

# Recurrent Transformer Variational Autoencoders for Multi-Action Motion Synthesis

Rania Briq<sup>1,2</sup>    Chuhan Zou<sup>2</sup>    Leonid Pishchulin<sup>2</sup>  
Chris Broaddus<sup>2</sup>    Jürgen Gall<sup>1</sup>  
<sup>1</sup>University of Bonn    <sup>2</sup>Amazon

## Abstract

*We consider the problem of synthesizing multi-action human motion sequences of arbitrary lengths. Existing approaches have mastered motion sequence generation in single-action scenarios, but fail to generalize to multi-action and arbitrary-length sequences. We fill this gap by proposing a novel efficient approach that leverages the expressiveness of Recurrent Transformers and generative richness of conditional Variational Autoencoders. The proposed iterative approach is able to generate smooth and realistic human motion sequences with an arbitrary number of actions and frames while doing so in linear space and time. We train and evaluate the proposed approach on PROX dataset [15] which we augment with ground-truth action labels. Experimental evaluation shows significant improvements in FID score and semantic consistency metrics compared to the state-of-the-art.*

## 1. Introduction

This paper addresses the problem of synthesizing multi-action human motion sequences of arbitrary lengths. This is a major research problem with a vast range of applications, including animating avatars and virtual characters [17, 21, 22, 29, 33, 38, 38], character visualization in stories [23], inpainting of missing frames in videos and completion of action sequences [4, 37], and synthesizing realistic motion under 3D scene constraints [14, 36]. Generating realistic motions of humans performing multiple action over longer periods of time is a challenging task, since it requires modeling dependencies between multiple possibly overlapping actions in the temporal domain, while ensuring that transition points between different actions are continuous and smooth. Recent approaches [13, 30] were able to considerably improve the fidelity of synthesized motion by building on conditional Variational Autoencoders (VAEs) [20] and Transformers [35]. However, both approaches are restricted to generating single action motion

sequences only. Furthermore, as motion sequences are generated in a single shot in a non-causal way, generation is typically limited to short sequences of a pre-defined length to avoid running out of memory, which often results in context fragmentation. On the other hand, Recurrent Transformers (RT) [6, 7, 9, 32] have been shown to combine the expressiveness of self-attention mechanism [35] to model long-range dependencies with the efficiency of recurrent architectures [16]. However, RTs lack conditioning and thus have not been used for motion sequence synthesis.

We propose a novel approach that takes an ordered set of action labels as input and synthesizes an arbitrary-length motion sequence of 3D human meshes performing these multiple potentially-overlapping actions. At the core of our approach lies a novel spatio-temporal formulation that marries Recurrent Transformers (RT) with conditional Variational Autoencoders (VAE) to generate multi-action sequences. We thus dub the proposed approach as RTVAE-Multi. The approach (cf. Fig. 1) is trained using sequences of SMPL [26, 28] 3D body meshes of people performing multiple actions, and associated ordered sets of action labels. During training, at each timestamp, we concatenate a 3D body mesh with input actions embeddings and feed them into the encoder together with the previous hidden state. Once the encoder has seen all frames, we sample multiple latent vectors each corresponding to the whole action subsequence. The decoder then uses the query sequence corresponding to the positional embedding of a given timestamp and stacked latent vectors to reconstruct the input poses. During inference, the length of the sequence that can be generated is arbitrary, and each frame is generated at constant time and memory.

To train RTVAE-Multi we require a dataset with both multi-action sequences annotated with GT action labels, and frame-wise SMPL [26, 28] 3D body mesh fittings. Since no such dataset is publicly available, we found PROX dataset [15] to be closest to our requirements, and have therefore augmented it with GT multi-action labels.

Our main contribution is a novel spatio-temporal formulation that combines the expressiveness and efficiency of Recurrent Transformers with generative richness of condi-

\* Work done during internship at Amazon.

