

MotionAug: Augmentation with Physical Correction for Human Motion Prediction

Takahiro Maeda Norimichi Ukita
 Toyota Technological Institute, Japan
 {sd21601, ukita}@toyota-ti.ac.jp

Abstract

This paper presents a motion data augmentation scheme incorporating motion synthesis encouraging diversity and motion correction imposing physical plausibility. This motion synthesis consists of our modified Variational AutoEncoder (VAE) and Inverse Kinematics (IK). In this VAE, our proposed sampling-near-samples method generates various valid motions even with insufficient training motion data. Our IK-based motion synthesis method allows us to generate a variety of motions semi-automatically. Since these two schemes generate unrealistic artifacts in the synthesized motions, our motion correction rectifies them. This motion correction scheme consists of imitation learning with physics simulation and subsequent motion debiasing. For this imitation learning, we propose the PD-residual force that significantly accelerates the training process. Furthermore, our motion debiasing successfully offsets the motion bias induced by imitation learning to maximize the effect of augmentation. As a result, our method outperforms previous noise-based motion augmentation methods by a large margin on both Recurrent Neural Network-based and Graph Convolutional Network-based human motion prediction models. The code is available at <https://github.com/meaten/MotionAug>.

1. Introduction

Human motion prediction, which forecasts future body poses based on past poses, is a crucial technique for human-robot interaction [22,23,28,29,41], autonomous driving [30], VR/AR applications [18], performance capture [8,32,33,43,44], etc. However, these applications are limited because of the lack of training motion data, which results in low prediction accuracy. This data insufficiency is caused by the enormous cost of motion data acquisition, such as motion capture equipment, recordings, post-processing, and denoising.

Such data insufficiency can be alleviated by Data

Augmentation (DA), for example, image recognition [25,39]. Compared with image data augmentation, however, DA for motion data is hard to address because simple numerical transformations (*e.g.*, additive noise) may generate physically-implausible motions such as too high velocities or floating motions.

This paper presents a novel motion data augmentation approach, including motion synthesis and motion correction. Our motion synthesis uses Variational AutoEncoder (VAE) [24] to exploit a training data distribution and Inverse Kinematics (IK) to exploit human knowledge, shown in “Motion Synthesis” in Fig. 1. Although most of our synthesized motions are physically-plausible, we observed some of them have unrealistic artifacts, which lead to the low accuracy of human motion prediction. These artifacts are corrected with our proposed motion correction method. This correction method uses (i) imitation learning with physics simulation to rectify these artifacts and (ii) subsequent motion debiasing to offset biases imposed by the mismatch between the bodies of a human and the character during imitation learning (“Motion Correction” in Fig. 1). Our contributions for motion diversity and physical plausibility are as follows:

1. **VAE-based human-motion synthesis:** Our generative model with adversarial training in sequencewise and framewise, and sampling-near-samples can generate plausible motions even with insufficient motion data.
2. **IK-based human-motion synthesis:** Compared with annotating IK target points in all frames as the standard IK does, our method requires less effort because only a target sampling space for a keyframe is manually given.
3. **PD-residual force:** We propose the PD-residual force that accelerates the training of imitation learning in a physics simulator to rectify the physical implausibility of synthesized motions.

