

NEURAL MARIONETTE: A Transformer-based Multi-action Human Motion Synthesis System

Weiqiang Wang, Xuefei Zhe, Huan Chen, Di Kang, Tingguang Li, Ruizhi Chen, and Linchao Bao

Abstract—We present a neural network-based system for long-term, *multi-action* human motion synthesis. The system, dubbed as NEURAL MARIONETTE, can produce high-quality and meaningful motions with smooth transitions from simple user input, including a sequence of action tags with expected action duration, and optionally a hand-drawn moving trajectory if the user specifies. The core of our system is a novel Transformer based motion generation model, namely MARIONET, which can generate diverse motions given action tags. Different from existing motion generation models, MARIONET utilizes contextual information from the past motion clip and future action tag, dedicated to generating actions that can smoothly blend into historical and future actions. Specifically, MARIONET first encodes target action tag and contextual information into an action-level latent code. The code is unfolded into frame-level control signals via a time unrolling module, which could be then combined with other frame-level control signals like the target trajectory. Motion frames are then generated in an auto-regressive way. By sequentially applying MARIONET, the system NEURAL MARIONETTE can robustly generate long-term, multi-action motions with the help of two simple schemes, namely “Shadow Start” and “Action Revision”. Along with the novel system, we also present a new dataset dedicated to the multi-action motion synthesis task, which contains both action tags and their contextual information. Extensive experiments are conducted to study the action accuracy, naturalism, and transition smoothness of the motions generated by our system.

Index Terms—Motion Synthesis, Character Control, Transformer, Autoregressive Model.



1 INTRODUCTION

HUMAN motion capturing, modeling, and synthesis have been widely studied for decades, due to the dramatically increasing demand from the entertainment industry. Although new motion capture equipment and the latest motion processing algorithms have made motion collecting more accurate and efficient, it is still a challenging task to synthesize meaningful long-term human motions in a controllable yet simple way. The observation that a meaningful long-term human motion can be easily described by a sequence of semantic action tags, naturally raises the following question, can we use a sequence of action tags together with corresponding action duration as control signals to generate a realistic and smooth human motion clip.

Benefiting from the rapid development of deep generative models, recently proposed deep neural network-based methods have demonstrated promising achievements in motion synthesis. Animation synthesis for game control is well studied in [1], [2], [3], MANN [2], which has shown

superior quality in the locomotion synthesis task. However, most locomotion synthesis methods are tailored to the gaming industry with limited model size to reduce response time for improving players’ experience, which constrains their ability to generate diverse and complicated motions. Moreover, these locomotion synthesis methods generate motion in a frame-by-frame scheme without having an overall scope of the complete action, which is another prominent limitation for producing non-periodical action such as “kick” or “throw”. On the other hand, two recent works, action2motion [4] and ACTOR [5], can generate diverse motions from a set of semantic action tags. However, both of them are elaborately designed to be conditioned on action tag, thus hard to incorporate other contextual information such as initial poses for a more controllable generation. Other works [6], [7], [8], [9], [10] can generate continuous motion based on previous poses. However, these methods are only suitable for short-term motion prediction and can not be used to generate a preferred action from a specified category.

We can see that existing motion synthesis methods as mentioned above still struggle to generate meaningful long-term motions from a set of action tags. More specifically, there are three major challenges: (1) The first challenge is that there is a lack of suitable datasets for this task. Early datasets, such as [4], [11], [12] only provide action tag and motion data without any contextual information, which makes it hard to generate continuous actions. The latest motion dataset named BABEL [13] provides frame-level motion annotations, which can be considered as containing contextual information such as previous motion and future tag. However, due to the lack of a clear definition of each action, the annotated tags in BABEL are very

- W. Wang and R. Chen are with the State Key Laboratory of Information Engineering in Surveying, Mapping, and Remote Sensing, Wuhan University, Wuhan 430079, China.
E-mail: {weiqking, ruizhi.chen}@whu.edu.cn.
- H. Chen is with Shenzhen International Graduate School, Tsinghua University, Shenzhen 518055, China.
E-mail: huan-che20@mails.tsinghua.edu.cn.
- X. Zhe, D. Kang and L. Bao are with Tencent AI Lab, Shenzhen 518054, China. E-mail: {zhexuefei, di.kang}@outlook.com, linchaobao@gmail.com.
- T. Li is with Tencent Robotics X, Shenzhen 518054, China.
E-mail: tgli0809@gmail.com.
- W. Wang and X. Zhe contributed equally to the work; L. Bao is the corresponding author; W. Wang and H. Chen did the work during their internships at Tencent AI Lab.

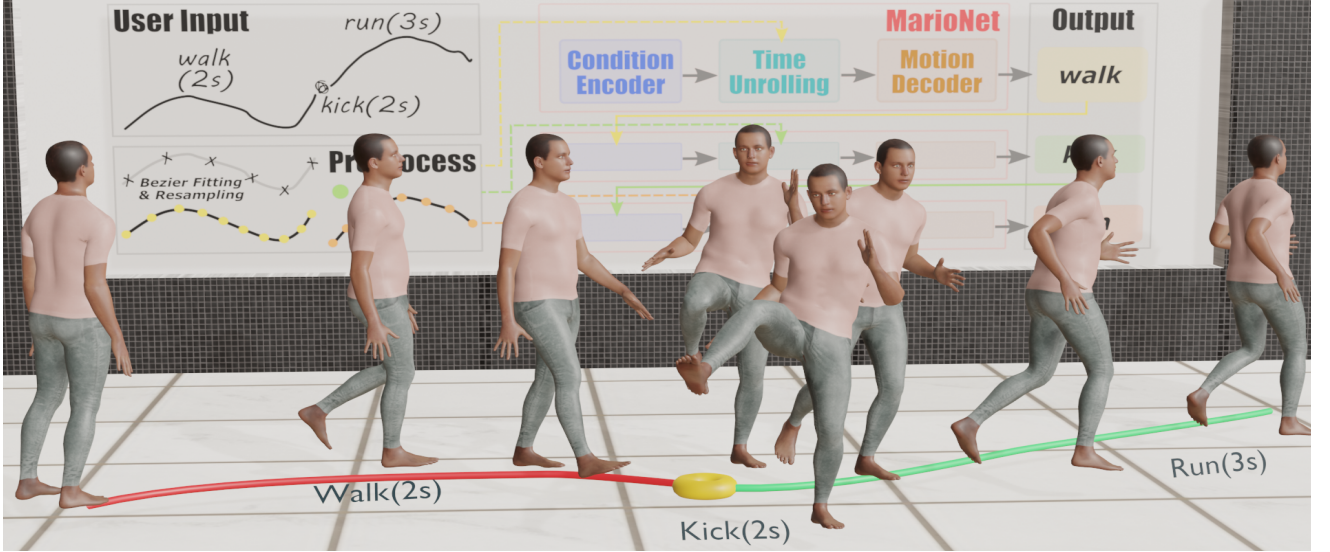


Fig. 1. Our system NEURAL MARIONETTE only needs a user to draw a trajectory line with the action tag for each segment and can produce a plausible motions with smooth transitions between actions. The core of the system is a novel Transformer based neural network named MARIONET.

noisy and ambiguous, making it hard to train a motion generative model. (2) The second challenge is how to formulate and encode contextual information into the long-term motion generation task. In practice, we have found it challenging to train a model which *directly* generates a long-term motion only with sequential action tags. One natural solution is to decompose this complicated action-tag sequence into a motion-sequence generative problem into single action generation blocks, where each action is generated with a smooth transition to the previous one. Then the expected long-term motion can be obtained by stitching these continuous actions together. In other words, a powerful generative model that takes not only semantic action tag but also contextual information is required. Unfortunately, there are few existing methods that can deal with the requirements. (3) The third challenge is that errors, or unnaturalness in synthesized motions can accumulate fast over time. Each separate action in multi-action generation needs additional contextual information to ensure a smooth transition between actions, compared to the original single-action generation which only considers target tag information. Nevertheless, the contextual information that comes from previously generated motion may not always be in the distribution seen in training. Thus, errors would gradually pass through the sequential process of multiple action generation if not proper correction strategies are adopted. It is quite similar to the common drifting problem in a temporal prediction model.

To address these three challenges, we propose our solution from three aspects. (1) First, we present a new dataset based on BABEL, named as *BABEL for Multi-Action Generation* (BABEL-MAG). We collect motion clips from the first 40 largest action categories from BABEL. Then we re-categorize them into 31 action groups following a well-defined and consistent labeling protocol. We also provide *contextual information* for each action, including initial motion and the tag of the next action, together with the current action tag. (2) We propose a single-action generative model based on Transformer [14] that takes both the current action tag and contextual information as input. Specifically,

our model first encodes the target action tag and contextual information into an action-level latent code Z . Then this action-level latent code is unfolded into time-varying frame-level codes with our time unrolling net. The frame-level codes provide better duration control and more stable single-action generation. Besides, other control signals such as trajectory can be combined together with frame-level codes during this process. Finally, The entire action can be obtained frame by frame with our auto-regressive motion decoder. With the help of contextual information, actions that are generated in sequential order can be stitched easily to produce a long-term motion with smooth transitions. (3) Lastly, to alleviate error accumulation, we propose a recognition-guaranteed hybrid pipeline to generate a continuous motion with multiple actions, consisting of a feature projection scheme named “*Shadow Start*” and a recognition-based scheme named “*Action Revision*”. As a result, we propose an effective and easy-to-use motion synthesis system that can generate accurate, natural human motions with smooth transitions from a sequence of action tags. Extensive experiments on both single- and multi-action generation are presented to demonstrate the superiority of our proposed method.

Contributions. Our contributions can be summarized as follows:

- We propose a novel single-action generation network, namely MARIONET, which utilizes both semantic tag and contextual information, to generate transition-smooth actions. Furthermore, MARIONET surpasses state-of-the-art single-action generative models in terms of action accuracy and naturalism.
- We present a neural network-based system NEURAL MARIONETTE which utilizes MARIONET iteratively for long-term, multi-action human motion synthesis. Two robust schemes, called “*Shadow Start*” and “*Action Revision*” are introduced to deal with the error accumulation problem, which significantly improves the capability of the system for generating long-term motion sequences.
- We introduce a new motion dataset BABEL-MAG for the multi-action motion generation task, which provides both

action tags and their contextual information.

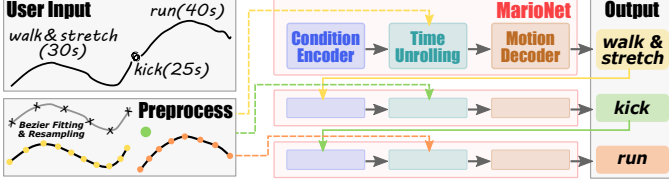


Fig. 2. **NEURAL MARIONETTE**. A user can specify a *sequence* of action tags, duration, and even motion trajectory by free-form hand drawings. Our method takes these inputs as control signals, then synthesizes a long-term human motion with smooth transitions by stitching sequential actions together.

2 RELATED WORK

2.1 Action-Conditioned Motion Synthesis

Action-conditioned human motion synthesis aims at generating realistic and diverse human motions corresponding to specified action tags. This is an under-explored field compared to other well-defined tasks including locomotion generation [1], [2], [8], [15], dance synthesis [16], [17], [18], [19], [20], [21], and gesture synthesis [22], [23], [24], [25]. Earlier works [4], [5] trained on exiting motion dataset [11], [12], [26] are able to generate different types of actions, including locomotion, exercises, and some daily actions. However, [4], [5] cannot be applied without modifications to generate a relatively complicated motion consisting of a sequence of different actions. Blending multiple independently generated actions cannot guarantee natural transitions and smooth motion trajectories. In order to generate consecutive actions with natural transitions, besides the action tag, our network is additionally conditioned on contextual information for motion synthesis. An earlier motion prediction paper [27] has a similar input setting to ours, which takes a 2-second previous motion and optionally the action tag as input to predict either a short-term (0.4s) or a long-term (1s) future motion. However, our network only requires one-tenth of previous motion (0.2s) compared to [27] while generating up to 3-second future motion, which is significantly more challenging.

Besides, another difference between our network and common motion prediction methods is that frame-level control signals like trajectory constraints can be effectively incorporated through the proposed time unrolling module, leading to motions that better satisfy users' expectations.

2.2 Control Signals in Motion Synthesis

There are several well-known motion synthesis subfields categorized according to the motion types, including locomotion synthesis, dance motion synthesis, co-speech gesture synthesis, etc. Here we focus on the *control signals* used in various motion synthesis tasks.

Some literature take *dense* control signal for every *frame* as input. For example, the locomotion synthesis task usually specifies dense signals for motion trajectory including turning, forward, and sideways velocity for every frame [8], [28]. Later, other control signals including foot contact [15], phase [1], or pace [29] are utilized to resolve motion ambiguity during the prediction, achieving better results. Dense control signals could help the network avoid

generating static poses or floating characters [8], [15], [30] in locomotion, which are major challenges for all types of long-term motion synthesis tasks. Similarly, in the dance motion synthesis task, dense acoustic features extracted from the input music such as beat, onset strength, chroma, MFCC and etc. are used to guide the motion generation [16], [17], [31], [32]. Most previous works [16], [17], [31], [33] only take music (features) as input to generate dance motions, formulating it as a cross-modal translation task. A recent paper [19] takes an additional 2-second seed motion along with the music features as input to generate the following dancing poses, which can be equivalently regarded as a motion prediction task conditioned on input music. [19] can generate better dance motions over the state-of-the-art methods that only take music as input since the presence of the seed motion significantly reduces the possible motion space for the following motions. We can observe that introducing proper dense and informative control signals can potentially improve the results of motion synthesis. Inspired by this observation, we utilize previous poses as additional control signals so that the generated actions can be concatenated to accomplish a complex task including a series of actions.