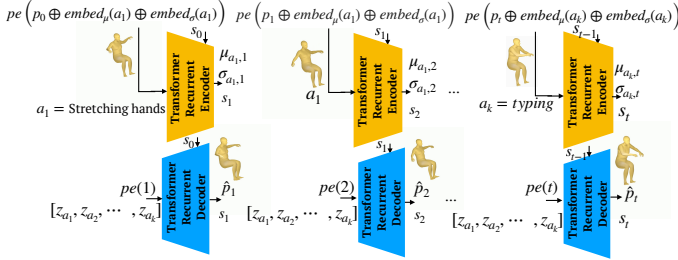


Figure 1. Overview of our approach. The encoder and decoder steps are unrolled to demonstrate the synthesis process of sequences conditioned on multiple actions. At iteration  $t$ , the input to the encoder is a SMPL pose denoted by  $p_t$ , concatenated with the embedding of the input action  $a_i$ , where  $a_i$  spans multiple frames, and  $pe$  is the positional embedding. The previous hidden state used to compute the output at  $t$  is denoted by  $s_{t-1}$ . For each action, we save the VAE distribution parameters  $\mu_{a_i, t_i}$  and  $\sigma_{a_i, t_i}$  where  $t_i$  denotes the end timestamp of action  $a_i$ . Once the encoder has seen all frames, we sample  $k$  latent vectors  $z_i \sim N(\mu_{a_i, t_i}, \sigma_{a_i, t_i})$  each corresponding to the entire sub-sequence matching  $a_i$ . The decoder transformer query sequence corresponds to the positional embedding of timestamp  $t$ , and  $[z_{a_1}, z_{a_2}, \dots, z_{a_k}]$  is a stacking of all the latent vectors samples from encoder output. The decoder is optimized to output a reconstruction of the input poses using an MSE loss term.

tional Variational Autoencoders to generate arbitrary-length multi-action motion sequences. In contrast to previous works [13, 30], we address a more challenging problem of multi-action motion synthesis, are able to generalize to arbitrary number of actions per generated sequence, and to generate arbitrary length sequences where space and time requirements grow linearly in the number of frames. In Sec. 4 we demonstrate significant improvements in motion synthesis over multi-action extension of [30], while performing on-par with [30] in single-action scenarios.

Our second contribution is the extension of PROX dataset [15] with GT multi-action labels, thus completing the dataset to contain multi-action motion sequences, per-frame SMPL [26, 28] fittings and GT action labels.

Our third contribution is a thorough experimental evaluation on the challenging task of synthesizing arbitrary-length motion sequences with an arbitrary number of actions and comparisons to the state-of-the-art and strong baselines. In Sec. 4 we demonstrate that the proposed design choices result in significant improvements in action recognition when evaluating a pre-trained off-the-shelf action recognition approach [41] on generated sequences.

## 2. Related Work

Motion synthesis is a widely researched problem in graphics and computer vision [2, 13, 25, 30, 36, 39]. [3, 25, 27] condition a future sequence on a few past frames and rely on autoregressive models such as LSTMs [16]. These models are inherently incapable of modeling long-term dependencies

between frames and their ability to generate realistic motion degenerates for longer sequences. Works such as [10] jointly predict future frames in action classification frameworks in order to learn action-specific motion representation. Deep generative models such as Variational Autoencoders (VAE) [20] and Generative Adversarial Networks (GANs) [12] have made breakthroughs in synthesizing high-fidelity visual content such as images, videos and human motion [24, 31, 34]. For example, Barsoum *et al.* [2] proposes a GAN-based Seq2seq model while [39] proposes a GAN-based model using graph convolutional network that samples a latent vector from an observed prior. Methods based on VAEs for motion synthesis conditioned on action labels have been proposed in [13], while [30] further proposed generating the full body mesh rather than the 3D pose. Other tasks include motion transfer or video completion in which generating human videos is guided by motion sequences that are generated based on a semantic label [4, 40]. Similarly, controllable characters using user-defined controls are used to generate an image sequence of a given person [11], or to animate characters and avatars [22, 33, 38]. In a related task [23] visualize a given story and its characters by generating a sequence of images at the sentence level and [1] synthesize motion sequences conditioned on natural language. Another group of approaches predicts human motion constrained by scene context or human-object interaction [5, 8, 14, 36].

