

Seamless Human Motion Composition with Blended Positional Encodings

German Barquero Sergio Escalera Cristina Palmero
 Universitat de Barcelona and Computer Vision Center, Spain
 {germanbarquero, sescalera}@ub.edu, crpalmec7@alumnes.ub.edu
<https://barquero german.github.io/FlowMDM/>




Figure 1. We present FlowMDM, a diffusion-based approach capable of generating seamlessly continuous sequences of human motion from textual descriptions (left). The whole sequence is generated simultaneously and it does not require any postprocessing. FlowMDM also makes strides in the challenging problem of extrapolating and controlling periodic motion such as walking, jumping, or waving (right).

Abstract

Conditional human motion generation is an important topic with many applications in virtual reality, gaming, and robotics. While prior works have focused on generating motion guided by text, music, or scenes, these typically result in isolated motions confined to short durations. Instead, we address the generation of long, continuous sequences guided by a series of varying textual descriptions. In this context, we introduce FlowMDM, the first diffusion-based model that generates seamless Human Motion Compositions (HMC) without any postprocessing or redundant denoising steps. For this, we introduce the Blended Positional Encodings, a technique that leverages both absolute and relative positional encodings in the denoising chain. More specifically, global motion coherence is recovered at the absolute stage, whereas smooth and realistic transitions are built at the relative stage. As a result, we achieve state-of-the-art results in terms of accuracy, realism, and smoothness on the Babel and HumanML3D datasets. FlowMDM excels when trained with only a single description per motion sequence thanks to its Pose-Centric Cross-Attention, which makes it robust against varying text descriptions at inference time. Finally, to address the limitations of existing HMC metrics, we propose two new metrics: the Peak Jerk and the Area Under the Jerk, to detect abrupt transitions.

1. Introduction

In the field of computer vision, recent progress has been made in developing photorealistic avatars [53] for applications like virtual reality, gaming, and robotics [60, 76]. Aside from looking visually realistic, avatars must also move in a convincing manner. This is challenging due to the intricate nature of human motion, strongly influenced by factors such as the environment, interactions, and physical contact [14]. Furthermore, complexity increases when attempting to control these motions. Recent advances include the generation of motion sequences from *control signals* like text descriptions or actions [106]; however, such methods only produce isolated, standalone motion. Therefore, these approaches fail to handle scenarios where a long motion is driven by distinct control signals on different time slices. Such capability is needed to provide full control over the sequence of desired actions and their duration. In these scenarios, the generated motion needs to feature seamless and realistic transitions between actions. In this work, we tackle this problem, which we refer to as generative Human Motion Composition (HMC). In particular, we focus on generating single-human motion from text (Fig. 1).

One of the primary obstacles in HMC is the lack of datasets that offer long motion sequences with diverse textual annotations. Existing datasets typically feature sequences of limited duration, often lasting only up to 10 seconds, and with just a single control signal governing the en-

tire sequence [26, 62]. This limitation calls for innovative solutions to address the inherent complexities of the task. Prior works have tackled this problem mostly with autoregressive approaches [4, 44, 47, 64, 101]. These methods iteratively create compositions by using the current motion as a basis to generate subsequent motions. However, they require datasets with multiple consecutive annotated motions, and tend to degenerate in very long HMC scenarios due to error accumulation [104]. Other recent works have leveraged the infilling capabilities of motion diffusion models to generate motion compositions [70, 100]. However, for these, a substantial portion of each motion sequence is generated independently from adjacent motions, and generating transitions requires computing redundant denoising steps. In this work, we propose a novel architecture designed to address these specific challenges. Our contributions are:

- We propose FlowMDM, the first diffusion-based model that generates seamless human motion compositions without any postprocessing or extra denoising steps. To accomplish it, we introduce Blended Positional Encodings (BPE), a new technique for diffusion Transformers that combines the benefits of both absolute and relative positional encodings during sampling. In particular, the denoising first exploits absolute information to recover the global motion coherence, and then leverages relative positions to build smooth and realistic transitions between actions. As a result, FlowMDM achieves state-of-the-art results in terms of accuracy, realism, and smoothness in the HumanML3D [26] and Babel [63] datasets.
- We introduce a new attention technique tailored for HMC: the Pose-Centric Cross-Attention (PCCAT). This layer ensures each pose is denoised based on its own condition and its neighboring poses. Consequently, FlowMDM can be trained on a dataset with only a single condition available per motion sequence and still generate realistic transitions when using multiple conditions at inference time.
- We reveal the lack of sensitivity of current HMC metrics to identify discontinuous or sharp transitions, and introduce two new metrics that help to detect them: the Peak Jerk (PJ) and the Area Under the Jerk (AUJ).

2. Related work

Conditional human motion generation. Recent studies in motion generation have shown notable progress in synthesizing movements conditioned on diverse modalities such as text [21, 26, 27, 35, 40, 61, 78, 79, 97–99], music [2, 17, 46, 74, 81, 93, 107], scenes [15, 30, 84–86, 94], interactive objects [1, 18, 41, 89], and even other humans’ behavior [9, 10, 28, 77, 90]. Traditionally, these approaches have been designed to generate motion sequences matching a single condition. The progress of this domain has been boosted by the release of big datasets including diverse modalities or manual annotations [12, 26, 28, 46, 50, 58, 62, 63]. Re-

search has also focused on problems like human motion prediction [3, 52, 56, 69, 75, 80, 83, 96] and motion infilling [29, 36, 39, 48, 49, 57, 65, 66, 72, 105], which do not rely on extensive manual annotations but rather on motion itself. Both tasks share a common challenge with HMC: the synthesized motion must not only be plausible but also integrate seamlessly with the neighboring behavior, ensuring fluidity and continuity. In this context, the utilization of human motion priors has been proven to be a successful technique to ensure any generated motion includes natural transitions [8, 45, 88]. In line with these approaches, our method learns a motion prior specifically tailored for HMC.

Autoregressive human motion composition. As in many other sequence modeling tasks, HMC was also first tackled with autoregressive methods. The gold standard has been pairing variational autoencoders with autoregressive decoders such as recurrent neural networks [101] or Transformers [4, 44, 47, 64]. Alternative approaches have introduced specialized reinforcement learning frameworks [51, 92, 102]. Autoregressive models rely on the availability of annotated motion transitions, a requirement that constrains the robustness of the models due to the scarcity of such data. To mitigate this issue, some methods include additional postprocessing steps like linear interpolations [4], or affine transformations [44]. However, these can distort the human motion dynamics and require a predetermined estimation of the transitions duration. Furthermore, autoregressive approaches generate motion solely based on the preceding motion. We argue that an accurate model should mimic the humans innate capacity to anticipate their next action and adapt their current behavior accordingly [24, 42].

Diffusion-based human motion composition. Diffusion models excel at conditional generation [20, 32, 71]. They also possess great zero-shot capabilities for image inpainting [67], and its equivalence in motion: motion infilling. DiffCollage [100], MultiDiffusion [7], and DoubleTake [70] introduced diffusion sampling processes that simultaneously generate temporally superimposed motions, and combine the overlapped regions so that an infilled transition emerges. DoubleTake extended such sampling with a refinement step in which the emerged transition undergoes further unconditional denoising steps. All these methods share two main limitations. First, they are constrained to modeling dependencies among neighboring motion sequences. This becomes a limitation when three or more consecutive actions share semantics and collectively represent a more comprehensive action. In this case, dependencies may extend beyond contiguous actions. Second, they need to set the number of frames that each transition takes between consecutive actions, leading to extra computations. Our work addresses these constraints with a solution able to model longer inter-sequence dynamics without extra computational burdens or predefined transition durations.

