

AAMDMD: Accelerated Auto-regressive Motion Diffusion Model

Tianyu Li
Georgia Tech

Calvin Qiao
UBC

Guanqiao Ren
BUAA

KangKang Yin
SFU

Sehoon Ha
Georgia Tech




Figure 1. We introduce the Accelerated Auto-regressive Motion Diffusion Model (AAMDMD), a novel framework designed to synthesize diverse and high-quality character motions at interactive rates.

Abstract

Interactive motion synthesis is essential in creating immersive experiences in entertainment applications, such as video games and virtual reality. However, generating animations that are both high-quality and contextually responsive remains a challenge. Traditional techniques in the game industry can produce high-fidelity animations but suffer from high computational costs and poor scalability. Trained neural network models alleviate the memory and speed issues, yet fall short on generating diverse motions. Diffusion models offer diverse motion synthesis with low memory usage, but require expensive reverse diffusion processes. This paper introduces the Accelerated Auto-regressive Motion Diffusion Model (AAMDMD), a novel motion synthesis framework designed to achieve quality, diversity, and efficiency all together. AAMDMD integrates Denoising Diffusion GANs as a fast Generation Module, and an Auto-regressive Diffusion Model as a Polishing Module. Furthermore, AAMDMD operates in a lower-dimensional embedded space rather than the full-dimensional pose space, which reduces the training complexity as well as further improves the performance. We show that AAMDMD outperforms existing methods in motion quality, diversity, and runtime efficiency, through comprehensive quantitative analyses and visual comparisons. We also demonstrate the effectiveness of each algorithmic component through ablation studies.

1. Introduction

The landscape of interactive motion synthesis, particularly in the realm of video games, has seen a notable expansion. Today’s AAA titles boast tens of thousands of unique characters in real-time, all needing to be contextually animated [18]. Therefore, the efficiency of motion synthesis has emerged as a critical focus of research in the field of computer animation. Motion Matching [25], a prevalent technique for industry-grade animation, was first developed by UbiSoft for the game “*For Honor*” [1]. The main objective of Motion Matching (MM) is to identify the most contextually suitable animation in a large dataset based on manually defined motion features. This approach, while capable of yielding responsive high-quality animations, is computationally intensive and scales poorly with respect to the size of the dataset.

Alternatively, trained neural networks have emerged to reduce the memory footprints and enhance runtime performance. However, these models possess their own challenges, such as unstable convergence at training time and compromised motion quality at testing time. Recently, diffusion-based generative models have revolutionized content creation, thanks to their power to create diverse high-quality content with lean memory demands. However, standard diffusion models are often impractical for time-critical applications, due to their poor run-time performance caused by expensive reverse diffusion processes.

We introduce the Accelerated Auto-regressive Motion

Diffusion Model (*AAMDM*), a novel framework crafted to generate diverse high-fidelity motion sequences without the need for prolonged reverse diffusion. Diffusion-based transition models naturally produce diverse multi-modal motion would be too slow for interactive applications. To overcome this challenge, our *AAMDM* framework mainly adopts two synergistic modules: a Generation Module, for rapid initial motion drafting using Denoising Diffusion GANs; and a Polishing Module, for quality improvements using an Auto-regressive Diffusion Model with just two additional denoising steps. Another distinctive feature of *AAMDM* is its operation in a learned lower-dimensional latent space rather than the traditional full pose space, further accelerating the training process.

We evaluate our algorithm on the LaFAN1 [13] dataset and demonstrate its capability of synthesizing diverse high-quality motions at interactive rates. Our method outperforms a number of baseline algorithms, such as LMM [25], MotionVAE [31], and AMDM [50], using various quantitative evaluation metrics. Furthermore, we conduct an analysis on an artificial multimodal dataset. This analysis confirms that our model can successfully capture the multi-modal transition model and is better suited for diverse and intricate motion synthesis tasks. Finally, we perform ablation studies to justify various design choices within our framework.

In summary, our primary contributions are as follows:

- We introduce *AAMDM*, a novel diffusion-based framework capable of generating extended motion sequences at interactive rates. The key idea is to combine the strengths of Denoising Diffusion GANs and Auto-regressive Diffusion Models in a compact embedded space.
- We conduct thorough comparative analyses between *AAMDM* and various established benchmarks using multiple metrics for measuring motion quality, diversity, and runtime efficiency. Together with our ablation studies, we provide a deep understanding of our algorithm with respect to alternative prior arts.
- We showcase novel high-quality multi-modal motions synthesized from our model, some impossible to achieve by previous methods, such as following a user-controlled root trajectory with diverse arm movements.

2. Related Work

2.1. Data-Driven Kinematic Motion Synthesis

The quest to create virtual characters that move naturally stands as a fundamental challenge in computer animation. Graph-based approach structures motion data into a graph and employs search algorithms to retrieve contextually appropriate animations [4, 22, 26, 27, 35, 46]. It offers high-fidelity motion, but its scalability is curtailed by substantial memory demands and search times.

Statistical methods have been devised to encapsulate motions within numerical models, such as linear, kernel-based, and neural network categories. Linear models represent poses with low-dimensional vectors, but they often fail to encompass the full spectrum of human movement [5, 21, 47]. Kernel-based models, including Radial Basis Functions (RBF) and Gaussian Processes (GP), embrace the non-linearity in motion data [11, 28, 36, 37, 40, 45, 59]. However, these methods are memory-intensive, especially when managing large covariance matrices.

The neural network paradigm has gained prominence for its scalability and efficiency at runtime [8, 12, 14, 31, 39, 41, 42, 54, 55, 60]. Innovative neural architectures have been proposed to better capture motion sequences within datasets, such as those adjusting weights according to a phase variable [19], employing gating mechanisms [64], and extracting periodic latent features [55]. Nevertheless, these models predominantly focus on locomotion and character’s leg movement, leaving room for broader exploration.

2.2. Generative Diffusion Model

Generative diffusion models are a groundbreaking class of algorithms that learn to replicate data distributions through the reverse of diffusion processes [16, 51, 52]. In conditional generation scenarios, innovations such as classifier-guided diffusion [7] and classifier-free guidance [15] have been introduced, offering fine-tuned control over the balance between diversity and fidelity. Applications of diffusion models span across image and video synthesis to robotics [2, 15–17, 20, 23, 38, 53, 58].

