

Morph: A Motion-free Physics Optimization Framework for Human Motion Generation

Zhuo Li^{*1}, Mingshuang Luo^{*2,3,4}, Ruibing Hou^{†2}, Xin Zhao⁵,

Hao Liu¹, Hong Chang^{2,4}, Zimo Liu³, Chen Li¹

¹WeChat, Tencent Inc, ²State Key Laboratory of AI Safety, Institute of Computing Technology, CAS, China

³Peng Cheng Laboratory, China, ⁴University of Chinese Academy of Sciences, China

⁵MoE Key Laboratory of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

Abstract

Human motion generation has been widely studied due to its crucial role in areas such as digital humans and humanoid robot control. However, many current motion generation approaches disregard physics constraints, frequently resulting in physically implausible motions with pronounced artifacts such as floating and foot sliding. Meanwhile, training an effective motion physics optimizer with noisy motion data remains largely unexplored. In this paper, we propose **Morph**, a **Motion-Free physics optimization framework**, consisting of a Motion Generator and a Motion Physics Refinement module, for enhancing physical plausibility without relying on expensive real-world motion data. Specifically, the motion generator is responsible for providing large-scale synthetic, noisy motion data, while the motion physics refinement module utilizes these synthetic data to learn a motion imitator within a physics simulator, enforcing physical constraints to project the noisy motions into a physically-plausible space. Additionally, we introduce a prior reward module to enhance the stability of the physics optimization process and generate smoother and more stable motions. These physically refined motions are then used to fine-tune the motion generator, further enhancing its capability. This collaborative training paradigm enables mutual enhancement between the motion generator and the motion physics refinement module, significantly improving practicality and robustness in real-world applications. Experiments on both text-to-motion and music-to-dance generation tasks demonstrate that our framework achieves state-of-the-art motion quality while improving physical plausibility drastically.

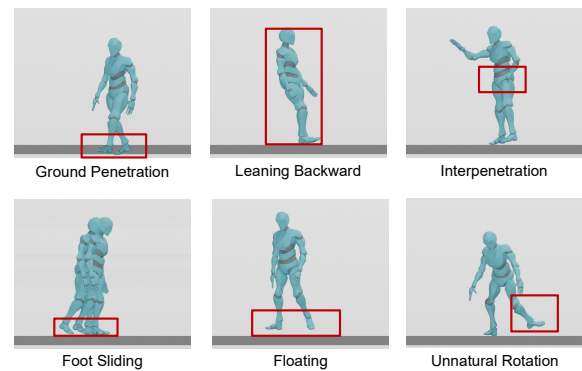


Figure 1. Examples of physical inconsistencies in generations.

1. Introduction

Accurate human motion synthesis plays a crucial role in applications such as robotics, video gaming, and virtual reality [1, 9, 38, 46, 47]. Recent advances in AI have enabled motion generation under various control conditions, including textual descriptions and music inputs. Current approaches can be categorized into three main paradigms: conditional diffusion models [15, 34, 35, 41, 43], conditional autoregressive models [12, 19, 26, 48, 49], and generative masked modeling [8, 27]. These models have significantly improved the quality of generated motion by better capturing the complex multimodal distribution of human movement.

However, most existing motion generation approaches overlook a fundamental aspect of human motion: the *laws of physical*. Although these generation models excel in capturing statistical distribution of human motions, they lack explicit mechanisms to enforce physical constraints. As a result, the generated motions frequently exhibit pronounced artifacts such as ground penetration, leaning backward, interpenetration, foot sliding, floating and unnatural rotation, as shown in Fig. 1. Given human sensitivity to even slight physical inconsistencies, these physically implausible mo-

*Equal contribution

†Corresponding author

tions hinder many real-world applications such as animation and virtual reality [11, 30, 41].

Recently, a few studies [10, 41] have attempted to enhance the physical plausibility of motion generation, making the generated motions more realistic. For example, PhysDiff [41] and Reindiffuse [10] incorporate physical constraints into the denoising diffusion process. However, these approaches face two main limitations. On the one hand, although these works [10, 41] introduce promising approaches for motion physics optimization, their generalizability and transferability across different types of motion generation models and tasks have not been fully demonstrated or explored. On the other hand, these works apply physics and diffusion iteratively, embedding physics optimization into multiple steps of diffusion process. The frequent execution of physics optimization significantly increases computational costs during inference. Additionally, prior studies [20–22] have explored physics-based human motion imitation using a physics simulator, typically relying on large-scale real 3D motion data as priors to enhance accuracy and naturalness in imitation learning. However, collecting large-scale realistic motion data is challenging. A natural question arises: *Is it possible to learn an efficient, model-agnostic physical optimizer without relying on real motion data?*

To achieve this, we propose a **Motion-free physics optimization framework**, namely **Morph**. As shown in Fig. 2, Morph consists of two main modules: a Motion Generator (**MG**) that can be any existing motion generator, and a Motion Physics Refinement (**MPR**) module for enhancing physical plausibility. Morph employs a two-stage optimization process. *In the first stage*, using large-scale noisy motion data produced by the motion generator, MPR module is optimized to project input motions into a physically-plausible space. Specifically, MPR module designs a motion imitator that controls a character agent to mimic the given noisy motions within a physics simulator. The simulator enforces multiple physical constraints, effectively reducing artifacts such as floating and foot sliding. To further ensure the naturalness of simulated motions, MPR module introduces a prior reward module to align the distribution of physics-refined motions with that of input motions. This prior reward module can provide smoother and more stable rewards. Feedback signals from both physics simulator and prior reward module guide the optimization of the motion imitator via reinforcement learning. *In the second stage*, we address the limitations of pretrained motion generators, which may not be trained on data that adheres to physical constraints and thus produce suboptimal results. To this end, we finetune the motion generator using the physic-refined motion data produced by MPR module, further enhancing its generation capability. During inference, the fine-tuned motion generator and MPR module work in

tandem to generate physics-plausible and high-quality motions.

By fostering the synergy between motion generation and motion physics refinement, morph provides a scalable, efficient, and generalizable solution to realistic human motion synthesis. We evaluate our morph framework on text-to-motion and music-to-dance tasks, using diffusion-based, autoregressive, and generative mask models. Extensive experiments demonstrate that morph achieves significant reduces physical errors, while maintaining competitive motion quality across different generators and tasks, despite not being trained on real motion data.

