. Common types of conditions include text [3, 4, 12, 20, 34, 41, 43, 56, 66, 76, 84], action [18, 55], music [2, 35, 37, 39], speech [5, 6, 22, 92], control signals [54, 71, 73, 81], action labels [18, 55, 82], incomplete motion sequences [14, 23], images [11, 25, 40, 60, 61] and scene [91]. The advent of Diffusion Models [26, 70] has given a big boost to text-driven motion synthesis. Kim et al. [34] develops a transformer-based diffusion model which could generate and edit motions well aligned with the given text. Motion Diffusion Model (MDM) [77] combines insights already well established in the motion generation domain and in each diffusion step predicts the sample rather than the noise. Motion Latent-based Diffusion model (MLD) [12] performs a diffusion process on the motion latent space. Besides the diffusion model, T2M-GPT [87] investigates a framework that combines VQ-VAE [79] and autoregressive model. However, since these models are trained only based on paired text-motion datasets such as HumanML3D [20], they cannot well handle unseen text prompts.

**Open-vocabulary Generation.** Instead of relying on pre-existing data, zero-shot text-driven generation leverages general knowledge learned during training to create novel content from text prompts. Reed et al. [59] uses DC-GAN architecture to steer for zero-shot text-to-image synthesis. CLIP [57], pre-trained on four hundred million image-text pairs, possesses the remarkable ability to effectively comprehend and generate meaningful associations between images and text. With CLIP’s strong ability, many works are able to generate high-quality zero-shot text-driven images [17, 51, 83] or 3D objects [31, 32, 49, 53, 65, 80]. In the motion synthesis field, some works have attempted to investigate open-vocabulary text-to-motion generation and achieved good results. MotionCLIP [76] utilizes a decoder to decode the CLIP motion embeddings. AvatarCLIP [28] synthesizes a key pose and then retrieves the closest motion from the database. Lin et al. [42] pre-train a motion

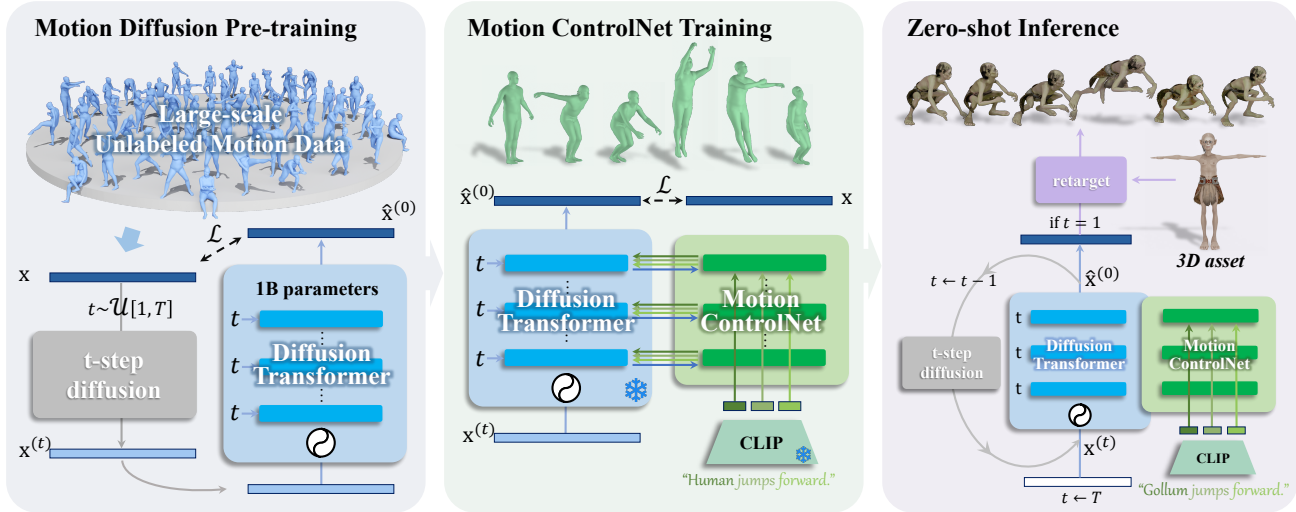


Figure 2. **Method overview.** We train our OMG model in two stages. First, we leverage large-scale unlabeled motion data to pre-train an unconditional diffusion model with up to 1B parameters (Sec. 3.1). Then, we adopt a conditional fine-tuning scheme called motion ControlNet to condition the pre-trained diffusion model on text prompts (Sec. 3.2). During inference, the pre-trained unconditional denoiser and the fine-tuned conditional denoiser are combined with classifier-free guidance, generating realistic motions with zero-shot text inputs.

generator that learns to reconstruct the full motion from the masked motion with the key pose. Make-An-Animation [8] pre-trains generative model with diverse in-the-wild text-pose pair. However, since these approaches are based on text-pose alignment, the generated motion is deficient in realism due to a lack of global rotation and translation. Recently, MotionGPT [33] pre-trains and fine-tunes the large language model with tokenized motion and text data. However, it still struggles to generate novel motion from open-vocabulary text prompts.

**Mixture-of-Experts.** In the deep learning field, Mixture-of-Experts (MoE) [16, 30, 67] in a neural network architecture divides specific model parameters into distinct groups, referred to as “experts”. Eigen et al. [15] use a different gating network at each layer in a multi-layer network, increasing the number of effective experts by introducing an exponential number of paths through different combinations of experts at each layer. Other research that regards the entire model as an expert [38] also achieves good results. Moreover, endeavors are underway to enhance the training and inference speed within the MoE framework [24, 36]. The MoE paradigm is applicable not only to language processing tasks but also extends to visual models [62], Mixture-of-Modality-Experts in multi-modal transformers [68], as well as motion synthesis. Holden et al. [27] develop phase-functioned experts blended by pre-defined blending weights to control character at a specific phase, and Starke et al. [72, 73], Zhang et al. [85] uses a gating network to predict the blending weights, achieving impressive results of character control. Inspired by them, we devise text-token-specific experts to control the corresponding sub-motions.