Recent adaptations of diffusion models for motion synthesis have been particularly promising, with efforts aimed at generating 3D human motion from textual descriptions [24, 56, 65]. Enhancements to these models have come through novel architectural designs [56], the integration of geometric losses [56], and the incorporation of physical guidance mechanisms [63]. Additionally, the synthesis of human dance motions from audio signals has been explored, with models using auditory cues to direct the generative process [3, 6, 34, 57]. However, the latency inherent in diffusion models, often taking considerable time to generate brief motion clips, precludes their application in real-time settings. The work of Shi et al. [50] represents a significant stride towards curtailing inference times through a reduced number of denoising steps.

2.3. Accelerating Diffusion Model

The typically slow sampling speeds of diffusion models are primarily attributed to the extensive series of denoising steps required. A range of strategies has been suggested to expedite this process, such as the application of knowledge distillation techniques [33], the employment of adaptive noise scheduling [48], and the design of single-step

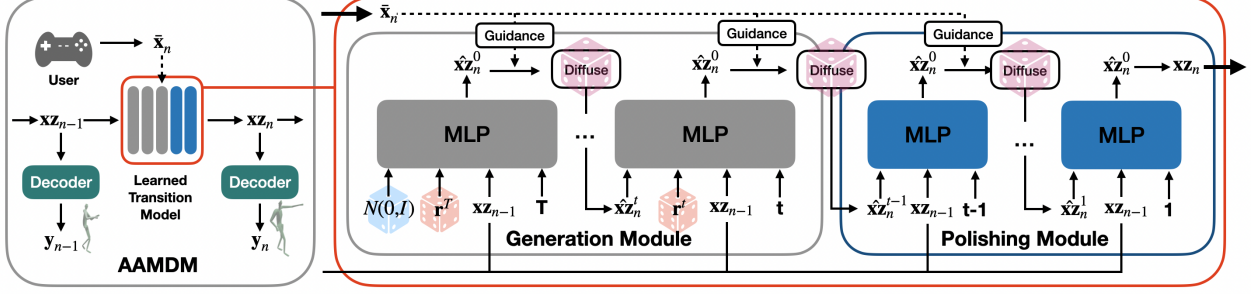


Figure 2. Overview of AAMDM. AAMDM incorporates three pivotal components for better motion quality and faster inference. Firstly, it models transitions within a low-dimensional embedded space $\mathbf{xz} \in \mathbf{XZ}$. Secondly, the framework features a *Generation* module, which employs Denoising Diffusion GANs. This module is responsible for efficiently generating initial drafts of motion sequences. Lastly, a *Polishing* module, which utilizes an Auto-regressive Diffusion Model, refines these initial drafts. A full-pose vector \mathbf{y}_n is then reconstructed from the corresponding embedded vector \mathbf{xz}_n using the learned decoder D^{AE} .

denoising distributions as conditional energy-based models [9]. Integrating reinforcement learning with diffusion models has also been proposed to decrease the number of reverse diffusion steps needed [50]. Nevertheless, such methods have often had to contend with either diminished sample quality or expensive multiple generation steps. The introduction of Denoising Diffusion GANs [62] is a notable innovation, integrating the strengths of diffusion models with Generative Adversarial Networks to concurrently address sample quality, generation speed, and mode coverage [10]. In this work, we have employed this technique to enhance the diffusion process for fast motion synthesis.

3. Method

The architecture of our Accelerated Auto-regressive Motion Diffusion Model (AAMDM) is illustrated in Figure 2. AAMDM incorporates three key components: transition in a low-dimensional embedded space, a *Generation* module with Denoising Diffusion GANs for efficient draft generation, and a *Polishing* module with Auto-regressive Diffusion Model for refining the draft.

In the following subsections, we will first explain the construction of the low-dimensional embedded space (Section 3.1). Then, we will describe the foundation of the Polishing (Section 3.2) and Generation modules (Section 3.3), namely the auto-regressive diffusion model and denoising diffusion GANs. Next, we will provide the design of the Generation and Polishing modules (Section 3.4), followed by an explanation of how the sampling procedure is guided to follow user’s commands (Section 3.5). Finally, we will provide a model representation (Section 3.6).

3.1. Constructing Embedded Space

Current learning-based methods for motion synthesis typically target to capture pose transitions in the full-body space, which often complicates learning and violates kinematic constraints intrinsic to the character’s morphology.

We introduce a compact embedded vector $\mathbf{xz} \in \mathbf{XZ}$ to replace a full-body pose $\mathbf{y} \in \mathbf{Y}$, where \mathbf{x} denotes an engineered feature and \mathbf{z} a learned latent vector.

An autoencoder is employed to learn the optimal embedded space, where an Encoder network $E^{AE}(\mathbf{y}) \rightarrow \mathbf{z}$ maps pose vectors to latent vectors, and a Decoder network $D^{AE}(\mathbf{xz}) \rightarrow \hat{\mathbf{y}}$ reconstructs poses from the encoded vectors. On top of learned features \mathbf{z} , we extract manual feature \mathbf{x} as well. The networks are trained jointly to minimize both the perceptual discrepancy losses, $L_{val}^{D,E}$ and $L_{vel}^{D,E}$, and a regularization loss $L_{reg}^{D,E}$:

$$L_{val}^{D,E} = \|\hat{\mathbf{y}} \ominus \mathbf{y}\| + \|F(\hat{\mathbf{y}}) \ominus F(\mathbf{y})\| \quad (1)$$

$$L_{vel}^{D,E} = \left\| \frac{F(\hat{\mathbf{y}}_0) \ominus F(\hat{\mathbf{y}}_1)}{\delta n} - \frac{F(\mathbf{y}_0) \ominus F(\mathbf{y}_1)}{\delta n} \right\| \quad (2)$$

$$L_{reg}^{D,E} = \|\mathbf{z}\|_2^2 \quad (3)$$

$$L_{D,E} = w_{val}^{D,E} L_{val}^{D,E} + w_{vel}^{D,E} L_{vel}^{D,E} + w_{reg}^{D,E} L_{reg}^{D,E} \quad (4)$$

Here, F indicates the forward kinematics function that converts joint rotations into joint positions, and the operator \ominus calculates the difference between two poses. \mathbf{y}_0 and \mathbf{y}_1 represent two consecutive frames of a motion sequence. δn denotes the time interval between frames. $w_{val}^{D,E}$, $w_{vel}^{D,E}$, $w_{reg}^{D,E}$ are weights for balancing different loss terms. Once we construct the embedded vector space, we can learn an embedded state transition model $S(\mathbf{xz}_{n-1}) \rightarrow \hat{\mathbf{xz}}_n$ instead of the full pose transition model $S(\mathbf{y}_{n-1}) \rightarrow \hat{\mathbf{y}}_n$.