Our contributions can be summarized as follows:

- We introduce Morph, a generalizable framework that enhances physical plausibility and generation quality across various motion generation models and tasks. Also, a prior reward module is designed to accelerate imitation learning and improve motion realism in physics simulation.
- We design a collaborative training paradigm, where the Motion Generator provides large-scale synthetic motion data to enhance the Motion Physics Refinement module, which, in turn, refines motion quality by feeding back physics-refined data.
- Extensive qualitative and quantitative experiments demonstrate that Morph achieves competitive physical plausibility and generation performance.

2. Related Work

Human Motion Generation. Motion generation is a long-history task that can be conditioned on various signals, such as text description, music and action [2, 4, 7, 8, 16, 17, 31, 32, 35, 39, 42, 45, 50]. Our work specifically focuses on text-to-motion [7, 8, 27, 34, 42, 44] and music-to-dance generation [15, 23, 31, 35]. Mainstream approaches can be roughly divided into three categories: diffusion-based methods [3, 15, 34, 35, 41, 43], autoregressive models [12, 18, 19, 31, 42, 48, 49] and generative masked modeling [8, 27]. For instance, MDM [34] uses a transformer-based diffusion model with conditional text representations extracted from CLIP [29]. T2M-GPT [42] designs a conditional autoregressive transformer model based on VQ-VAE [36] and GPT [28, 37]. MoMask [8] employs residual vector quantization and generative masked transformers to iteratively generate motions. For music-to-dance generation, Bailando [31] predicts discrete token sequences conditioned on music and uses an autoregressive transformer to regenerate the dance sequence. However, existing motion generation models often produce physically implausible motions, as they overlook physical laws during training. In contrast, our proposed MPR module effectively enforces physical constraints, significantly enhancing the physical plausibility of generated motions.

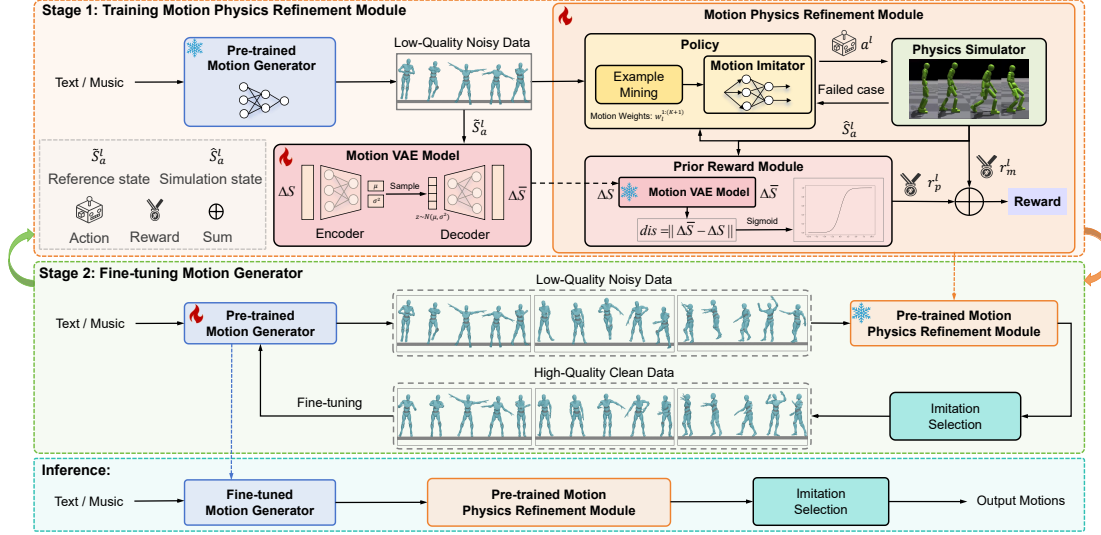


Figure 2. An overview of the Morph framework. Morph comprises a Motion Generator and a Motion Physics Refinement module. Morph employs a two-stage training process: Motion Physics Refinement module training and Motion Generator fine-tuning. And a Imitation Selection Operation is employed to ensure the motion quality after physics refinement. The solid curved arrows on the left and right (in orange and green) represent the iterative, collaborative optimization between Stage 1 and Stage 2.

Physically Plausible Motion Generation. Physical plausibility refers to the degree to which generated motions adhere to physical rules, such as foot sliding, foot-ground contact, and body leaning. Generating physically-plausible motions is crucial for many real-world applications, such as animation and virtual reality. Recently, a few works have attempted to address this challenge [10, 35, 41]. For example, [10, 41] integrate physical constraints into the diffusion process, iteratively applying physics and diffusion to maintain alignment with motion data distribution while enhancing physical realism. EDGE [35] introduces auxiliary losses to align specific aspects of physical realism, specifically for generating more physically plausible dances. However, these methods rely on real motion data and have not fully demonstrated generalizability across different motion generation models and tasks, limiting their flexibility and scalability.

3. Method

To enhance physical plausibility of motion generation without relying on real motion data, we propose a motion-free physics optimization framework, namely morph. As illustrated in Fig. 2, morph consists of two key modules: a Motion Generator (MG) that could be any existing pre-trained motion generation model, and a Motion Physics Refinement module (MPR) independent of Motion Generator. Morph employs a two-stage training process. In the first stage, the Motion Generator produces large-scale noisy motion data (Sec. 3.1), which is subsequently used to train the Physics Refinement module to map input motions into a physically plausible space (Sec. 3.2). In the second stage, these

physics-refined motions, in turn, are used to fine-tune Motion Generator, further enhancing its capabilities (Sec 3.3). Through these two-stage training process, morph drastically reduces physical errors while improving the overall quality of generated motions.

3.1. Noisy Motion Data Generation

Different from [10, 41] that rely on real motion data, our goal is to develop a *motion-free* physical optimization framework. To achieve this, we utilize an existing pre-trained motion generation model to produce large-scale, noisy motion data for training the physics refinement module. As shown in Fig. 2, given control signals c (e.g., text or music), the pre-trained motion generator is employed to generate a motion sequence $\tilde{\mathbf{x}}^{1:L}$, as follows:

$$\tilde{\mathbf{x}}^{1:L} = f_{\xi}(c), \text{ where, } \tilde{\mathbf{x}}^{1:L} = \{\tilde{\mathbf{x}}^l = [\boldsymbol{\theta}^l, \mathbf{p}^l]\}_{l=1}^L, \quad (1)$$

where f_{ξ} represents the motion generator with parameters ξ , $\tilde{\mathbf{x}}^l$ denotes the l^{th} pose of the synthetic motion sequence, represented by the joint rotations $\boldsymbol{\theta}$ and positions \mathbf{p} . Notably, our framework, morph, is agnostic to the specific instantiation of f_{ξ} , allowing it to be generally compatible with various pretrained motion generators.