### 3. Methods

We aim to generate realistic and diverse human motions that are conditioned on text prompts, which capture complex and abstract motion characteristics with zero-shot open-vocabulary descriptions. To this end, we adopt pretrain-then-finetune paradigm to enhance the capability of our model. For the unconditional denoiser (Fig. 2 left), we leverage a large-scale unlabeled motion dataset for motion diffusion pre-training and scale the model size up (Sec. 3.1). For the conditional denoiser (Fig. 2 middle), we devise a specific conditional fine-tuning scheme called motion ControlNet, including a novel conditioning design called Mixture-of-Controllers, to exploit the zero-shot capability of the CLIP text embeddings (Sec. 3.2). During inference (Fig. 2 right), we further use classifier-free guidance to combine the unconditional denoiser and the conditional one.

#### 3.1. Motion Diffusion Pre-training

**Large-scale Unlabeled Motion Dataset.** To facilitate large-scale pre-training in the motion domain, we collected a large amount of high-quality human motion data totaling over 20M frames from various sources, such as character animation datasets, marker-based, IMU-based, and multi-view markerless mocap datasets. We then standardize the frame rate and re-target them to the unified parametric skeleton of SMPL body model [45] using off-the-shelf re-targeting algorithms [86]. After that, following the previous successful works [19, 77] on motion generation, we enrich the SMPL skeletons with an over-parameterized motion representation that facilitates network learning.

Model Name	$n_{layers}$	$d_{model}$	$n_{heads}$	$d_{head}$	$n_{params}$
OMG-Base	8	1024	8	128	88M
OMG-Large	12	1280	10	128	201M
OMG-Huge	16	1664	13	128	405M
OMG-Giant	24	2048	16	128	1B

Table 1. Sizes and architectures of our 4 models.

**Model Scaling Up.** We design our pre-trained diffusion model with minimalism, only preserving the essential modules that are scalable and efficient to adapt to expanding data. To this end, we adopt the Diffusion Transformer (DiT) [52] architecture as our network backbone, due to its scalability and impressive performance with increasing input tokens and training data. The only difference is that we use rotary positional embedding [74] to encourage the model to capture the relative temporal relations among the frames. To study the dependence of performance on model size, we train 4 different sizes of our OMG model, ranging from 88 million parameters to 1 billion parameters. Tab. 1 shows the configuration details of our 4 models. Here  $n_{layers}$  is the total number of layers,  $d_{model}$  is the number of units in each bottleneck layer (we always have the feed-forward layer two times the size of the bottleneck layer,  $d_{ff} = 2d_{model}$ ), and  $d_{head}$  is the dimension of each attention head, and  $n_{params}$  is the total number of trainable parameters.

**Training.** We train our unconditional denoiser  $\mathcal{D}_u$  to predict the clean motion  $\hat{\mathbf{x}}^{(0)}$  given diffusion timestep  $t$  with the simple objective:

$$\mathcal{L}_{simple} = \mathbb{E}_{\mathbf{x}, t, \epsilon} [\lambda_t \|\mathbf{x} - \mathcal{D}_u(\mathbf{x}^{(t)}, t)\|_2^2], \quad (1)$$

where  $\mathbf{x}^{(t)} = \mathbf{x} + \sigma_t \epsilon$  is  $t$ -step noised motion, noise  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ,  $\lambda_t$  is the loss weighting factor. Furthermore, we introduce geometric losses, analogous to MDM [77], that contain velocity loss  $\mathcal{L}_{vel}$ , and foot contact loss  $\mathcal{L}_{foot}$  to enforce physical properties and anti-sliding. Overall, the unconditional denoiser  $\mathcal{D}_u$  is thus trained with the following objective:

$$\mathcal{L} = \mathcal{L}_{simple} + \lambda_{vel} \mathcal{L}_{vel} + \lambda_{foot} \mathcal{L}_{foot}. \quad (2)$$

In our experiments,  $\lambda_{vel} = 30$  and  $\lambda_{foot} = 30$ . The number of diffusion timesteps  $T$  is set to 1,000 with cosine noise level schedule [50]. We train all sizes of the model for 1M iterations with a batch size of 256 using the AdamW [46] optimizer with a weight decay of 0, a maximum learning rate of  $10^{-4}$ , and a cosine LR schedule with 10K linear warmup steps. The models are trained using Pytorch with ZeRO [58] memory redundancy strategy on 8 NVIDIA A100 GPUs. The largest model OMG-G takes 1500 GPU hours for pre-training.

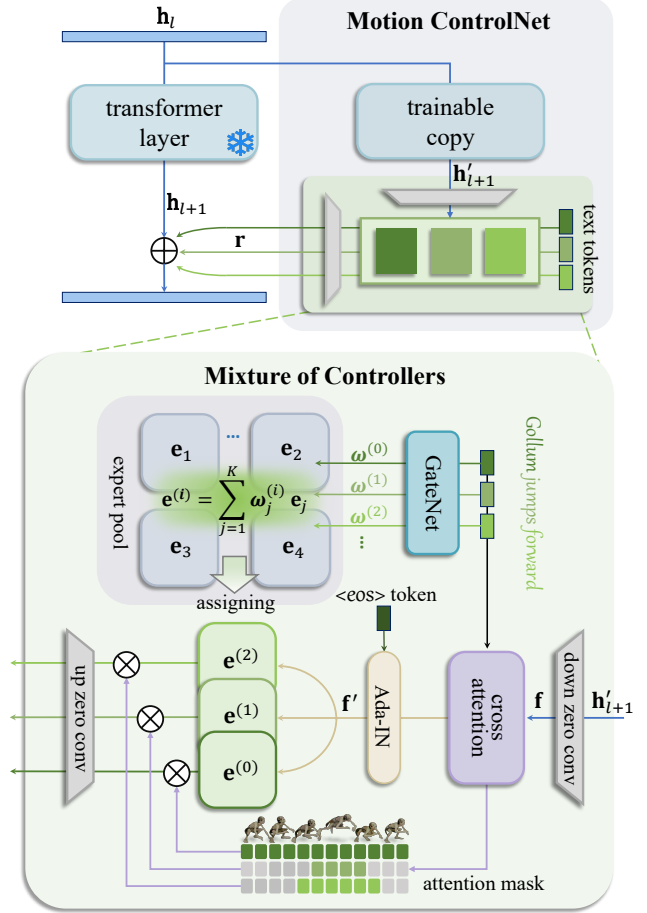


Figure 3. **Motion ControlNet (top)** freezes the parameters of the pre-trained transformer layer and combines a trainable copy of the layer with the proposed **Mixture-of-Controllers (bottom)** block. The MoC block first fuses the text features and motion features and simultaneously determines the sub-motion ranges for each text token with the cross-attention mechanism. Then it performs fine-grained control of sub-motions with text-token-specific experts.