**Variational Autoencoder Transformer.** [30] proposed the currently state-of-the-art approach to synthesizes motion sequences conditioned on a given action label. The underlying generative method is a Transformer-based [35] VAE [20] which learns the latent distribution of the observed pose sequences. [30] is limited to single action sequence generation, where each sequence is generated in one shot in a non-causal way. This bounds the length of the sequence by the resources available on a machine, a limitation further amplified by quadratic complexity attention computation, which requires splitting long sequences into several parts thereby causing context fragmentation.

**Recurrent Transformers.** Transformer [35] is an expressive network architecture that relies on the self-attention mechanism for weighing different elements of the input. Several approaches have been proposed to alleviate the quadratic complexity needed to compute the attention [6, 7, 9, 32]. One prominent work is a recurrent approach based on formulating the attention using kernel functions [19], making it possible to calculate the attention in linear time.

## 3. Multi-action Recurrent Transformer Variational Autoencoder

Given an ordered set of action labels  $A = \{a_1, a_2, \dots, a_k\}$ , we propose a model that synthesizes

a motion sequence of 3D human meshes performing these actions. The ability to generate long sequences that are plausible and coherent requires allowing information to flow persistently across the temporal domain. To that end, we propose using a recurrent model with a hidden state that maintains compact yet long-range information to help it reason about the sequence.

**Approach.** In order to overcome the aforementioned limitations of [30], we propose to extend their approach to generate sequences of an arbitrary length conditioned on an arbitrary number of action labels. An overview of the approach is shown in Fig. 1. The input during training is a sequence consisting of multiple possibly overlapping actions and an ordered set of corresponding action labels. The transformer network additionally receives its last hidden state as part of the input, and outputs an encoding of the input representing the current pose and distribution parameters  $\mu$  and  $\sigma$ . We save the distribution parameters corresponding to the last frame of every action label, since we hypothesize that the encoder learns a better representation as it accumulates more knowledge of an action sub-sequence. We sample  $k$  latent vectors  $z_i \in \mathbb{R}^d$  each corresponding to a sub-sequence of one action. In order to guarantee a continuous and uniform representation for the entire sequence, especially at the cross points where action  $a_i$  ends and  $a_{i+1}$  starts, we stack the latent vectors of all actions. In the decoding phase, we do not rely on the timestamps to decide the number of frames synthesized for each action, and instead, we let the decoder reason about this number. This makes it possible to train the decoder in a weakly supervised manner, as it learns to relate an action-specific latent vector to a generated sub-sequence, but that also entails a limitation where during inference, it is not straightforward to specify the length of the sub-sequence corresponding to an action.

## 4. Experiments

**Dataset.** Existing motion datasets such as Human-Act12 [13] and Human3.6M [18] consist of single-action sequences, making it hard to train and evaluate the proposed model. We thus extended PROX dataset [15] which originally contains 3D pose annotations and SMPL [26, 28] fittings by adding GT action labels. We identified 50 action labels in 60 various recorded scenarios in 12 different 3D scenes in which 20 subjects are interacting with the scene. The number of frames is 100K and we split the dataset into training and test sets, where the training set consists of 27 scenes with subjects performing actions in indoor environments such as in the library, cafeteria and office. The test set consists of 3 scenes, some of which have only ground truth (GT) action annotations without SMPL fittings. The scenes consist of 1000 – 3000 frames and the actions performed can be intricate and are not repetitive or simple in nature.