3. Methodology

Problem definition. Our goal consists in generating a motion sequence of N frames, with the capability of conditioning the generated motion inside non-overlapping intervals $[0, \tau_1), [\tau_1, \tau_2), \dots, [\tau_j, N)$, with $0 < \tau_1 < \dots < \tau_j < N$. We will refer to the motion inside these intervals as *motion subsequences*, or $\mathcal{S}_i = \{x_{\tau_i}, \dots, x_{\tau_{i+1}-1}\}$, each driven by its corresponding condition c_i , and with a maximum length of L . It is essential that consecutive subsequences, influenced by different control signals, transition seamlessly and realistically. In particular, we aim at the even more challenging case where motion sequences containing several pairs of (\mathcal{S}_i, c_i) are not necessarily available in our dataset.

In this section, we present FlowMDM, an architecture with strong inductive biases that promote the emergence of a **robust translation-invariant motion prior**. Such *motion prior* is learned with a diffusion model equipped with a bidirectional (i.e., encoder-only) Transformer, similar to prior works [70, 79]. With it, we overcome the main limitations of autoregressive methods (Sec. 3.1). However, previous works are constrained in terms of motion duration. We could arguably provide extrapolation capabilities to the diffusion model by replacing the absolute positional encoding with a relative alternative, thus making the denoising of each pose *translation invariant*. However, this technique would fail to build complex compositional semantics that require knowledge about the start and end of each subsequence. For example, when generating the motion composition $\mathcal{S}_i \rightarrow \mathcal{S}_{i+1}$ with $c_i = \text{'walking'}$ and $c_{i+1} = \text{'walk and sit down'}$, \mathcal{S}_{i+1} might only feature the action *'sit down'* because, with only relative positional information, the Transformer cannot know if the partially denoised *'walking'* motion preceding the beginning of \mathcal{S}_{i+1} belongs to \mathcal{S}_i or \mathcal{S}_{i+1} . To combine the benefits of both relative and absolute positional encodings, we introduce BPE (Sec. 3.2). This novel technique exploits the iterative nature of diffusion models to promote intra-subsequence global coherence in earlier denoising stages, while making later denoising stages translation invariant, ensuring that realistic and plausible transitions naturally emerge between subsequences. Still, during training, the condition remains unchanged throughout all ground truth motion sequences. In order to make our denoising model *robust* to having multiple conditions per sequence at inference, we introduce a new attention paradigm called PCCAT (Fig. 3.3). As a result, FlowMDM is able to simultaneously generate very long compositions of human motion subsequences, all in harmony and fostering plausible transitions between them, without explicit supervision on transitions generation.

3.1. Bidirectional diffusion

The cumulative nature of errors in autoregressive models often results in a decline in performance when generating

long sequences [104]. This is exacerbated in HMC, where transitions are scarce or even missing in the training corpus, and the model needs to deal with domain shifts at inference. Another limitation of autoregressive methods is that the generated \mathcal{S}_i only depends on $\{\mathcal{S}_j\}_{j < i}$. We discussed in Sec. 2 why this is a suboptimal solution for HMC. Thus, an appropriate model for HMC should also be able to anticipate the following motion, \mathcal{S}_{i+1} , and possibly adapt \mathcal{S}_i so that the transition is feasible. We argue that the iterative paradigm of diffusion models provides very appropriate inductive biases for naturally mimicking such ability: the partially denoised \mathcal{S}_i and \mathcal{S}_{i+1} are refined later in successive denoising steps. By choosing a bidirectional Transformer as our denoising function [38], we enable the modeling of both past and future dependencies. Therefore, we design our framework as a bidirectional motion diffusion model, similar to MDM [79]. We refer the reader to [91] for more details on the theoretical aspects of diffusion models.

3.2. Blended positional encodings

Diffusion models can learn strong motion priors that ensure any motion generated is realistic and plausible [70]. In fact, they can also generate smooth transitions between subsequences [7, 70, 100]. However, these capabilities stem from inference-time motion infilling techniques, which we argue do not exploit the full potential of human motion priors. In fact, building a prior that extrapolates well to sequences longer than those observed during training is very challenging. The field of natural language processing has made progress in sequence extrapolation techniques, notably by substituting absolute positional encoding (APE) with a relative (RPE) counterpart [37]. By only providing information regarding how far tokens are between them, they achieve sequence-wise translation invariance and, therefore, can extrapolate their modeling capabilities to longer sequences. Yet, the absolute positions of poses within a motion, including their distances to the start and end of the action, are necessary to build the global semantics of the motion, as exemplified at the beginning of this section.

Here, we propose BPE, a novel positional encoding scheme designed for diffusion models that enables motion extrapolation while preserving the global motion semantics. Our BPE is inspired by the observation that in motion, high frequencies encompass local fine details, whereas low frequencies capture global structures. Similar insights have been drawn for images [59]. Diffusion models excel at decomposing the generation process into recovering lower frequencies, and gradually transitioning to higher frequencies. Fig. 2 shows how at early denoising phases, motion diffusion models prioritize global inter-frame dependencies, shifting towards local relative dependencies as the process unfolds. The proposed BPE harmonizes these dynamics *during inference*: at early denoising stages, our de-




Figure 2. Attention scores of a single query pose (current frame) as a function of the pose attended to (x-axis) in a diffusion-based motion generation model with a sinusoidal absolute positional encoding. Curves show the scores at each denoising step. We observe that, whereas early steps show strong global dependencies (blue), later denoising stages exhibit a clearly local behavior (red).

noising model is fed with an APE and, towards the conclusion, with an RPE. A scheduler guides this transition. As a result, intra-subsequence global dependencies are recovered at the beginning of the denoising, and intra- and inter-subsequences motion smoothness and realism are promoted later. To make the model understand APE and RPE at inference, we expose it to both encodings by randomly alternating them during training. As a result, the BPE schedule can be tuned at inference time to balance the intra-subsequence coherence and the inter-subsequence realism trade-off.

Rotary Position Encoding (RoPE). Our choice for RPE is rotary embeddings [73]. RoPE integrates a position embedding into the queries and keys, ensuring that after dot-product multiplication, the attention scores’ positional information reflects only the relative pairwise distance between queries and keys. Specifically, let W_q and W_k be the projection matrices into the d -dimensional spaces of queries and keys. Then, RoPE encodes the absolute positions m and n of a pair of query ($q_m = W_q x_m$) and key ($k_n = W_k x_n$), respectively, as d -dimensional rotations R_m^d, R_n^d over the projected poses x_m, x_n . The rotation angles are parameterized by m and n so that the attention formulation becomes:

$$q_m^T k_n = (R_m^d W_q x_m)^T (R_n^d W_k x_n) = x_m^T W_q R_{n-m}^d W_k x_n. \quad (1)$$

Note that the resulting rotation R_{n-m}^d only depends on the distance between n and m , and any absolute information about n or m is removed. RoPE is a natural choice for our RPE due to its simplicity and convenient injection before the attention takes place. As a result, RoPE is compatible with faster attention techniques like FlashAttention [22, 23].

Sinusoidal Position Encoding. Our APE is the classic sinusoidal position encoding [82], which leverages sine and cosine functions to inject positional information. It is added to the queries, keys, and values of the attention layers.

Note that for APE, attention is limited to each subsequence, while for RPE, attention spans all frames up to the attention horizon $H < L < N$. Since L defines the maximum range of motion dynamics learned during RPE training, there is no advantage in setting $H \geq L$ (Tabs. D/E in supp.

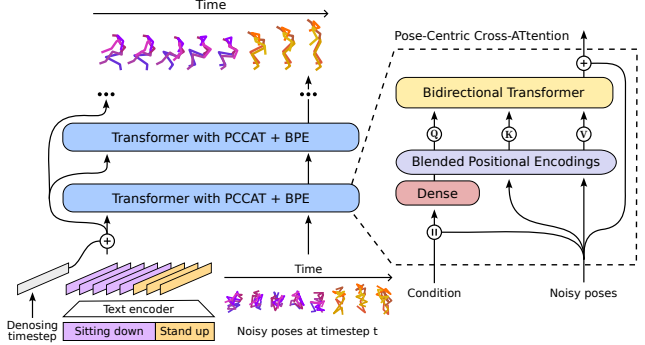


Figure 3. **Pose-centric cross-attention.** Our attention minimizes the entanglement between the control signal (e.g., text, objects) and the noisy motion by feeding the former only to the query. Consequently, our model denoises each frame’s noisy pose only leveraging its own condition, and the neighboring noisy poses.

material). Leveraging both APE and RPE constraints ensures quadratic complexity over the maximum subsequence length L in both memory and computation [11]. As a result, FlowMDM’s complexity is equivalent to that of other Transformer-based motion diffusion models [70, 100].