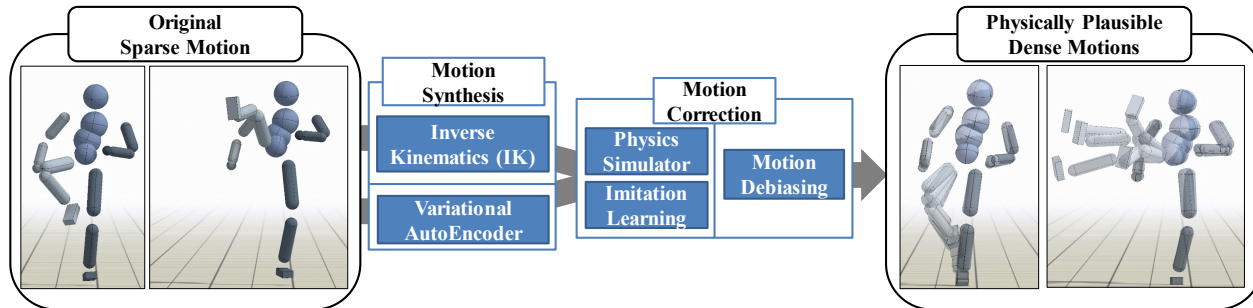


Figure 1. **Overview of our proposed motion data augmentation.** Original motions are augmented independently using our VAE- and IK-based syntheses. We also propose (i) motion correction where the synthesized motions are modified to be physically-plausible using imitation learning, a physics simulator, and (ii) subsequent motion debiasing.

4. **Motion debiasing:** Our motion debiasing successfully offsets the motion bias induced by the imitation learning to maximize the effect of our data augmentation.

2. Related Work

2.1. Human Motion Prediction and Augmentation

From the releases of large-scale motion capture sequence datasets [10, 20], many deep learning-based human motion prediction methods were proposed. Most approaches [6, 12, 13, 15, 21, 32, 35, 45, 51] are built upon Recurrent Neural Networks (RNN) to model sequence-to-sequence relationships between past and future motions. Recently, Graph Convolutional Network (GCN)-based approaches [31] achieved a better performance than RNN-based models by encoding motions with Discrete Cosine Transform. Along with improving model architectures, the stochasticity of human motion is addressed by using generative models such as Generative Adversarial Networks (GAN) [3, 26], VAE [1, 46], and Flow-based models [47].

These approaches assume a large-scale motion dataset that is too expensive to obtain in real-world tasks. Despite this difficulty, motion data augmentation approaches are almost ignored. Fragkiadaki *et al.* [12] proposed corrupting input motions with zero-mean Gaussian noise for motion data augmentation. While this simple additive noise improves the variety of input motions, the augmented motions might lose motion contexts and defy the laws of physics.

2.2. Data augmentation with generative models

In image classification tasks, generative models such as GAN are used for data augmentation by generating within-class images [7, 19, 42]. This approach is applicable to other tasks, including image segmentation [38] and person re-identification [50]. However, generative models for human motion might synthesize

several kinds of physically-implausible motions because it is difficult to learn such physical plausibility from a limited number of training motion data, especially in data insufficient settings.

2.3. Inverse Kinematics (IK)

IK modifies the pose of a whole body so that key points in the body reach their target positions. IK can also modify a motion by providing the target positions in all frames of the motion [14]. Although IK can significantly modify each pose and potentially be helpful for motion augmentation, it is impractical to manually annotate the target positions in all frames included in a training dataset for augmenting all motions [5, 17].

2.4. Motion Synthesis with Physics Simulation

Motion prediction models should be trained on physically-plausible motions for better accuracy and reliability. Therefore, physics simulation might improve the quality of augmented motions. Recent deep reinforcement learning enables a physically-simulated character to imitate various motions [4, 27, 36, 48, 49]. However, these methods often require more than one day to converge for imitating only one motion. To incorporate this, we need to reduce the vast computational cost to augment a large number of motions.

3. Proposed Motion Augmentation

We propose two independent motion synthesis approaches with VAE and IK, described in Secs. 3.1 and 3.2, respectively. Furthermore, a method for motion correction is also proposed for rectifying the artifacts of synthesized motions in Sec. 3.3. Finally, we propose motion debiasing to offset the bias imposed by dynamic mismatch, as presented in Sec. 3.4.