3.2. Physics-Based Motion Refinement

Since most existing motion generation models lack explicit physical constraints, they often produce motions with noticeable artifacts, as shown in Fig. 1. To address this, the Motion Physics Refinement (MPR) module is tasked with projecting the generated motion $\tilde{\mathbf{x}}^{1:L}$, which disregards the

laws of physics, into a physically-plausible motion $\hat{\mathbf{x}}^{1:L}$. As shown in Fig. 2, the MPR module consists of three components: a motion imitator, a physics simulator and a motion prior reward model. Specifically, the motion imitator controls a simulated character to mimic the input motion $\tilde{\mathbf{x}}^{1:L}$ within the physics simulator. The resulting motion $\hat{\mathbf{x}}^{1:L}$ from the physics simulator is considered physically plausible, as it adheres to the laws of physics. Additionally, the prior reward model penalize the non-human motions further ensuring the naturalness and realism of the physics-refined motions.

Motion Imitator Learning. Motion Imitator Learning can be formulated as a Markov decision process, represented by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R})$ of states, actions, transition dynamics and reward function. Formally, the motion imitator is described by a policy $\pi(\mathbf{a}^l | \mathbf{s}^l)$, which specifies the probability distribution of selecting an action $\mathbf{a}^l \in \mathcal{A}$ given the current state $\mathbf{s}^l \in \mathcal{S}$. The physic simulator, in turn, defines the transition dynamics $\mathcal{T}(\mathbf{s}^{l+1} | \mathbf{s}^l, \mathbf{a}^l)$, which determines the next state \mathbf{s}^{l+1} based on current state \mathbf{s}^l and action \mathbf{a}^l . Specifically, starting from an initial state \mathbf{s}^1 , a character agent acts in the physic simulator according to policy $\pi(\mathbf{a}^l | \mathbf{s}^l)$, iteratively sampling action \mathbf{a}^l . Then, the physic simulator, governed by the transition dynamics $\mathcal{T}(\mathbf{s}^{l+1} | \mathbf{s}^l, \mathbf{a}^l)$, generates the next state \mathbf{s}^{l+1} , from which the simulated pose $\hat{\mathbf{x}}^{l+1}$ is derived. By running the policy for L steps, we can obtain the simulated motion sequence $\hat{\mathbf{x}}^{1:L}$. In implementation, we adopt a standard RL algorithm Proximal Policy Optimization (PPO) to train the motion imitator π , where a reward is assigned based on how well the simulated motion $\hat{\mathbf{x}}^{1:L}$ aligns with input motion $\tilde{\mathbf{x}}^{1:L}$. In the following, we elaborate on the design of states, rewards, policy and actions.

States. The simulation state \mathbf{s}^l consists of the input motion’s next pose, along with the differences between the input *next* pose and the *current* simulated pose across multiple aspects, including joint rotation, position, velocity and angular velocity. This difference information informs the policy of the pose residuals that require compensation, enabling the system to better align the simulated motion with the input motion. Formally, the state is defined as $\mathbf{s}^l = [\tilde{\boldsymbol{\theta}}^{l+1}, \tilde{\mathbf{p}}^{l+1}, \tilde{\boldsymbol{\theta}}^{l+1} - \hat{\boldsymbol{\theta}}^l, \tilde{\mathbf{p}}^{l+1} - \hat{\mathbf{p}}^l, \tilde{\mathbf{v}}^{l+1} - \hat{\mathbf{v}}^l, \tilde{\boldsymbol{\omega}}^l - \hat{\boldsymbol{\omega}}^l]$. Here, $\boldsymbol{\theta}$, \mathbf{p} , \mathbf{v} and $\boldsymbol{\omega}$ represent the joint rotation, position, joint velocity, and angular velocity, respectively, $[\cdot]/[\hat{\cdot}]$ denote the quantities of the input/simulated motions, and $[\cdot]^l$ denotes the quantities at the l^{th} timestep.

Actions. We use the target joint angles of proportional derivative (PD) controllers as the action representation, where the action \mathbf{a}^l specifies the PD target to enable robust motion imitation.

Policy. Following [10, 22, 41], we employ a parameterized Gaussian policy $\pi(\mathbf{a}^l | \mathbf{s}^l) = \mathcal{N}(\mu_\phi(\mathbf{s}^l), \Sigma)$, where

the mean action $\mu_\phi(\mathbf{s}^l)$ is the output by our motion imitator, a simple multi-layer perceptron network with parameter ϕ , and Σ is a fixed diagonal covariance matrix.

Rewards. The reward function is designed to encourage the simulated motion to match the input motion. At each timestep l , the reward r^l consists of a mimic reward r_m^l and a energy penalty r_e^l and a prior reward r_p^l , formulating as $r^l = r_m^l + r_e^l + r_p^l$. The mimic reward r_m^l includes joint rotation $\boldsymbol{\theta}$, position \mathbf{p} , joint velocity \mathbf{v} and angular velocity $\boldsymbol{\omega}$, as follows:

$$r_m^l = w_\theta \exp[-\alpha_\theta \|\tilde{\boldsymbol{\theta}}^l - \hat{\boldsymbol{\theta}}^l\|] + w_p \exp[-\alpha_p \|\tilde{\mathbf{p}}^l - \hat{\mathbf{p}}^l\|] + w_v \exp[-\alpha_v \|\tilde{\mathbf{v}}^l - \hat{\mathbf{v}}^l\|] + w_\omega \exp[-\alpha_\omega \|\tilde{\boldsymbol{\omega}}^l - \hat{\boldsymbol{\omega}}^l\|], \quad (2)$$

where $w_\theta, w_p, w_v, w_\omega$, are weighting factors, $\alpha_\theta, \alpha_p, \alpha_v, \alpha_\omega$ are scaling factors, $[\cdot]/[\hat{\cdot}]$ denote the quantities of the input/simulated pose, and $\|\cdot\|$ is the L1 norm.