Some other works take *sparse* control signal that stays the same for each frame in one action. For example, [4], [5] take as input the target action tag and generate diverse motions as expected. This setting with limited input information is severely challenging since adapting dense-signal conditioned methods aforementioned to handle the sparse control signal will quickly result in generating static poses as demonstrated in [4]. In order to alleviate this difficulty and generate diverse motions belonging to the same category, [4], [5] adopt a probabilistic VAE framework that can take a random sample as input during inference. However, these methods can only synthesize a single action at one time and cannot generate long and consecutive motions like locomotion, dance, and co-speech motion synthesis methods [1], [17], [23], [25].

2.3 Deep Learning in Motion Synthesis

Traditional motion synthesis methods usually rely on a search-and-blend framework to *concatenate* short motion clips from an existing motion database. Representative methods include motion graph-based methods [34], which formulate motion synthesis as a path searching problem in a graph constructed based on motion similarity for all the motion clips in the database. Various constraints can be incorporated into the path searching process, making this framework very flexible. For example, a recent graph-based method [35] has incorporated deeply learned high-level style and rhythm features of choreography-music to search for the most suitable candidates, achieving convincing dance choreography.

However, relying on neural networks to directly synthesize motion instead of searching-and-blending short motions from an existing database is becoming more and more popular since network-based methods can encode database information into their weights instead of storing a huge motion database [3] and it is possible to generate novel and realistic motions that have not been seen in the training

set [5]. For example, early deep learning methods mostly adopt the deterministic framework for motion synthesis. [36] proposed a representative Encoder-Recurrent-Decoder (ERD) framework to jointly learn spatial representations (with the encoder and decoder) and their dynamics (with a recurrent network) to predict the next frame pose. [27] predicts both short-term and long-term future motion by GRU layers [37] with a novel sampling-based loss. [15] uses the MLPs for mapping the control signal to a hidden motion space which is later modeled by the CNN autoencoder to generate corresponding locomotion.

Recently, probabilistic framework gains more and more attention since it can better avoid generating frozen motion and generating diverse results by sampling [4], [5], [8], [17], [25], [30]. [17] proposes a dance unit VAE to encode and decode the dance units and a music-to-movement GAN to generate dance movement conditioned on the input music. [25] utilizes a TCN-VAE framework and split the latent space into two parts to explicitly model the one-to-many mapping between speech and gesture, producing more diverse and realistic results. [30] proposes to model the conditional probability distribution of the next frame pose using normalizing flows that takes previous poses and control signals as input. [4] proposes a GRU-VAE and [5] proposes Transformer-VAE to generate realistic and diverse motions belonging to the given action category. We also adopt a Transformer-based probabilistic framework but introduce several key designs in order to synthesize complex motions containing multiple actions.

3 BABEL-MAG DATASET

To synthesize high-quality and continual human actions, we need a large-scale tag-labeled motion dataset. A good dataset for continuous action synthesis should satisfy the following criteria: (1) The motion sequences in the dataset should contain multiple actions that can provide contextual information for continual action generation. (2) The action tags should be *precise and unambiguous*. For example, the same tag should not be used to describe two actions with different movement patterns.

Existing datasets such as [4], [12], [26] can only support action-level synthesis (i.e. synthesize a motion with respect to a single tag one time). A recent dataset, BABEL [13], provides motion annotations to the large-scale MoCap dataset AMASS [38]. Its motion annotation includes frame-level tags in continual motion sequences, which can be potentially used for multi-action synthesis. However, the action tags in BABEL do not fully satisfy the criterion (2), which poses difficulties in the training motion generative models. As we observe (also mentioned in the original paper [13]), BABEL has some unneglectable issues: (1) The action tags are *unbalanced* and have a long-tailed distribution, which makes it difficult for network training. (2) Since the action tags are abstracted from sentences proposed by different annotators, some action tags are *ambiguous*, which can hardly be used for motion synthesis. For example, the tag "turn" is used for describing a body turn-around, walking in circles, and even head-turning; the tag "step" is used for both strolling and climbing stairs. (3) The action tags can be composed but no clear *composition rules* are discussed or presented, which

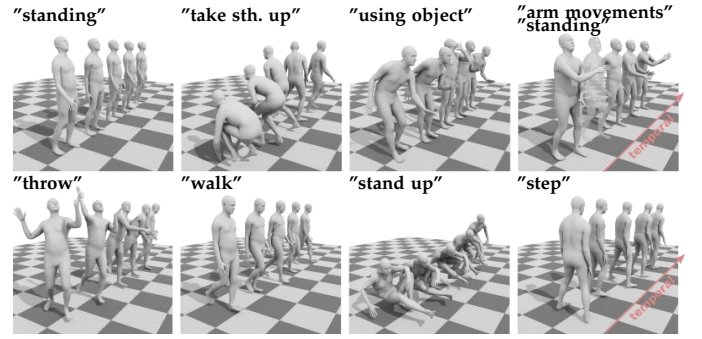


Fig. 3. **Labeling error in BABEL.** The descriptions from BABEL can be ambiguous. For example, the upper actions miss the most prominent movement and are labeled as (1) "arm movements&raising body part", (2) "turn&head movements", (3) "hand movements", and (4) "stand" in BABEL respectively while lower actions are all labeled as "stand" despite of significantly different movements. For comparison, our tagging rules for BABEL-MAG can easily separate different types of actions and the corrected tags are annotated above the images.

makes it hard for synthesizing diverse and complicated actions in real life. See Fig. 3 for some qualitative illustrations.

To avoid these issues above, we present a new dataset, BABEL-MAG, by relabeling actions from the first 40 *largest* and as *meaningful* as possible categories in BABEL, containing over 3 hour motion data for "Multi-Action Generation" task. First of all, we abstract **31** action tags from the 40 BABEL motion category names, including **24** tags for describing action intentions, **3** for body states, **4** for body part movements. For intention-describing tags, we further divided them into *root-motion* tags and *in-place* tags, depending on whether the common trajectories of these motions change significantly. As a result, we successfully resolve the tagging ambiguity in the chosen **7643** actions from BABEL. See Fig. 3 for some examples with corrected tags in BABEL-MAG. Besides action tags, we also provide initial motion and future action tags describing the original following motion as action contextual information. That is, for every item, it includes: (1) k frames of previous motion ($\{P\}_{-k}^{-1}$); (2) current action tag (A^{cur}); (3) current action duration (T); (4) next action tag (A^{next}). And each motion P is represented as joint rotation with root translation per frame as [38], [39]. See Sec. 2 in Supplementary for more details about BABEL-MAG labeling.

4 METHODOLOGY

4.1 Problem formulation

Given a sequence of n action tag annotations $\{A_i\}_{i=1\dots n}$ with corresponding expected action duration $\{T_i\}_{i=1\dots n}$ and optionally user defined trajectories $\{U_i\}_{i=1\dots n}$ (we set U_i to be zero-value if not specified), our goal is to generate a *sequence of actions* $\{M_i\}_{i=1\dots n}$ where the action $M_i = \{P_1, \dots, P_{T_i}\}$ can be accurately described by the input label A_i , and follow the input trajectory U_i ($U_i \in \mathbf{R}^2$ if specified). More importantly, the motion sequence is expected to be realistic with smooth transitions between any two actions M_i and M_{i+1} .

However, it is challenging to directly synthesize long-term human motion with deep generative models. Instead, we first propose a single action generative model, named as MARIONET, then demonstrate a hybrid system named

as NEURAL MARIONETTE synthesize multi-action motion by sequentially using MARIONET.

4.2 MARIONET: Single-action Generation Model

In this part, we discuss the design details of MARIONET, which is the main building block of our motion sequence generating pipeline (illustrated in Fig. 4). MARIONET can synthesize a single action that accurately corresponds to the input action tag and smoothly transits from initial motion while taking the future action tag into consideration.

4.2.1 Network Structure

MARIONET is an auto-regressive attention-based model *conditioned* on (1) initial motion P^{ini} , (2) current action tag A^{curr} , and (3) the following action tag A^{next} , which contains three main building blocks with Transformer layers, namely the Condition Encoder, Time Unrolling Module, and Motion Decoder.

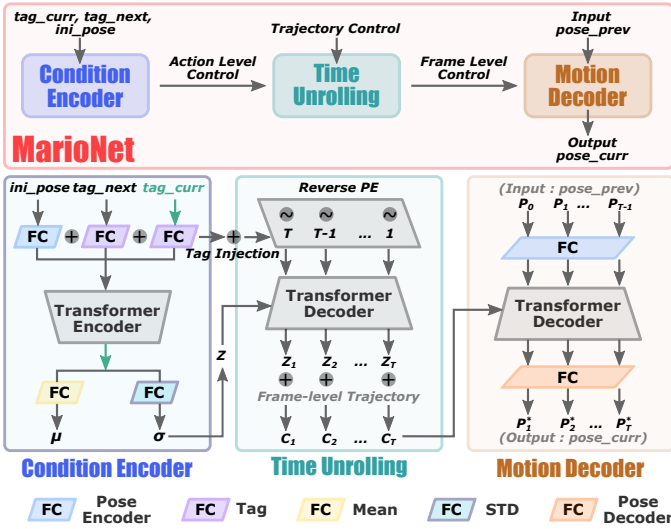


Fig. 4. **Network Overview of MARIONET.** The key component of our method is the MARIONET trainable block that generates new poses from given poses. MARIONET contains three major building blocks: Condition Encoder block, Time Unrolling block, and Motion Decoder block. The fully connected (FC) layers with the same color share the same weights.

Condition Encoder. The Condition Encoder projects inputs (initial poses, current and next action tags) to an *action-level* latent distribution parameterized by mean μ and variance σ . Specifically, we first embed pose and tag information to feature space with different FC layers. Then embedded P^{ini} , A^{curr} , A^{next} are concatenated together and fed into a Transformer encoder. We are particularly interested in the embedded A^{curr} (dubbed as *action-token*). Similar to ACTOR [5], only the action token is used for pooling. The output of the Transformer encoder that corresponds to the action token is used to regress the distribution parameters (μ, σ) via two independent FC layers. Sampling from the obtained distribution leads to an *action-level* latent code Z .

Time Unrolling Module. This module learns a *frame-level* latent space by unrolling the action-level latent code and potentially combining user frame-level input information (e.g. trajectory), which we call frame-level controls as a whole. Specifically, the *action-level* latent code Z from previous step is firstly unrolled into *frame-level* latent codes $[z_1, \dots, z_T]$

via a classic Transformer Decoder. The Transformer Decoder takes the action-level latent code Z as Key K and Value V , and takes the *reverse* positional encoding (Rev-PE) with *Tag Injection* operation (described in the next paragraph) as Query Q for cross-attention. As a result, we get the corresponding frame-level latent codes $[z_1, \dots, z_T]$ in \mathbf{R}^d . At this stage, we can also append other frame-level information such as motion trajectories to better guide or control the motion synthesis. To achieve this, the latent codes $[z_1, \dots, z_T]$ is first projected into $\mathbf{R}^{\tilde{d}}$ via a FC layer, while the user input is projected into $\mathbf{R}^{d-\tilde{d}}$ via another FC layer. Combining them together leads to the finalized frame-level controls $[c_1, \dots, c_T]$ in \mathbf{R}^d . In this way, the frame-level trajectory can be concatenated with frame-level action latent space, resulting in more powerful frame-level controls. Note that, unlike ACTOR in which standard PE is used, we adopt Rev-PE to encode positions for time unrolling inspired by the time-to-arrival (TTA) embeddings in [40], which is beneficial for motion generation with a pre-defined number of frames since the action termination information is encoded in the Rev-PE.

Tag Injection. Due to different signals being mixed up in the condition encoder for an action-level representation, the control ability of the current action tag A^{curr} which we expect to be dominant is inevitably weakened. Therefore, we propose a tag injection operation to enhance its control ability, leading to generated action more exactly corresponding to the specified tag. Specifically, we add the *embedded* A^{curr} to every frame of Rev-PE, then use them as the Query Q in Time Unroll Module for cross-attention.

Motion Decoder. The Motion Decoder takes a sequence of poses $P = [P_0, \dots, P_{T-1}]$, and the frame-level controls $C = [c_1, \dots, c_T]$ as input, and outputs the synthesized action $P^* = [P_1^*, \dots, P_T^*]$, i.e., a set of new poses shifted 1 frame right temporally from the input poses. Specifically, the input poses P_t are firstly projected into the same pose latent space as the one discussed in the Condition Encoder via a FC layer. Then the Transformer Decoder takes the embedded input poses with positional encodings as Query Q , and takes the frame-level controls c_t as Key K and Value V for cross-attention. The outputs of the Transformer Decoder are then recovered into pose space by linear projections with FC. During the training phase, the Transformer Decoder with attention-mask can process the ground-truth input poses in parallel. During the test, the complete action with T frames is generated iteratively. Specifically, at each iteration, one new frame is firstly generated and then feed to the Transformer Decoder again to synthesize a new pose in the following frame.