**Sliding Random Windows.** During training, we propose sliding random windows that iterate over each motion state frame  $x_i$  in the motion dataset to sample motion sequences starting from  $x_i$ . Specifically, for the  $i^{th}$  frame, we randomly crop a sub-sequence  $\mathbf{x}$  of length  $l$  with a random window  $[i, i + l)$ , where  $l \sim \mathcal{U}[1, L]$  is a uniform variable, and  $L$  is the hyper-parameter that determines the maximum length of the generated motion sequence of our model. This pushes the model to learn the relations among both temporal frames and spatial features of a single motion state to process arbitrary lengths of motions even a single keyframe, to facilitate the following fine-grained control of sub-motions of arbitrary lengths. In our experiments, for balancing resolution and duration, the maximum length  $L$  of sliding random windows is set to 300 with the framerate set to 30.

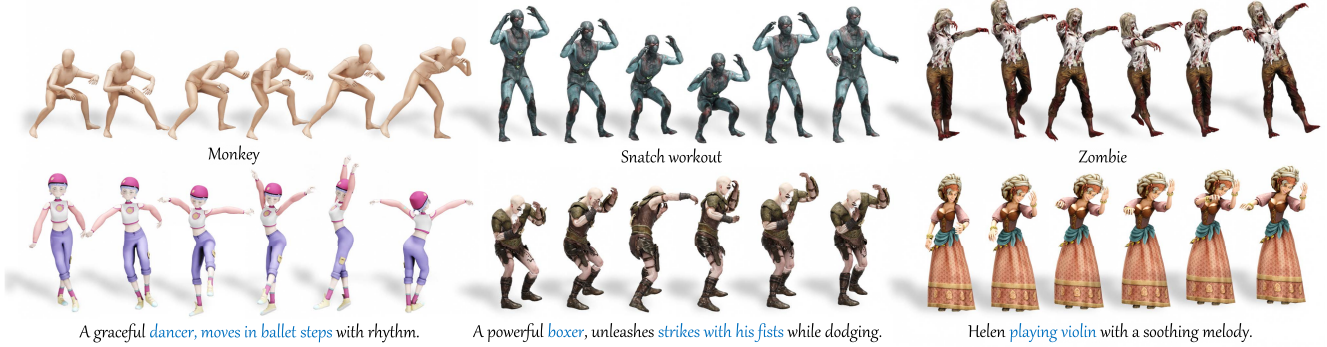


Figure 4. **Qualitative results** generated by our model given various unseen text prompts. Our model effectively captures the motion characteristics from either a single phrase or longer natural sentences.

### 3.2. Motion ControlNet

To incorporate the text prompts  $\mathbf{c} = \{\mathbf{c}_i\}_{i=1}^n$  into the pre-trained motion model, We adopt a fine-tuning scheme we call motion ControlNet to condition on text prompts, inspired by image ControlNet [88], including a trainable copy of original transformer layers and the proposed conditioning blocks called Mixture-of-Controllers (MoC).

As illustrated in Fig. 3 top, we freeze the parameters of the pre-trained transformer layer and combine a trainable copy of the layer with a MoC block that injects  $n$  residuals corresponding to  $n$  text tokens into the output of the original transformer layer. Specifically, the frozen parameters retain the pre-trained denoising ability, while the trainable copy reuses the large-scale pre-trained model as a strong initial backbone to learn to extract semantic motion features from intermediate representation  $\mathbf{h}_l$  of the last layer. And based on that, the MoC block is employed to predict the  $n$  conditional residuals  $\mathbf{r}$  corresponding to  $n$  text tokens, all of which are then added to the original output  $\mathbf{h}_{l+1}$ . To do so, a pre-trained CLIP text encoder  $\mathcal{E}(\mathbf{c})$  is employed to extract the text token embeddings.

**Mixture-of-Controllers.** In the Mixture-of-Controllers block (Fig. 3 bottom), our key idea is to control the sub-motion sequences separately with different expert controllers in an MoE fashion [16, 67], so as to better align them to the corresponding text token embeddings in CLIP space. For instance, for the given text prompt “Gollum jumps forward”, the “Gollum” token corresponds to the entire sequence while “jumps” corresponds only to the sub-sequence when he is jumping.

To this end, we introduce two cooperative designs. The first one is a **cross-attention mechanism** to fuse the text features and motion features and simultaneously determine the sub-motion ranges for each text token. The second one is the **text-token-specific experts** selected from an expert pool to perform fine-grained control of sub-motions.

For the **cross-attention mechanism**, a sequence of mo-

tion features  $\mathbf{f} \in \mathbb{R}^{l \times d_m}$  is input into a cross-attention layer along with text embeddings  $\mathcal{E}(\mathbf{c}) \in \mathbb{R}^{n \times d_c}$ , where  $l$  denotes the length of the sequence,  $d_m$  represents the dimension of the motion features,  $n$  denotes the number of the text tokens, and  $d_c$  represents the dimension of text embedding. The cross-attention layer projects  $\mathbf{f}$  to a query matrix  $\mathbf{Q} = \mathbf{f} \cdot \mathbf{W}_m^Q$ , where  $\mathbf{W}_m^Q \in \mathbb{R}^{d_m \times d_m}$  is the motion projection matrix. And  $\mathcal{E}(\mathbf{c})$  is projected to a key matrix  $\mathbf{K} = \mathcal{E}(\mathbf{c})\mathbf{W}_c^K$  and a value matrix  $\mathbf{V} = \mathcal{E}(\mathbf{c})\mathbf{W}_c^V$ , where  $\mathbf{W}_c^K, \mathbf{W}_c^V \in \mathbb{R}^{d_c \times d_m}$  are the text projection matrices. Then the text values  $\mathbf{V}$  are distributed into the motion sequence with the attention mechanism:

$$\mathbf{f}' = \mathbf{f} + \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V}. \quad (3)$$

Simultaneously an attention map between text tokens and motion sequence is produced, denoted by  $\mathbf{A} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right) \in \mathbb{R}^{l \times n}$ . Within this attention map, the  $i^{\text{th}}$  column  $\mathbf{A}_{*,i}$  indicates a correspondence between a sub-motion sequence and the  $i^{\text{th}}$  text token. Besides, we employ the adaptive instance normalization (Ada-IN) conditioning on  $\langle eos \rangle$  (end of sequence) token to normalize  $\mathbf{f}'$  before sending them to the experts. Since CLIP  $\langle eos \rangle$  token summarizes the entire text, we intuitively use it to unify the distribution of the entire motion sequence.