3.3. Pose-centric cross-attention

In order to make motion generation with diffusion models efficient, we would like to *simultaneously* generate very long sequences. In motion Transformers, the generation is conditioned at a sequence level by injecting the condition as a token [79], or as a sequence-wise transformation in intermediate layers [99]. Therefore, they cannot be conditioned on multiple signals in different subsequences. For this reason, diffusion-based methods for HMC opted for individually generating sequences and then merging them [70, 100]. To enable such simultaneous heterogeneous conditioning without any extra postprocessing, we propose to inject the condition at every frame. However, we still need to deal with a challenge: the condition never varies at training time. Therefore, at inference time, attention scores are computed with the embeddings E_{x_m, c_m} and E_{x_n, c_n} of the pose-condition pairs (x_m, c_m) and (x_n, c_n) as:

$$q_m^T k_n = (W_q E_{x_m, c_m})^T (W_k E_{x_n, c_n}) = E_{x_m, c_m}^T W_q^T W_k E_{x_n, c_n}. \quad (2)$$

When $c_m \neq c_n$, $q_m^T k_n$ was never encountered during training. If instead of injecting the condition at every frame, we used cross-attention layers, distinct conditions would also be temporally mixed, and we would face the same problem. To reduce the presence and impact of such training-inference misalignment, we introduce PCCAT, see Fig. 3, which aims at minimizing the entanglement between conditions and noisy poses. Specifically, PCCAT combines every frame’s noisy pose and condition into queries, while using only noisy poses as keys and values. Thus, Eq. 2 becomes:

$$q_m^T k_n = (W_q E_{x_m, c_m})^T (W_k E_{x_n}) = E_{x_m, c_m}^T W_q^T W_k E_{x_n}. \quad (3)$$

With PCCAT, the attention output for pose m becomes a weighted average of the value projections of its neighboring noisy poses. A residual connection adds the PCCAT output to the noisy poses. With comprehensive coverage of the motion spectrum in the training dataset, the network observes various poses preceding and following each pose, particularly within its local neighborhood. Therefore, local relationships do not suffer from unseen intermediate representations. Still, there is an obstacle to address: long-range dependencies. However, as discussed in Sec. 3.2, their importance is mostly confined to the initial stages of denoising. There, the network is exposed to very noisy motion data, thus becoming robust to such unseen combinations of poses. In the latest denoising stages, when the network deals with almost clean input sequences, global dependencies have already been developed and attention is short-ranged (Fig. 2).

4. Experiments

4.1. Experimental setup

Datasets. Our experiments are conducted on the Babel [63] and HumanML3D [26] datasets, with their train and test splits. HumanML3D features multiple textual descriptions of each motion sequence, but lacks explicit transition annotations, making supervised learning infeasible for transition generation. Babel, on the other hand, provides finely-grained textual descriptions at an atomic level, including transitions, which facilitates more precise and dynamic motion control but also presents a greater challenge due to fast and short transitions. To demonstrate the flexibility of FlowMDM, we employ the standard motion representations provided with each dataset. HumanML3D utilizes a 263D pose vector that includes joint coordinates, angles, velocities, and feet contact. By contrast, Babel uses the global position and orientation and a 6D rotation representation [103] of the SMPL model joints [13], as in [61].

Evaluation. Our evaluation uses the metrics established by [26], and later refined for this task in [47, 70, 92]. More specifically, motion sequences are synthesized as compositions of 32 pairs of textual descriptions and their durations. The 32 subsequences and the 31 transitions between S_{i-1} and S_i pairs are evaluated independently. In particular, each transition is defined as the set of consecutive poses $\{x_{\tau_i-L_{tr}/2}, \dots, x_{\tau_i+L_{tr}/2-1}\}$, sharing $\frac{L_{tr}}{2}$ frames with S_{i-1} and S_i . The transition duration L_{tr} is set to 30/60 frames for Babel/HumanML3D (1/3 seconds). The top-3 R-precision (R-prec), and the multimodal distance (MM-Dist) are used to evaluate how well the subsequences’ motion matches their textual description [26]. The FID score and the average pairwise distance among all motion embeddings (diversity) assess the quality and variety of both subsequences and transitions, respectively [26, 31].

Closing the gap: the Jerk. Generative models are hard

to evaluate [19, 68, 91]. The FID score [31] has proven to be a very reliable metric in quantifying the similarity between distributions of generated and real motion data while being sensitive to motion artifacts or noise [54]. Nevertheless, only relying on perceptual metrics like FID for assessing transition quality can be misleading due to their insensitivity to motion anomalies such as abrupt accelerations [8], or foot skating [55]. To complement the FID, our work introduces two novel metrics built upon the concept of *jerk* (i.e., time derivative of acceleration), which is indicative of motion smoothness and proven sensitive to kinetic irregularities [5, 6, 16, 25, 34, 43, 95]. Given that natural human motion typically exhibits constrained jerk due to relatively consistent acceleration patterns [25, 43], our metrics are tailored to highlight *persistent deviations* from this norm on transitions. Firstly, we compute the Peak Jerk (PJ) as the maximum jerk value throughout the transition motion over all joints. While this measure captures extreme fluctuations, it favors models that unnaturally smooth transitions across several wider peaks of jerk. To measure this undesirable effect, we introduce the Area Under the Jerk (AUJ), calculated as the sum of L1-norm differences between a method’s instantaneous jerk and the dataset’s average jerk value. This measure serves as an aggregate indicator of motion smoothness, quantifying the cumulative deviation from natural human movement across the entire transition. The PJ and AUJ of a transition are formally defined as follows:

$$PJ = \max_{\substack{1 \leq i \leq K \\ 1 \leq \tau \leq L_{tr}}} |j_i(\tau)|_1, \quad AUJ = \sum_{\tau=1}^{L_{tr}} \max_{1 \leq i \leq K} |j_i(\tau) - j_{avg}|_1, \quad (4)$$

where $j_i(\tau)$ is the jerk at time τ for joint i , K is the number of joints, and j_{avg} is the average joints-wise maximum jerk across the dataset.

Baselines. We compare our method to publicly released related works that can generate sequential motions from text: the autoregressive TEACH [4], and the diffusion sampling techniques DoubleTake [70], DiffCollage [100], and MultiDiffusion [7]. Sampling techniques are evaluated with PCCAT and APE for a fairer comparison. Additionally, we evaluate TEACH with its spherical linear interpolation over transitions turned off (TEACH.B), and DoubleTake with MDM, as originally proposed (DoubleTake*). TEACH and TEACH.B cannot be trained for HumanML3D due to the lack of pairs of consecutive actions and textual descriptions.