The energy penalty is computed as:

$$r_e^l = -0.0005 \cdot \|\hat{\mathbf{v}}^l \hat{\boldsymbol{\omega}}^l\|_2^2, \quad (3)$$

where $\hat{\mathbf{v}}^l$ and $\hat{\boldsymbol{\omega}}^l$ denotes the joint torque and angular velocity of the simulated pose in the l^{th} timestep, and $\|\cdot\|_2$ is the L2 norm.

Prior Reward Module. Relying solely on physical simulators and imitation learning often results in non-human-like actions, characterized by a mechanical feel and unrealistic rotations. Previous works [22, 25] have attempted to address these issues by employing adversarial-based methods trained on high-quality motion data to incorporate human motion knowledge, improving the naturalness of the motion. However, this adversarial-based method often suffers from training instability and produces rewards that are not smooth enough, which also prolongs the overall imitation learning training process. To address this issue, we propose a Motion VAE Reward model as a prior to constrain the motion generated by the physics simulator, making it more natural. As shown in Fig. 2, to accelerate the training speed of motion imitation learning, we first train a lightweight Motion VAE model using motion data generated by the motion generator. This allows the model to effectively learn the distribution of the generated data. The objective function for optimizing the Motion VAE model is as follows:

$$\mathcal{L}(\theta, \phi; \mathbf{s}) = \mathbb{E}_{q_\phi(\mathbf{z} | \Delta \mathbf{s})} [\log p_\theta(\Delta \mathbf{s} | \mathbf{z})] - D_{KL}(q_\phi(\mathbf{z} | \Delta \mathbf{s}) \| p(\mathbf{z})) \quad (4)$$

where θ and ϕ denote the parameters of VAE, $\Delta \mathbf{s}$ refer to the difference in state between two consecutive time steps for reference motion. $D_{KL}(\cdot, \cdot)$ refers to the KL divergence.

As shown in Fig. 2, after the Motion VAE model is trained, the pre-trained Motion VAE model is applied to compute rewards as follows:

$$r_p^l = 1 - \frac{1}{1 + e^{-|\Delta \bar{\mathbf{s}} - \Delta \mathbf{s}|}}, \quad (5)$$

where Δs and $\Delta \bar{s}$ refer to the difference in state between two consecutive time steps for simulated motion and the reconstructed simulated motion, respectively. $\|\cdot\|$ is the L1 norm.

Initialization and Early Termination. Similar to [24], a starting point is randomly selected from a motion clip for imitation. To improve training efficiency and accelerate convergence, we terminate the episode when the MPJPE (Mean Per Joint Position Error) between the state of input pose and that of simulated pose exceeds 0.5 meter.

Hard Negative Mining. As training progresses, the motion imitator gradually learns to imitate simple motion sequences. However, more challenging examples in the large-scale motion dataset may be overlooked, limiting the model’s ability to handle difficult samples. To address this, we implement a Hard Negative Mining process that identifies motions where the physical simulator fails to imitate as hard samples. Specifically, a *dynamic* weight is assigned to each motion sequence for sampling, doubling whenever imitation fails. This process progressively increases the focus on challenging samples, guiding the imitator to effectively learn from difficult examples.

3.3. Motion Generator Fine-tuning

Existing motion generators are often trained with limited real motion data, which may limit their capabilities. However, collecting large-scale real 3D motion data is costly. So, we leverage a large volume of synthetic, physically-plausible motion data provided by MPR module, to further enhance the capabilities of the motion generator.

As shown in Fig. 2, in the second stage, we further fine-tune the motion generator using the physics-refined motions produced by the MPR module. Notably, since the physical simulator cannot replicate non-grounded motions (*e.g.*, sitting on a chair or swimming), such simulated motions may deviate from the true data distribution. To this end, we apply an *Imitation Selection Operation* to filter out simulated data of non-grounded motions. Specifically, we calculate the average per-joint position error (MPJPE) between the samples before and after physical optimization. A threshold τ is set to determine whether to accept the physically refined motion $\hat{\mathbf{x}}^{1:L}$ (with $\text{MPJPE} < \tau$) or input motions $\tilde{\mathbf{x}}^{1:L}$ (with $\text{MPJPE} > \tau$). The selected data is then paired with the original condition signals (*e.g.*, text or music). Ultimately, a large-scale, physically plausible motion data is constructed for fine-tuning the motion generator.

Following [8, 34], we use the mean squared error to optimize the motion generator. Denote the selected motion sequence paired with the condition as $(\mathbf{x}^{1:L}, c)$, the objective function for motion generator f_ξ is defined as:

$$\mathcal{L}_{\text{MG}}(\xi) = \mathbb{E} \left[\left\| \mathbf{x}^{1:L} - f_\xi(c) \right\|_2^2 \right] \quad (6)$$

where ξ is the parameters of the motion generator.

Multi-Round Iterative Optimization. In Fig. 2, we define the complete two-stage process as a single training round. To maximize the synergy between the MG and MPR modules, we repeat this process multiple times, using the model weights from each round as initialization for the next, thus creating a multi-round iterative training paradigm. We present the results of this iterative training in the Appendix, demonstrating its effectiveness and potential.

Inference. As shown in Fig 2, after the two-stage optimization, the finetuned motion generator (Stage 2) and the trained MPR module (Stage 1) are combined to perform inference. To mitigate simulation error in non-grounded motions, the *imitation selection operation* described above is used as a post-process step.

4. Experiment

Extensive experiments evaluate the performance of our Morph across multiple motion generation tasks and datasets. Specifically, we assess Morph on two motion generation tasks, text-to-motion and music-to-dance. As Morph is agnostic to specific instantiation of generation models, we combine Morph with three types of generation models for text-to-motion, *i.e.*, diffusion-based models, including MotionDiffuse [42] and MDM [33], autoregressive-based models, including T2M-GPT [42] and generative make modeling, including MoMask [8]. For music-to-dance, we combine Morph with the diffusion-based model EDGE [35] and autoregressive model Bailando [31]. The Appendix provides additional visualizations and experimental results showcasing Morph’s performance on varying amounts of noisy motion data, different and threshold τ in imitation selection operation, and multi-round optimization of the MPR module and motion generator.

4.1. Datasets and Evaluation Metrics

HumanML3D. HumanML3D [6] is a large-scale 3D human motion-language dataset, consisting of 14,616 motion clips and 44,970 motion descriptions annotations. The total motion duration is 28.59 hours, with each motion clip downsampled to 20 FPS and accompanied by 3-4 textual descriptions. The dataset is split into training, validation and test sets in an 80%, 5% and 15% ratios.