3.2. Auto-regressive Diffusion Model (ADM)

Character animations are intrinsically multi-modal. For a given pose, there may be multiple follow-up poses at the next moment. The transition from $S(\mathbf{xz}_{n-1}) \rightarrow \hat{\mathbf{xz}}_n$ is essentially a many-to-many mapping. Neural network models that use a Mean Square Error (MSE) based loss to train, such as Learned Motion Matching [25], are unable to capture these many-to-many transitions, since MSE losses work on one-to-one mappings. Therefore we employ a diffusion model as our backbone model. Our diffusion model

follows the structure of DDPM [16]. For each forward diffusion step, a small noise vector is added on top of the future embedded vector \mathbf{xz}_n :

$$q(\mathbf{xz}_n^t | \mathbf{xz}_n^{t-1}) = N(\sqrt{\alpha^t} \mathbf{xz}_n^{t-1}, (1 - \alpha^t)I) \quad (5)$$

The reverse diffusion phase $p(\mathbf{xz}_n^{t-1} | \mathbf{xz}_n^t)$ generates embedded vector $\hat{\mathbf{xz}}_n$ by gradually removing the noise on top of \mathbf{xz}_n . In our setting, the reverse diffusion model G^{ADM} follows the formulation of [43, 56] and directly predicts the embedded vector rather than the added noise as in [16, 50]. The previous vector \mathbf{xz}_{n-1} is used as a condition term:

$$\hat{\mathbf{xz}}_n^0 = G^{ADM}(\mathbf{xz}_n^t, \mathbf{xz}_{n-1}, t) \quad (6)$$

The predicted $\hat{\mathbf{xz}}_n^0$ is then used as a condition $\mathbf{xz}_{n'-1} = \hat{\mathbf{xz}}_n^0$ for generating the next $\mathbf{xz}_{n'}$, where $n' = n + 1$.

To ensure high-quality generation over a long horizon, the loss for training G^{ADM} measures the difference between the auto-regressively generated h -length embedded vector and the ground truth value. Specifically, to generate an embedded vector sequence, we start with the trajectory $\mathbf{xz}_{0:h}$ and add forward diffusion noise. This can be done in an auto-regressive manner using Equation 6 starting from the initial condition of $(\mathbf{xz}_0, \mathbf{xz}_1^{t_1}, t_1)$ until $(\hat{\mathbf{xz}}_{h-1}^0, \mathbf{xz}_h^{t_h}, t_h)$. The loss is designed as:

$$L_{val}^{ADM} = \|\hat{\mathbf{xz}}_{1:h}^0 - \mathbf{xz}_{1:h}\| \quad (7)$$

$$L_{vel}^{ADM} = \left\| \frac{(\hat{\mathbf{x}}_{1:h}^0 - \hat{\mathbf{x}}_{0:h-1}^0)}{h * \delta n} - \frac{(\mathbf{x}_{1:h} - \mathbf{x}_{0:h-1})}{h * \delta n} \right\| \quad (8)$$

$$L_{G^{ADM}} = w_{val}^{ADM} L_{val}^{ADM} + w_{vel}^{ADM} L_{vel}^{ADM} \quad (9)$$

Here L_{val}^{ADM} encourages the reconstruction of the trajectory, and L_{vel}^{ADM} aims to imitate the velocities.

Although this basic diffusion model can produce high-quality samples and achieve improved mode coverage, the sampling process is time consuming primarily due to the iterative nature of diffusion and denoising.

3.3. Fast Generation via Denoising Diffusion GANs

The Diffusion Model typically involves multiple steps to generate solid predictions. This is based on the assumption that the denoising follows a Gaussian distribution [61]. However, this assumption is only valid when a small amount of noise is eliminated at each denoising step. As a result, it takes numerous steps to generate a high-quality prediction from pure noise. To minimize the number of steps in the reverse process and therefore accelerating the generating process, an alternative approach is to utilize a non-Gaussian multimodal distribution.

Our *AAMDM* utilizes Denoising Diffusion-GANs (DD-GANs) as proposed by Xiao et al. [61]. This method formulates the reverse diffusion process using a multimodal distribution. It achieves this by parameterizing the reverse diffusion process as conditional GANs. The reverse diffusion

generator, denoted as G^{GAN} , takes an additional latent variable \mathbf{r}^t as a conditional term, in addition to $(\mathbf{xz}_n^t, \mathbf{xz}_{n-1}, t)$:

$$\hat{\mathbf{xz}}_n^0 = G^{GAN}(\mathbf{xz}_n^t, \mathbf{xz}_{n-1}, \mathbf{r}^t, t). \quad (10)$$

We use $G^{GAN}(\sim)$ as an abbreviation of $G^{GAN}(\mathbf{xz}_n^t, \mathbf{xz}_{n-1}, \mathbf{r}^t, t)$, which can be trained by minimizing the KL divergence between two distributions: $D_{KL}(p(\mathbf{xz}_n^{t-1} | \mathbf{xz}_n^t, \mathbf{xz}_{n-1}) || q(\mathbf{xz}_n^{t-1} | \mathbf{xz}_n^t, \mathbf{xz}_{n-1}))$:

$$L_{GAN} = -\mathbb{E}_{p(\mathbf{xz}_n^{t-1} | \mathbf{xz}_n^t, \mathbf{xz}_{n-1})}[\log(D^{GAN}(\sim))]. \quad (11)$$