For the **text-token-specific experts**, we define an expert controller corresponding to  $i^{\text{th}}$  text token  $\mathcal{E}(\mathbf{c}_i)$  as a two-layer feed-forward network (denoted by  $\mathcal{F}$ ) with parameters  $\mathbf{e}^{(i)} = \{\mathbf{W}_0 \in \mathbb{R}^{d_m \times 2d_m}, \mathbf{W}_1 \in \mathbb{R}^{2d_m \times d_m}, \mathbf{b}_0 \in \mathbb{R}^{2d_m}, \mathbf{b}_1 \in \mathbb{R}^{d_m}\}$ . Furthermore, the parameters  $\mathbf{e}^{(i)}$  are generated by blending  $K$  expert parameters  $\{\mathbf{e}_1, \dots, \mathbf{e}_K\}$  in an expert pool, each of which is in a form of the same parameter configuration:

$$\mathbf{e}^{(i)} = \sum_j^K \omega_j^{(i)} \mathbf{e}_j, \quad (4)$$

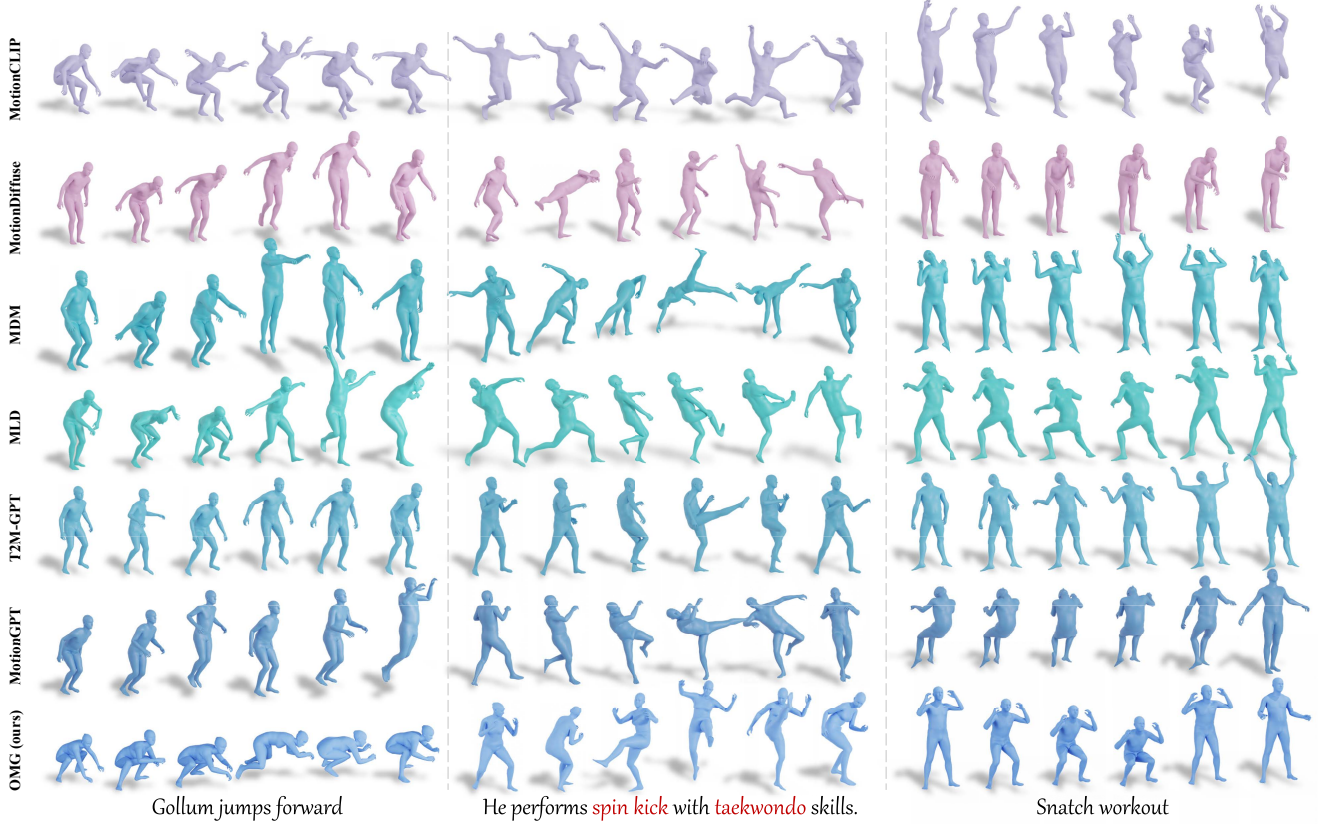


Figure 5. **Qualitative comparison.** Our method can generate high-quality human motions that better align with text prompts than previous state-of-the-art methods.

where  $j$  indicates the index of expert pool,  $K$  is expert pool size that can be adjusted, and  $\omega^{(i)} = \{\omega_j^{(i)}\}_{j=1}^K \in \mathbb{R}^K$  is the blending weights vector controlled by the text token  $c_i$  with a three-layer fully-connected gating network  $\mathcal{G}$ :

$$\omega^{(i)} = \text{softmax}(\mathcal{G}(\mathcal{E}(c_i))). \quad (5)$$

Then, the  $n$  experts take motion features  $\mathbf{f}'$  as input and output  $n$  conditional masked residuals  $\mathbf{r} = \{\mathbf{r}_i\}_{i=1}^n$ :

$$\mathbf{r}_i = \mathbf{M}_{*,i} \circ \mathcal{F}(\mathbf{f}' | \mathbf{e}^{(i)}), \quad (6)$$

where  $\mathbf{M}_{*,i} = \text{sigmoid}(\gamma(\mathbf{A}_{*,i} - \beta \max(\mathbf{A}_{*,i})))$  is an attention mask ranging from 0 to 1, and  $\gamma$  and  $\beta$  are hyper-parameters to control the sharpness and threshold respectively.