AIST++. AIST++ [14] is a human dance-music dataset, which includes 992 pieces of high-quality dance sequence across ten dance genres, each dance paired with a corresponding music. The dance durations range from 7.4 to 48 seconds, with a frame rate of 60 FPS. The dataset is split into training and evaluation sets, with allocations of 952 and 40, respectively.

Evaluation Metrics. We evaluate motion generation from two perspectives: common generation metrics and physi-

Table 1. Ablation study on Morph-MoMask (combined with MoMask [8] generator) for text-to-motion task on HumanML3D dataset. IS: imitation selection operation; Prior: using prior reward training MPR module; Adv: using adversarial reward training MPR module; Energy: using energy reward training MPR module; Real Data: using real motion data training MPR module; FT: fine-tuning Motion Generator with physics-refined motions. The arrows (\uparrow / \downarrow) indicate that higher/smaller values are better.

Model Number	Methods						Common Generation Metrics		Physical Plausibility Metrics				
	IS	Prior	Adv	Energy	Real Data	FT	RTOP-3 \uparrow	FID \downarrow	PFC \downarrow	Penetrate \downarrow	Float \downarrow	Skate \downarrow	IFR \downarrow
A	Baseline (only motion generator MoMask [8])						0.807	0.045	1.058	23.152	10.660	5.262	-
B		✓		✓			0.800	0.172	0.821	0.000	2.233	0.016	0.0152
C			✓	✓			0.792	0.194	0.852	0.000	2.272	0.020	0.0155
D	✓			✓			0.782	0.276	0.735	0.000	2.554	0.032	0.0362
E	✓	✓					0.795	0.163	0.711	0.000	2.362	0.024	0.0270
F	✓	✓		✓			0.802	0.074	0.669	0.000	2.268	0.011	0.0153
G	✓	✓		✓	✓		0.792	0.179	0.734	0.000	2.420	0.015	0.0327
H	✓	✓		✓		✓	0.816	0.041	0.647	0.000	2.141	0.010	0.0149

cal plausibility metrics. First, for text-to-motion generation, we follow [7, 44] and use three standard metrics: *Frechet Inception Distance (FID)*, which measures the distance between generated and ground-truth motion distributions; *R-Precision (RTOP-1, RTOP-2, RTOP-3)*, which computes retrieval accuracy of generated motions based on input text; and *Diversity (Div)*, which measures the diversity of generated motions. For music-to-dance generation, following [15, 31, 40], we use five metrics: FID_k and Div_k for kinetic features (denoted ‘ k ’); FID_g and Div_g for geometric features (denoted ‘ g ’); and *Beat Align Score (BAS)*, which measures the alignment between input music and generated dances.

Next, we present the physical plausibility metrics. Following [35, 41], we use four physic-based metrics to assess the physical plausibility of generated motions: *Penetrate* that measures ground penetration; *Float* that measures floating; *Skate* that measures foot sliding; *Physical Foot Contact score (PFC)* that measures the realism of foot-ground contact. Additionally, we introduce an imitation selection metric, *imitation failure rate (IFR)*, which calculates the failure rate of the physics refinement module in imitating motion.

4.2. Implementation Details

Training Setup. We implement Morph based on PyTorch. The NVIDIA’s Isaac Gym is used as the physics simulator. During the MPR module’s training phase (Stage 1), we employ the Adam optimizer, with a batch size of 64 and a learning rate of 4×10^{-5} . During the finetuning of motion generator (Stage 2), we follow the original training setup, only adjusting the learning rate to 1×10^{-5} . For hyperparameters setting, we set w_θ , w_q , w_v and w_w (Eq. 2) to 0.5, 0.3, 0.1 and 0.1, α_θ , α_q , α_v and α_w (Eq. 2) to 100, 10, 0.1 and 0.1 and the imitation selection threshold τ is set to 0.5. All our experiments are conducted on 8 Tesla V100 GPUs. Further details on implementation and hyperparameter analysis are provided in the Appendix.

Data Preprocessing. In text-to-motion generation, some motion sequences involve environmental interactions, such

as sitting down, climbing stairs, and swimming, which cannot be simulated by current physical simulator. Therefore, we annotate each text description in HumanML3D [6] based on its semantic content. Text Descriptions involving interaction-based motions are labeled as ‘0’, while other descriptions are labeled as ‘1’. Text labeled as ‘1’ will be used to train the MPR module in Stage 1. We will release this annotation on HumanML3D in the final version.

4.3. Ablation Studies

In this section, we conduct ablation studies to validate the effectiveness of each part of our method. We use MoMask [8] as the motion generator in Morph, denoted as Morph-MoMask, throughout the ablation studies. The ablation results are shown in Tab. 1.

Effectiveness of training MPR module using only generated data. We investigate the effectiveness of using only synthetic motion data from two perspectives: its overall effectiveness and its performance relative to training with real motion data. First, we compare the baseline model A (only the MoMask generator) and model F which combines MoMask with the trained MPR module. In Tab. 1, model F outperforms model A on all physical metrics (0.669/1.058 on PFC, 0.0/23.152 on Penetrate, 2.268/10.66 on Float, 0.011/5.262 on Skate), while maintaining competitive performance on generation metrics. These results confirm the feasibility of training the MPR module solely with synthetic data. Next, we compare model F with model G, which uses real motion data to train the MPR module. In Tab. 1, model G achieves inferior performance compared to model F. This can be attributed to the significant domain gap between the training data (real motions) and test data (generated motions). This gap hinders the MPR module trained on real data from effectively adapting to the generated motion data, leading to performance decline.