This objective can then be converted to training a diffusion-step-dependent discriminator network $D^{GAN}(\mathbf{xz}_n^{t-1}, \mathbf{xz}_n^t, \mathbf{xz}_{n-1}, t)$ to distinguish if \mathbf{xz}_n^{t-1} is diffused from the original data \mathbf{xz}_n or generated fake data $\hat{\mathbf{xz}}_n^0$, and the generator is trained to disguise the discriminator. We use $D^{GAN}(\sim)$ as an abbreviation of $D^{GAN}(\mathbf{xz}_n^{t-1}, \mathbf{xz}_n^t, \mathbf{xz}_{n-1}, t)$.

$$L_{D^{GAN}} = -\mathbb{E}_{q(\mathbf{xz}_n^{t-1} | \mathbf{xz}_n^t, \mathbf{xz}_{n-1})}[\log(D^{GAN}(\sim))] - \mathbb{E}_{p(\mathbf{xz}_n^{t-1} | \mathbf{xz}_n^t, \mathbf{xz}_{n-1})}[\log(1 - D^{GAN}(\sim))] \quad (12)$$

DD-GANs offer high sampling speed while maintaining excellent mode coverage and output quality in a single motion step. However, when used in an autoregressive generation setting, we have observed that DD-GANs often lead to unstable training, resulting in deteriorated motion quality. To address this issue, we propose combining ADM and DD-GANs to achieve fast and high-quality sampling.

3.4. Combining ADM and DD-GANs

The combination of ADM and DD-GANs is based on the insight that the diffusion process transitions from generating samples from noise at early stages to making small adjustments in the prediction at late stages. To achieve higher quality output, the generation of single motion steps is divided into two sub-steps: Generation and Polishing. The Generation module utilizes DD-GANs to generate a draft prediction in a few steps, while the Polishing module refines the output from the Generation module using ADM.

The process begins with a random noise input, $\mathbf{xz}_n^T \sim N(0, I)$, and the Generation module goes through $T^{GAN} = 3$ reverse diffusion steps using G^{GAN} for $\mathbf{xz}_n^{T^{AA}-T^{GAN}}$. The generated $\mathbf{xz}_n^{T^{AA}-T^{GAN}}$ is then passed to the Polishing module, where G^{ADM} refines the result using $T^{ADM} = 2$ steps. The total number of generation steps, $T^{AA} = T^{GAN} + T^{ADM} = 5$. Finally, the generated $\hat{\mathbf{xz}}_n^0$ replaces \mathbf{xz}_{n-1}^0 for the next step prediction. The Generation Module and Polishing Module are trained separately using Equation 7, 12, and 11. For a more detailed sampling and training procedure, refer to the supplementary materials.

3.5. Motion Control with User Commands

To generate the motions that follow the user’s commands, for single pose transition, *AAMDM* guides the motion generation process through a guided diffusion method proposed by Rempe et al. [44]. Given the user’s query $\bar{\mathbf{x}}_n$, at each diffusion steps with noise vector \mathbf{xz}_n^t , we perturb the generated vector $\hat{\mathbf{xz}}_n^0$ to obtain the guided vector $\hat{\mathbf{xz}}_n^{0,*}$:

$$\hat{\mathbf{xz}}_n^{0,*} = \hat{\mathbf{xz}}_n^0 - \epsilon \alpha^t \nabla_{\mathbf{xz}_n^t} J(\hat{\mathbf{x}}_n^0, \bar{\mathbf{x}}_n). \quad (13)$$

Here, J is an objective function measuring distance between the generated feature vector and user’s query. ϵ is a step parameter, α^t is the noise parameter in diffusion model.

3.6. Model Representation

Character Representation The pose vector, denoted as \mathbf{y} , captures all the character’s pose information in a single frame of the animation. It is defined as $\mathbf{y} = \{\mathbf{y}^t, \mathbf{y}^r, \dot{\mathbf{y}}^t, \dot{\mathbf{y}}^r, \dot{\mathbf{r}}^t, \dot{\mathbf{r}}^r\}$, where \mathbf{y}^t and \mathbf{y}^r represent joint local translations and rotations, $\dot{\mathbf{y}}^t$ and $\dot{\mathbf{y}}^r$ represent joint local translational and rotational velocities, and $\dot{\mathbf{r}}^t$ and $\dot{\mathbf{r}}^r$ represent root translational and rotational velocities. The total dimension of \mathbf{y} is 338. Additionally, we define $\mathbf{x} = \{\mathbf{t}^t, \mathbf{t}^d\}$, where $\mathbf{t}^t \in \mathbb{R}^6$ and $\mathbf{t}^d \in \mathbb{R}^6$ represents the 2D future trajectory positions and facing direction projected on the ground, 20, 40, and 60 frames in the future local to the character. The latent vector \mathbf{z} has a dimension of 52, and thus $\mathbf{xz} \in \mathbb{R}^{64}$.

Neural Network Structure The encoder network E^{AE} , the decoder network D^{AE} , ADM generator G^{GAN} , DD-GANs generator G^{GAN} and DD-GANs discriminator D^{GAN} are all fully connected neural network. The details are presented in the supplementary materials.

4. Experiments

We conducted a series of experiments to evaluate the performance of the proposed method, *AAMDM*. Firstly, *AAMDM* is quantitatively compared against several baseline methods using different evaluation metrics. Subsequently, we conducted additional experiments on an artificial multi-modal dataset for detailed discussion. Lastly, we performed ablation studies to justify design choices. Overall, the results demonstrate that *AAMDM* can efficiently generate high-quality motions with long horizons auto-regressively. The motions can be seen in the supplementary video.

Implementation Details We implemented our motion generation framework in Pytorch and conducted experiments on a PC equipped with an NVIDIA GeForce RTX 3080 Ti and AMD Ryzen 9 3900X 12-Core Processor. For all networks, training was performed for 1M iterations using the RAdam optimizer [32] with a batch size of 64

and a learning rate of 0.0001. We trained the Encoder E^{AE} and Decoder D^{AE} first to construct the embedded vector space \mathbf{XZ} , then we trained the Polishing module and the Generation module. Both Polishing and Generation modules were trained with a window size of 10 frames. The total training procedure took around 20 hours.

Dataset We utilized the Ubisoft LaForge Animation Dataset (“LaFAN1”) [13] for evaluation. LaFAN1 is a collection of high-quality human character animations, encompassing a wide range of motions. Our dataset comprised 25 motion clips from LAFAN1, featuring 100,000 pose transitions, and had a total duration of 26.67 minutes.