4.2.2 Training

In this part, we discuss in detail how we train MARIONET on BABEL-MAG and how we construct the training loss for the motion synthesis task.

Scheduled Sampling for Transformer. There is a severe train-test discrepancy called “Exposure Bias” in the teacher-forcing training scheme for sequence-related generation [31], [41], [42]. The Scheduled Sampling strategy is proposed to alleviate this dilemma in sequence generation [43]. It is relatively easy to adapt Scheduled Sampling

to motion generation for models with RNN architecture [31], since they generate a motion sequence frame by frame during the training phase. However, Transformer processes the input and outputs the whole sequence in parallel by the attention-mask mechanism during training, making it hard to take a previously predicted pose as input in a frame-by-frame fashion [44], [45].

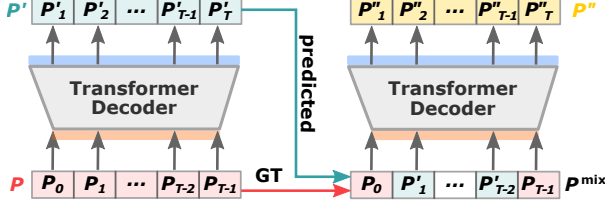


Fig. 5. **Scheduled Sampling Strategy.** *Left:* the Motion Decoder takes the ground-truth poses $\{P_t\}_{t=0 \dots T-1}$ as input and output the synthesized poses $\{P'_t\}_{t=1 \dots T}$. *Right:* $\{P_t\}$ and $\{P'_t\}$ are mixed and feed into the Motion Decoder again to synthesize new poses $\{P''_t\}_{t=1 \dots T}$. Both P'_t and P''_t are used for loss construction. For simplicity, we use an orange/blue box to represent the pose encoding/decoding FC layer.

In this work, we successfully adapt the Scheduled Sampling to motion generation with Transformer. Specifically, as illustrated in Fig. 5, two weights-sharing Motion Decoders with Transformer Decoder are involved. The first Decoder (left) takes the ground truth motion sequence $P = [P_0, P_1, \dots, P_{T-1}]$ as input, and outputs the predicted motion sequence $P' = [P'_1, P'_2, \dots, P'_T]$ that is one frame later than the input. Then the predicted motion P' is mixed with the ground truth $P = [P_0, P_1, \dots, P_{T-1}]$ which we call the mixed motion P^{mix} . Note that, the ratio (ξ) of poses from the predicted motion can vary in P^{mix} during the training. The mixed motion P^{mix} is then fed into the second Decoder (right), which outputs a new predicted motion sequence $P'' = [P''_1, P''_2, \dots, P''_T]$. The reconstruction loss takes both P' and P'' into consideration when this strategy is adopted, which helps to bridge the gap between training and inference. Note that, to avoid error accumulation of these predicted high-biased poses, our Scheduled Sampling strategy is activated only when the predicted motion P' is close enough to the input ground-truth P . During the training, the ratio ξ is increased linearly from 0% (i.e., $P^{\text{mix}} = P$) to 100% (i.e., $P^{\text{mix}} = P'$) between the *Activate* and *Accomplish* epoch. This strategy alleviates the train-test discrepancy, making our generated motion sequence both fluent and continuous at the testing phase.

Training Loss. We use the following loss to measure the reconstruction quality:

$$L_{\text{Recon}}(\hat{P}) = \lambda_R L_R + \lambda_D L_D + \lambda_J L_J, \quad (1a)$$

$$\text{where } L_R = \frac{1}{T} \sum_{t=1}^T \|R_t - \hat{R}_t\|_2^2 + \lambda_V \frac{1}{T-1} \sum_{t=2}^T \|R'_t - \hat{R}'_t\|_2^2, \quad (1b)$$

$$L_D = \frac{1}{T} \sum_{t=1}^T \|D_t - \hat{D}_t\|_2^2 + \lambda_V \frac{1}{T-1} \sum_{t=2}^T \|D'_t - \hat{D}'_t\|_2^2, \quad (1c)$$

$$L_J = \frac{1}{T} \sum_{t=1}^T \|J_t - \hat{J}_t\|_2^2 + \lambda_V \frac{1}{T-1} \sum_{t=2}^T \|J'_t - \hat{J}'_t\|_2^2 \quad (1d)$$

where the rotation \hat{R}_t , root translation \hat{D}_t , and joint position \hat{J}_t are associated with the synthesized pose \hat{P}_t ; $\hat{R}'_t = \hat{R}_t - \hat{R}_{t-1}$ give the rotation velocity of that pose (similar for the root translation and joint position). Specifically, for a given predicted pose \hat{P}_t , which is represented by the rotations \hat{R}_t and a root translation \hat{D}_t , we can then fit a SMPL model [46]

with mean shape parameters to get the 3D positions for the joints \hat{J}_t of the pose \hat{P}_t . After this, we can then measure the reconstruction error of the motion sequence \hat{P} by computing the per-frame differences and velocity difference of rotation, root translation, and joint position between the predicted motion and the ground-truth motion. Therefore, we have the *Reconstruction Loss on Rotation* L_R defined in (Eq. (1b)); the *Reconstruction Loss on Translation* L_D defined in (Eq. (1c)); and the *Reconstruction Loss on Local Joint Positions* L_J defined in (Eq. (1d)). We also adopt the Kullback-Leibler (KL) divergence to regularize the distribution of the action-level embedding in our Condition Encoder. In summary, we have the following loss function for training:

$$L = \lambda_1 L_{\text{Recon}}(P') + \lambda_2 L_{\text{Recon}}(P'') + \lambda_{\text{KL}} L_{\text{KL}} \quad (2)$$

Note that, $\lambda_R, \lambda_D, \lambda_J, \lambda_V, \lambda_1, \lambda_2$ are weights used for balancing the realism and diversity by manual setting.

4.2.3 Inference

During the inference phase of generating an action with tag A , the current tag A^{curr} (set as A), the initial motion P^{ini} that corresponding to tag A , and the following tag A^{next} are firstly fed into the Condition Encoder. A random sample from the action-level latent space is then combined with T frames reverse PE (corresponding to specified action duration), and trajectory U constraints (set to zero-value for in-place actions) to generate the frame-level controls. Finally, the Motion Decoder runs in a frame-by-frame fashion to auto-regressively generate the motion sequence that corresponds to the expected action tag A from the frame-level controls. For the first iteration, we use the last frame of P^{ini} as the Begin of Poses (BOP). We terminate this process when we obtain T frames of predicted poses as requested.

4.3 Generating Multi-action Motion

Previous parts introduce how to generate a single action conditioned on both tag and contextual information by MARIONET. In this section, we discuss how to synthesize a long-term multi-action motion by sequentially applying the MARIONET. We first introduce a pure generative pipeline, which only relies on the generative model such as MARIONET. Then based on this pipeline, we derived NEURAL MARIONETTE with two simple but effective correct schemes for a more robust multi-action generation.

4.3.1 Pure Generative Pipeline

A simple way to synthesize multi-action motion is to generate multiple single-action motions with an action generative model, then connect these single-action clips temporally. Specifically, we first normalize initial motion (x frames) $P^{\text{ini}} = \{P_t\}_{t=-x}^{t=-1}$ to facing $-y$ direction (zero-rotation of T-pose) then moved to origin point both based on its ending pose P_{-1} , where pose P_{-1} , which is the last frame pose of previously generated action or a given initial motion, has a global rotation $R_z = (0, 0, \gamma)$ on z-axis, ($R = (\alpha, \beta, \gamma)$ in x-y-z form of Euler angle) and a global translation $D_t(x, y, z)$. For each frame of normalized initial motion:

$$\{P_t^*\}_{t=-x}^{t=-1} = R_z^{-1} \{P_t\}_{t=-x}^{t=-1} - D_t \quad (3)$$

After obtaining the normalized initial motion, an action is generated by applying the MARIONET as described in Sec. 4.2.3:

$$\{\tilde{P}_t\}_{t=1}^{t=T} = \text{MARIONET}(\{P_t^*\}_{t=-x}^{t=-1}, A^{\text{curr}}, A^{\text{next}}, T, U) \quad (4)$$

For stitching with the previous action, the generated action needs to be rotated back by

$$\{\hat{P}_t\}_{t=1}^{t=T} = R_z \{\tilde{P}_t\}_{t=1}^{t=T} + D_t \quad (5)$$

Then the following actions can be generated by sequentially applying these steps.

4.3.2 NEURAL MARIONETTE: a Hybrid Pipeline

During testing, we notice that the pure generative pipeline may fail sometime, especially when generating long multi-action motions, i.e., motions including more than 10 actions. Through further digging into this problem, we find two main reasons. The first reason is *insufficient* connected pairs of actions for training. By counting connection actions tags, we find that, except a few tag pairs (e.g. turn-walk), most of the tag pairs have only limited connected samples in the training data. (See more detail in Supp. Fig 1.) A single-action network trained on data with such a sparse connection matrix tends to face a severe out-of-distribution problem when generating continuous actions.

A straightforward solution is to obtain more connected motion capture data. However, it is not feasible because the desirable data grows with the tag number by the magnitude of the square. The second reason is the error accumulation issue, which is a major challenge faced by many temporal generative models when generating long sequences. To tackle the above challenges, we propose a hybrid system named “NEURAL MARIONETTE”, which includes two important parts: “Shadow Start” and “Action Revision”.

Projection based “Shadow Start” Inspired by solutions for long-term tracking [47], [48] and the manifold projection in DeepFaceDrawing [49], we proposed a pose embedding projection method to improve the stability of long-term generation, named “Shadow Start”. The main idea of Shadow Start is to project the feature embedding of Begin of Pose (BOP) P_0 , as mentioned in Sec. 4.2.3, to a hyperplane supported by some embeddings derived from BOPs in the training set. More specifically, the BOP embedding feature, f_0 is firstly extracted by $f_0 = FC(P_0)$, where FC is the pose embedding layer in Fig. 4. Then, N most similar embeddings extracted from BOPs $\{\tilde{P}^i\}_N$ belonging to the current action tag A^{curr} in the training dataset are retrieved according to their $L2$ distance to f_0 in the embedding space. Here, pose embeddings of $\{\tilde{P}^i\}_N$ are denoted as $\tilde{F} = [\tilde{f}^i]_N$. Following this, the pose embedding f_0^* of desired projected P_0^* can be found by solving the linear combination coefficients $\Phi = [\varphi_i]_N$ as follows:

$$\Phi^* = \underset{\Phi}{\text{argmin}} \|f_0 - \Phi \cdot \tilde{F}\|_2^2, \quad f_0^* = \Phi^* \cdot \tilde{F} \quad (6)$$

After obtaining f_0^* , the new action can be generated as described in Sec. 4.2.3 only by replacing origin BOP embedding f_0 with f_0^* .

Since the actual BOP P_0^* corresponding to embedding f_0^* is unknown, this method is named “Shadow Start”. By doing so, it is expected that the modified feature embedding

of BOP f_0^* is less out-of-distribution than the original one. Meanwhile, the resulting start pose of generated action would be closer to the end pose of the previous action compared to naively replacing the original embedding with the most similar one.

Recognition based “Action Revision” Although Shadow Start can improve the robustness (in terms of action recognition accuracy) of long-term motion generation, it also causes jittering effects due to the modification of the BOP feature. Since our method can directly synthesize correct actions in most cases. therefore, a more practical way is to only use Shadow Start when MARIONET fails to generate the correct actions. Specifically, we first generate one action by MARIONET, then an action classifier denoted as F_c is used to classify whether the motion is the desired action. If the result is not, we regenerate this action only once with the aforementioned Shadow Start operation. We repeat this process during long-term motion generation and call it an “Action Revision” scheme. With above two parts, we summarize the whole NEURAL MARIONETTE in Algorithm 1.

Algorithm 1: NEURAL MARIONETTE

Input : $\{A_i, T_i, U_i\}_N, \{P_t\}_{t=-x}^{t=-1}$
for $i \leftarrow 1$ **to** N **do**
 R_z and $D_t \leftarrow P_{-1}$;
 Normalized initial pose
 $\{P_t^*\}_{t=-x}^{t=-1} \leftarrow R_z^{-1} \{P_t\}_{t=-x}^{t=-1} - D_t$;
 Generate i -th action $\{\tilde{P}_t\}_{t=1}^{t=T_i} \leftarrow$
 MARIONET($\{P_t^*\}_{t=-x}^{t=-1}, A_i, A_{i+1}, T_i, U_i$);
 if $F_c(\{\tilde{P}_t\}_{t=1}^{t=T_i})$ **is not correct** **then**
 $f_0^* \leftarrow$ “Shadow Start” ;
 Re-generate i -th action with f_0^* ;
 end
 Rotated i -th action $\{\hat{P}_t\}_{t=1}^{t=T_i} \leftarrow R_z \tilde{P}_t + D_t$;
 Update initial pose sequence
 $\{P_t\}_{t=-x}^{t=-1} \leftarrow \{\hat{P}_t\}_{t=T_{i-1}-x+1}^{t=T_i}$;
end

5 EXPERIMENTS

We conduct experiments for two main tasks: single-action generation and multi-action generation. Although the main purpose of NEURAL MARIONETTE is to synthesize multi-action human motion, the main building block of our system is our MARIONET, which is a single-action motion generative model. So we first compare MARIONET with other baselines under the single-action generation setting. Besides, we further analyze component effectiveness and controllability of MARIONET. For the multi-action task, we first compare different generative models under a purely generative pipeline. Both objective and subjective measurements are provided for performance evaluation. Then we show how our NEURAL MARIONETTE can further improve the performance of multi-action synthesis.