Implementation details. We tune the hyperparameters of all models with grid search. The attention horizon for RPE, H , is set to 100/150 for Babel/HumanML3D. The number of diffusion steps is 1K for all experiments. Our model is trained with the x_0 parameterization [87], and minimizes the L2 reconstruction loss. During training, RPE and APE are alternated randomly at a frequency of 0.5. We use classifier-free guidance with weights 1.5/2.5 [33]. We use

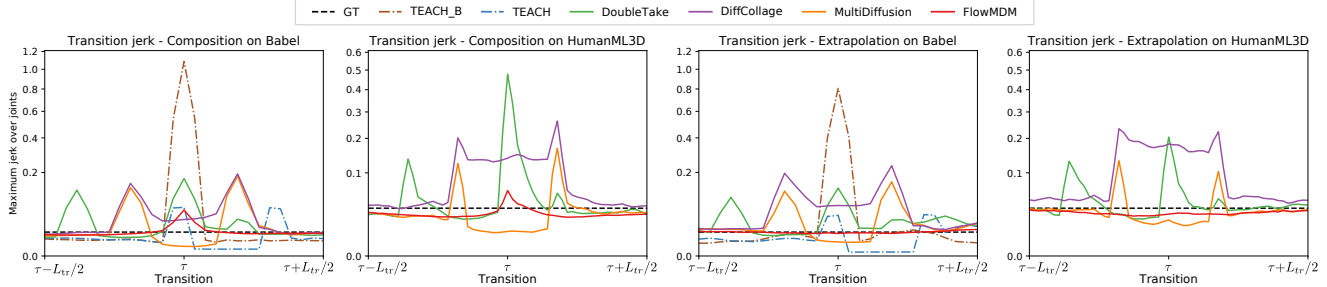


Figure 4. **Transitions smoothness.** Average maximum jerk over joints at each frame of the transitions for both motion composition (left) and extrapolation (right) tasks. While other methods show severe smoothness artifacts in the beginning and end of their transition refinement processes, FlowMDM’s jerk curve has the shortest peak for composition, and an absence of peaks for extrapolation.

	R-prec \uparrow	Subsequence			FID \downarrow	Transition		
		FID \downarrow	Div \rightarrow	MM-Dist \downarrow		Div \rightarrow	PJ \rightarrow	AUJ \downarrow
GT	0.715 \pm 0.003	0.00 \pm 0.00	8.42 \pm 0.15	3.36 \pm 0.00	0.00 \pm 0.00	6.20 \pm 0.06	0.02 \pm 0.00	0.00 \pm 0.00
TEACH_B	0.703 \pm 0.002	1.71 \pm 0.03	8.18 \pm 0.14	3.43 \pm 0.01	3.01 \pm 0.04	6.23 \pm 0.05	1.09 \pm 0.00	2.35 \pm 0.01
TEACH	0.655 \pm 0.002	1.82 \pm 0.02	7.96 \pm 0.11	3.72 \pm 0.01	3.27 \pm 0.04	6.14 \pm 0.06	0.07 \pm 0.00	0.44 \pm 0.00
DoubleTake*	0.596 \pm 0.005	3.16 \pm 0.06	7.53 \pm 0.11	4.17 \pm 0.02	3.33 \pm 0.06	<u>6.16</u> \pm 0.05	0.28 \pm 0.00	1.04 \pm 0.01
DoubleTake	0.668 \pm 0.005	<u>1.33</u> \pm 0.04	7.98 \pm 0.12	3.67 \pm 0.03	3.15 \pm 0.05	<u>6.14</u> \pm 0.07	0.17 \pm 0.00	0.64 \pm 0.01
MultiDiffusion	0.702 \pm 0.005	1.74 \pm 0.04	8.37 \pm 0.13	3.43 \pm 0.02	6.56 \pm 0.12	5.72 \pm 0.07	0.18 \pm 0.00	0.68 \pm 0.00
DiffCollage	0.671 \pm 0.003	1.45 \pm 0.05	7.93 \pm 0.09	3.71 \pm 0.01	4.36 \pm 0.09	6.09 \pm 0.08	0.19 \pm 0.00	0.84 \pm 0.01
FlowMDM	<u>0.702</u> \pm 0.004	0.99 \pm 0.04	<u>8.36</u> \pm 0.13	3.45 \pm 0.02	2.61 \pm 0.06	6.47 \pm 0.05	0.06 \pm 0.00	0.13 \pm 0.00

Table 1. Comparison of FlowMDM with the state of the art in Babel. Symbols \uparrow , \downarrow , and \rightarrow indicate that higher, lower, or values closer to the ground truth (GT) are better, respectively. Evaluation is run 10 times and \pm specifies the 95% confidence intervals.

	R-prec \uparrow	Subsequence			FID \downarrow	Transition		
		FID \downarrow	Div \rightarrow	MM-Dist \downarrow		Div \rightarrow	PJ \rightarrow	AUJ \downarrow
GT	0.796 \pm 0.004	0.00 \pm 0.00	9.34 \pm 0.08	2.97 \pm 0.01	0.00 \pm 0.00	9.54 \pm 0.15	0.04 \pm 0.00	0.07 \pm 0.00
DoubleTake*	0.643 \pm 0.005	<u>0.80</u> \pm 0.02	9.20 \pm 0.11	3.92 \pm 0.01	1.71 \pm 0.05	8.82 \pm 0.13	0.52 \pm 0.01	2.10 \pm 0.03
DoubleTake	0.628 \pm 0.005	1.25 \pm 0.04	9.09 \pm 0.12	4.01 \pm 0.01	4.19 \pm 0.09	8.45 \pm 0.09	0.48 \pm 0.00	1.83 \pm 0.02
MultiDiffusion	0.629 \pm 0.002	1.19 \pm 0.03	9.38 \pm 0.08	4.02 \pm 0.01	4.31 \pm 0.06	8.37 \pm 0.10	<u>0.17</u> \pm 0.00	<u>1.06</u> \pm 0.01
DiffCollage	0.615 \pm 0.005	1.56 \pm 0.04	8.79 \pm 0.08	4.13 \pm 0.02	4.59 \pm 0.10	8.22 \pm 0.11	0.26 \pm 0.00	2.85 \pm 0.09
FlowMDM	0.685 \pm 0.004	0.29 \pm 0.01	9.58 \pm 0.12	3.61 \pm 0.01	1.38 \pm 0.05	8.79 \pm 0.09	0.06 \pm 0.00	0.51 \pm 0.01

Table 2. Comparison of FlowMDM with the state of the art in HumanML3D.

a binary step function to guide the BPE sampling, yielding 125/60 initial APE steps. The minimum/maximum lengths for training subsequences are set to 30/200 and 70/200 frames (i.e., 1/6.7s and 3.5/10s). For Babel, training subsequences include consecutive ground truth motions with distinct textual descriptions in order to increase the motions variability, and make the network explicitly robust to multiple conditions. The ablation study includes two conditioning baselines: 1) concatenating each frame’s condition and noisy pose, and replacing the PCCAT with vanilla self-attention (SAT), and 2) injecting the condition with cross-attention layers (CAT). More details in supp. material A.

4.2. Quantitative analysis

Comparison with the state of the art on HMC. Tables 1 and 2 show the comparison of FlowMDM with current state-of-the-art models in Babel and HumanML3D datasets, respectively. In HumanML3D, our model outperforms by a fair margin the other methods in terms of subsequence accuracy-wise metrics (R-prec and MM-Dist), and FID. In

Babel, it matches the state of the art in accuracy and excels in FID score. FlowMDM produces transitions of higher quality and smoothness on both datasets, as indicated by FID, PJ, and AUJ metrics. The lack of correlation between the FID score and the AUJ underscores the importance of the latter as a complementary metric for assessing smoothness. Fig. 4-left shows the average jerk values across the generated transitions. We observe that state-of-the-art methods exhibit severe smoothness artifacts. During TEACH’s spherical linear interpolation, the jerk quickly reaches values near zero. By contrast, DiffCollage leans toward higher-than-average jerk values, while MultiDiffusion exhibits the opposite trend. DoubleTake shows three peaks, caused by their two-stage noise estimation process. In comparison, FlowMDM successfully minimizes peak jerk values, producing the smoothest transitions between subsequences. See supp. material Sec. C for in-depth analyses.