**Baseline.** Since we are the first to address multi-action motion synthesis, we create a strong baseline by extending single-action ACTOR [30] approach to multi-actions. For that, we fix the maximum number of actions to be  $M$ , and divide latent vector  $z$  into  $M$  equal-sized sub-vectors each of size  $z_i \in \mathbb{R}^{[d/M]}$  to encode one action sub-sequence. We call this approach ACTOR-Multi.

**Standard evaluation metrics.** We follow [13] and train an action recognition model [41] on the GT training data and evaluate on the synthesized sequences using Fréchet Inception Distance (*FID*), *accuracy* of predicted action labels, *multimodality* and *diversity*, where the latter two refer to the variance between the different actions classes and variance within the same class.

**Semantic consistency metrics.** We evaluate semantic consistency between GT conditioning labels and synthesized sequences. Specifically, we calculate the proportion of sequences whose conditioning labels match the labels of the closest sequence in GT training set. We calculate cross-product distances between each 3D pose in generated sequence and each 3D pose in GT sequence, and find the sequence-to-sequence matching using the Hungarian algorithm.

### 4.1. Evaluation

**Using standard metrics.** Tab. 1 shows the performance of RTVAE-Multi and comparison to ACTOR-Multi, while varying the sequence length (60, 80, 120) for sequences with at most 8 actions. *tr* denotes training set, *gen* refers to synthesized sequences, and *gt* denotes GT data. RTVAE-Multi significantly improves over ACTOR-Multi, with a performance gap that rises with an increased sequence length. This is due to small capacity of the latent vector which does not suffice to represent significant variations in actions when synthesizing multi-action sequences. Tab. 2 provides direct comparison to [30] in single-action use case, with RTVAE-Single performing on-par with ACTOR [30].

**Using semantic consistency metrics.** RTVAE-Multi significantly outperforms ACTOR-Multi when fixing number of actions to four and varying the sequence length (2 (a)). RTVAE-Single matches ACTOR [30] in single-action use cases (2 (b)). Increasing the number of actions for a fixed sequence length degrades RTVAE-Multi accuracy much more gracefully compared to ACTOR-Multi (2 (c)).

**Qualitative evaluation.** Sample multi-action generated results are shown in Fig. 3. We observe continuity in the spatial and temporal domain.

### 4.2. Ablative studies

Ablative studies performed using *standard metrics* (Tab. 3) and *semantic consistency metrics* (Tab. 4).

**Last frame distribution parameters.** RTVAE-Multi uses distribution parameters of the last frame of each action la-

seq. length	method	accuracy <sub>tr<sub>gen</sub></sub>	accuracy <sub>tr<sub>gen</sub></sub> ↑	accuracy <sub>test<sub>gen</sub></sub> ↑	FID <sub>tr<sub>gen</sub></sub> ↓	diversity <sub>tr<sub>gen</sub></sub>	multimodality <sub>tr<sub>gen</sub></sub>	diversity <sub>tr<sub>gen</sub></sub>	multimodality <sub>tr<sub>gen</sub></sub>
60	ACTOR-Multi	98.1	74.3 ± 1.26	60.0 ± 6.1	680 ± 11.46	86.57	26.96	65.67 ± 0.51	23.66 ± 0.46
	RTVAE-Multi		<b>82.8 ± 0.84</b>	<b>61.8 ± 2.62</b>	<b>590 ± 6.23</b>			73.0 ± 0.39	20.5 ± 0.36
80	ACTOR-Multi	97.1	56.2 ± 0.92	41.5 ± 1.54	1197.4 ± 33.43	89.2	35.16	67.58 ± 1.32	26.69 ± 0.35
	RTVAE-Multi		<b>74.3 ± 1.11</b>	<b>49.4 ± 1.72</b>	<b>996 ± 23.12</b>			75.5 ± 0.64	25.85 ± 0.39
120	ACTOR-Multi	99.1	38.3 ± 0.47	27.9 ± 3.07	2938.44 ± 35.6	130.47	47.86	97.24 ± 0.82	38.95 ± 0.51
	RTVAE-Multi		<b>50.8 ± 0.71</b>	<b>37.3 ± 1.49</b>	<b>2584 ± 9.7</b>			100.14 ± 0.8	36.81 ± 0.54