5.1 Single-Action Generation

5.2 Experiment setup

Train/test split. For every action, we randomly select 20% to 100 samples as testing set and use rest data for training.

Compared methods. We consider three state-of-the-art methods, which are carefully adapted and trained on BABEL-MAG, including: **MANN** [2], **Action2motion (A2M)** [4] and **ACTOR** [5]

5.2.1 Evaluation Metric

We first consider the commonly used *action-level* metrics to measure the recognition accuracy, realism, and diversity of the synthesized single-action motions [4], [5], [17]. Specifically, a Transformer-based action recognition network \mathcal{N} pre-trained on BABEL-MAG is used to extract action features. Besides, another pose-level metric is also included to provide a more comprehensive evaluation.

- **Recognition Accuracy (Acc).** We use the pre-trained network \mathcal{N} [4] to classify the synthesized actions and report the recognition accuracy for each method.
- **Frechet Inception Distance (FID).** FID [50] measures the similarity between the distributions of action features (extracted via \mathcal{N}) from real-world samples and those from synthetic samples. A *lower* FID suggests a better result when comparing different methods.
- **Diversity (Div).** We measure the *variance* of the features f (extracted via \mathcal{N}) of the generated actions across all action categories. We evenly split the generated actions into two subsets $\{f_i\}$ and $\{\tilde{f}_i\}$, and approximate the variance by $\frac{1}{K} \sum_{i=1}^K \|f_i - \tilde{f}_i\|_2$ as in [4]. A *higher* value means the test motions are more diverse.
- **Multimodality (Multi-mod).** Multimodality measures the diversity *within each action category*. Specifically, we first compute *variance* of the motion features (via \mathcal{N}) of the generated actions sharing the same action tag, then average the variances over different action categories. A *higher* value suggests a larger inter-class diversity.
- **Joint Position Error (Δ_{JPE}).** We measure the joint position error by calculating the L2 distance between the real motion and the generated one in per-frame and per-joint levels. A *lower* value means the generated motions are closer to the real ones.

our method can generate motions accurately corresponding to the given action tag. As for diversity measurements, three methods considering contextual information achieve similar scores while our method obtains the highest diversity score with **9.9** which is even more diverse than the real data marginally. MANN outperforms the other two methods on multimodality metric with the value **2.08**, while our method obtains a very closing value of **2.07**. In addition to these action-level metrics, MARIONET also shows the best performance in terms of Joint Position Error among all methods.

We could easily find from Table 1 that ACTOR, as the state-of-the-art action-conditioned generative model, is still hard to generate tag-desired and high-quality motions without using any contextual information in our dataset. Note that, Diversity and Multimodality are meaningful only when high-quality and tag-preferred motions are provided. Therefore, the highest Multimodality score achieved by ACTOR is not necessarily the inter-class diversity that we need. In summary, our method outperforms other baselines on overall quality and action accuracy with a substantial margin and has a similar performance to other baselines on diversity and multimodality.

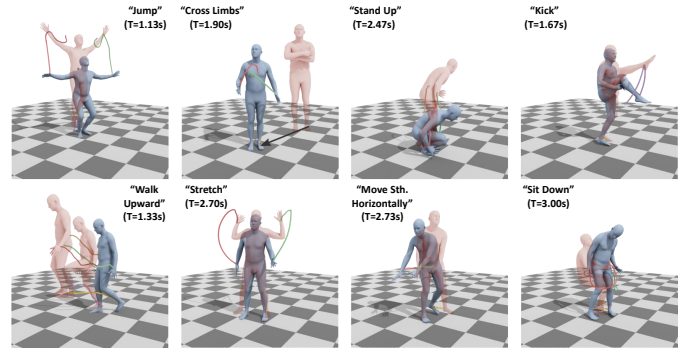


Fig. 6. We show eight synthesized actions with different tags using our MARIONET, including the first pose (blue mesh), pose in keyframes (red meshes), and joint trajectories (lines in different colors).

TABLE 1

Comparison on single-action synthesis. The unit of Δ_{JPE} is cm.

Metrics	FID ↓	Acc ↑	Div ↑	Multi-m ↑	Δ_{JPE} ↓
Real Test Data	/	0.894	9.865	1.956	/
MANN [2]	3.406	0.838	9.789	2.079	5.09
A2M [4]	3.152	0.826	9.738	2.048	4.84
Ours	1.115	0.874	9.904	2.068	4.58
ACTOR [5]	9.558	0.744	9.424	2.356	8.39

5.2.2 Comparison

In this task, a method is required to generate actions according to the provided action tag. MANN, A2M, and ours use both action tag and contextual information as input while ACTOR takes only the action tag. Results of all metrics are presented in Table 1. Our method obtains a significantly lower (better) FID score (**1.1**) among all methods, indicating better motion quality. A visual comparison result is presented in Fig. 6. For the recognition evaluation, our method achieves the accuracy of **0.87**, which is higher than the others and approaches the level of real data. It shows that

5.2.3 Ablation Study

We ablate several important components of our method to verify their effectiveness, including Schedule Sampling strategy, Time Unrolling (*TU*), and Tag Injection (*TI*). All the ablative experiments (see Table 2) are conducted under the single-action generation setting. First of all, MARIONET trained without the scheduled sampling suffers from larger Joint Position Error inevitably while our proposed training strategy greatly narrows the train-test discrepancy.

Besides, we find that both Time Unrolling and Tag Injection contribute to improving the quality of generated motions in terms of all metrics except multimodality. At the same time, the former has a larger impact than the latter on FID, accuracy, and Joint Position Error, which shows that unrolling latent code Z to time-varying sequential features plays a key role in our generative model and can significantly improve naturalness and accuracy of the generated actions.

5.2.4 Qualitative results of model controllability

As a conditional generative model, MARIONET can produce diverse actions corresponding to different inputs. In this

TABLE 2
Ablation study of key designs of MARIONET.

Metrics	FID ↓	Acc ↑	Div ↑	Multi-mod ↑	Δ_{JPE} ↓
Ours (MARIONET)	1.115	0.874	9.904	2.068	4.58
<i>w/o Schedule Sampling</i>	<u>1.274</u>	<u>0.871</u>	9.846	1.985	<u>6.99</u>
<i>w/o TU and w/o TI</i>	11.294	0.638	9.706	<u>2.404</u>	9.32
<i>w/o TU</i>	11.522	0.676	<u>9.867</u>	2.442	9.42
<i>w/o TI</i>	5.040	0.819	9.458	2.284	7.84

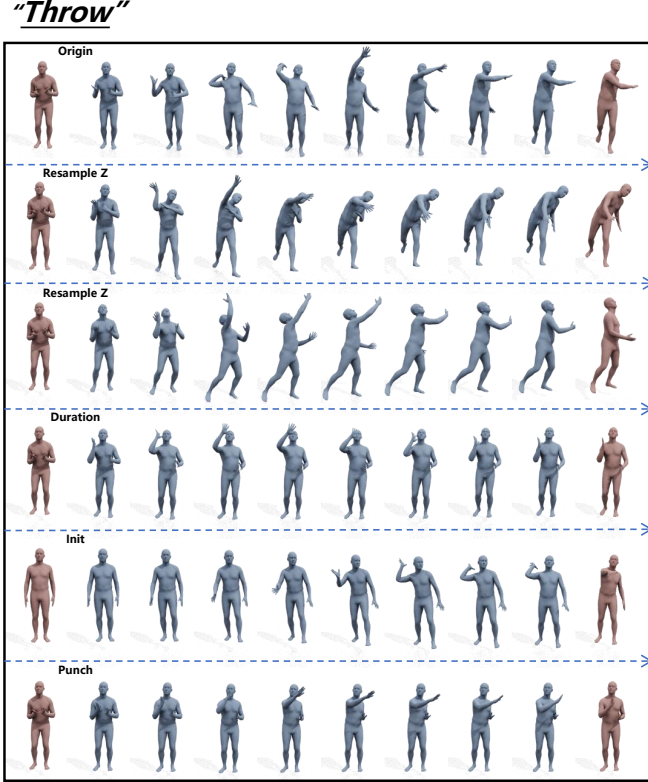


Fig. 7. Effectiveness of different control signals. The input of the 1st row is from the real action of “throw”. Compared to the 1st row: (2),(3) sampling condition latent code Z ; (4) changes duration (5) changes initial pose (6) changes action tag to “Punch”. It can be clearly seen that actions are different from each other.

part, we present two visual examples in Fig. 7 to demonstrate the controllability of MARIONET. The first row of both examples is base generated actions, of which all input conditions are from original real data including action tag (“throw” and “run” respectively), duration, initial motion, and trajectory. The latent code Z is directly set to the predicted μ . The second and the third rows show actions generated with different re-sampling latent code Z . It can be seen that these three actions (1st row to 3rd row) are significantly different, which demonstrates that, even with the same input, our model still can generate diverse motions belonging to the same action category. For each row from the fourth to the sixth row, we show generated actions with only one different input control signal compared to the first row. It can be easily found that all these actions are different as expected. Especially, for the last row, we change action tags (change “throw” to “punch” for Fig 7). Even with the same initial motion, MARIONET can generate different action types according to the input action tags.

5.3 Multi-Action Generation

In this part, we are focusing on evaluating multi-action generation performance. For conducting more comprehensive experiments, we firstly construct two multi-action testing sets from the BABEL dataset by considering different connection frequencies between actions. Based on testing sets, we compare our MARIONET with two adapted state-of-the-art action generation methods under the pure generative pipeline. Both subjective and objective evaluations are provided to discuss the performance of multi-action synthesis. Lastly, we demonstrate how the hybrid system, NEURAL MARIONETTE can further improve performance compared to the pure generative pipeline.

5.3.1 Construction of Multi-action Evaluation Sets

For this task, we need to prepare a sequence of action tags and initial motion as inputs. Directly selecting an animation from BABEL may include extra actions that are not labeled in our dataset. Instead, we create multi-action animations by connecting actions from the single-action task test set according to the transition matrix where the number of connections between different action categories in the training set is calculated in advance (see Fig. 1 in Supp.). Please note these multi-action animations are not continual motions in the original data, and we only use their conditional information as inputs, including action tags, duration, and trajectories (for “root motion” only).

In order to analyze the influence of the number of transitions between action categories that existed in training data, we set up two kinds of multi-action generation evaluation sets, namely “overall” and “sufficient” testing. The “overall” set is obtained by randomly connected action tags which are at least connected once in the training set. The “sufficient” set is only included action tag sequences that have frequently connected in the training set. Every set has **764** multi-action samples that contains **20** actions. (Detailed construction steps of these two sets are provided in Supplementary.).

5.3.2 Evaluation Metrics

It has not been well established that the criteria to evaluate the multi-action motion generation task. In this paper, we suggest evaluating multi-action animation from three aspects:

- **Per-action Quality:** we measure action recognition accuracy and FID (See Sec. 5.2.1) for action clips which are split from generated results. These two metrics are presented to jointly demonstrate the quality of synthetic actions with the same sequential index.
- **Transition Smoothness:** we evaluate transition smoothness between two actions based on pose and velocity gaps. For *pose gap* (Δ_{pose}), we use the L2 distance of per-joint position between the end pose of the previous action and the start one in the current action. A lower value usually indicates a smoother transition. For *velocity gap* (Δ_{vel}), we also measure the difference of per-joint velocity between the end of the previous action and the start of the current action, which evaluates the dynamic smoothness of transitions.

- **Qualification Ratio:** We define action-level and motion-level qualification ratio (QR) as follows. For action-level QR, we calculate the percentage of accurately classified actions among all independent actions in a generated motion sequence. Given an example, a motion including 3 accurate actions among total 5 actions, then its action-level QR is 60%. We report the average of all testing data as action-level QR for every method. For motion-level QR, we define the animation as a qualified one only if all actions are recognized correctly. Then the Motion QR can be calculated simply as:

$$\text{Motion QR} = \frac{\text{number of qualified animations}}{\text{number of all animations}}. \quad (7)$$

Motion-level QR represents the quality of the full motion sequence, which is a much more strict measurement than action-level QR.