Human motion extrapolation. In single text-to-motion, the duration of the generated motion is limited to the longest subsequence length L available in the training set. Extrapolation

Cond.	Train. PE	Inf. PE	Subsequence				Transition				
			R-prec \uparrow	FID \downarrow	Div \rightarrow	MM-Dist \downarrow	FID \downarrow	Div \rightarrow	PJ \rightarrow	AUJ \downarrow	
GT	-	-	0.715 \pm 0.003	0.00 \pm 0.00	8.42 \pm 0.15	3.36 \pm 0.00	0.00 \pm 0.00	6.20 \pm 0.06	0.02 \pm 0.00	0.00 \pm 0.00	
PCCAT	A	A	0.699 \pm 0.004	1.34 \pm 0.04	8.36 \pm 0.12	3.40 \pm 0.02	4.26 \pm 0.07	5.98 \pm 0.06	1.81 \pm 0.01	3.73 \pm 0.01	
PCCAT	R	R	0.635 \pm 0.006	1.28 \pm 0.03	8.05 \pm 0.11	4.02 \pm 0.02	2.18 \pm 0.07	6.14 \pm 0.08	<u>0.03</u> \pm 0.00	0.20 \pm 0.00	
PCCAT	B	A	0.716 \pm 0.006	1.20 \pm 0.04	8.31 \pm 0.14	3.32 \pm 0.02	3.01 \pm 0.06	6.35 \pm 0.07	1.78 \pm 0.01	3.66 \pm 0.02	
PCCAT	B	R	0.635 \pm 0.004	0.85 \pm 0.02	8.25 \pm 0.12	3.98 \pm 0.02	<u>2.14</u> \pm 0.04	6.44 \pm 0.09	0.04 \pm 0.00	0.15 \pm 0.00	
SAT	B	B	0.681 \pm 0.004	1.52 \pm 0.04	8.22 \pm 0.11	3.61 \pm 0.02	1.91 \pm 0.03	6.41 \pm 0.07	0.06 \pm 0.00	<u>0.12</u> \pm 0.00	
CAT	B	B	0.719 \pm 0.004	1.29 \pm 0.02	8.16 \pm 0.13	3.27 \pm 0.02	2.57 \pm 0.08	<u>6.06</u> \pm 0.07	0.02 \pm 0.00	0.07 \pm 0.00	
PCCAT	B	B	0.702 \pm 0.004	<u>0.99</u> \pm 0.04	8.36 \pm 0.13	3.45 \pm 0.02	2.61 \pm 0.06	6.47 \pm 0.05	0.06 \pm 0.00	0.13 \pm 0.00	

Table 3. Ablation study in Babel. Cond. indicates the conditioning scheme, Train./Inf. PE specify the positional encodings (PE) used at training/inference time, and A, R, and B refer to absolute, relative, and blended PE, respectively. \uparrow , \downarrow , and \rightarrow indicate that higher, lower, or values closer to the ground truth (GT) are better, respectively. Evaluation is run 10 times and \pm specifies the 95% confidence intervals.

Cond.	Train. PE	Inf. PE	Subsequence				Transition				
			R-prec \uparrow	FID \downarrow	Div \rightarrow	MM-Dist \downarrow	FID \downarrow	Div \rightarrow	PJ \rightarrow	AUJ \downarrow	
GT	-	-	0.796 \pm 0.004	0.00 \pm 0.00	9.34 \pm 0.08	2.97 \pm 0.01	0.00 \pm 0.00	9.54 \pm 0.15	0.04 \pm 0.00	0.07 \pm 0.00	
PCCAT	A	A	0.689 \pm 0.005	0.66 \pm 0.02	9.73 \pm 0.12	3.63 \pm 0.02	3.90 \pm 0.12	8.29 \pm 0.08	1.50 \pm 0.01	3.40 \pm 0.02	
PCCAT	R	R	0.531 \pm 0.005	1.75 \pm 0.07	8.71 \pm 0.10	4.80 \pm 0.03	2.53 \pm 0.12	8.62 \pm 0.08	<u>0.03</u> \pm 0.00	0.58 \pm 0.01	
PCCAT	B	A	0.699 \pm 0.005	0.61 \pm 0.02	9.76 \pm 0.10	<u>3.54</u> \pm 0.02	2.42 \pm 0.09	8.39 \pm 0.09	1.40 \pm 0.01	3.29 \pm 0.02	
PCCAT	B	R	0.554 \pm 0.007	1.06 \pm 0.06	9.02 \pm 0.11	4.54 \pm 0.02	1.12 \pm 0.04	9.00 \pm 0.10	0.05 \pm 0.00	0.53 \pm 0.01	
SAT	B	B	0.692 \pm 0.004	<u>0.49</u> \pm 0.02	<u>9.08</u> \pm 0.09	3.51 \pm 0.01	3.19 \pm 0.08	8.09 \pm 0.11	0.04 \pm 0.00	0.36 \pm 0.02	
CAT	B	B	0.622 \pm 0.005	1.27 \pm 0.04	8.86 \pm 0.15	4.10 \pm 0.01	3.93 \pm 0.14	8.23 \pm 0.10	0.04 \pm 0.00	<u>0.49</u> \pm 0.02	
PCCAT	B	B	0.685 \pm 0.004	0.29 \pm 0.01	9.58 \pm 0.12	3.61 \pm 0.01	<u>1.38</u> \pm 0.05	<u>8.79</u> \pm 0.09	0.06 \pm 0.00	0.51 \pm 0.01	

Table 4. Ablation study in HumanML3D.

olating periodic actions into sequences longer than those in the ground truth presents a notable challenge. Achieving this through HMC requires the harmonization of periodicity across adjacent subsequences. Common strategies that combine independently generated subsequences often disrupt the periodicity of the motion. To assess our model’s capabilities in addressing this issue, we construct an evaluation set comprising 32 consecutive repetitions of 32 different extrapolatable actions such as ‘walk forward’, ‘jumping’, or ‘playing the guitar’, extracted from the Babel and HumanML3D test sets (more details in supp. material Sec. B). Fig. 4-right displays the motion jerk across transitions for all models on this task. We observe that, while other models exhibit smoothness anomalies similar to those in HMC, FlowMDM closely mirrors the ground truth jerk. This observation indicates that the jerk peak noted in FlowMDM for the composition task is likely attributed to smoothness irregularities in more complex transitions.

Ablation study. The effectiveness of BPE and PCCAT is presented in Tables 3 and 4. Reasonably, the baseline model trained solely with APE fails to generate smooth transitions. Conversely, a model trained only with RPE, despite producing the smoothest transitions, struggles to model global motion dependencies and accurately reflect the corresponding textual descriptions. Interestingly, training with BPE improves the performance of both APE- and RPE-only samplings. Sampling with BPE combines the best of both worlds by preserving the excellent AUJ values of the RPE models and reaching the state-of-the-art accuracy and FID

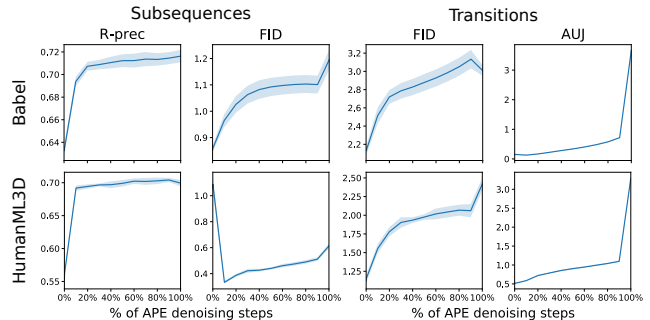


Figure 5. **BPE trade-offs.** Increasing the number of APE steps undergone during BPE sampling improves the correspondence between motion and textual description (R-prec), but reduces the transition realism and smoothness (FID and AUJ). The best balance is reached around 10% of APE denoising steps.

scores of the APE models. Fig. 5 illustrates this balance. Specifically, increasing the number of APE steps enhances the motion’s congruence with the textual description, at the cost of reducing the smoothness and realism of the transitions. In HumanML3D, the SAT and CAT conditioning schemes lead to worse transitions in terms of FID and diversity. This is caused by the coexistence of different conditions in the local neighborhood of the transition at inference, which never happens during training. Our PCCAT conditioning technique effectively solves this problem. In Babel, such effect is not present because the training motion sequences include several subsequences, thus increasing the model’s robustness to transitions with varying conditions.