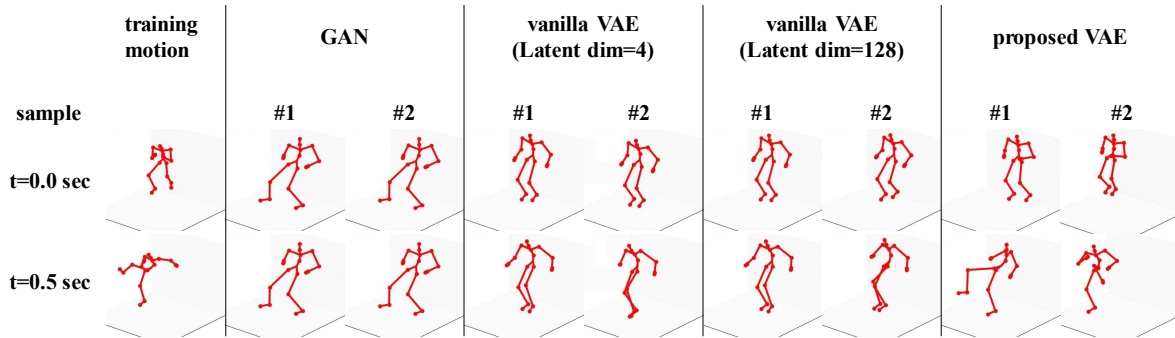


Figure 2. **Synthesized motions from GAN, vanilla VAE, and our VAE.** Despite the dynamic training motions, GAN produces static motions due to data insufficiency. The vanilla VAE produces non-diverse motions regardless of the dimension of the latent space. Our VAE with adversarial training and sampling-near-samples successfully synthesizes dynamic motions different from training motions.




Figure 3. **Proposed VAE-based network with adversarial training.** The synthesized motions and training motions are discriminated framewise and sequencewise.

3.1. DA with VAE

Although GAN is widely used as a generative model, we found that, for motion synthesis, GAN often produces only static motions where all poses are almost identical due to data insufficiency and training instability of GAN (i.e., mode collapse). Instead, we propose a VAE-based model that is free from these problems. Our VAE-based model described below successfully generates various motions despite insufficient data using adversarial training and sampling-near-samples, as shown in Fig. 2.

Adversarial training: Our proposed network is shown in Fig. 3. The encoder produces a mean μ and a variance σ^2 in the latent space from an input motion $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ where each \mathbf{x}_t denotes a pose vector in t -th frame. The latent representation \mathbf{z} is sampled from the normal distribution $\mathcal{N}(\mu, \sigma^2)$. The decoder reconstructs a motion $\hat{\mathbf{X}}$ from \mathbf{z} . Frame-wise and sequence-wise discriminators (denoted by Dis^f and Dis^s , respectively, in Fig. 3) discriminate \mathbf{X} from $\hat{\mathbf{X}}$ for improving $\hat{\mathbf{X}}$ in terms of the fidelity of poses and motion dynamics. We validated that this VAE with adversarial training can suppress mode collapse and generate more realistic motions than the vanilla VAE.

Sampling-near-samples in the latent space: In

the inference of the vanilla VAE, a latent representation \mathbf{z} is sampled from a normal distribution with zero mean and unit variance $\mathcal{N}(\mathbf{0}, \mathbf{I})$. This normal distribution should be represented well by all training samples in the latent space for the better quality of synthesized motions, as shown in Fig. 4a.

However, the dimension of the latent space should not be set too low so that insufficient data could cover the whole normal distribution because too low-dimensional latent space has a too narrow bottleneck and generates inaccurate motions that lack motion details. On the other hand, a high-dimensional representation leads to the sparsity of training data, making it difficult to sample realistic data from the learned regions, as shown in Fig. 4b. Therefore, we have a tradeoff between motion details and sampling easiness.