4.1. Baseline Comparison

We compared *AAMDM* with the following baselines:

- **Learned Motion Matching (LMM)**: LMM is an interactive motion synthesis method proposed by Kolsi et al. [25]. Similar to our method, LMM uses an embedded vector space. It comprises three networks: Projector that maps the human input vector $\bar{\mathbf{x}}$ to the embedded vector \mathbf{xz} for addressing user’s command, Decompressor that reproduces the pose vector \mathbf{y} from \mathbf{xz} , and Stepper that maps \mathbf{xz}_{n-1} to \mathbf{xz}_n for learning the pose transition. Both Stepper and Projector are trained using MSE based loss. Unlike LMM that treats user commands and pose transition separately, *AAMDM* fuses these two requirement using the guided diffusion process.
- **Motion VAE (MVAE)**: MVAE [31] is based on an autoregressive conditional variational autoencoder. Given the current pose, MVAE predicts a distribution of possible next poses, as it is conditioned on a set of stochastic latent variables. The key distinction between *AAMDM* and MVAE is that the former models transitions using a diffusion-based model, whereas the latter employs a VAE.
- **Autoregressive Motion Diffusion Model (AMDM)**: AMDM [50] is an autoregressive diffusion model-based framework for motion synthesis. There are three main differences between *AAMDM* and AMDM. First, AMDM accelerates the diffusion process by simply taking fewer reverse diffusion steps, while *AAMDM* leverages DD-GANs. Second, AMDM operates in the full pose space, whereas *AAMDM* learns transitions in an embedded space. Third, AMDM predicts noise at each reverse diffusion step, while *AAMDM* directly predicts the target vector as Ramesh et al. [43]. We implemented two versions of AMDM, named AMDM5 and AMDM200, to indicate the use of 5 and 200 diffusion steps, respectively.

4.1.1 Evaluation of Random Motion Synthesis

We first evaluated the performance of these methods in random motion generation over the following metrics:

	DIV \rightarrow	FID \downarrow	FFR \downarrow	FPS \uparrow	TE-UC \downarrow	FID-UC \downarrow	FFR-UC \downarrow
Dataset	14.533	0.000	0.113	N/A	N/A	0.000	0.113
AAMDMM(Our)	11.574	14.051	0.131	173	0.034	15.367	0.143
LMM [25]	5.374	49.706	0.194	812	0.021	55.674	0.233
MVAE [31]	7.223	22.981	0.312	703	0.027	47.453	0.349
AMDM5 [50]	8.134	18.741	0.214	192	0.054	25.943	0.256
AMDM200 [50]	11.165	12.132	0.129	4.72	0.012	14.254	0.133

Table 1. In our quantitative analysis, we demonstrate that the *AAMDMM* framework is capable of generating motions of a quality comparable to that of *AMDM200*, while significantly outperforming other methods in both random sampling and user control scenarios. Meanwhile, the result also indicates that *AAMDMM* is approximately 40 times faster than *AMDM200*.




Figure 3. Comparison between motions generated by LMM (top) and *AAMDMM* (Bottom). Starting from a similar character pose, LMM is unable to generate diverse motions while *AAMDMM* can reproduce diverse complex motions.

- **Diversity (DIV):** Diversity measures the distributional spread of the generated motions in the character pose space. This metric, adopted from several previous works [29, 30, 56, 57], assesses how well the generated motion matches the distribution of the ground truth dataset. We follow the implementation used in MDM [56], computing Diversity using 1,000 frames from each generated motion clip. A good Diversity score should be close to that of the motion dataset.
- **Frechet Inception Distance (FID):** FID evaluates the difference between the distributions of generated and ground truth motions. FID serves as an indicator of the overall quality of generated motions in many prior works [56, 57].
- **Footskating Frame Ratio (FFR):** FFR quantifies the realism of generated motion, particularly focusing on foot-ground contact. We measured foot skating artifacts as described in Zhang et al. [64]. A lower FFR score indicates better physical plausibility of the generated motions.
- **Frames Per Second (FPS):** FPS is a measure of the efficiency of motion generation methods in creating new frames. Higher FPS values indicate faster frame generation rates, essential for interactive applications.

The qualitative results are summarized in Table 1. Notably, *AAMDMM* achieve similar performance as *AMDM200*

while surpasses the other baselines in all motion quality metrics with more than 40 times faster than *AMDM200*. This demonstrates *AAMDMM*'s capability to efficiently generate high-quality character animations.

LMM's motion quality was generally found to be inferior to *AAMDMM*, as reflected in the FID and DIV metrics. This discrepancy is likely due to LMM's training with MSE loss, presupposing a one-to-one mapping. However, this assumption may not be valid in datasets with multiple possible transitions from a single pose. Figure 3 provides an example. A more detailed discussion on this aspect will be presented in a subsequent section. However, LMM showed a higher FPS score, attributed to its single feed-forward operation, compared to *AAMDMM*'s five feed-forward operations.

In motion quality evaluation, MVAE slightly outperformed LMM with scores of 7.223 and 22.981 in DIV and FID respectively. MVAE's better quality can be linked to its use of VAE for handling multiple mappings in pose transitions. Although MVAE offered improved training stability and performance, *AAMDMM* still outperformed in these metrics. MVAE also exhibited faster performance than *AAMDMM* due to its single-step feedforward process.

In comparison between *AMDM5* and *AAMDMM*, both methods used 5 diffusion steps which led to similar FPS scores (173 vs 192). However, the diffusion steps in *AMDM5* were modeled using a Gaussian distribution, which is typically effective when the total number of denoising steps is in the order of hundreds. As *AMDM5* utilized only five steps, this assumption did not hold and it led to compromised motion quality. On the other hand, *AAMDMM* leveraged DD-GANs to model multimodal transitions, which reduced the number of steps required for generating a new frame without sacrificing motion quality.