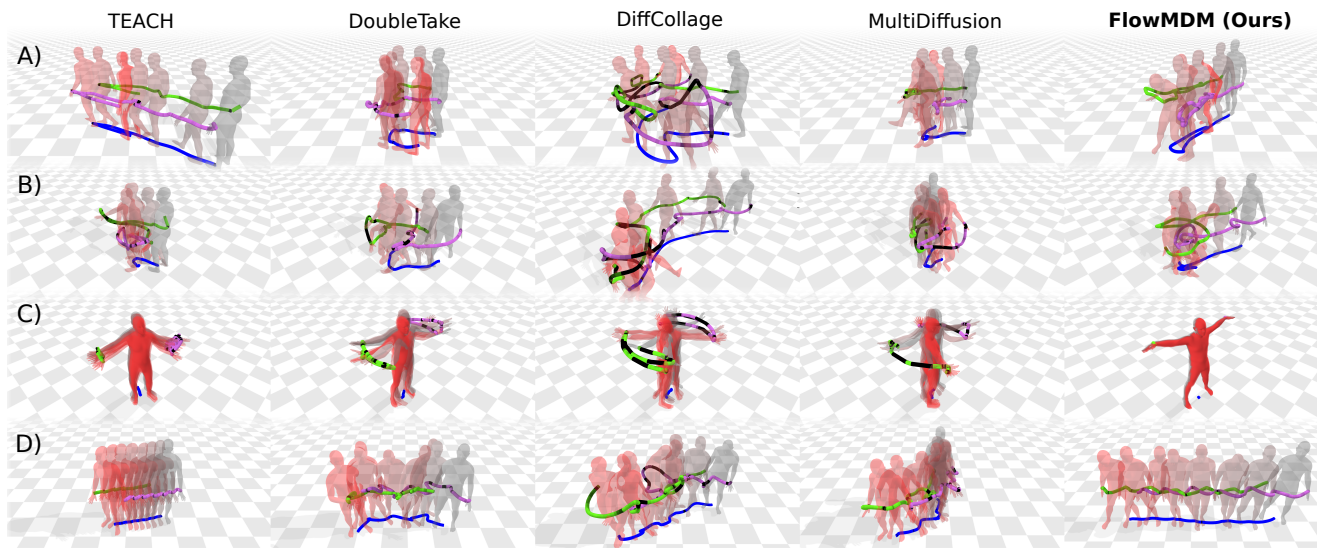


Figure 6. **Qualitative analysis (Babel)**. A) and B) show compositions of 3 motions (‘walk straight’→‘side steps’→‘walk backward’, and ‘walk’→‘turn around’→‘sit on the bench’, respectively), and C) and D) illustrate extrapolations that repeat 6 times a static (‘t-pose’) and a dynamic (‘step to the right’) action, respectively. Solid curves match the trajectories of the global position (blue) and left/right hands (purple/green). Darker colors indicate instantaneous jerk deviations from the median value, saturating at twice the jerk’s standard deviation in the dataset (black segments). Abrupt transitions manifest as black segments amidst lighter ones. FlowMDM exhibits the most fluid motion and preserves the staticity or periodicity of extrapolated actions, in contrast to other methods that show spontaneous high jerk values and fail to keep the motion coherence in extrapolations.

On the efficiency of FlowMDM. Diffusion-based state-of-the-art methods such as MultiDiffusion and DiffCollage denoise poses from the transition more than once in order to harmonize it with the adjacent motions. DoubleTake’s transitions undergo an additional denoising process, which adds computational burden and can not be parallelized. Oppositely, FlowMDM does not apply redundant denoising steps to any pose. In particular, our model goes through 47.1%, 28.4%, and 16.5% less pose-wise denoising steps than DoubleTake, DiffCollage, and MultiDiffusion, respectively.

4.3. Qualitative results

Fig. 6 illustrates how our quantitative findings translate into visual outcomes on the human motion composition and extrapolation tasks. First, as anticipated by Fig. 4, we confirm that state-of-the-art methods produce short intervals of jerk peaks around transitions. These do not typically match long-range motion scenarios, where such jerks might be contextually appropriate. Contrarily, FlowMDM produces motion that is realistic, accurate, and smooth. Particularly, we notice that DiffCollage’s bias toward producing constantly high jerk values around transitions is perceived as an overall chaotic motion. Due to the independent generation of their subsequences, DoubleTake, DiffCollage, and MultiDiffusion are unable to maintain the static or periodic nature of actions when extrapolating them. Only TEACH and FlowMDM are able to successfully extrapolate a static ‘t-pose’, and ours is the only one capable of extrapolating a ‘step to the right’ sequence realistically. Finally, FlowMDM

also inherits the trajectory control capabilities of motion diffusion models as shown in Fig. 1-right.

5. Conclusion

We presented FlowMDM, the first approach that generates human motion compositions simultaneously, without undergoing postprocessing or redundant denoising diffusion steps. We also introduced the blended positional encodings to combine the benefits of absolute and relative positional encodings during the denoising chain. Finally, we presented the pose-centric cross-attention, a technique that improves the generation of transitions when training with only a single condition per motion sequence.

Limitations and future work. The absolute stage of BPE does not model relationships between subsequences. Consequently, their low-frequency spectrum is generated independently. This limitation could be addressed in future work by incorporating an intention planning module. Finally, our method learns a strong motion prior that generates transitions between combinations of actions never seen at training time. Such capability could theoretically be used with different models leveraging different control signals, assuming they all are trained under the same framework. Future work will experimentally validate this hypothesis.

Acknowledgements. This work has been partially supported by the Spanish projects PID2022-136436NB-I00, TED2021-131317B-I00, and PDC2022-133305-I00 and by ICREA under the ICREA Academia programme.

References

- [1] Vida Adeli, Mahsa Ehsanpour, Ian Reid, Juan Carlos Niebles, Silvio Savarese, Ehsan Adeli, and Hamid Rezaatoughi. Tripod: Human trajectory and pose dynamics forecasting in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13390–13400, 2021. 2
- [2] Simon Alexanderson, Rajmund Nagy, Jonas Beskow, and Gustav Eje Henter. Listen, denoise, action! audio-driven motion synthesis with diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–20, 2023. 2
- [3] Sadeqh Aliakbarian, Microsoft Fatemeh Saleh ACRV, Stephen Gould ACRV, and Anu Mathieu Salzmann CVLab. Contextually plausible and diverse 3d human motion prediction. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [4] Nikos Athanasiou, Mathis Petrovich, Michael J Black, and Gül Varol. Teach: Temporal action composition for 3d humans. In *2022 International Conference on 3D Vision (3DV)*, pages 414–423. IEEE, 2022. 2, 5
- [5] Sivakumar Balasubramanian, Alejandro Melendez-Calderon, and Etienne Burdet. A robust and sensitive metric for quantifying movement smoothness. *IEEE transactions on biomedical engineering*, 59(8):2126–2136, 2011. 5
- [6] Sivakumar Balasubramanian, Alejandro Melendez-Calderon, Agnes Roby-Brami, and Etienne Burdet. On the analysis of movement smoothness. *Journal of neuroengineering and rehabilitation*, 12(1):1–11, 2015. 5
- [7] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. Multidiffusion: Fusing diffusion paths for controlled image generation. In *International Conference on Machine Learning*, pages 1737–1752. PMLR, 2023. 2, 3, 5
- [8] German Barquero, Sergio Escalera, and Cristina Palmero. Belfusion: Latent diffusion for behavior-driven human motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2317–2327, 2023. 2, 5
- [9] German Barquero, Johnny Núñez, Sergio Escalera, Zhen Xu, Wei-Wei Tu, Isabelle Guyon, and Cristina Palmero. Didn't see that coming: a survey on non-verbal social human behavior forecasting. In *Understanding Social Behavior in Dyadic and Small Group Interactions*, pages 139–178. PMLR, 2022. 2
- [10] German Barquero, Johnny Núñez, Zhen Xu, Sergio Escalera, Wei-Wei Tu, Isabelle Guyon, and Cristina Palmero. Comparison of spatio-temporal models for human motion and pose forecasting in face-to-face interaction scenarios. In *Understanding Social Behavior in Dyadic and Small Group Interactions*, pages 107–138. PMLR, 2022. 2
- [11] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020. 4
- [12] 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. 2
- [13] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 561–578. Springer, 2016. 5
- [14] Paulo Vinicius Koerich Borges, Nicola Conci, and Andrea Cavallaro. Video-based human behavior understanding: A survey. *IEEE transactions on circuits and systems for video technology*, 23(11):1993–2008, 2013. 1
- [15] Zhe Cao, Hang Gao, Karttikeya Mangalam, Qi-Zhi Cai, Minh Vo, and Jitendra Malik. Long-term human motion prediction with scene context. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 387–404. Springer, 2020. 2
- [16] Angela Castillo, Maria Escobar, Guillaume Jeanneret, Albert Pumarola, Pablo Arbeláez, Ali Thabet, and Arsiom Sanakoyeu. Bodiffusion: Diffusing sparse observations for full-body human motion synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4221–4231, 2023. 5
- [17] Kang Chen, Zhipeng Tan, Jin Lei, Song-Hai Zhang, Yuan-Chen Guo, Weidong Zhang, and Shi-Min Hu. Choreomaster: choreography-oriented music-driven dance synthesis. *ACM Transactions on Graphics (TOG)*, 40(4):1–13, 2021. 2
- [18] Enric Corona, Albert Pumarola, Guillem Alenya, and Francesc Moreno-Noguer. Context-aware human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6992–7001, 2020. 2
- [19] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1):53–65, 2018. 5
- [20] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2
- [21] 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
- [22] Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. In *The Twelfth International Conference on Learning Representations*, 2023. 4
- [23] Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359, 2022. 4
- [24] David A Engström, JA Scott Kelso, and Tom Holroyd. Reaction-anticipation transitions in human perception-action patterns. *Human movement science*, 15(6):809–832, 1996. 2
- [25] Philipp Gulde and Joachim Hermsdörfer. Smoothness met-