To solve this tradeoff, we propose a novel sampling method robust to sparsity, specifically sampling \mathbf{z} from only learned regions that are appropriately represented by training data in the latent space, as shown in Fig. 4c. In this method, each motion in the training data is encoded into mean μ and variance σ^2 . We apply k-means clustering to all training motions based on μ to make n_c clusters. Given $\bar{\mu}$ and $\bar{\sigma}^2$ that are respectively the mean of μ and the mean of σ^2 over randomly-sampled n_s training motions from each cluster, the latent representation \mathbf{z} is drawn from $\mathcal{N}(\bar{\mu}, \bar{\sigma}^2)$, and \mathbf{z} is fed into the decoder for generating \mathbf{X}_{aug} , as expressed by $\mathbf{X}_{\text{aug}} = \text{Dec}(\mathbf{z})$ and $\mathbf{z} \sim \mathcal{N}(\bar{\mu}, \bar{\sigma}^2)$. We sampled motion subsets from each cluster for efficiency.

3.2. DA with IK

The IK-based motion editing needs target positions in all frames of a motion. To achieve this semi-automatically, we present an effortless IK-based motion synthesis that only requires a user to provide a target sampling space \mathbb{P} for the pose of the keyframe \mathbf{x}_{key} on

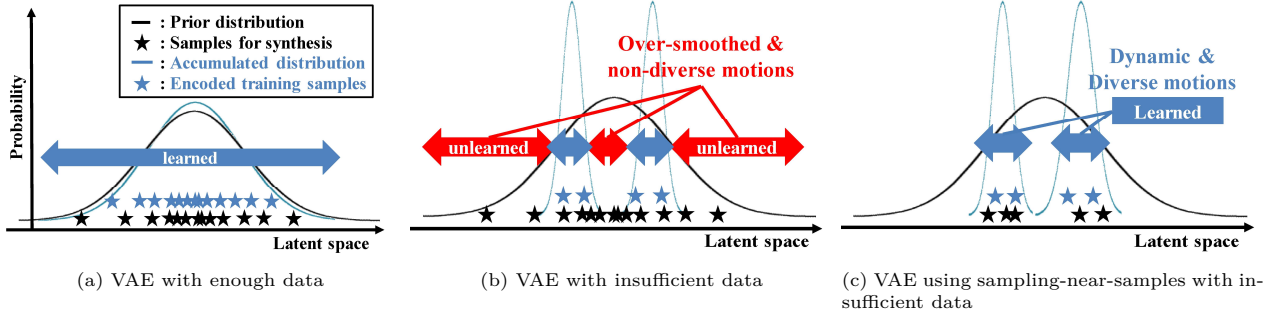


Figure 4. **Our sampling-near-samples for insufficient data.** The prior distribution matches the accumulated distribution that aggregates the distributions of encoded training data with enough data (a). Synthesized motions are mostly sampled from learned regions (blue arrows). With insufficient data (b), the accumulated distribution (learned regions) gets sparse, and the prior distribution often samples from the unlearned regions (red arrows), which leads to over-smoothed and non-diverse motions. We propose sampling-near-samples (c), which samples only from learned regions by sampling latent representations using the clusters of training motions.

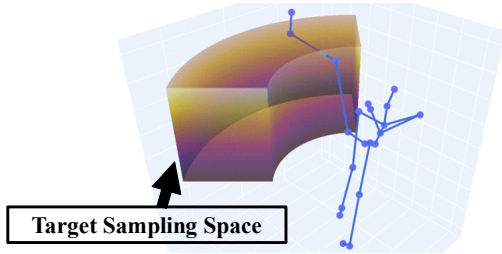


Figure 5. **Target sampling space for the action class *kick*.** The target space is a fan-shaped one that a foot end-effector may reach. We uniformly sample targets $\mathbf{p}_{t_{\text{key}}}^{\text{sample}}$ from this space for the keyframe t_{key} .

each action class. Examples of a kick class are shown in Figs. 5 and 6. The user determines the target sampling space as shown in Fig. 5. Then, the keyframe for the kick class is defined as the frame where a kicking foot reaches the farthest position from the body.