Additionally, a down-projection and up-projection 1-D convolution pair is used to process input and output respectively to reduce the control latent dimension and thus the trainable parameters. Moreover, similar to image ControlNet [88], we use zero-initialized convolution parameters to protect the trainable copy from the harmful gradient noises in the initial training steps.

**Training and Inference.** The conditional denoiser  $\mathcal{D}_c$  is also trained using the same objective  $\mathcal{L}$  (Eq. (2)) with trainable parameters in motion ControlNet. In our experiments, we employ the frozen *CLIP-VIT-L/14* [57] text encoder to extract text embeddings at the final layer normalization with dimension  $d_c = 768$  and we truncate the text tokens with the maximum token number 77. We set the down-projection latent dimension  $d_m$  of MoC blocks to 256, the expert pool size  $K = 12$ , and attention mask sharpness  $\gamma = 24$  and threshold  $\beta = 0.25$ . We train all the motion ControlNets for 500 epochs with a batch size of 64 using the AdamW optimizer with a weight decay of  $1 \times 10^{-5}$ , a maximum learning rate of  $3 \times 10^{-5}$ , and a cosine LR schedule with 1K linear warmup steps. In the training process, we randomly replace  $\langle eos \rangle$  token with the empty token with 50% probability to increase the ability to capture global semantics in text tokens as a replacement.

During inference, classifier-free guidance is used to combine unconditional denoiser and conditional denoiser:

$$\hat{\mathbf{x}}^{(0)} = (1 - s) \cdot \mathcal{D}_u(\mathbf{x}^{(t)}, t) + s \cdot \mathcal{D}_c(\mathbf{x}^{(t)}, t, \mathbf{c}). \quad (7)$$

In our experiments, DDIM [69] sampling strategy with 200 timesteps is employed and the guidance strength  $s = 4.5$ .

Methods	HumanML3D [20]			Mixamo [1] (zero-shot)		
	FID↓	R-Precision↑	Diversity →	FID ↓	CLIP-score↑	Diversity →
Real	0.002±.000	0.797±.002	9.503±.065	0.106±.003	0.648±.001	2.665±.022
MotionCLIP [76]	-	-	-	2.542±.012	0.511±.004	2.205±.012
MAA [8]	0.774±.002	0.676±.001	8.230±.064	-	-	-
MotionDiffuse [89]	0.630±.001	<u>0.782</u> ±.001	9.410±.049	2.363±.010	0.505±.002	2.411±.016
MDM [77]	0.544±.044	0.611±.007	<u>9.559</u> ±.086	1.297±.004	0.536±.004	<u>2.594</u> ±.011
MLD [12]	0.473±.013	0.772±.002	9.724±.082	<u>1.220</u> ±.004	<u>0.556</u> ±.003	2.583±.018
T2M-GPT [87]	<b>0.116</b> ±.004	0.775±.002	9.844±.095	1.420±.003	0.541±.002	2.590±.022
MotionGPT [33]	<u>0.232</u> ±.008	0.778±.002	<b>9.528</b> ±.071	1.365±.003	0.552±.002	2.589±.018
OMG (ours)	0.381±.008	<b>0.784</b> ±.002	9.657±.085	<b>1.164</b> ±.009	<b>0.588</b> ±.002	<b>2.632</b> ±.021

Table 2. Comparison of text-to-motion generation on HumanML3D [20] and Mixamo [1] test set. We ran all the evaluations 20 times, with the average reported alongside a 95% confidence interval. The right arrow  $\rightarrow$  means the closer to real motion the better. **Bold** and underline indicate the best and the second best result. The term (Zero-shot) implies that the dataset contains unseen open-vocabulary texts.

## 4. Experiments

We show qualitative results in Fig. 4. Our method enables fine-grained control of complex and abstract motion trait descriptions. We encourage the reader to appreciate more qualitative results provided in supplements and the video.

### 4.1. Datasets and Evaluation Metrics

**Training Datasets.** At the pre-training stage, we utilize various publicly available human motion datasets, such as artist-created datasets [23, 48], marker-based [7, 29, 44, 47, 75], IMU-based [40, 78] and multi-view markerless [10, 39, 41, 90] motion capture datasets, totaling over 20 million frames. In the subsequent conditional fine-tuning stage, we train our motion ControlNet using the text-motion HumanML3D [20] dataset, for fair comparisons with previous methods.

**Evaluation Datasets.** We test on two benchmarks, the HumanML3D [20] and the Mixamo [1] test set. The HumanML3D test set evaluates the in-domain performance. And the Mixamo dataset consists of abundant artist-created animations and human-annotated descriptions, offering a wide variety of diverse and dynamic motions, utilized to compare the zero-shot performance across the domains.

**Evaluation Metrics** are summarized as four parts. (1) Frechet Inception Distance (FID) is our principal metric to evaluate the feature distributions between the generated and real motions. (2) Motion-retrieval precision (R-Precision) calculates the text and motion Top 3 matching accuracy under feature space. (3) CLIP-score. Borrowing from text-to-image synthesis [64], we use CLIP-score to evaluate zero-shot text-motion consistency by measuring cosine similarity in CLIP space. To do so, we train a motion encoder using the training set of HumanML3D and Mixamo to extract motion features that are aligned to text embeddings in CLIP space as same as [76]. (4) Diversity is assessed by calculating variance through features.

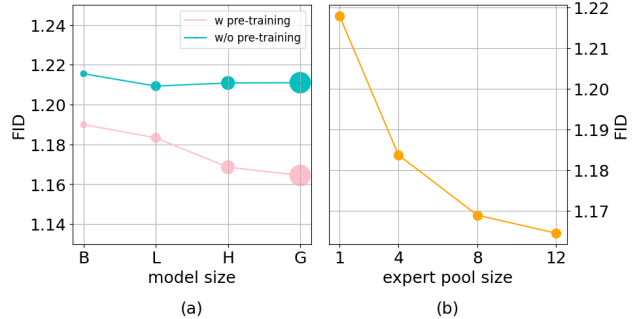


Figure 6. Quantitative evaluation on pre-training, model size, and expert pool size. (a) Models *w pre-training* show consistently improved performance over *w/o pre-training*, and *w pre-training* models, which benefit from large-scale motion data, improve with increasing model size. (b) Larger expert pool sizes improve the performance.