- rics in complex movement tasks. *Frontiers in neurology*, 9:615, 2018. 5
- [26] 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, 5
- [27] Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *European Conference on Computer Vision*, pages 580–597. Springer, 2022. 2
- [28] Wen Guo, Xiaoyu Bie, Xavier Alameda-Pineda, and Francisc Moreno-Noguer. Multi-person extreme motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13053–13064, 2022. 2
- [29] Félix G Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. *ACM Transactions on Graphics (TOG)*, 39(4):60–1, 2020. 2
- [30] Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael J Black. Stochastic scene-aware motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11374–11384, 2021. 2
- [31] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 5
- [32] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2
- [33] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. 5
- [34] Neville Hogan and Dagmar Sternad. Sensitivity of smoothness measures to movement duration, amplitude, and arrests. *Journal of motor behavior*, 41(6):529–534, 2009. 5
- [35] 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. 2
- [36] Manuel Kaufmann, Emre Aksan, Jie Song, Fabrizio Pece, Remo Ziegler, and Otmar Hilliges. Convolutional autoencoders for human motion infilling. In *2020 International Conference on 3D Vision (3DV)*, pages 918–927. IEEE, 2020. 2
- [37] Amirhossein Kazemnejad, Inkit Padhi, Karthikeyan Natesan Ramamurthy, Payel Das, and Siva Reddy. The impact of positional encoding on length generalization in transformers. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [38] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019. 3
- [39] Jihoon Kim, Taehyun Byun, Seungyoun Shin, Jungdam Won, and Sungjoon Choi. Conditional motion in-betweening. *Pattern Recognition*, 132:108894, 2022. 2
- [40] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. Flame: Free-form language-based motion synthesis & editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 8255–8263, 2023. 2
- [41] Nilesh Kulkarni, Davis Rempe, Kyle Genova, Abhijit Kundu, Justin Johnson, David Fouhey, and Leonidas Guibas. Nifty: Neural object interaction fields for guided human motion synthesis. *arXiv preprint arXiv:2307.07511*, 2023. 2
- [42] Wilfried Kunde, Katrin Elsner, and Andrea Kiesel. No anticipation–no action: the role of anticipation in action and perception. *Cognitive Processing*, 8:71–78, 2007. 2
- [43] Caroline Larboulette and Sylvie Gibet. A review of computable expressive descriptors of human motion. In *Proceedings of the 2nd International Workshop on Movement and Computing*, pages 21–28, 2015. 5
- [44] Taeryung Lee, Gyeongsik Moon, and Kyoung Mu Lee. Multiact: Long-term 3d human motion generation from multiple action labels. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 1231–1239, 2023. 2
- [45] Jiaman Li, Ruben Villegas, Duygu Ceylan, Jimei Yang, Zhengfei Kuang, Hao Li, and Yajie Zhao. Task-generic hierarchical human motion prior using vaes. In *2021 International Conference on 3D Vision (3DV)*, pages 771–781. IEEE, 2021. 2
- [46] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13401–13412, 2021. 2
- [47] Shuai Li, Sisi Zhuang, Wenfeng Song, Xinyu Zhang, Hejia Chen, and Aimin Hao. Sequential texts driven cohesive motions synthesis with natural transitions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9498–9508, 2023. 2, 5
- [48] Weiyu Li, Xuelin Chen, Peizhuo Li, Olga Sorkine-Hornung, and Baoquan Chen. Example-based motion synthesis via generative motion matching. *ACM Transactions on Graphics (TOG)*, 42(4):1–12, 2023. 2
- [49] Yunhao Li, Zhenbo Yu, Yucheng Zhu, Bingbing Ni, Guangtao Zhai, and Wei Shen. Skeleton2humanoid: Animating simulated characters for physically-plausible motion in-betweening. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1493–1502, 2022. 2
- [50] Jing Lin, Ailing Zeng, Shunlin Lu, Yuanhao Cai, Ruimao Zhang, Haoqian Wang, and Lei Zhang. Motion-x: A large-scale 3d expressive whole-body human motion dataset. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 2
- [51] Zhengyi Luo, Jinkun Cao, Kris Kitani, Weipeng Xu, et al. Perpetual humanoid control for real-time simulated avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10895–10904, 2023. 2
- [52] Hengbo Ma, Jiachen Li, Ramtin Hosseini, Masayoshi Tomizuka, and Chiho Choi. Multi-objective diverse human motion prediction with knowledge distillation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pat-*