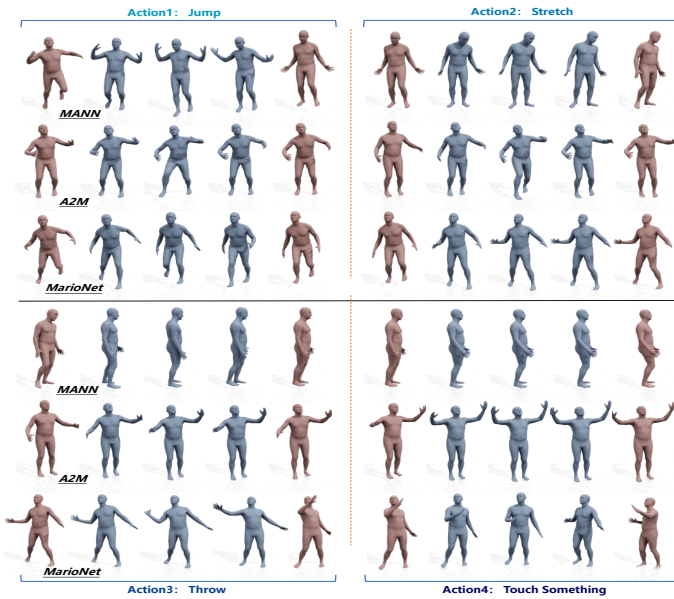


Fig. 8. Action visualization of multi-action generation by different methods. Severe abrupt pose changes between actions could be easily observed in A2M, while MANN does not generate correct desired actions. Our MARIONET successfully generates expected actions with a smooth transition.

5.3.3 Comparison

We firstly compare our MARIONET with the other two methods in a purely generative pipeline. Then we demonstrate how the proposed “Shadow Start” and “Action Revision” schemes can further improve the performance of multi-action generation. For the purely generative pipeline, the whole motions are only synthesized by generative models sequentially, detail algorithm can be found in Supplementary Algorithm.1. Generative methods are evaluated with both overall and sufficient testing. Results are presented in Table 3 and Fig. 9. Besides, we present a qualitative result of multi-action generation by the above three methods in Fig. 8. Note that the *averaged* scores in the table are calculated from the *second* action to the last one, since the initial motion of the first action is from BABEL-MAG, which is different from later action generation that uses previously generated poses as input.

First, we compare per-action quality and qualification ratio metrics among different methods. Under the “overall

testing” setting, our method obtains the highest Recognition Accuracy, Action QR, Motion QR, and lowest FID (i.e. most realistic motions). Fig. 9 shows how these metrics change as more actions are sequentially generated by different methods. As expected, the accuracy of all methods decreases, since the ending poses in the previously generated action may have not been seen in the training set. However, our method consistently performs best on all metrics within 8 actions. After 8 actions, our method is significantly better in terms of FID and generates multiple actions with smooth transitions while performing similarly with A2M on other metrics, indicating that the realism of our generated actions remains high quality.

Next, we compare the Transition Smoothness by two pose-level metrics among different methods (see Table 3). For both settings, our method has a slightly larger pose gap than MANN but obtains a smoother dynamical transition with a lower velocity gap. A2M reports the worst result in terms of the two Transition Smooth metrics over both settings, which is verified in Fig. 8 where A2M easily produces abrupt pose changes.

Finally, we can see choosing the next action restricted on the transition connection matrix (Fig. 1 in Supp.) rankings can improve the performance by 8% on both Accuracy and Action QR, reducing/increasing the FID/Motion QR by nearly half (“sufficient” vs “overall”). This phenomenon supports that the quality of multi-action generation can be further improved if there are more combinations available between every two action categories during training.

TABLE 3
Comparison on multi-action synthesis for pure generative methods.
 Δ_{pose} and Δ_{vel} are in cm and cm/frame respectively.

Method	Overall				Sufficient			
	Acc \uparrow	FID \downarrow	$\Delta_{pose}\downarrow$	$\Delta_{vel}\downarrow$	Acc \uparrow	FID \downarrow	$\Delta_{pose}\downarrow$	$\Delta_{vel}\downarrow$
MANN	0.723	12.785	1.47	0.42	0.746	14.744	1.56	0.37
A2M	0.775	15.796	3.41	1.92	0.830	7.933	4.12	1.73
MARIONET	0.786	7.854	2.03	0.40	0.862	4.624	2.53	0.35

5.3.4 Subjective Evaluation

Along with the above objective experiments, we also conduct a user study to evaluate the quality of actions purely generated by generative models. The quality is judged from three aspects: action accuracy, transition smoothness, and naturalism. The user study has two parts including 8 videos independently. In the first part, each video contains two actions of human motions from the real data and those generated by MANN, A2M, and our method. We use the initial motion, action tags, and duration from the compared real data to generate motions. Based on the video contents, users are requested to verify whether the first and the second actions match to given action tags respectively. Then they give scores for transition smoothness and naturalism.

The second part does *not* include real motions. And every motion contains 3 ~ 5 actions generated by different methods with the same control signals. In addition to the scores of transitions smoothness and naturalism as in the previous part, users are asked to select the number of accurate actions corresponding to action tags for each motion. Every user should finish both parts. We collect 35 valid

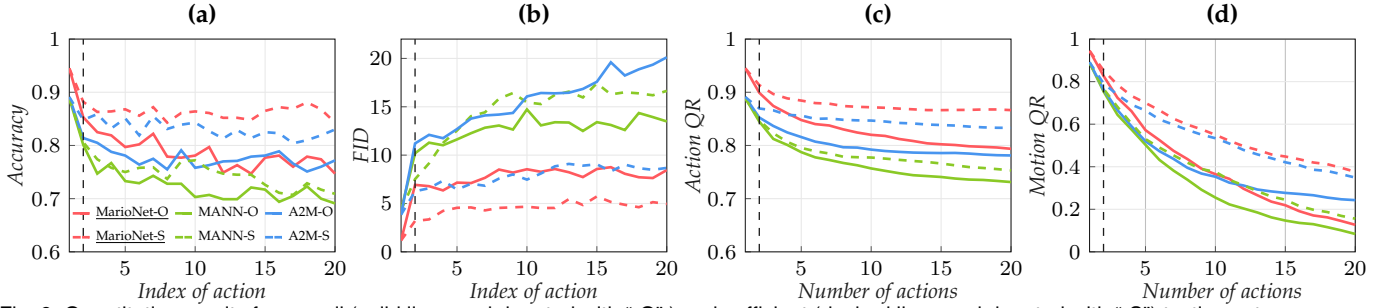


Fig. 9. Quantitative results for overall (solid lines and denoted with "-O") and sufficient (dashed lines and denoted with "-S") testing sets.

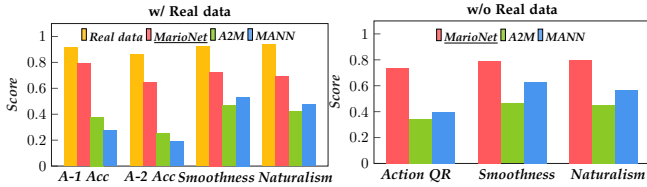


Fig. 10. Results of subject evaluation. "A-1 Acc" and "A-2 Acc" in the left figure indicate the action recognition of the first and the second actions respectively. Both figures plot average scores which range from 0 \sim 1, and 1 represents the best.

evaluation statistics from subjects varying in age, gender, and country.

The average results of the first part are presented in Fig. 10 denoted as "W/ real data". "A1-Acc" and "A2-Acc" are the percentages of how many users believe actions match their tags for the first and the second actions respectively. It can be found that our method has the closest accuracy performance to real data and a significantly higher score compared to MANN and A2M. Besides, the accuracy of the second action has an obvious decrease contrasted to the first one, which is consistent with the objective evaluation in Fig. 9. We conclude this as the challenges of connection sparsity and error accumulation described in Sec. 4.3.2. As for smoothness and naturalism, our method achieves the best performance among generative models but is still lower than real data. Although no real motion is provided in the second part, we observe similar results that our method obtain much higher scores compared to MANN and A2M. Especially for Action QR, our method achieves 73.9%, which is more than 30% higher than the scores of A2M and MANN. This improvement is significantly larger than that in the objective study in Fig. 9.

In summary, the subjective study demonstrates the consistent results with objective experiments that our method surpasses two compared motion synthesis methods in terms of action accuracy, transition smoothness, and naturalism. However, more effort should be made in future work that can help to decrease the gap between our synthetic motions and the real ones.

5.3.5 Performance improvement of "Shadow Start" and "Action Revision"

In order to alleviate the phenomenon of BOP out-of-distribution sometimes caused by the common drifting problem in temporal auto-regressive models, we introduce "Shadow Start" denoted as SS, and "Action Revision" denoted as AR, to the motion synthesis pipeline.

- MARIONET is a pure generative pipeline with our model.

- MARIONET+SS indicates always using MARIONET with Shadow Start operation.
- MARIONET+SS+AR is the hybrid method with a classifier to decide when to use the Shadow Start operation.

Quantitative results on *overall* testing in Table 4 shows the effectiveness of the two proposed schemes. "SS" considers candidate BOPs in training data *every* step to correct the drifting BOP. As expected, it improves the quality of synthetic animation on Accuracy and FID while suffering from a more dramatic gap in smoothness inevitably. With the help of a recognition network, we could determine when is suitable to call the "Shadow Start", which extremely relieves undesired abrupt changes between sequential actions caused by projection operation. We present a visual comparison example in Fig 5.3.5. It can be seen that, directly generating (the 1st row) action with MARIONET leading to a "frozen" motion. With our "SS", "walking backward" action can be synthesized correctly. The last row is a "waling backward" action we search from training set that has most similar starting pose to the end one of previous action. By comparing these three starting pose to the end pose of the previous action, we can find that "SS" provide a good compromise between end poses of previous action, which may be our-of-distribution, and end poses belong to desired action in training set.

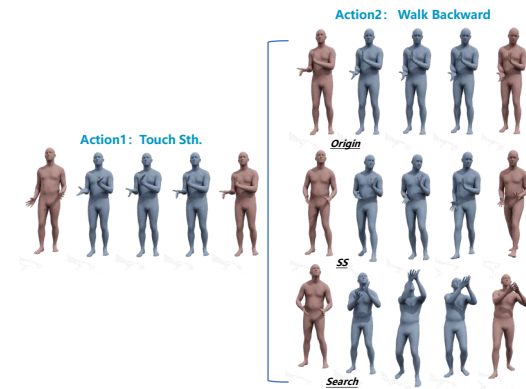


Fig. 11. Action synthesis w/o and w/ SS. The 1st and 2nd rows are action generated w/o and w/ SS, respectively. W/o SS, motion in the 1st row is "frozen", and w/ SS, the second row shows the synthesis action is correct. The last row is action with most similar starting pose to the end pose of previous action.

Since an action may fail only when both normal generation and re-generation using a projected BOP do not work at the same time, we get the best Accuracy by 85.3% for over 14000 actions in total. To conclude, recognition-based "Ac-

tion Revision” makes NEURAL MARIONETTE more robust in high-quality motion generation while providing comparable performance in terms of Transition Smoothness compared to the pure generative pipeline.

TABLE 4
Comparison of SS and AR.

Method	Acc \uparrow	FID \downarrow	$\Delta_{pose} \downarrow$	$\Delta_{vel} \downarrow$
MARIONET	0.786	7.854	2.03	0.40
MARIONET+SS	0.819	6.576	4.57	0.47
MARIONET+SS+AR	0.853	6.887	2.63	0.42

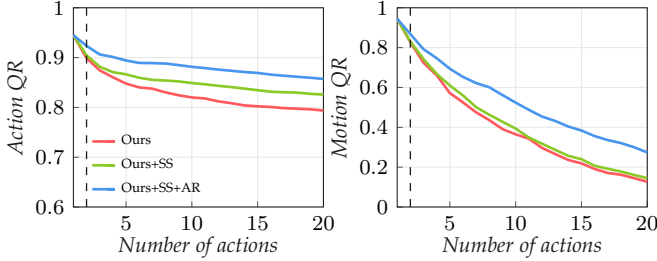


Fig. 12. Action QR (left) and Motion QR (Right) of Our variants.

5.3.6 Failure case

Although, from the above comparison, it can be seen that our NEURAL MARIONETTE achieves significantly higher performance. It has to point out that, even with “shadow start”, our method may fail to generate correct actions sometimes. Failure cases usually appear when the end pose of previous action is significantly different from the start pose of current action in training sets. Fig 5.3.6 present an example to illustrate this phenomenon, which we expect to generate a “kick-walking” motion clip. The left part is a “kick” action while the first and second rows of the right part are “walk” actions generated without and with “SS” respectively. However, both of them are not correct “walk” actions obviously. The last row is another action with the “walk” tag which is retrieved from the training dataset with the closest starting pose to the end pose of the previous action. Nonetheless, this starting pose still has a significant difference from the previous end one, making our MARIONET hard to generate correct action, even with feature projection.

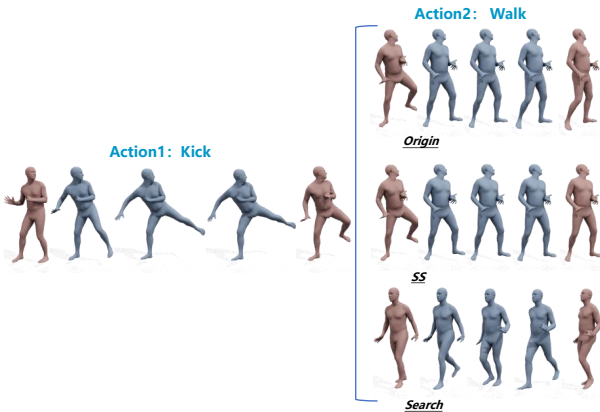


Fig. 13. A failure case of generating the expected motion containing two actions “kick-walk”.