AMDM200 with more diffusion steps is better aligned with the Gaussian distribution assumption, which is connected to highly improved motion quality metrics. However, this increase in diffusion steps comes at the cost of efficiency. As the number of steps rises, the generation speed decreases. This trade-off highlights the balance between motion quality and generative efficiency, with *AMDM200*

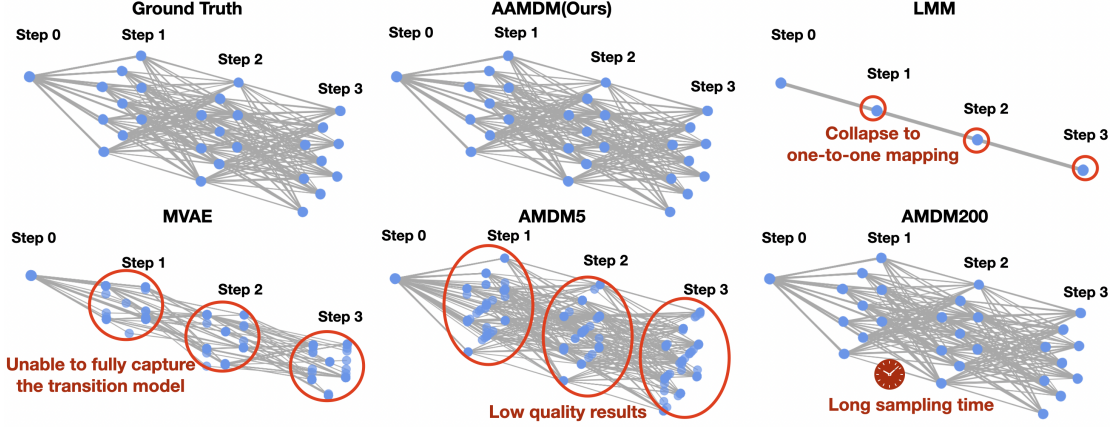


Figure 4. Visualization of the learned transition results of an artificial Squ-9-Gaussian experiment in 2D. We show that *AAMD* outperforms baseline methods in learning the many-to-many distribution mapping in sequential scenarios.

favoring the former at the expense of the latter.

4.1.2 Evaluation of Interactive Synthesis

We evaluated the performance of these methods in an interactive motion synthesis scenario. The experiment involved interactively controlling the character’s moving direction while allowing the arms to move freely. Our evaluation employed the following metrics, with ‘-UC’ denoting ‘Under Control’:

- **Tracking Error (TE-UC):** The TE-UC metric assesses the method’s ability to follow user commands \bar{x} . It is defined as the discrepancy between the user’s command and the generated motion $|\bar{x} - \hat{x}|$. A lower TE-UC value signifies better alignment with user input, reflecting superior performance.
- **Frechet Inception Distance (FID-UC):** The FID-UC is used to measure the similarity between the motion dataset and the generated trajectories. A lower FID-UC indicates a higher quality of the generated motion.
- **Footskating Frame Ratio (FFR-UC):** This metric evaluates the realism of the motion when the character is under user control. It assesses aspects such as naturalness and adherence to physical constraints. Lower FFR-UC scores suggest more physically plausible and realistic motion generation.

In user control scenarios, our results demonstrate that the *AAMD* framework consistently outperformed baseline methods across nearly all metrics evaluated. Compared with Learned Motion Matching (LMM), *AAMD* addresses several key issues inherent in the LMM’s approach. LMM employs a projector network trained with an MSE loss to interpret user commands, which leads to two primary issues. Firstly, multiple candidate poses could potentially match the user command, but LMM’s projector network struggles to handle multi-modal transitions. Secondly, the projector network often ignores the character’s current

pose, necessitating blending techniques to ensure smooth transitions. MVAE faces challenges in training to capture all the possible transitions, resulting in a quality of motion that does not match that of *AAMD*. Similarly, AMDM5’s reduces the number of diffusion steps, which breaks the Gaussian distribution assumption and consequently downgrades the motion quality. Although AMDM200 provides higher-quality generation due to more diffusion steps, its low speed (4.72 FPS) is not suitable for any interactive applications.

4.2. Additional Studies on Artificial Dataset

In addition to the previous experiment, we conducted an additional study to analyze the effectiveness of various methods on a many-to-many transition dataset. For this purpose, we created a 2D “Squ-9” dataset characterized by its multi-modal dynamics where any given point in three by three Gaussian distributions can transit to any other Gaussian distributions in the next time step. By learning this dataset, we evaluated the effectiveness of each method to capture this many-to-many dynamics. The comparative results are visually depicted in Figure 4.

In our results, our *AAMD* captured all the possible modes while preserving sample quality. In contrast, LMM struggled to represent the dataset’s many-to-many vector transitions, resulting in a singular vector cluster at each step. MVAE showed an improvement in mode coverage, yet it cannot illustrate all possible modes. Among other diffusion model-based approaches, AMDM5 exhibited better transitions but their qualities are still worse than *AAMD*. Although AMDM200 produced results of comparable quality to *AAMD*, it required 40 times more inference time.

4.3. Ablation Studies

We provided additional insights of *AAMD* by conducting three ablation studies summarized in Table 2.

T^{ADM}	T^{GAN}	DIV \rightarrow	FID \downarrow	FPS \uparrow
Dataset		14.533	0.000	N/A
0	3	N/A	N/A	311
1	3	10.612	16.332	215
2*	3*	11.574	14.051	173
10	3	12.041	11.779	47
2	2	12.415	28.476	211
2*	3*	11.574	14.051	173
2	4	11.313	14.534	135
2	10	9.775	16.312	58
wo/ Emb		56.341	128.412	146

Table 2. Ablation study results. *The default parameters.

Polishing Steps In our study, we investigated the impact of the number of polishing steps (T^{ADM}) on the generation process. Specifically, we denote $T^{ADM} = 0$ as the scenario where no polishing module is used, meaning the output from the generation module is directly utilized for future frame generation. In our experiments, settings with $T^{ADM} > 0$ exhibited significant performance enhancements compared to the $T^{ADM} = 0$ scenario as when $T^{ADM} = 0$, the framework was unable to generate reliable long-horizon trajectory due to the diverges of the character’s pose. This suggests that relying solely on denoising diffusion GANs may not yield high-quality outputs for long-horizon generation. In contrast, additional polishing steps markedly improved the output quality, making it more suitable for long-horizon predictions. Furthermore, results indicate a positive correlation between the number of polishing steps T^{ADM} and the output quality. However, it is important to note that increasing T^{ADM} also leads to longer sampling times.