- tern Recognition*, 2022. 2
- [53] Shugao Ma, Tomas Simon, Jason Saragih, Dawei Wang, Yuecheng Li, Fernando De La Torre, and Yaser Sheikh. Pixel codec avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 64–73, 2021. 1
- [54] Antoine Maiorca, Youngwoo Yoon, and Thierry Dutoit. Evaluating the quality of a synthesized motion with the fréchet motion distance. In *ACM SIGGRAPH 2022 Posters*, pages 1–2, 2022. 5
- [55] Antoine Maiorca, Youngwoo Yoon, and Thierry Dutoit. Validating objective evaluation metric: Is fréchet motion distance able to capture foot skating artifacts? In *Proceedings of the 2023 ACM International Conference on Interactive Media Experiences*, pages 242–247, 2023. 5
- [56] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. Generating smooth pose sequences for diverse human motion prediction. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [57] Boris N Oreshkin, Antonios Valkanas, Félix G Harvey, Louis-Simon Ménard, Florent Bocquet, and Mark J Coates. Motion in-betweening via deep delta-interpolator. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 2
- [58] Cristina Palmero, German Barquero, Julio CS Jacques Junior, Albert Clapés, Johnny Núñez, David Curto, Sorina Smeureanu, Javier Selva, Zejian Zhang, David Saeteros, et al. Chalearn lap challenges on self-reported personality recognition and non-verbal behavior forecasting during social dyadic interactions: Dataset, design, and results. In *Understanding Social Behavior in Dyadic and Small Group Interactions*, pages 4–52. PMLR, 2022. 2
- [59] Zizheng Pan, Jianfei Cai, and Bohan Zhuang. Fast vision transformers with hilo attention. *Advances in Neural Information Processing Systems*, 35:14541–14554, 2022. 3
- [60] Sang-Min Park and Young-Gab Kim. A metaverse: Taxonomy, components, applications, and open challenges. *IEEE access*, 10:4209–4251, 2022. 1
- [61] Mathis Petrovich, Michael J Black, and Gül Varol. Temos: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision*, pages 480–497. Springer, 2022. 2, 5
- [62] Matthias Plappert, Christian Mandery, and Tamim Asfour. The kit motion-language dataset. *Big data*, 4(4):236–252, 2016. 2
- [63] Abhinanda R Punnakkal, Arjun Chandrasekaran, Nikos Athanasiou, Alejandra Quiros-Ramirez, and Michael J Black. Babel: Bodies, action and behavior with english labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 722–731, 2021. 2, 5
- [64] Yijun Qian, Jack Urbanek, Alexander G Hauptmann, and Jungdam Won. Breaking the limits of text-conditioned 3d motion synthesis with elaborative descriptions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2306–2316, 2023. 2
- [65] Jia Qin, Youyi Zheng, and Kun Zhou. Motion in-betweening via two-stage transformers. *ACM Transactions on Graphics (TOG)*, 41(6):1–16, 2022. 2
- [66] Tianxiang Ren, Jubo Yu, Shihui Guo, Ying Ma, Yutao Ouyang, Zijiao Zeng, Yazhan Zhang, and Yipeng Qin. Diverse motion in-betweening with dual posture stitching. *arXiv preprint arXiv:2303.14457*, 2023. 2
- [67] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2
- [68] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. *Advances in neural information processing systems*, 31, 2018. 5
- [69] Tim Salzmann, Marco Pavone, and Markus Ryll. Motron: Multimodal probabilistic human motion forecasting. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [70] Yoni Shafir, Guy Tevet, Roy Kapon, and Amit Haim Bermano. Human motion diffusion as a generative prior. In *The Twelfth International Conference on Learning Representations*, 2023. 2, 3, 4, 5
- [71] 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
- [72] 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
- [73] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 4
- [74] Guofei Sun, Yongkang Wong, Zhiyong Cheng, Mohan S Kankanhalli, Weidong Geng, and Xiangdong Li. Deepdance: music-to-dance motion choreography with adversarial learning. *IEEE Transactions on Multimedia*, 23:497–509, 2020. 2
- [75] Jiarui Sun and Girish Chowdhary. Towards globally consistent stochastic human motion prediction via motion diffusion. *arXiv preprint arXiv:2305.12554*, 2023. 2
- [76] Ryo Suzuki, Adnan Karim, Tian Xia, Hooman Hedayati, and Nicolai Marquardt. Augmented reality and robotics: A survey and taxonomy for ar-enhanced human-robot interaction and robotic interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–33, 2022. 1
- [77] Julian Tanke, Linguang Zhang, Amy Zhao, Chengcheng Tang, Yujun Cai, Lezi Wang, Po-Chen Wu, Juergen Gall, and Cem Keskin. Social diffusion: Long-term multiple human motion anticipation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9601–9611, 2023. 2
- [78] Guy Tevet, Brian Gordon, Amir Hertz, Amit H Bermano, and Daniel Cohen-Or. Motionclip: Exposing human motion generation to clip space. In *European Conference on Computer Vision*, pages 358–374. Springer, 2022. 2
- [79] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel

- Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*, 2022. 2, 3, 4
- [80] Sib0 Tian, Minghui Zheng, and Xiao Liang. Transfusion: A practical and effective transformer-based diffusion model for 3d human motion prediction. *arXiv preprint arXiv:2307.16106*, 2023. 2
- [81] 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
- [82] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4
- [83] Jacob Walker, Kenneth Marino, Abhinav Gupta, and Martial Hebert. The pose knows: Video forecasting by generating pose futures. *Proceedings of the IEEE international conference on computer vision*, 2017. 2
- [84] Jingbo Wang, Yu Rong, Jingyuan Liu, Sijie Yan, Dahua Lin, and Bo Dai. Towards diverse and natural scene-aware 3d human motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20460–20469, 2022. 2
- [85] Jiashun Wang, Huazhe Xu, Jingwei Xu, Sifei Liu, and Xiaolong Wang. Synthesizing long-term 3d human motion and interaction in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9401–9411, 2021.
- [86] Jingbo Wang, Sijie Yan, Bo Dai, and Dahua Lin. Scene-aware generative network for human motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12206–12215, 2021. 2
- [87] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion GANs. In *International Conference on Learning Representations (ICLR)*, 2022. 5
- [88] Jiachen Xu, Min Wang, Jingyu Gong, Wentao Liu, Chen Qian, Yuan Xie, and Lizhuang Ma. Exploring versatile prior for human motion via motion frequency guidance. In *2021 International Conference on 3D Vision (3DV)*, pages 606–616. IEEE, 2021. 2
- [89] Sirui Xu, Zhengyuan Li, Yu-Xiong Wang, and Liang-Yan Gui. Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14928–14940, 2023. 2
- [90] Sirui Xu, Yu-Xiong Wang, and Liangyan Gui. Stochastic multi-person 3d motion forecasting. In *The Eleventh International Conference on Learning Representations*, 2022. 2
- [91] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 2022. 3, 5
- [92] Zhao Yang, Bing Su, and Ji-Rong Wen. Synthesizing long-term human motions with diffusion models via coherent sampling. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3954–3964, 2023. 2, 5
- [93] Zijie Ye, Haozhe Wu, Jia Jia, Yaohua Bu, Wei Chen, Fanbo Meng, and Yanfeng Wang. Choreonet: Towards music to dance synthesis with choreographic action unit. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 744–752, 2020. 2
- [94] Hongwei Yi, Chun-Hao P Huang, Shashank Tripathi, Lea Hering, Justus Thies, and Michael J Black. Mime: Human-aware 3d scene generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12965–12976, 2023. 2
- [95] Xinyu Yi, Yuxiao Zhou, and Feng Xu. Transpose: Real-time 3d human translation and pose estimation with six inertial sensors. *ACM Transactions on Graphics (TOG)*, 40(4):1–13, 2021. 5
- [96] Ye Yuan and Kris Kitani. Dlow: Diversifying latent flows for diverse human motion prediction. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16*, pages 346–364. Springer, 2020. 2
- [97] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16010–16021, 2023. 2
- [98] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Yong Zhang, Hongwei Zhao, Hongtao Lu, Xi Shen, and Ying Shan. Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14730–14740, 2023.
- [99] 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, 4
- [100] Qinsheng Zhang, Jiaming Song, Xun Huang, Yongxin Chen, and Ming-Yu Liu. Diffcollage: Parallel generation of large content with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10188–10198, 2023. 2, 3, 4, 5
- [101] Yan Zhang, Michael J Black, and Siyu Tang. Perpetual motion: Generating unbounded human motion. *arXiv preprint arXiv:2007.13886*, 2020. 2
- [102] Yan Zhang and Siyu Tang. The wanderings of odysseus in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20481–20491, 2022. 2
- [103] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019. 5
- [104] Yi Zhou, Zimo Li, Shuangjiu Xiao, Chong He, Zeng Huang, and Hao Li. Auto-conditioned recurrent networks for extended complex human motion synthesis. In *International Conference on Learning Representations*, 2018. 2, 3
- [105] Yi Zhou, Jingwan Lu, Connelly Barnes, Jimei Yang, Sitao Xiang, et al. Generative tweening: Long-term inbetweening

of 3d human motions. *arXiv preprint arXiv:2005.08891*, 2020. [2](#)

- [106] 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. [1](#)
- [107] Wenlin Zhuang, Congyi Wang, Jinxiang Chai, Yangang Wang, Ming Shao, and Siyu Xia. Music2dance: Dancenet for music-driven dance generation. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 18(2):1–21, 2022. [2](#)