6 APPLICATION: INTERACTIVE MULTI-ACTION GENERATION

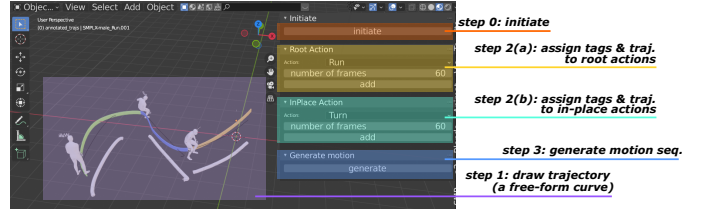


Fig. 14. Screenshot of our trajectory annotation GUI.

In this section, we demonstrate that by using our NEURAL MARIONETTE, multi-action can be controllably generated in an interactive fashion. We develop an interactive multi-action system as a blender add-on and the parametric human model is from SMPLX [39]. Besides with our NEURAL MARIONETTE, a simple spherical linear interpolation [51] is used to further smooth the transition when “Shadow Start” is used. The interface is shown in Fig. 14. Users can generate a multi-action motion by firstly plotting a free hand curve as Fig 6, then annotates this curve with multiple action tags, duration, see Fig. 15. After finishing per-action annotating, user can click “generating” button, then our NEURAL MARIONETTE starts to generate actions and display on screen.

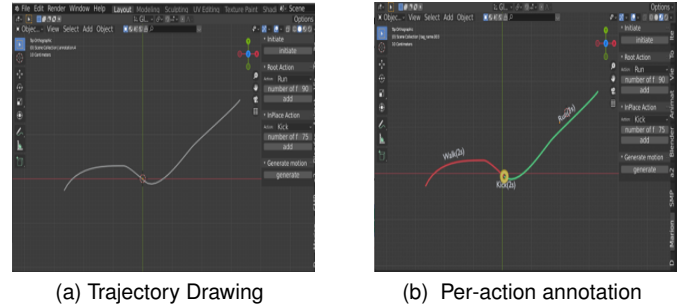


Fig. 15. Interactive steps of multi-action synthesis.

In summary, with our interactive animation synthesis system, human animation creation is simplified as drawing a trajectory curve and annotating with a few mouse clicks. It is not only significantly improving efficiency but also lowers the bar of making human animation.

7 CONCLUSION AND LIMITATION

In conclusion, we propose a new scheme, NEURAL MARIONETTE, for the multi-action human motion synthesis task. MARIONET, the core part of our method, can fully utilize action tag and contextual information to generate high-quality single-action motions with smooth transitions which can be later stitched together to produce a much longer multi-action motion. Experiments on both conditional single-action and multi-action generation tasks demonstrate that our MARIONET and NEURAL MARIONETTE achieve superior motion synthesis performance in terms of action accuracy, naturalism, and transition smoothness over multiple baseline methods. We further demonstrate that our method can be used for human animation creation in an interactive

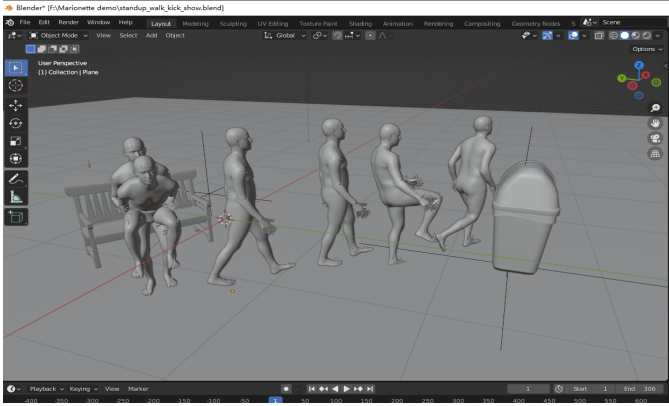


Fig. 16. A user case of our interactive motion synthesis GUI. It demonstrate a man *stands up*, then *walks* to a trash bin then *kick* it, and finally *runs away*.

way via a Blender add-on integrated with our NEURAL MARIONETTE.

Although our method has achieved state-of-the-art performance among generative models, the result of subject evaluation suggests that there is still a gap between the real human motions and our synthesized motions in the terms of Realism and Transition Smoothness. Another limitation of our NEURAL MARIONETTE is that it does not support interactive frame-editing during generation which limits the usage in real animation creation where more strict requirements have to be satisfied. Though the whole generated motion can be edited afterward, it would be more convenient for technical artists to generate desired motion if frame-editing is available during the current interactive workflow. Therefore, in the future, we would like to investigate more powerful techniques to narrow the gap between synthetic motions and real ones. Meanwhile, frame-editing, this meaningful function would be inserted into our system to make animation creation more controllable and efficient.

REFERENCES

- [1] D. Holden, T. Komura, and J. Saito, "Phase-functioned neural networks for character control," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, pp. 1–13, 2017.
- [2] H. Zhang, S. Starke, T. Komura, and J. Saito, "Mode-adaptive neural networks for quadruped motion control," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–11, 2018.
- [3] D. Holden, O. Kanoun, M. Perepichka, and T. Popa, "Learned motion matching," *ACM Transactions on Graphics*, vol. 39, no. 4, pp. 53–1, 2020.
- [4] C. Guo, X. Zuo, S. Wang, S. Zou, Q. Sun, A. Deng, M. Gong, and L. Cheng, "Action2motion: Conditioned generation of 3d human motions," in *ACM International Conference on Multimedia*, 2020, pp. 2021–2029.
- [5] M. Petrovich, M. J. Black, and G. Varol, "Action-conditioned 3d human motion synthesis with transformer vae," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 985–10 995.
- [6] Y. Yuan and K. Kitani, "Dlow: Diversifying latent flows for diverse human motion prediction," in *European Conference on Computer Vision*, 2020, pp. 346–364.
- [7] E. Barsoum, J. Kender, and Z. Liu, "Hp-gan: Probabilistic 3d human motion prediction via gan," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1418–1427.
- [8] I. Habibie, D. Holden, J. Schwarz, J. Yearsley, and T. Komura, "A recurrent variational autoencoder for human motion synthesis," in *British Machine Vision Conference*, 2017.
- [9] Z. Wang, J. Chai, and S. Xia, "Combining recurrent neural networks and adversarial training for human motion synthesis and control," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 1, pp. 14–28, 2021.
- [10] H. Wang, E. S. L. Ho, H. P. H. Shum, and Z. Zhu, "Spatio-temporal manifold learning for human motions via long-horizon modeling," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 1, pp. 216–227, 2021.
- [11] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, "NTU RGB+D: A large scale dataset for 3d human activity analysis," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016, pp. 1010–1019.
- [12] J. Liu, A. Shahroudy, M. Perez, G. Wang, L.-Y. Duan, and A. C. Kot, "NTU RGB+D 120: A large-scale benchmark for 3d human activity understanding," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 42, no. 10, pp. 2684–2701, 2019.
- [13] A. R. Punnaakkal, A. Chandrasekaran, N. Athanasiou, A. Quiros-Ramirez, and M. J. Black, "Babel: Bodies, action and behavior with english labels," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 722–731.
- [14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [15] D. Holden, J. Saito, and T. Komura, "A deep learning framework for character motion synthesis and editing," *ACM Transactions on Graphics*, vol. 35, no. 4, pp. 1–11, 2016.
- [16] T. Tang, J. Jia, and H. Mao, "Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis," in *ACM International Conference on Multimedia*, 2018, pp. 1598–1606.
- [17] H.-Y. Lee, X. Yang, M.-Y. Liu, T.-C. Wang, Y.-D. Lu, M.-H. Yang, and J. Kautz, "Dancing to music," *Advances in neural information processing systems*, vol. 32, 2019.
- [18] J. Li, Y. Yin, H. Chu, Y. Zhou, T. Wang, S. Fidler, and H. Li, "Learning to generate diverse dance motions with transformer," *arXiv preprint arXiv:2008.08171*, 2020.
- [19] R. Li, S. Yang, D. A. Ross, and A. Kanazawa, "AI Choreographer: Music Conditioned 3D Dance Generation with AIST++," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 13 401–13 412.
- [20] A. Aristidou, A. Yiannakidis, K. Aberman, D. Cohen-Or, A. Shamir, and Y. Chrysanthou, "Rhythm is a dancer: Music-driven motion synthesis with global structure," *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2022.
- [21] J. Chen, C. Fan, Z. Zhang, G. Li, Z. Zhao, Z. Deng, and Y. Ding, "A music-driven deep generative adversarial model for guzheng playing animation," *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2021.
- [22] S. Levine, P. Krähenbühl, S. Thrun, and V. Koltun, "Gesture controllers," in *ACM Transactions on Graphics*, 2010, pp. 1–11.
- [23] S. Ginosar, A. Bar, G. Kohavi, C. Chan, A. Owens, and J. Malik, "Learning individual styles of conversational gesture," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 3497–3506.
- [24] Y. Yoon, B. Cha, J.-H. Lee, M. Jang, J. Lee, J. Kim, and G. Lee, "Speech gesture generation from the trimodal context of text, audio, and speaker identity," *ACM Transactions on Graphics*, 2020.
- [25] J. Li, D. Kang, W. Pei, X. Zhe, Y. Zhang, Z. He, and L. Bao, "Audio2gestures: Generating diverse gestures from speech audio with conditional variational autoencoders," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 11 293–11 302.
- [26] Y. Ji, F. Xu, Y. Yang, F. Shen, H. T. Shen, and W.-S. Zheng, "A large-scale RGB-D database for arbitrary-view human action recognition," in *ACM International Conference on Multimedia*, 2018, pp. 1510–1518.
- [27] J. Martinez, M. J. Black, and J. Romero, "On human motion prediction using recurrent neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2891–2900.
- [28] Y.-H. Wen, Z. Yang, H. Fu, L. Gao, Y. Sun, and Y.-J. Liu, "Autoregressive stylized motion synthesis with generative flow," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 13 612–13 621.
- [29] D. Pavilo, D. Grangier, and M. Auli, "Quaternet: A quaternion-based recurrent model for human motion," in *British Machine Vision Conference*, 2018.

- [30] G. E. Henter, S. Alexanderson, and J. Beskow, "Moglow: Probabilistic and controllable motion synthesis using normalising flows," *ACM Transactions on Graphics*, vol. 39, no. 6, pp. 1–14, 2020.
- [31] R. Huang, H. Hu, W. Wu, K. Sawada, M. Zhang, and D. Jiang, "Dance revolution: Long-term dance generation with music via curriculum learning," in *bp-iclr*, 2021.
- [32] R. Bellini, Y. Kleiman, and D. Cohen-Or, "Dance to the beat: Synchronizing motion to audio," *Computational Visual Media*, vol. 4, no. 3, pp. 197–208, 2018.
- [33] W. Zhuang, C. Wang, J. Chai, Y. Wang, M. Shao, and S. Xia, "Music2dance: Dancenet for music-driven dance generation," *ACM Transactions on Multimedia Computing, Communications, and Applications*, vol. 18, no. 2, pp. 1–21, 2022.
- [34] L. Kovar, M. Gleicher, and F. Pighin, "Motion Graphs," *ACM Transactions on Graphics*, 2002.
- [35] K. Chen, Z. Tan, J. Lei, S.-H. Zhang, Y.-C. Guo, W. Zhang, and S.-M. Hu, "Choreomaster: choreography-oriented music-driven dance synthesis," *ACM Transactions on Graphics*, vol. 40, no. 4, pp. 1–13, 2021.
- [36] K. Fragkiadaki, S. Levine, P. Felsen, and J. Malik, "Recurrent network models for human dynamics," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2015, pp. 4346–4354.
- [37] K. Cho, B. Van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder-decoder approaches," *arXiv preprint arXiv:1409.1259*, 2014.
- [38] N. Mahmood, N. Ghorbani, N. F. Troje, G. Pons-Moll, and M. J. Black, "Amass: Archive of motion capture as surface shapes," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 5442–5451.
- [39] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. A. Osman, D. Tzionas, and M. J. Black, "Expressive body capture: 3D hands, face, and body from a single image," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10 975–10 985.
- [40] F. G. Harvey, M. Yurick, D. Nowrouzezahrai, and C. Pal, "Robust motion in-betweening," *ACM Transactions on Graphics*, vol. 39, no. 4, pp. 60:1–60:12, 2020.
- [41] M. Ranzato, S. Chopra, M. Auli, and W. Zaremba, "Sequence level training with recurrent neural networks," *arXiv preprint arXiv:1511.06732*, 2015.
- [42] W. Zhang, Y. Feng, F. Meng, D. You, and Q. Liu, "Bridging the gap between training and inference for neural machine translation," *arXiv preprint arXiv:1906.02448*, 2019.
- [43] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer, "Scheduled sampling for sequence prediction with recurrent neural networks," *Advances in neural information processing systems*, vol. 28, 2015.
- [44] T. Mihaylova and A. F. T. Martins, "Scheduled sampling for transformers," *arXiv preprint arXiv:1906.07651*, 2019.
- [45] Y. Li, Y. Pan, T. Yao, J. Chen, and T. Mei, "Scheduled sampling in vision-language pretraining with decoupled encoder-decoder network," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 10, 2021, pp. 8518–8526.
- [46] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, "Smpl: A skinned multi-person linear model," *ACM Transactions on Graphics*, vol. 34, no. 6, pp. 1–16, 2015.
- [47] Z. Kalal, K. Mikolajczyk, and J. Matas, "Tracking-learning-detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 34, no. 7, pp. 1409–1422, 2011.
- [48] Z. Zhu, Q. Wang, B. Li, W. Wu, J. Yan, and W. Hu, "Distractor-aware siamese networks for visual object tracking," in *European Conference on Computer Vision*, September 2018.
- [49] S.-Y. Chen, W. Su, L. Gao, S. Xia, and H. Fu, "Deepfacedrawing: Deep generation of face images from sketches," *ACM Transactions on Graphics*, vol. 39, no. 4, pp. 72–1, 2020.
- [50] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," *Advances in neural information processing systems*, vol. 30, 2017.
- [51] K. Shoemake, "Animating rotation with quaternion curves," pp. 245–254, 1985.