Generation Steps In our second study, we examined the effects of the number of generation steps. Theoretically, increasing T^{GAN} should reduce the amount of noise that needs to be removed at each denoising step, potentially simplifying the training process. However, our results show that a specific value of T^{GAN} , $T^{GAN} = 3$ and $T^{GAN} = 4$, yielded the highest overall motion quality, yet $T^{GAN} = 3$ was more efficient. Although $T^{GAN} = 2$ achieved the best performance in the DIV metric, we observed a few cases of divergence in the motion, which resulted in worse performance in FID compared to $T^{GAN} = 3$. When $T^{GAN} = 10$, the learning task should be easier since the distance between each diffusion step is smaller. However, our results show that the performance was the worst. We hypothesize that this is because we utilized a simple MLP network; thus, it may not be adequately equipped to handle

larger values of T^{GAN} for effectively training the denoising diffusion GANs.

Importance of Embedded Transition Space In this analysis, we explored the advantages of learning transitions in an embedded space \mathbf{XZ} as opposed to the full pose space \mathbf{Y} . Our results suggest that utilizing the full pose space yields inferior outcomes compared to an embedded space. We attribute this finding to two primary factors. Firstly, learning in a higher-dimensional space, like the full pose space, is inherently more challenging than in a lower-dimensional space, particularly under multimodal distribution conditions. Secondly, as discussed in the previous section, *AAMDM* does not employ a complex neural network architecture or specialized techniques for constructing the latent space in Denoising Diffusion GANs. MLP networks used in our framework may not be sufficiently robust to capture transitions in larger spaces effectively. This limitation further supports the advantage of using an embedded space for learning transitions.

5. Discussion and future work

We have introduced a novel framework for motion synthesis: Accelerated Auto-regressive Motion Diffusion Model (*AAMDM*). *AAMDM* is designed to efficiently generate high-quality animation frames for interactive user engagement. This is achieved by several technical components: the use of a low-dimensional embedded space for compact representation, Denoising Diffusion GANs for fast approximations, and the Diffusion Model for robust and accurate long-horizon synthesis. Our benchmarking of *AAMDM* against various baseline methods has demonstrated its superior capabilities in motion synthesis. We have also investigated the nuances of different autoregressive motion synthesis methods, providing valuable insights into this domain. Additionally, our ablation studies have validated the design choices made for *AAMDM* and identified the influence of various hyperparameters on the overall system performance.

In the future, we plan to explore several research directions. One notable challenge is the trade-off between the motion quality and the computational cost. Future work could explore advanced techniques such as parallel computing and the use of temporal information to accelerate the generation process. In addition, the model performance can be further improved by introducing more sophisticated methods to structure the latent space of the denoising diffusion GANs, such as a structured matrix-Fisher distribution [49]. Finally, it will be interesting to improve the controllability of the framework by introducing a learning-based control mechanism rather than relying on gradient-based sampling guidance.

References

- [1] For honor. <https://www.ubisoft.com/en-us/game/for-honor>. Accessed: November 13, 2023. **1**
- [2] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision-making? *arXiv preprint arXiv:2211.15657*, 2022. **2**
- [3] Simon Alexanderson, Rajmund Nagy, Jonas Beskow, and Gustav Eje Henter. Listen, denoise, action! audio-driven motion synthesis with diffusion models. *arXiv preprint arXiv:2211.09707*, 2022. **2**
- [4] Okan Arikan and David A Forsyth. Interactive motion generation from examples. *ACM Transactions on Graphics (TOG)*, 21(3):483–490, 2002. **2**
- [5] Jinxiang Chai and Jessica K Hodgins. Performance animation from low-dimensional control signals. In *ACM SIGGRAPH 2005 Papers*, pages 686–696. 2005. **2**
- [6] 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. **2**
- [7] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. **2**
- [8] Katerina Fragkiadaki, Sergey Levine, and Jitendra Malik. Recurrent network models for kinematic tracking. *CoRR*, abs/1508.00271, 1(2):4, 2015. **2**
- [9] Ruiqi Gao, Yang Song, Ben Poole, Ying Nian Wu, and Diederik P Kingma. Learning energy-based models by diffusion recovery likelihood. *arXiv preprint arXiv:2012.08125*, 2020. **3**
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. **3**
- [11] Keith Grochow, Steven L Martin, Aaron Hertzmann, and Zoran Popović. Style-based inverse kinematics. In *ACM SIGGRAPH 2004 Papers*, pages 522–531. 2004. **2**
- [12] Félix G Harvey and Christopher Pal. Recurrent transition networks for character locomotion. In *SIGGRAPH Asia 2018 Technical Briefs*, pages 1–4. 2018. **2**
- [13] Félix G. Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. 39(4), 2020. **2, 5**
- [14] Gustav Eje Henter, Simon Alexanderson, and Jonas Beskow. Moglow: Probabilistic and controllable motion synthesis using normalising flows. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020. **2**
- [15] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. **2**
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. **2, 4**
- [17] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. **2**
- [18] Daniel Holden. Character control with neural networks and machine learning. *Proc. of GDC 2018*, 1:2, 2018. **1**
- [19] Daniel Holden, Taku Komura, and Jun Saito. Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017. **2**
- [20] Tobias Höppe, Arash Mehrjou, Stefan Bauer, Didrik Nielsen, and Andrea Dittadi. Diffusion models for video prediction and infilling. *arXiv preprint arXiv:2206.07696*, 2022. **2**
- [21] Nicholas Howe, Michael Leventon, and William Freeman. Bayesian reconstruction of 3d human motion from single-camera video. *Advances in neural information processing systems*, 12, 1999. **2**
- [22] Kyunglyul Hyun, Kyungho Lee, and Jehee Lee. Motion grammars for character animation. In *Computer Graphics Forum*, pages 103–113. Wiley Online Library, 2016. **2**
- [23] Michael Janner, Yilun Du, Joshua B Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022. **2**
- [24] 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**
- [25] Marko Kolsi, Mikko Mononen, and Joonas Javanainen. Learned motion matching. In *Proceedings of the 19th ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 6:1–6:10. ACM, 2018. **1, 2, 3, 5, 6**
- [26] Lucas Kovar, Michael Gleicher, and Frédéric Pighin. Motion graphs. In *ACM SIGGRAPH 2008 classes*, pages 1–10. 2008. **2**
- [27] Jehee Lee, Jinxiang Chai, Paul SA Reitsma, Jessica K Hodgins, and Nancy S Pollard. Interactive control of avatars animated with human motion data. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 491–500, 2002. **2**
- [28] Sergey Levine, Jack M Wang, Alexis Haraux, Zoran Popović, and Vladlen Koltun. Continuous character control with low-dimensional embeddings. *ACM Transactions on Graphics (TOG)*, 31(4):1–10, 2012. **2**
- [29] Tianyu Li, Hyunyoung Jung, Matthew Gombolay, Yong Kwon Cho, and Sehoon Ha. Crossloco: Human motion driven control of legged robots via guided unsupervised reinforcement learning. *arXiv preprint arXiv:2309.17046*, 2023. **6**
- [30] Tianyu Li, Jungdam Won, Alexander Clegg, Jeonghwan Kim, Akshara Rai, and Sehoon Ha. Ace: Adversarial correspondence embedding for cross morphology motion retargeting from human to nonhuman characters. *arXiv preprint arXiv:2305.14792*, 2023. **6**
- [31] Hung Yu Ling, Fabio Zinno, George Cheng, and Michiel Van De Panne. Character controllers using motion vaes. *ACM Transactions on Graphics (TOG)*, 39(4):40–1, 2020. **2, 5, 6**
- [32] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019. **5**