Effectiveness of prior reward and energy reward. When training the MPR module, we add an prior reward and energy reward alongside the common imitation task reward. In Tab. 1, comparing model F with model D, which omits

Table 2. Comparison results for text-to-motion task on HumanML3D dataset. Morph is combined with different types of motion generators. MG: Motion Generator; MPR: Motion Physics Refinement module; FT: fine-tuning motion generator with the physics-refined motion data. † denotes Morph without fine-tuning the motion generator (only Stage 1 training)

MG	MPR	FT	Common Generation Metrics				Physical Plausibility Metrics				
			RTOP-1 ↑	RTOP-3 ↑	FID ↓	Diversity ↑	PFC ↓	Penetrate ↓	Float ↓	Skate ↓	IFR ↓
PhysDiff w/ MD [41]	-	-	-	0.780	0.551	-	-	0.898	1.368	0.423	-
PhysDiff w/ MDM [41]	-	-	-	0.631	0.433	-	-	0.998	2.601	0.512	-
Reindiffuse [10]	-	-	-	0.622	0.385	-	-	<u>0.000</u>	0.711	0.058	-
MDM [33]	-	-	0.455	0.749	0.489	9.920	0.811	17.384	17.502	3.540	-
Morph-MDM†	✓	-	0.449	0.748	0.496	9.736	0.715	0.000	2.271	0.017	0.0176
Morph-MDM	✓	✓	0.471	0.757	0.482	<u>9.865</u>	0.697	0.000	2.258	0.016	0.0153
MotionDiffuse(MD) [42]	-	-	0.491	0.782	0.630	9.410	0.533	16.670	5.698	3.419	-
Morph-MD†	✓	-	0.488	0.782	0.649	9.523	<u>0.433</u>	0.000	1.162	0.028	0.0181
Morph-MD	✓	✓	0.494	0.788	0.568	9.493	0.414	0.000	<u>0.853</u>	0.029	0.0163
T2M-GPT [42]	-	-	0.491	0.775	0.116	9.761	0.998	72.250	8.918	7.801	-
Morph-T2M-GPT†	✓	-	0.490	0.774	0.133	9.648	0.764	0.000	2.702	0.042	0.0205
Morph-T2M-GPT	✓	✓	0.498	0.784	0.105	9.771	0.748	0.000	2.700	0.039	0.0195
MoMask [8]	-	-	<u>0.521</u>	<u>0.807</u>	<u>0.045</u>	9.641	1.058	23.152	10.660	5.262	-
Morph-MoMask†	✓	-	0.516	0.802	0.074	9.578	0.669	0.000	2.268	<u>0.011</u>	0.0153
Morph-MoMask	✓	✓	0.525	0.816	0.041	9.689	0.645	0.000	2.141	0.010	0.0149

Table 3. Comparison results on common generation metrics for text-to-motion on HumanML3D dataset.

Methods	RTOP-1 ↑	RTOP-2 ↑	RTOP-3 ↑	FID ↓	MM-Dist ↓	Diversity ↑	MModality ↑
MDM [33]	0.455	0.645	0.749	0.489	3.330	9.920	2.290
MotionDiffuse (MD) [44]	0.491	0.681	0.782	0.630	3.113	9.410	1.553
MLD [3]	0.481	0.673	0.772	0.473	3.196	9.724	<u>2.413</u>
T2M-GPT [42]	0.491	0.680	0.775	0.116	3.118	9.761	1.856
AttT2M [49]	0.499	0.690	0.786	0.112	3.038	9.700	2.452
MMM [27]	0.515	0.708	0.804	0.089	2.926	9.577	1.226
MoMask [8]	0.521	0.713	0.807	<u>0.045</u>	2.958	9.641	1.241
BAMM [26]	0.525	<u>0.720</u>	<u>0.814</u>	0.055	2.919	9.717	1.687
Morph-MDM	0.471	0.664	0.757	0.482	3.103	<u>9.865</u>	2.179
Morph-MD	0.494	0.691	0.788	0.568	3.000	9.494	1.582
Morph-T2M-GPT	0.498	0.690	0.784	0.105	2.994	9.771	1.910
Morph-MoMask	0.525	0.726	0.816	0.041	<u>2.920</u>	9.689	1.655

the prior reward, we observe a performance decline in both generation and physical metrics on model D. The significant drop in FID metric indicates that, without the prior reward, the simulated motions lack the distribution constraints of the input motions, leading to a deviation from the intended motion distribution. The B, which leverages prior rewards, significantly outperforms model C, which uses adversarial rewards, across various generation and physics-based metrics. This highlights the superior effectiveness of the proposed prior reward over the adversarial reward. We compare model F with E, which omits the energy reward. As shown in Tab. 1, model F outperforms model E in generation metrics. We argue that introducing energy reward helps suppress high-frequency jitter during the humanoid robot control process, resulting in more natural motions.

Effectiveness of imitation selection strategy. Like [10, 41], the proposed Morph focuses on addressing physical inconsistencies in non-interactive and ground-contact human

motions. In training and testing data, there are some interactive and non-grounded motions, such as sitting on a chair, swimming, or climbing stairs. To ensure the quality of the motions after refining, we design an imitation selection operation to filter out the simulations of such motions. In Tab. 1, we compare model F with B, which omits the imitation selection operation. As shown, model F outperforms B in both generation and physic metrics, validating the effectiveness of the imitation selection strategy.

Effectiveness of finetuning motion generator using simulated motions. Furthermore, we compare model F and model H, which finetunes the motion generator using the physically refined motions produced by the MPR module. As shown in Tab. 1, model H achieves further gains in both generation and physic metrics, with RTOP-3 increasing by 1.4% and FID improving by 0.033. These results suggest that the proposed MPR module can, in turn, enhance the performance of the motion generator.

Table 4. Comparison results for music-to-dance on AIST++ dataset. † denotes Morph without fine-tuning motion generator.

MG	MPR	FT	Common Generation Metrics					Physical Plausibility Metrics				
			FID _k ↓	FID _g ↓	Div _k ↑	Div _g ↑	BAS ↑	PFC ↓	Penetrate ↓	Float ↓	Skate ↓	IFR ↓
FACT [13]	-	-	35.35	22.11	5.94	6.18	0.2209	-	-	-	-	-
TM2D [5]	-	-	19.01	20.09	9.45	6.36	0.2049	-	-	-	-	-
EDGE [35]	-	-	42.16	22.12	3.96	4.61	0.2334	0.610	76.490	60.982	7.906	-
Bailando [31]	-	-	28.16	<u>9.62</u>	7.83	6.34	0.2332	0.074	31.183	12.650	4.466	-
Morph-EDGE†	✓	-	43.01	23.69	4.30	5.18	0.2229	0.302	0.000	3.720	<u>0.011</u>	0.0092
Morph-EDGE	✓	✓	39.24	19.82	4.93	5.36	0.2329	0.295	0.000	3.519	0.009	0.0082
Morph-Bailando†	✓	-	35.48	12.03	7.70	<u>6.49</u>	<u>0.2409</u>	<u>0.044</u>	0.000	<u>2.076</u>	0.023	<u>0.0064</u>
Morph-Bailando	✓	✓	<u>26.25</u>	9.38	<u>7.88</u>	6.62	0.2412	0.043	0.000	2.057	0.021	0.0055