Table 1. Evaluation on PROX dataset with a maximum of 8 actions/sequence, while varying sequence length.

seq. length	method	accuracy <sub>tr<sub>gen</sub></sub>	accuracy <sub>tr<sub>gen</sub></sub> ↑	accuracy <sub>test<sub>gen</sub></sub> ↑	FID <sub>tr<sub>gen</sub></sub> ↓	diversity <sub>tr<sub>gen</sub></sub>	multimodality <sub>tr<sub>gen</sub></sub>	diversity <sub>tr<sub>gen</sub></sub>	multimodality <sub>tr<sub>gen</sub></sub>
60	ACTOR [30]	97.3	<b>74.4 ± 0.69</b>	<b>65.8 ± 3.1</b>	733 ± 6.85	61.64	19.14	49.0 ± 0.45	14.14 ± 0.38
	RTVAE-Single		72.7 ± 0.75	64.8 ± 2.34	<b>725 ± 5.56</b>			48.4 ± 0.42	13.72 ± 0.3
80	ACTOR [30]	97.8	70.7 ± 1.0	61.0 ± 1.98	1696 ± 5.46	77.19	26.33	57.72 ± 0.37	26.33 ± 0.28
	RTVAE-Single		<b>71.1 ± 0.62</b>	<b>61.5 ± 2.87</b>	<b>1631 ± 9.84</b>			58.67 ± 0.52	17.77 ± 0.34
120	ACTOR [30]	98.4	64.2 ± 0.53	58.2 ± 2.5	5311 ± 14.73.6	105.22	39.31	81.84 ± 0.63	25.16 ± 0.57
	RTVAE-Single		<b>65.6 ± 1.08</b>	58.2 ± 2.69	<b>4885 ± 23.5</b>			81.75 ± 0.82	25.34 ± 0.56

Table 2. Evaluation on PROX dataset when using single action/sequence, while varying sequence length.

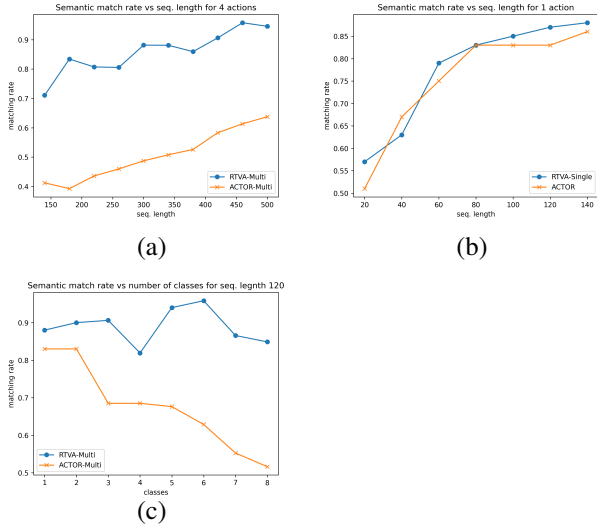


Figure 2. Match rate of action labels between synthesized and ground-truth samples while varying the sequence length for single action (a) and four actions (b), and varying the number of actions for constant sequence length (c).

method	accuracy <sub>tr<sub>gen</sub></sub> ↑	accuracy <sub>test<sub>gen</sub></sub> ↑	FID <sub>tr<sub>gen</sub></sub> ↓
RTVAE-Multi	<b>74.3 ± 1.11</b>	<b>49.4 ± 1.72</b>	<b>996 ± 23.12</b>
Average stats	72.8 ± 0.47	46.9 ± 1.98	1129 ± 9.56
All diff-latent	62.0 ± 0.4	38.1 ± 1.35	1249.86 ± 4.54
W/o look-b.-a.	29.1 ± 1.36	16.5 ± 2.0	2161.42 ± 21.46

Table 3. Ablative evaluations on PROX using *standard metrics* (max 80 frames- and max 8 actions per sequence).

bel. We evaluate the performance when using averaged parameters over all frames in an action sub-sequence. We observe deteriorated performance in this case, indicating that the learned distribution becomes better as the model accumulates more context information about the sequence (see *Average stats* in Tab. 3 and 4).