- [33] Eric Luhman and Troy Luhman. Knowledge distillation in iterative generative models for improved sampling speed. *arXiv preprint arXiv:2101.02388*, 2021. 2
- [34] Jianxin Ma, Shuai Bai, and Chang Zhou. Pretrained diffusion models for unified human motion synthesis. *arXiv preprint arXiv:2212.02837*, 2022. 2
- [35] Jianyuan Min and Jinxiang Chai. Motion graphs++ a compact generative model for semantic motion analysis and synthesis. *ACM Transactions on Graphics (TOG)*, 31(6):1–12, 2012. 2
- [36] Tomohiko Mukai. Motion rings for interactive gait synthesis. In *Symposium on Interactive 3D Graphics and Games*, pages 125–132, 2011. 2
- [37] Tomohiko Mukai and Shigeru Kuriyama. Geostatistical motion interpolation. In *ACM SIGGRAPH 2005 Papers*, pages 1062–1070. 2005. 2
- [38] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 2
- [39] Soohwan Park, Hoseok Ryu, Seyoung Lee, Sunmin Lee, and Jehee Lee. Learning predict-and-simulate policies from unorganized human motion data. *ACM Transactions on Graphics (TOG)*, 38(6):1–11, 2019. 2
- [40] Sang Il Park, Hyun Joon Shin, and Sung Yong Shin. On-line locomotion generation based on motion blending. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 105–111, 2002. 2
- [41] Dario Pavllo, David Grangier, and Michael Auli. Quaternion: A quaternion-based recurrent model for human motion. *arXiv preprint arXiv:1805.06485*, 2018. 2
- [42] Dario Pavllo, Christoph Feichtenhofer, Michael Auli, and David Grangier. Modeling human motion with quaternion-based neural networks. *International Journal of Computer Vision*, 128:855–872, 2020. 2
- [43] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 4, 5
- [44] Davis Remppe, Zhengyi Luo, Xue Bin Peng, Ye Yuan, Kris Kitani, Karsten Kreis, Sanja Fidler, and Or Litany. Trace and pace: Controllable pedestrian animation via guided trajectory diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13756–13766, 2023. 5
- [45] Charles Rose, Michael F Cohen, and Bobby Bodenheimer. Verbs and adverbs: Multidimensional motion interpolation. *IEEE Computer Graphics and Applications*, 18(5):32–40, 1998. 2
- [46] Alla Safonova and Jessica K Hodgins. Construction and optimal search of interpolated motion graphs. In *ACM SIGGRAPH 2007 papers*, pages 106–es. 2007. 2
- [47] Alla Safonova, Jessica K Hodgins, and Nancy S Pollard. Synthesizing physically realistic human motion in low-dimensional, behavior-specific spaces. *ACM Transactions on Graphics (ToG)*, 23(3):514–521, 2004. 2
- [48] Robin San-Roman, Eliya Nachmani, and Lior Wolf. Noise estimation for generative diffusion models. *arXiv preprint arXiv:2104.02600*, 2021. 2
- [49] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Hierarchical kinematic probability distributions for 3d human shape and pose estimation from images in the wild. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11219–11229, 2021. 8
- [50] Yi Shi, Jingbo Wang, Xuekun Jiang, and Bo Dai. Controllable motion diffusion model. *arXiv preprint arXiv:2306.00416*, 2023. 2, 3, 4, 5, 6
- [51] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015. 2
- [52] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2
- [53] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020. 2
- [54] Paul Starke, Sebastian Starke, Taku Komura, and Frank Steinicke. Motion in-betweening with phase manifolds. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 6(3):1–17, 2023. 2
- [55] 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
- [56] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 2, 4, 6
- [57] 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. 2, 6
- [58] Vikram Voleti, Alexia Jolicœur-Martineau, and Chris Pal. Mcvd-masked conditional video diffusion for prediction, generation, and interpolation. *Advances in Neural Information Processing Systems*, 35:23371–23385, 2022. 2
- [59] Jack M Wang, David J Fleet, and Aaron Hertzmann. Gaussian process dynamical models for human motion. *IEEE transactions on pattern analysis and machine intelligence*, 30(2):283–298, 2007. 2
- [60] Zhiyong Wang, Jinxiang Chai, and Shihong Xia. Combining recurrent neural networks and adversarial training for human motion synthesis and control. *IEEE transactions on visualization and computer graphics*, 27(1):14–28, 2019. 2
- [61] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion gans. *arXiv preprint arXiv:2112.07804*, 2021. 4
- [62] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion gans. *arXiv preprint arXiv:2112.07804*, 2021. 3
- [63] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. *arXiv preprint arXiv:2212.02500*, 2022. 2

- [64] 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. [2](#), [6](#)
- [65] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. [2](#)