Supplementary for NEURAL MARIONETTE: A Transformer-based Multi-action Human Motion Synthesis System

Weiqliang Wang, Xuefei Zhe, Huan Chen, Di Kang, Tingguang Li, Ruizhi Chen, and Linchao Bao,



1 TERMINOLOGY

In this section, we briefly provide some terminologies used in our paper for better understanding.

1.1 Motion and Action

We define a *motion* as a sequence of human poses with different movements, which can be long and complicated. As a comparison, an *action* is a relatively short motion clip, which only contains a single type of movement pattern. Therefore, a motion can be decomposed into multiple sequential actions. For example, a motion sequence that has a human walking for a few seconds and then running can be decomposed into two actions, "Walk" and "Run".

1.2 Action Description

We first define \mathcal{A} as a collection of action *tags* that are as independent of each other as possible. As discussed above, an action can then be described by tags from \mathcal{A} , i.e., $A = \{a_1, \dots, a_k\}$, where $a_i \in \mathcal{A}, \forall i = 1 \dots k$. We call this action has *Tag Annotation A*. Note that, multi-tag annotation happens only for non-intention actions

1.3 Action Representation

An action M can be represented by a sequence of human poses, i.e., $M = \{P_1, \dots, P_t, \dots, P_T\}$, where P_t denotes a single pose at t -th frame and T is the total number of frames in the clip. A human pose P_t can be characterized by a *root translation* $D_t \in \mathbf{R}^3$ and the *rotations* of all the joints

Algorithm 1: Pure generative multi-action pipeline

Input : $\{A_i, T_i, U_i\}_N, \{P_t\}_{t=-x}^{t=-1}$
for $i \leftarrow 1$ **to** N **do**
 R_z and $D_t \leftarrow P_{-1}$;
 Normalized initial pose
 $\{P_t^*\}_{t=-x}^{t=-1} \leftarrow R_z^{-1} \{P_t\}_{t=-x}^{t=-1} - D_t$;
 Generate i -th action $\{\hat{P}_t\}_{t=1}^{t=T_i} \leftarrow$
 $G(\{P_t^*\}_{t=-x}^{t=-1}, A_i, A_{i+1}, T_i, U_i)$;
 Rotated i -th action $\{\hat{P}_t\}_{t=1}^{t=T_i} \leftarrow R_z \hat{P}_t + D_t$;
 Update initial pose sequence
 $\{P_t\}_{t=-x}^{t=-1} \leftarrow \{\hat{P}_t\}_{t=T_{i-1}-x+1}^{t=T_i}$;
end

$R_t \in \mathbf{R}^{n_J \times 6}$, where n_J is the number of joints and the continuous 6D rotation representation is used [1]. Given a pose $P_t = (R_t, D_t)$ and certain body shape parameter, we can recover the *body mesh* using SMPL [2] or SMPLX [3].

1.4 Action Trajectory

The trajectory U of an action M can be described by *root translation* $D_t (t = 1, \dots, T)$. We denote this as a function, $F : [1, \dots, T] \rightarrow \mathbf{R}^2$, which assigns a translation vector in the ground plane to a specified frame. Thus, one way to directly manipulate the trajectory of the synthesized motion is by drawing a curve on the ground in the world coordinate system and then sampling discretely. More details are provided in Sec. 5.2.

2 DATASET

2.1 BABEL-MAG Dataset

In this part, we show more details about the BABEL-MAG. To clean the annotations of the collected motion data, we propose an intention-first tagging protocol, including a simple tagging pipeline with consist of several tagging rules, which instruct how to compose tags to consistently and thoroughly describe an action. Specifically, we first annotate those actions with a clear intention such as "punch" and "sit down". These actions are tagged according to their final goals. For example, if an action demonstrates a man *walks* for *moving something*, we will tag this action with "moving

- W. Wang and R. Chen are with the State Key Laboratory of Information Engineering in Surveying, Mapping and Remote Sensing, Wuhan University, Wuhan 430079, China.
E-mail: {weiqking, ruizhi.chen}@whu.edu.cn.
- H. Chen is with Shenzhen International Graduate School, Tsinghua University, Shenzhen 518055, China.
E-mail: huan-che20@mails.tsinghua.edu.cn.
- Z. Xue, D. Kang and B. Lin are with Tencent AI Lab, Shenzhen 518054, China. E-mail: {zhexuefei, di.kang}@outlook.com, linchaobao@gmail.com.
- T. Li is with Tencent Robotics X, Shenzhen 518054, China.
E-mail: tgli0809@gmail.com.
- W. Wang and X. Zhe contributed equally to the work; B. Lin is the corresponding author; W. Wang and H. Chen did the work during their internships at Tencent AI Lab.

§. These two authors contribute equally to the work.

something” because his final goal is to move something and walking is helping to achieve this goal. If an action does not have a clear intention, we then use the corresponding body state (including “stand”, “bend”, and “sit”) and body part (including “head”, “arm”, “waist”, and “leg”) movement to describe it. Note that only non-intention actions are allowed to be labeled with multiple tags. Our action tags have simple and unambiguous definitions, and they are as independent as possible from each other, greatly simplifying the action annotation process.

The statistic and type of every action tag used in BABEL-MAG are presented in Table 1. Besides, we calculate the action pairs between every two tags and show this by the Transition Adjacent Matrix via a heatmap in Fig. 1.

TABLE 1
Action Tags for BABEL-MAG. Note that, five action types — “Intention”, “Root”, “In-place”, “Status”, “Body parts” are correspondingly abbreviated as “I”, “R”, “P”, “S”, “B” for simplification.

Tag ID	Tag Type	Tag Name	Segment Number
0	I & R	Walk	1370
1	I & R	Walk Backward	160
2	I & R	Walk Sideway	105
3	I & R	Walk Upward/Downward	272
4	I & R	Step	536
5	I & R	Run	198
6	I & R	Jump	168
7	I & P	Turn	1063
8	I & P	Sit Down	86
9	I & P	Stand Up	123
10	I & P	Using Object	227
11	I & P	Take Something Up	375
12	I & P	Place Something Down	261
13	I & P	Move Something	102
14	I & P	Throw	143
15	I & P	Touch Something	301
16	I & P	Kick	91
17	I & P	Punch	69
18	I & P	Wave Hand	29
19	I & P	Stretch	410
20	I & P	Crossing Limbs	33
21	I & P	Touch Body Part	200
22	I & P	Using Object and Sit	39
23	I & P	Step and Touch Body Part	54
24	S	Stand	754
25	S	Bend	149
26	S	Sit	325
27	B	Head Movement	199
28	B	Arm Movement	859
29	B	Waist Movement	30
30	B	Leg Movement	193

3 IMPLEMENTATION

We implemented MARIONET with Pytorch, and NEURAL MARIONETTE as an add-on in Blender. All motion sequences are firstly downsampled to 30 Hz, and the initial motion provided in each contextual information is fixed to 6 frames, while tag annotations of BABEL-MAG are saved in the same format as BABEL. The detailed settings for MARIONET is shown in Table 2. And we trained it using Adam [4] with a batch size of 30 and a learning rate of 0.0001, for 1000 epochs where Activate epoch is 250 and Accomplish epoch is 750. Besides, we empirically set different weights of training loss with $\lambda_{KL} = 10^{-8}$, $\lambda_D = 0.5$, $\lambda_R = 1$, $\lambda_J = 2$, $\lambda_V = 2$. When the scheduled sampling is activated, $\lambda_1 = 0.1$ and $\lambda_2 = 0.9$;

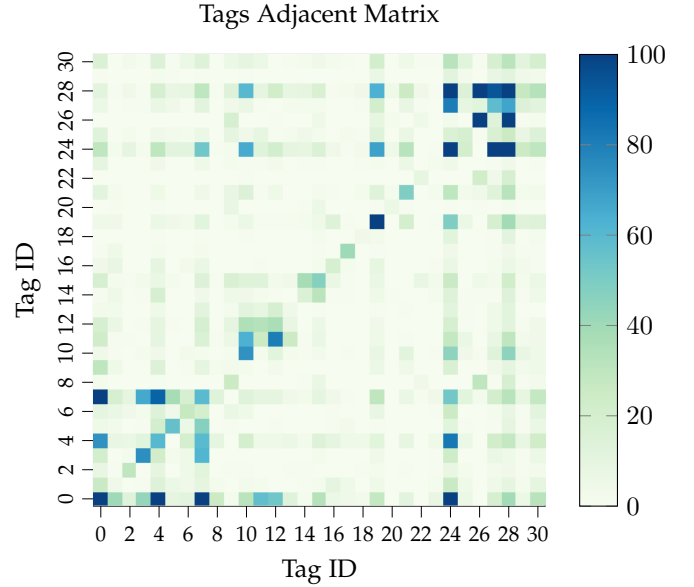


Fig. 1. We report the number of action pairs that are adjacent to each other between every two tags via a heatmap.

$\lambda_1 = 1$ and $\lambda_2 = 0$ otherwise. We will release our code and BABEL-MAG to reproduce all the results.

4 MORE EXPERIMENTS DETAILS

4.1 Details of baselines and their adaption

- **MANN** [5] is a mixture-of-expert (MoE) based method for motion generation, consisting of K expert network(s) and a gating network to output the blend weights of the experts. Although MANN is designed for quadruped characters, its learned blending strategy takes complicated dynamic states into consideration automatically and outperforms PFNN [6] which uses both a predefined blending strategy and manually labeled phase labels.
- **Action2motion (A2M)** [7] is a temporal variational auto-encoder (VAE) which can generate 3D human motion based on the input action tag. From the view of input/output, A2M is the closest method to ours. For a fair comparison, we made two adjustments to A2M including (1) replacing lie algebra representation of the joint rotations with 6D continuous representation [1]; (2) providing A2M with the same initial pose as our setting rather than a rest pose.
- **ACTOR** [8] is a conditional variational auto-encoder based on Transformer, which is also used in our MARIONET. Note that ACTOR is included only as a reference since it cannot take as input a beginning pose which makes it unsuitable to generate continuous actions one by one.

4.2 Construction for Multi-action Evaluation Set

We first use a trained motion classification net to eliminate those test samples with lower confidence scores than 50% to obtain action candidates set of 764 high-quality action clips. In the overall testing set, we firstly use any action with tag A_0 as the first action. Then a next tag A_1 is randomly selected from all tags except ones that are not connected

TABLE 2
Network architecture for MARIONET.

		Condition Encoder				Time Unrolling		Motion Decoder	
Transformer	d_model	256				32		256	
	n_head	8				1		8	
	dim_feedforward	1024				64		1024	
	num_layers	2				1		2	
	dropout	0.1				0.001		0.01	
		gelu				relu		relu	
FC		tag_embd	pose_encode	mean_embd	std_embd	traj_embd	unroll_embd	pose_encode	pose_decode
	in_size	31	135	256	256	2	256	135	256
	out_size	256	256	256	256	32	224	256	135
		relu				relu		relu	

by the current tag A_0 . Finally, the next action is randomly chosen from actions with A_1 in the action candidates set. A multi-action test sample is formed by **20** actions selected in this way. By going through the action candidates set, we obtain an overall testing dataset with **764** multi-action samples.

As for sufficient testing, we consider building multi-action animation with sufficient connections between actions. The difference from the overall testing is how the next action tag is selected. We restrict the next candidate action tag derived from the following tags:

- 1) the top 4 tags connected by current action A_0 in transition matrix;
- 2) the current action tag A_0
- 3) the next action tag of the current action sample

Then next tag is randomly selected from these candidate tags. We generated the same number of sufficient multi-action testing samples as the overall. Under such a setting, we expect to study the performance of training data with sufficient action connections.

4.3 Additional Results

In this section, we provide another example to show the effectiveness of control signals in Fig. 2 and single-action visualization samples by sequences of dense frames in Fig. 4 as complements to the main text. Besides, more qualitative results of multi-action generation are included in supplementary videos for comprehensive evaluation.

5 APPLICATION

5.1 System Workflow

Our interactive motion generation system is a client-server application.(see Fig. 5). It includes three main parts: (1) a user interface that allows users to a draw free-form curve to specify motion trajectory, select multiple action tags with the corresponding duration, and support displaying multiple generated motions; (2) a pre-processing step to convert the hand-drawn trajectory into proper inputs of our network; (3) a multi-action synthesis step utilizing NEURAL MARIONETTE to generate the whole motion and return data to the interface. A user can follow the steps below to generate the expected motion.