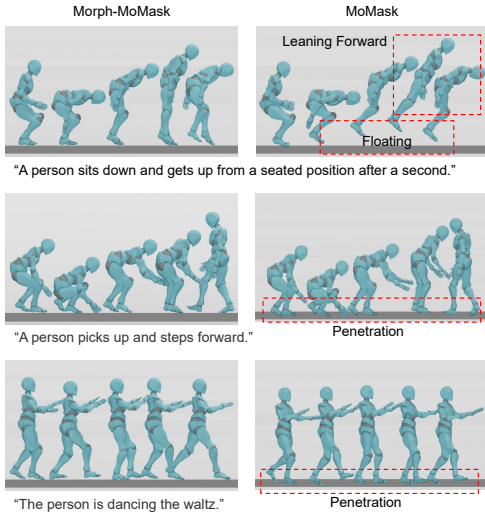


Figure 3. Qualitative comparison between our Morph-MoMask and MoMask in text-to-motion task. Morph-MoMask significantly reduces physical artifacts such as leaning forward, floating and penetration.

4.4. Evaluation with Different Motion Generators

We evaluate the adaptability of our Morph with different motion generators, including MDM [33], Motiondiffuse [44], T2M-GPT [42], and MoMask [8]. As shown in Tab. 2, Morph significantly enhances physical fidelity while maintaining competitive generation metrics across different motion generation models. For example, after integrating Morph with MDM (Morph-MDM), RTOP-1 metric improved by 0.016. In terms of physical plausibility metrics, penetration dropped to zero, float decreased from 17.502 to 2.258, and skate reduced from 3.540 to 0.016. Similar performance improvements are observed with other motion generators. The consistent gains across different generators highlight the versatility of our Morph framework.

4.5. Comparisons with State-of-the-arts

Main results on Text-to-Motion Generation. On the text-to-motion dataset, HumanML3D, we compare Morph framework with other state-of-the-art methods. The com-

parison results are shown in Tab. 2 and Tab. 3. As shown, Morph achieves significant gains in physical metrics while maintaining competitive performance in generation metrics, demonstrating its capability to enhance physical plausibility. Fig. 3 illustrates the generated motions from both our Morph-MoMask and MoMask. As shown, motions generated by MoMask often exhibit physically unrealistic artifacts, such as penetration, floating, and leaning forward. In contrast, Morph-MoMask effectively reduces these artifacts, producing motions that are both physically plausible and realistic.

Main results on Music-to-Dance Generation. On the music-to-dance dataset, AIST++, we compare Morph combined with EDGE [35] and Bailando [31] generators against other state-of-the-art methods. As shown in Tab. 4, Morph significantly improves physical metrics compared to existing methods. For generation metrics, our Morph-EDGE model achieves the best results in multiple metrics, including FID_g, Div_g, and BAS, while also demonstrating competitive performance in FID_k. These results validate the superiority of our framework in music-to-dance generation.

5. Conclusion

In this paper, we present Morph, a model-agnostic physical optimization framework designed to enhance physical plausibility in motion generation without relying on costly real-world data. We begin by using a pretrained motion generator (MG) to synthesize large-scale noisy motion data. Then, we introduce a Motion Physics Refinement (MPR) module that leverages this synthesized data to train a motion imitator enforcing physical constraints. These physically refined motions are used to fine-tune the motion generator, further enhancing its capabilities. To stabilize the physics optimization process, we introduce a prior reward module. This collaborative training paradigm enables mutual enhancement between MG and MPR module, improving practicality and robustness. Our framework is compatible with various motion generation models for both text-to-motion and music-to-dance tasks. Extensive experiments show that Morph substantially improves physical plausibility while achieving competitive generation quality.

References

- [1] Bharat Lal Bhatnagar, Xianghui Xie, Ilya A Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. BEHAVE: Dataset and method for tracking human object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15935–15946, 2022. 1
- [2] Zhi Cen, Huaijin Pi, Sida Peng, Zehong Shen, Minghui Yang, Shuai Zhu, Hujun Bao, and Xiaowei Zhou. Generating human motion in 3d scenes from text descriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1855–1866, 2024. 2
- [3] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18000–18010, 2023. 2, 7
- [4] Sisi Dai, Wenhao Li, Haowen Sun, Haibin Huang, Chongyang Ma, Hui Huang, Kai Xu, and Ruizhen Hu. Interfusion: Text-driven generation of 3d human-object interaction. *arXiv preprint arXiv:2403.15612*, 2024. 2
- [5] Kehong Gong, Dongze Lian, Heng Chang, Chuan Guo, Zihang Jiang, Xinxin Zuo, Michael Bi Mi, and Xinchao Wang. Tm2d: Bimodality driven 3d dance generation via music-text integration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9942–9952, 2023. 8
- [6] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5152–5161, 2022. 5, 6
- [7] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5152–5161, 2022. 2, 6
- [8] Chuan Guo, Yuxuan Mu, Muhammad Gohar Javed, Sen Wang, and Li Cheng. Momask: Generative masked modeling of 3d human motions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1900–1910, 2024. 1, 2, 5, 6, 7, 8
- [9] Vladimir Guzov, Aymen Mir, Torsten Sattler, and Gerard Pons-Moll. Human POSEitioning system (HPS): 3D human pose estimation and self-localization in large scenes from body-mounted sensors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4318–4329, 2021. 1
- [10] Gaoge Han, Mingjiang Liang, Jinglei Tang, Yongkang Cheng, Wei Liu, and Shaoli Huang. Reindiffuse: Crafting physically plausible motions with reinforced diffusion model. *arXiv preprint arXiv:2410.07296*, 2024. 2, 3, 4, 7
- [11] Ludovic Hoyet, Rachel McDonnell, and Carol O’Sullivan. Push it real: Perceiving causality in virtual interactions. *ACM Transactions on Graphics (TOG)*, 31(4):1–9, 2012. 2
- [12] 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, 2
- [13] 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. 8
- [14] Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa. Learn to dance with aist++: Music conditioned 3d dance generation, 2021. 5
- [15] Ronghui Li, YuXiang Zhang, Yachao Zhang, Hongwen Zhang, Jie Guo, Yan Zhang, Yebin Liu, and Xiu Li. Lodge: A coarse to fine diffusion network for long dance generation guided by the characteristic dance primitives. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1524–1534, 2024. 1, 2, 6
- [16] Yong Li, Menglin Liu, Zhen Cui, Yi Ding, Yuan Zong, Wenming Zheng, Shiguang Shan, and Cuntai Guan. Decoupled doubly contrastive learning for cross domain facial action unit detection. *IEEE Transactions on Image Processing*, 2025. 2
- [17] Yong Li, Yi Ren, Xuesong Niu, Yi Ding, Xiu-Shen Wei, and Cuntai Guan. Beyond overfitting: Doubly adaptive dropout for generalizable au detection. *IEEE Transactions on Affective Computing*, 2025. 2
- [18] Shunlin Lu, Ling-Hao Chen, Ailing Zeng, Jing Lin, Ruimao Zhang, Lei Zhang, and Heung-Yeung Shum. Humantomato: Text-aligned whole-body motion generation. In *Forty-first International Conference on Machine Learning*, 2024. 2
- [19] Mingshuang Luo, Ruibing Hou, Zhuo Li, Hong Chang, Zimo Liu, Yaowei Wang, and Shiguang Shan. M³gpt: An advanced multimodal, multitask framework for motion comprehension and generation. *Advances in Neural Information Processing Systems*, 37:28051–28077, 2024. 1, 2
- [20] Zhengyi Luo, Ryo Hachiuma, Ye Yuan, and Kris Kitani. Dynamics-regulated kinematic policy for egocentric pose estimation. *Advances in Neural Information Processing Systems*, 34:25019–25032, 2021. 2
- [21] Zhengyi Luo, Jinkun Cao, Josh Merel, Alexander Winkler, Jing Huang, Kris Kitani, and Weipeng Xu. Universal humanoid motion representations for physics-based control. *arXiv preprint arXiv:2310.04582*, 2023.
- [22] Zhengyi Luo, Jinkun Cao, Alexander W. Winkler, Kris Kitani, and Weipeng Xu. Perpetual humanoid control for real-time simulated avatars. In *International Conference on Computer Vision (ICCV)*, 2023. 2, 4
- [23] Matthew Marchellus and In Kyu Park. M2c: Concise music representation for 3d dance generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3126–3135, 2023. 2
- [24] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic. *ACM Transactions on Graphics*, page 1–14, 2018. 5
- [25] 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*, page 1–20, 2021. 4