### 4.2. Comparison

We compare our approach with various state-of-the-art methods. Specifically, we apply conditional motion diffusion models, including MDM [77], MLD [12], MotionDiffuse [89], VAE-based model MotionCLIP [76], and autoregressive model T2M-GPT [87]. Besides, we also apply MAA [8] based on the text-pose alignment method, and MotionGPT [33] utilizing motion-language pre-training. Some results on HumanML3D are borrowed from their own paper. The quantitative results are presented in Tab. 2. Our method demonstrates the best text-to-motion alignment (R-Precision) in the in-domain evaluation. Moreover, among all diffusion-based methods, our model achieves the best FID. In terms of zero-shot performance, our approach achieves the best FID and CLIP-score, outperforming previous methods. This suggests a superior capability for high-quality motion generation and effective matching with zero-shot text prompts. As depicted in Fig. 5 of the qualitative results, our model demonstrates the capability to generate motions that show more realism, and better alignment with each motion characteristic description, whether in the sentence or phrase.

### 4.3. Ablation Study

To examine the specific contributions of our novel OMG model architecture, we conduct a series of ablation studies focusing on the roles of pre-training, model scale, and multiple expert controllers. Additionally, we undertake an in-depth analysis of our MoC Block. These evaluations utilize the out-of-domain dataset Mixamo [1], providing a robust testbed to ascertain the effectiveness of the technical designs on the zero-shot performance of our model.

**Effect of Pre-training and Model Scale.** Our investigation into the impact of pre-training involves training several vari-

Methods	FID ↓	CLIP-score ↑	Diversity →
Real	0.106±.003	0.648±.001	2.665±.022
(1) Cross-attn + FFN	1.252±.006	0.552±.003	2.576±.025
(2) w/o Zero Conv	1.339±.005	0.535±.004	2.695±.014
(3) w/o Attention Mask	1.246±.008	0.557±.003	<b>2.647±.026</b>
ours (complete)	<b>1.164±.009</b>	<b>0.588±.002</b>	2.632±.021

Table 3. Quantitative evaluation on MoC block. The damping performance of the three variants of our model highlights the effectiveness of our MoC block technical designs.

ant models without pre-training (*w/o\_pre-training*) across four distinct scales (see Tab. 1). These models are then compared against counterparts with the pre-training process (*w\_pre-training*). As illustrated in Fig. 6a, *w\_pre-training* consistently outperforms *w/o\_pre-training*, achieving lower FID across all sizes. This indicates that our model’s performance on zero-shot motion generation is effectively enhanced through pre-training. Meanwhile, we observe that the performance of *w\_pre-training* improves with increasing model size. Qualitative analysis, as shown in Fig. 7a, reveals that the “Gollum” characteristic of the generated motions becomes more pronounced with larger models. This suggests that larger model sizes enhance the overall quality and alignment of the generated motions.

**Effect of Expert Pool Size.** To systematically analyze the impact of multiple experts, we train four distinct variants of the OMG model, sweeping over different expert pool sizes ( $K = 1, 4, 8,$  and  $12$  respectively). The results, as depicted in Fig. 6b, demonstrate a discernible decrease in FID as expert pool size increases. This trend suggests that the larger expert hypothesis space enhances the experts’ ability to align sub-motions with their respective text embeddings.

**Evaluation of the MoC Block.** In order to assess the effectiveness of our proposed MoC block, we explore its performance through several variant configurations: (1) Cross-attn + FFN, where the architecture is pruned to include only a cross-attention layer and a feed-forward network; (2) w/o Zero Conv, substituting zero convolutions with standard convolutional layers initialized using Gaussian distributions; and (3) w/o Attention Mask, omitting the multiplication of attention masks with the output of the expert layer. These variant models are rigorously compared against our original OMG model. Quantitative results, as shown in Tab. 3, demonstrate that the exclusion of specific components within the MoC Block leads to a worse result in both the FID and CLIP-score. This observation underscores the integral contributions of these components to the motion generation process. Qualitative analysis, presented in Fig. 7b, indicates noticeable deficiencies in the motion generation of all three variants: variant (1) is not able to clearly depict “spin” motions; variant (2) exhibits less expressive

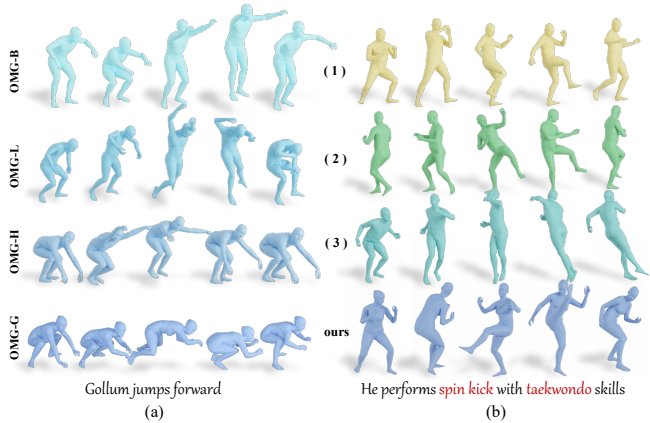


Figure 7. Qualitative evaluation on model sizes (a) and MoC block (b). Models with larger sizes effectively comprehend richer out-of-domain motion features to present better motion expressions. Besides, our technical designs effectively improve the alignment with the input texts.

motion features; and variant (3) inadequately captures the “kick” motion. In contrast, our OMG model demonstrates a superior ability to more accurately align motion with the text prompt.

## 5. Conclusion

In this paper, we present a novel text-to-motion generation framework, OMG, that combines the advantages of conditional generative models and text-pose alignment methods. It carefully tailors pretrain-then-finetune paradigm into text-to-motion generation. The pre-training stage leverages a large amount of unlabeled motion data to train a powerful unconditional diffusion model that ensures the realism and diversity of the generated motions. The fine-tuning stage introduces motion ControlNet including the proposed novel text conditioning block called Mixture-of-Controllers. With the cross-attention mechanism and text-token-specific experts, it adaptively aligns the sub-motion features to the text embeddings in the CLIP space in an MoE fashion. The extensive experiments demonstrate that OMG achieves state-of-the-art zero-shot performance on text-to-motion generation. We believe it is a significant step towards open-vocabulary motion generation of human characters, with wide potential applications in movies, games, robotics, and VR/AR.