- 1) Scene Preparation (Optional). Prepare or import a scene to display the expected motion.
- 2) Trajectory Drawing. Click the "Initiate" button, then draw a free-form curve by clicking on the canvas, which

"Run"

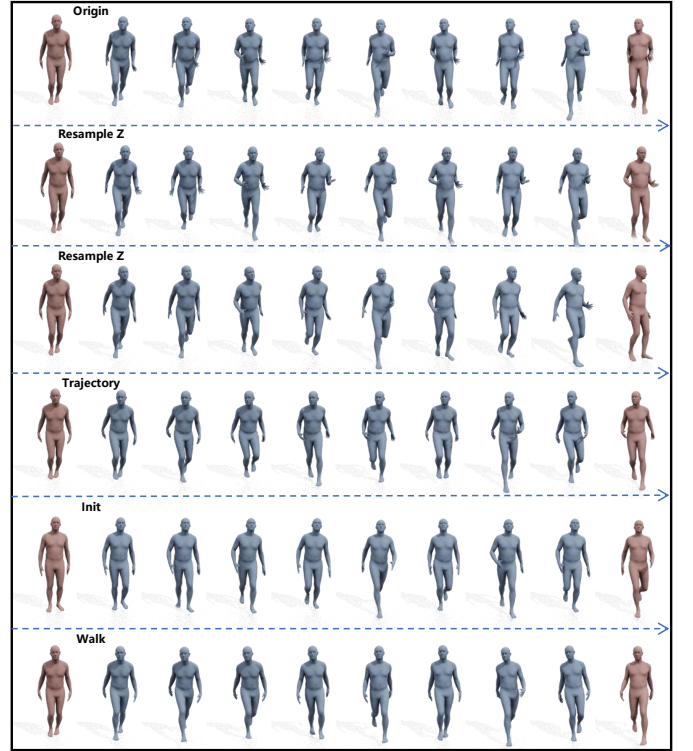


Fig. 2. Effectiveness of different control signals. Note that, "Run" in each row is downsampled to the same number of frames including the original start and end poses, and they are marked as the same color *red* for better visualization.

will be taken as the motion trajectory. We show this step in Fig. 15(a) in the main text.

- 3) Per-action annotation. If necessary, one can insert points into the curve to divide it into segments for different actions by clicking. Then select a curve segment (or a point) to assign tags and duration for a root action (or an in-place action) and click "add" to finish the current assignment. Note that, a path (or a circle) will be displayed to show the added root (or in-place) action. One can repeat this step to finish multiple action assignments. As for the in-place actions at the same circle, the action order in the synthesized motion is the same as the annotation order. Fig. 14(b) in the main text gives an example of annotating three actions.
- 4) Generation. Click the "Generate" button to start the motion synthesis and check the results displayed on

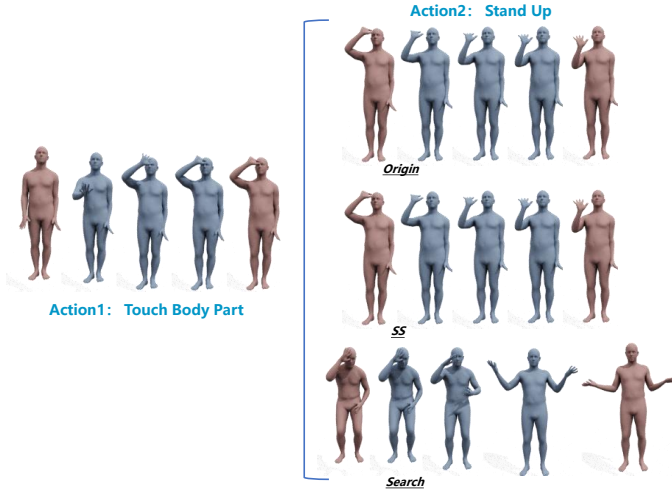


Fig. 3. failure case.

the canvas.

A user case of human motion generated in the 3D scene is demonstrated in Fig. 15 in the main text.

5.2 Annotation Pre-processing

Firstly, the initial motion of the first action is either provided by the user or randomly selected from the training set conditioned on the same current tag. As for that in the following action, it is obtained by collecting the last 6 frames from the previously generated action. The current tag, the following tag, and the duration for each action can be obtained from the user annotations directly. Note that, the ‘following tag’ is simply assigned as the corresponding ‘current tag’ for the last action by default.

The trajectory for an in-place action is set to zero value since no global translation of the root node is expected. For a root action, we fit a cubic Bézier curve. The positions of the samples are used as trajectory locations. Specifically, we have the *trajectory location* for the n -th frame as follows:

$$\mathbf{B}(t) = (1-t)^3 P_0 + 3(1-t)^2 t P_1 + 3(1-t) t^2 P_2 + t^3 P_3 \quad (\text{A.1})$$

where $t = \frac{n}{T}$, T is the total number of frames, and (P_0, P_1, P_2, P_3) are the positions of the Bézier points during the fitting processing. Finally, the body orientation of normalized BOP and the direction of processed trajectory for the first 6 frames would be checked whether their angle is out of the range of those in the training dataset. If an invalid angle appears, the body orientation would be automatically rotated to obey the rule.

REFERENCES

- [1] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, “On the continuity of rotation representations in neural networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5745–5753.
- [2] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, “Smpl: A skinned multi-person linear model,” *ACM Transactions on Graphics*, vol. 34, no. 6, pp. 1–16, 2015.
- [3] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. A. Osman, D. Tzionas, and M. J. Black, “Expressive body capture: 3D hands, face, and body from a single image,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10975–10985.
- [4] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [5] H. Zhang, S. Starke, T. Komura, and J. Saito, “Mode-adaptive neural networks for quadruped motion control,” *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–11, 2018.
- [6] D. Holden, T. Komura, and J. Saito, “Phase-functioned neural networks for character control,” *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, pp. 1–13, 2017.
- [7] C. Guo, X. Zuo, S. Wang, S. Zou, Q. Sun, A. Deng, M. Gong, and L. Cheng, “Action2motion: Conditioned generation of 3d human motions,” in *ACM International Conference on Multimedia*, 2020, pp. 2021–2029.
- [8] M. Petrovich, M. J. Black, and G. Varol, “Action-conditioned 3d human motion synthesis with transformer vae,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10985–10995.

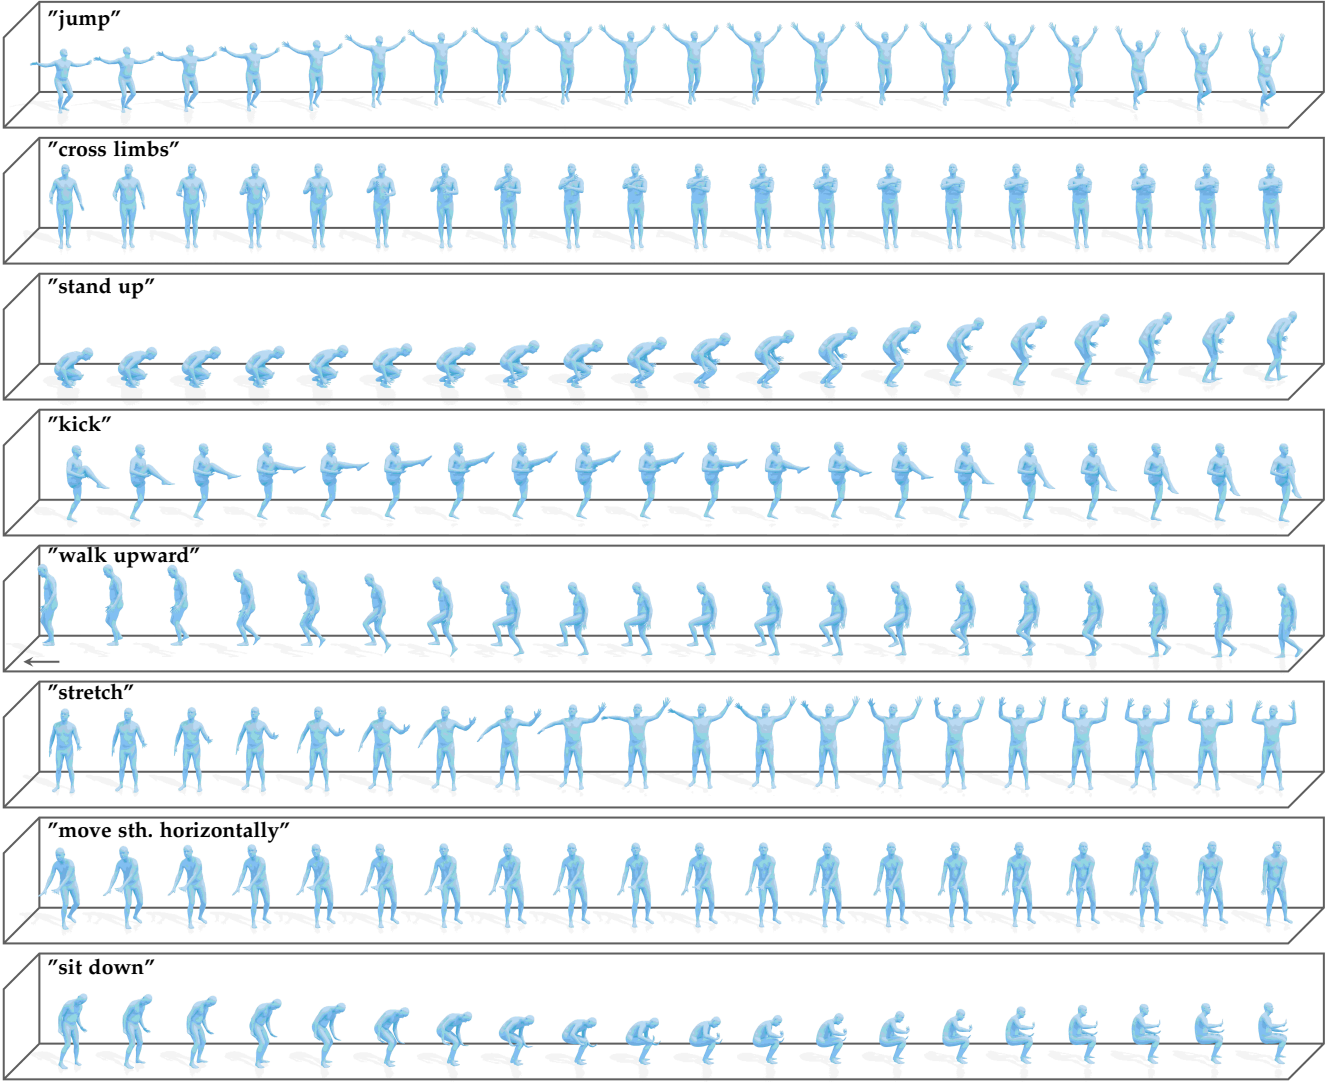


Fig. 4. Dense-frames visualization for eight synthesized actions using MARIONET. Each action is down-sampled to the same number of frames for better visualization.

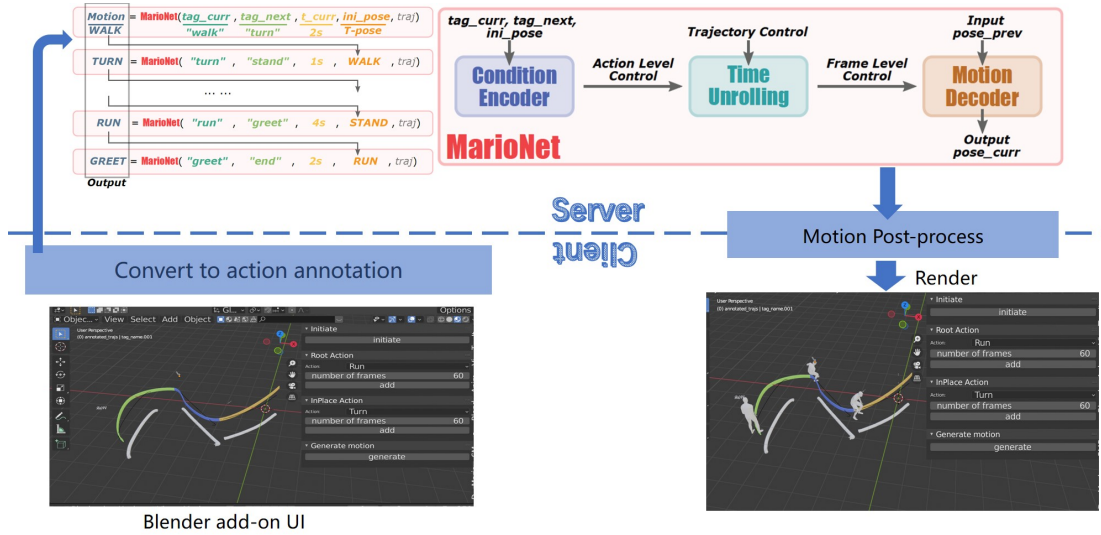


Fig. 5. Application structure. Our interactive multi-action generation system is a client-server application. The user creates action annotations on the client side and sends them to the server while the server runs NEURAL MARIONETTE to synthesize motions and send them back to display on the screen.