- [26] Ekkasit Pinyoanuntapong, Muhammad Usama Saleem, Pu Wang, Minwoo Lee, Srijan Das, and Chen Chen. Bamm: Bidirectional autoregressive motion model. *arXiv preprint arXiv:2403.19435*, 2024. 1, 7
- [27] Ekkasit Pinyoanuntapong, Pu Wang, Minwoo Lee, and Chen Chen. Mmm: Generative masked motion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1546–1555, 2024. 1, 2, 7
- [28] Alec Radford. Improving language understanding by generative pre-training. 2018. 2
- [29] 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, 2021. 2
- [30] Paul SA Reitsma and Nancy S Pollard. Perceptual metrics for character animation: sensitivity to errors in ballistic motion. In *ACM SIGGRAPH 2003 Papers*, pages 537–542. 2003. 2
- [31] Li Siyao, 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 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11050–11059, 2022. 2, 5, 6, 8
- [32] Mikihiro Tanaka and Kent Fujiwara. Role-aware interaction generation from textual description. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15999–16009, 2023. 2
- [33] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H. Bermano. Human motion diffusion model, 2022. 5, 7, 8
- [34] 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*, 2023. 1, 2, 5
- [35] Jonathan Tseng, Rodrigo Castellon, and Karen Liu. Edge: Editable dance generation from music. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 448–458, 2023. 1, 2, 3, 5, 6, 8
- [36] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 2
- [37] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 2
- [38] Timo Von Marcard, Roberto Henschel, Michael J Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3D human pose in the wild using imus and a moving camera. In *Proceedings of the European Conference on Computer vision*, pages 601–617, 2018. 1
- [39] Hongwei Yi, Hualin Liang, Yifei Liu, Qiong Cao, Yandong Wen, Timo Bolkart, Dacheng Tao, and Michael J Black. Generating holistic 3d human motion from speech. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 469–480, 2023. 2
- [40] Jiashuo Yu, Yaohui Wang, Xinyuan Chen, Xiao Sun, and Yu Qiao. Long-term rhythmic video soundtracker. In *International Conference on Machine Learning*, pages 40339–40353. PMLR, 2023. 6
- [41] 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 (ICCV)*, pages 16010–16021, 2023. 1, 2, 3, 4, 6, 7
- [42] 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. *arXiv preprint arXiv:2301.06052*, 2023. 2, 5, 7, 8
- [43] Mingyuan Zhang, Xinying Guo, Liang Pan, Zhongang Cai, Fangzhou Hong, Huirong Li, Lei Yang, and Ziwei Liu. Remodiffuse: Retrieval-augmented motion diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 364–373, 2023. 1, 2
- [44] 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. 2, 6, 7, 8
- [45] Mingyuan Zhang, Daisheng Jin, Chenyang Gu, Fangzhou Hong, Zhongang Cai, Jingfang Huang, Chongzhi Zhang, Xinying Guo, Lei Yang, Ying He, et al. Large motion model for unified multi-modal motion generation. *arXiv preprint arXiv:2404.01284*, 2024. 2
- [46] Siwei Zhang, Yan Zhang, Qianli Ma, Michael J Black, and Siyu Tang. PLACE: Proximity learning of articulation and contact in 3D environments. In *International Conference on 3D Vision*, pages 642–651, 2020. 1
- [47] 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*, pages 6194–6204, 2020. 1
- [48] Yaqi Zhang, Di Huang, Bin Liu, Shixiang Tang, Yan Lu, Lu Chen, Lei Bai, Qi Chu, Nenghai Yu, and Wanli Ouyang. Motiongpt: Finetuned llms are general-purpose motion generators. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7368–7376, 2024. 1, 2
- [49] Chongyang Zhong, Lei Hu, Zihao Zhang, and Shihong Xia. Att2m: Text-driven human motion generation with multi-perspective attention mechanism. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 509–519, 2023. 1, 2, 7
- [50] Wentao Zhu, Xiaoxuan Ma, Dongwoo Ro, Hai Ci, Jinlu Zhang, Jiaxin Shi, Feng Gao, Qi Tian, and Yizhou Wang. Human motion generation: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2