## Acknowledgement

This work was supported by National Key R&D Program of China (2022YFF0902301), Shanghai Local college capacity building program (22010502800). We also acknowledge support from Shanghai Frontiers Science Center of Human-centered Artificial Intelligence (ShangHAI).



## References

- [1] Adobe. <https://www.mixamo.com/>. Accessed:2023-9-28. **2, 7**
- [2] Gunjan Aggarwal and Devi Parikh. Dance2music: Automatic dance-driven music generation. *arXiv preprint arXiv:2107.06252*, 2021. **2**
- [3] Hyemin Ahn, Timothy Ha, Yunho Choi, Hwiyeon Yoo, and Songhwa Oh. Text2action: Generative adversarial synthesis from language to action. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5915–5920. IEEE, 2018. **2**
- [4] 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**
- [5] Tenglong Ao, Qingzhe Gao, Yuke Lou, Baoquan Chen, and Libin Liu. Rhythmic gesticulator: Rhythm-aware co-speech gesture synthesis with hierarchical neural embeddings. *ACM Transactions on Graphics (TOG)*, 41(6):1–19, 2022. **2**
- [6] Tenglong Ao, Zeyi Zhang, and Libin Liu. Gesturediffuclip: Gesture diffusion model with clip latents. *ACM Transactions on Graphics (TOG)*, 42(4):1–18, 2023. **2**
- [7] Joao Pedro Araújo, Jiaman Li, Karthik Vetrivel, Rishi Agarwal, Jiajun Wu, Deepak Gopinath, Alexander William Clegg, and Karen Liu. Circle: Capture in rich contextual environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21211–21221, 2023. **7**
- [8] Samaneh Azadi, Akbar Shah, Thomas Hayes, Devi Parikh, and Sonal Gupta. Make-an-animation: Large-scale text-conditional 3d human motion generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15039–15048, 2023. **1, 3, 7**
- [9] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators. 2024. **1**
- [10] Zhongang Cai, Daxuan Ren, Ailing Zeng, Zhengyu Lin, Tao Yu, Wenjia Wang, Xiangyu Fan, Yang Gao, Yifan Yu, Liang Pan, et al. Humman: Multi-modal 4d human dataset for versatile sensing and modeling. In *European Conference on Computer Vision*, pages 557–577. Springer, 2022. **7**
- [11] 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**
- [12] 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. **1, 2, 7**
- [13] Rishabh Dabral, Muhammad Hamza Mughal, Vladislav Golyanik, and Christian Theobalt. Mofusion: A framework for denoising-diffusion-based motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9760–9770, 2023. **1**
- [14] 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**
- [15] David Eigen, Marc’Aurelio Ranzato, and Ilya Sutskever. Learning factored representations in a deep mixture of experts. *arXiv preprint arXiv:1312.4314*, 2013. **3**
- [16] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research*, 23(1):5232–5270, 2022. **2, 3, 5**
- [17] Kevin Frans, Lisa Soros, and Olaf Witkowski. Clipdraw: Exploring text-to-drawing synthesis through language-image encoders. *Advances in Neural Information Processing Systems*, 35:5207–5218, 2022. **2**
- [18] 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**
- [19] 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. **1, 2, 3**
- [20] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5152–5161, 2022. **2, 7**
- [21] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. In *The Twelfth International Conference on Learning Representations*, 2023. **1**
- [22] Ikhsanul Habibie, Mohamed Elgharib, Kripasindhu Sarkar, Ahsan Abdullah, Simbarashe Nyatsanga, Michael Neff, and Christian Theobalt. A motion matching-based framework for controllable gesture synthesis from speech. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–9, 2022. **2**
- [23] 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, 7**
- [24] Jiaao He, Jiezhong Qiu, Aohan Zeng, Zhilin Yang, Jidong Zhai, and Jie Tang. Fastmoe: A fast mixture-of-expert training system. *arXiv preprint arXiv:2103.13262*, 2021. **3**
- [25] Yannan He, Anqi Pang, Xin Chen, Han Liang, Minye Wu, Yuexin Ma, and Lan Xu. Challencap: Monocular 3d capture of challenging human performances using multi-modal references. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11400–11411, 2021. **2**
- [26] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. **2**

- [27] Daniel Holden, Taku Komura, and Jun Saito. Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017. 3
- [28] Fangzhou Hong, Mingyuan Zhang, Liang Pan, Zhongang Cai, Lei Yang, and Ziwei Liu. Avatarclip: Zero-shot text-driven generation and animation of 3d avatars. *ACM Transactions on Graphics (TOG)*, 41(4):1–19, 2022. 1, 2
- [29] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. 7
- [30] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991. 3
- [31] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 867–876, 2022. 2
- [32] Nikolay Jetchev. Clipmatrix: Text-controlled creation of 3d textured meshes. *arXiv preprint arXiv:2109.12922*, 2021. 2
- [33] Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. Motiongpt: Human motion as a foreign language. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 3, 7
- [34] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. Flame: Free-form language-based motion synthesis & editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8255–8263, 2023. 2
- [35] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. Dancing to music. *Advances in neural information processing systems*, 32, 2019. 2
- [36] Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. In *International Conference on Learning Representations*, 2020. 3
- [37] 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
- [38] Margaret Li, Suchin Gururangan, Tim Dettmers, Mike Lewis, Tim Althoff, Noah A Smith, and Luke Zettlemoyer. Branch-train-merge: Embarrassingly parallel training of expert language models. In *First Workshop on Interpolation Regularizers and Beyond at NeurIPS 2022*, 2022. 3
- [39] 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, 7
- [40] Han Liang, Yannan He, Chengfeng Zhao, Mutian Li, Jingya Wang, Jingyi Yu, and Lan Xu. Hybridcap: Inertia-aid monocular capture of challenging human motions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1539–1548, 2023. 2, 7
- [41] Han Liang, Wenqian Zhang, Wenxuan Li, Jingyi Yu, and Lan Xu. Intergen: Diffusion-based multi-human motion generation under complex interactions. *arXiv preprint arXiv:2304.05684*, 2023. 2, 7
- [42] Junfan Lin, Jianlong Chang, Lingbo Liu, Guanbin Li, Liang Lin, Qi Tian, and Chang-wen Chen. Being comes from not-being: Open-vocabulary text-to-motion generation with wordless training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 23222–23231, 2023. 1, 2
- [43] Xiao Lin and Mohamed R Amer. Human motion modeling using dvgs. *arXiv preprint arXiv:1804.10652*, 2018. 2
- [44] Haiyang Liu, Zihao Zhu, Naoya Iwamoto, Yichen Peng, Zhengqing Li, You Zhou, Elif Bozkurt, and Bo Zheng. Beat: A large-scale semantic and emotional multi-modal dataset for conversational gestures synthesis. In *European Conference on Computer Vision*, pages 612–630. Springer, 2022. 7
- [45] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866. 2023. 3
- [46] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. 4
- [47] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. 1, 7
- [48] Ian Mason, Sebastian Starke, and Taku Komura. Real-time style modelling of human locomotion via feature-wise transformations and local motion phases. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 5(1), 2022. 7
- [49] Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. Text2mesh: Text-driven neural stylization for meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13492–13502, 2022. 2
- [50] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. 4
- [51] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2085–2094, 2021. 2
- [52] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 2, 4
- [53] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human