**Different latent vector per frame.** Using different latent vector for each generated frame also adversely affects performance. This indicates that a single latent vector representing the whole action sub-sequence provides a smoother frame generation, as similar frames should be close together in the latent space (see *All diff-latent* in Tab. 3 and 4).




Figure 3. Sample motion sequences synthesized by RTVAE-Multi.

**Look-back-ahead.** Stacking all latent vectors during the decoding phase is necessary for the synthesized sequence to avoid discontinuity between different action sub-sequences. We call it *look-back-ahead* strategy as it makes the decoder aware of both previous and future frames while reasoning about current frame. Using a separate latent vector from only the current frame results in a dramatic performance drop (see *W/o look-b.-a.* in Tab. 3 and 4).

method	correctly matched sequences %
RTVAE-Multi	<b>78.74</b>
Average stats	67.44
All diff-latent	67.77
W/o look-back-ahead	28.8

Table 4. Ablative evaluations on PROX using *semantic consistency metrics* (max 80 frames- and max 8 actions per sequence).

## 5. Conclusion

We addressed a challenging problem of synthesizing multi-action human motion sequences. We proposed a novel spatio-temporal formulation combining the expressiveness and efficiency of Recurrent Transformers with generative richness of conditional VAEs. The proposed approach generalizes to an arbitrary number of actions and frames per generated sequence, with space and time requirements growing linearly in the number of frames. Experimental evaluation showed significant improvements in multi-action motion synthesis over the state of the art.

## References

- [1] Chaitanya Ahuja and Louis-Philippe Morency. Language2pose: Natural language grounded pose forecasting. In *2019 International Conference on 3D Vision (3DV)*, pages 719–728. IEEE, 2019. 2
- [2] Emad Barsoum, John Kender, and Zicheng Liu. Hpgan: Probabilistic 3d human motion prediction via gan. 2018 ieeecv. In *CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1499–149909, 2017. 2
- [3] Judith Butepage, Michael J Black, Danica Kragic, and Hedvig Kjellstrom. Deep representation learning for human motion prediction and classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6158–6166, 2017. 2
- [4] Haoye Cai, Chunyan Bai, Yu-Wing Tai, and Chi-Keung Tang. Deep video generation, prediction and completion of human action sequences. In *Proceedings of the European conference on computer vision (ECCV)*, pages 366–382, 2018. 1, 2
- [5] Zhe Cao, Hang Gao, Karttikeya Mangalam, Qi-Zhi Cai, Minh Vo, and Jitendra Malik. Long-term human motion prediction with scene context. In *European Conference on Computer Vision*, pages 387–404. Springer, 2020. 2
- [6] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019. 1, 2
- [7] Krzysztof Choromanski, Valerii Likhoshesterov, Davidohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*, 2020. 1, 2
- [8] Enric Corona, Albert Pumarola, Guillem Alenya, and Francesc Moreno-Noguer. Context-aware human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6992–7001, 2020. 2
- [9] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019. 1, 2
- [10] Ali Diba, Vivek Sharma, Luc Van Gool, and Rainer Stiefelhagen. Dynamonet: Dynamic action and motion network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6192–6201, 2019. 2
- [11] Oran Gafni, Lior Wolf, and Yaniv Taigman. Vid2game: Controllable characters extracted from real-world videos. *arXiv preprint arXiv:1904.08379*, 2019. 2
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [13] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2021–2029, 2020. 1, 2, 3
- [14] Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael J Black. Stochastic scene-aware motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11374–11384, 2021. 1, 2
- [15] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J. Black. Resolving 3D human pose ambiguities with 3D scene constraints. In *International Conference on Computer Vision*, pages 2282–2292, Oct. 2019. 1, 2, 3
- [16] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 1, 2
- [17] Daniel Holden, Taku Komura, and Jun Saito. Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017. 1
- [18] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, jul 2014. 3
- [19] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, pages 5156–5165. PMLR, 2020. 2
- [20] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1, 2
- [21] Jehee Lee, Jinxiang Chai, Paul SA Reitsma, Jessica K Hodgins, and Nancy S Pollard. Interactive control of avatars animated with human motion data. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 491–500, 2002. 1
- [22] Jehee Lee and Kang Hoon Lee. Precomputing avatar behavior from human motion data. *Graphical models*, 68(2):158–174, 2006. 1, 2
- [23] Yitong Li, Zhe Gan, Yelong Shen, Jingjing Liu, Yu Cheng, Yuexin Wu, Lawrence Carin, David Carlson, and Jianfeng Gao. Storygan: A sequential conditional gan for story visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6329–6338, 2019. 1, 2
- [24] Zhiheng Li, Martin Renqiang Min, Kai Li, and Chenliang Xu. Style2i: Toward compositional and high-fidelity text-to-image synthesis. *arXiv preprint arXiv:2203.15799*, 2022. 2
- [25] Zimo Li, Yi Zhou, Shuangjiu Xiao, Chong He, Zeng Huang, and Hao Li. Auto-conditioned recurrent networks for extended complex human motion synthesis. *arXiv preprint arXiv:1707.05363*, 2017. 2
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. 1, 2, 3
- [27] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2891–2900, 2017. 2
- [28] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and

- Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. [1](#), [2](#), [3](#)
- [29] Xue Bin Peng, Angjoo Kanazawa, Jitendra Malik, Pieter Abbeel, and Sergey Levine. Sfv: Reinforcement learning of physical skills from videos. *ACM Trans. Graph.*, 37(6), Nov. 2018. [1](#)
- [30] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10985–10995, 2021. [1](#), [2](#), [3](#), [4](#)
- [31] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *International conference on machine learning*, pages 1060–1069. PMLR, 2016. [2](#)
- [32] Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. Efficient content-based sparse attention with routing transformers. *Transactions of the Association for Computational Linguistics*, 9:53–68, 2021. [1](#), [2](#)
- [33] Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. Neural state machine for character-scene interactions. *ACM Trans. Graph.*, 38(6):209–1, 2019. [1](#), [2](#)
- [34] Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Xiao-Yuan Jing, Fei Wu, and Bingkun Bao. Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis. *arXiv preprint arXiv:2008.05865*, 2020. [2](#)
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [1](#), [2](#)
- [36] Jiashun Wang, Huazhe Xu, Jingwei Xu, Sifei Liu, and Xiaolong Wang. Synthesizing long-term 3d human motion and interaction in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9401–9411, 2021. [1](#), [2](#)
- [37] Yonatan Wexler, Eli Shechtman, and Michal Irani. Space-time video completion. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I. IEEE, 2004. [1](#)
- [38] Jungdam Won, Deepak Gopinath, and Jessica Hodgins. A scalable approach to control diverse behaviors for physically simulated characters. *ACM Transactions on Graphics (TOG)*, 39(4):33–1, 2020. [1](#), [2](#)
- [39] Sijie Yan, Zhizhong Li, Yuanjun Xiong, Huahan Yan, and Dahua Lin. Convolutional sequence generation for skeleton-based action synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4394–4402, 2019. [2](#)
- [40] Ceyuan Yang, Zhe Wang, Xinge Zhu, Chen Huang, Jianping Shi, and Dahua Lin. Pose guided human video generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 201–216, 2018. [2](#)
- [41] Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875*, 2017. [2](#), [3](#)