- bodies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14314–14323, 2021. [2](#)
- [54] 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](#)
- [55] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10985–10995, 2021. [2](#)
- [56] Mathis Petrovich, Michael J Black, and Gül Varol. Temos: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision*, pages 480–497. Springer, 2022. [1](#), [2](#)
- [57] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [1](#), [2](#), [6](#)
- [58] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–16. IEEE, 2020. [4](#)
- [59] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *International conference on machine learning*, pages 1060–1069. PMLR, 2016. [2](#)
- [60] 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](#)
- [61] Yiming Ren, Chengfeng Zhao, Yannan He, Peishan Cong, Han Liang, Jingyi Yu, Lan Xu, and Yuexin Ma. Lidar-aid inertial poser: Large-scale human motion capture by sparse inertial and lidar sensors. *IEEE Transactions on Visualization and Computer Graphics*, 29(5):2337–2347, 2023. [2](#)
- [62] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts. *Advances in Neural Information Processing Systems*, 34:8583–8595, 2021. [3](#)
- [63] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. [1](#)
- [64] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. [7](#)
- [65] Aditya Sanghi, Hang Chu, Joseph G Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshah. Clip-forged: Towards zero-shot text-to-shape generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18603–18613, 2022. [2](#)
- [66] Yoni Shafir, Guy Tevet, Roy Kapon, and Amit Haim Bermano. Human motion diffusion as a generative prior. In *The Twelfth International Conference on Learning Representations*, 2023. [2](#)
- [67] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations*, 2016. [2](#), [3](#), [5](#)
- [68] Sheng Shen, Zhewei Yao, Chunyuan Li, Trevor Darrell, Kurt Keutzer, and Yuxiong He. Scaling vision-language models with sparse mixture of experts. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023. [3](#)
- [69] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2020. [6](#)
- [70] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2020. [2](#)
- [71] 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](#)
- [72] Sebastian Starke, Yiwei Zhao, Taku Komura, and Kazi Zaman. Local motion phases for learning multi-contact character movements. *ACM Transactions on Graphics (TOG)*, 39(4):54–1, 2020. [3](#)
- [73] 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](#), [3](#)
- [74] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. [4](#)
- [75] Omid Taheri, Nima Ghorbani, Michael J Black, and Dimitrios Tzionas. Grab: A dataset of whole-body human grasping of objects. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 581–600. Springer, 2020. [7](#)
- [76] Guy Tevet, Brian Gordon, Amir Hertz, Amit H Bermano, and Daniel Cohen-Or. Motionclip: Exposing human motion generation to clip space. In *European Conference on Computer Vision*, pages 358–374. Springer, 2022. [1](#), [2](#), [7](#)
- [77] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*, 2022. [1](#), [2](#), [3](#), [4](#), [7](#)
- [78] Matthew Trumble, Andrew Gilbert, Charles Malleon, Adrian Hilton, and John Collomosse. Total capture: 3d human pose estimation fusing video and inertial sensors.

- In *Proceedings of 28th British Machine Vision Conference*, pages 1–13, 2017. 7
- [79] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 2
- [80] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3835–3844, 2022. 2
- [81] Yiming Xie, Varun Jampani, Lei Zhong, Deqing Sun, and Huaizu Jiang. Omnicontrol: Control any joint at any time for human motion generation. In *The Twelfth International Conference on Learning Representations*, 2023. 2
- [82] Liang Xu, Ziyang Song, Dongliang Wang, Jing Su, Zhicheng Fang, Chenjing Ding, Weihao Gan, Yichao Yan, Xin Jin, Xiaokang Yang, et al. Actformer: A gan-based transformer towards general action-conditioned 3d human motion generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2228–2238, 2023. 2
- [83] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In *European Conference on Computer Vision*, pages 736–753. Springer, 2022. 2
- [84] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16010–16021, 2023. 2
- [85] 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
- [86] Jiaxu Zhang, Junwu Weng, Di Kang, Fang Zhao, Shaoli Huang, Xuefei Zhe, Linchao Bao, Ying Shan, Jue Wang, and Zhigang Tu. Skinned motion retargeting with residual perception of motion semantics & geometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13864–13872, 2023. 3
- [87] 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. 1, 2, 7
- [88] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 1, 2, 5, 6
- [89] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 1, 7
- [90] Siwei Zhang, Qianli Ma, Yan Zhang, Zhiyin Qian, Taein Kwon, Marc Pollefeys, Federica Bogo, and Siyu Tang. Ego-body: Human body shape and motion of interacting people from head-mounted devices. In *European Conference on Computer Vision*, pages 180–200. Springer, 2022. 7
- [91] Yan Zhang, Mohamed Hassan, Heiko Neumann, Michael J. Black, and Siyu Tang. Generating 3d people in scenes without people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [92] Qingcheng Zhao, Pengyu Long, Qixuan Zhang, Dafei Qin, Han Liang, Longwen Zhang, Yingliang Zhang, Jingyi Yu, and Lan Xu. Media2face: Co-speech facial animation generation with multi-modality guidance. *arXiv preprint arXiv:2401.15687*, 2024. 2