Given the sampling space for the keyframe, an IK target position $\mathbf{p}_{t_{\text{key}}}^{\text{sample}} \in \mathbb{P}$ for the keyframe is randomly sampled. Target positions $\mathbf{p}_t^{\text{target}}$ for all frames are determined by propagating the difference between $\mathbf{p}_{t_{\text{key}}}^{\text{sample}}$ and the end-effector position at the keyframe $\mathbf{p}_{t_{\text{key}}}$ to backward and forward, in a linearly-decreasing manner, as shown in Fig. 6 and expressed as follows:

$$\begin{aligned} \mathbf{p}^{\text{diff}} &= \mathbf{p}_{t_{\text{key}}}^{\text{sample}} - \mathbf{p}_{t_{\text{key}}} \\ \mathbf{p}_t^{\text{target}} &= \mathbf{p}_t + \mathbf{p}^{\text{diff}} \cdot f(t_{\text{key}}, t) \\ f(t_{\text{key}}, t) &= \begin{cases} \frac{t}{t_{\text{key}}} & \text{if } t \leq t_{\text{key}} \\ \frac{T-t}{T-t_{\text{key}}} & \text{if } t > t_{\text{key}} \end{cases} \\ \mathbf{X}_{\text{aug}} &= \{\text{IK}(\mathbf{x}_1, \mathbf{p}_1^{\text{target}}), \dots, \text{IK}(\mathbf{x}_T, \mathbf{p}_T^{\text{target}})\} \end{aligned}$$

where $\hat{\mathbf{x}} = \text{IK}(\mathbf{x}, \mathbf{p})$ is an IK function, and \mathbf{x} and \mathbf{p} denote a pose vector and a 3D position, respectively. We apply IK with automatically obtained targets $\mathbf{p}_t^{\text{target}}$ to all frames for obtaining a synthesized motion \mathbf{X}_{aug} where the end-effector smoothly reaches $\mathbf{p}_{t_{\text{key}}}^{\text{sample}}$.

3.3. Motion Correction with Imitation Learning using Physics Simulation

Although most synthesized motions generated by our method are physically realistic, some of them are not. For example, footskating by VAE, mutual penetrations between body parts by IK, and unstable poses by VAE and IK are empirically observed, as shown in the middle column of Fig. 7.

DeepMimic [36] is an imitation learning scheme that allows a physically-simulated character to mimic various motions. Given a goal motion (*e.g.*, motion measured by a motion capture system), imitation learning trains a policy that modifies a character pose at $t+1$ from its body status at t so that the sequence of the modified poses gets close to the goal motion. Then, to physically control the character toward the modified pose at each moment, a Proportional-Differential (PD) controller suggests torques given to the character at t . We can obtain the modified motion where the physical character performs by repeating this scheme.

While, in DeepMimic [36], the motion is modified for compensating dynamic mismatch (*i.e.*, the difference between the bodies of the goal motion and the character), the goal motion is already physically-plausible because a motion capture system measures it. On the other hand, we apply this imitation learning to rectify a physically-implausible motion produced by our method. This physically-implausible motion makes our problem more challenging because the policy must rec-




Figure 6. **Overview of our sequential IK scheme.** Given an IK target and the body pose in the keyframe, body poses in all other frames are automatically calculated by IK. IK target positions are automatically determined by propagating the positional difference on the keyframe in a linearly decreasing manner.




Figure 7. **Examples of our motion correction.** Upper: Legs penetrate each other. Lower: Unstable pose is observed. Synthesized motions (b) are generated from original motions (a) and then rectified (c).

tify physical implausibility and the dynamic mismatch. To cope with this more challenging problem, our imitation learning scheme employs Residual Force Control (RFC) [48] maintaining physical stability such as fall prevention. With RFC, learnable additional external forces given to the root joint of the character achieve physical stability. The rectified motion of the character is still physically-plausible because additional external forces are minimized in training while the pose similarity between a goal motion and the character is maximized.

While DeepMimic using RFC allows us to generate stable motions, the convergence of the training process usually takes more than one day for rectifying one motion with several CPU threads. This cost is a critical problem when we augment a large number of motions. The dominant cost in convergence time is on reinforcement learning of the policy network that requires exploration in the policy action space, specifically a target character pose and an additional external force. Al-

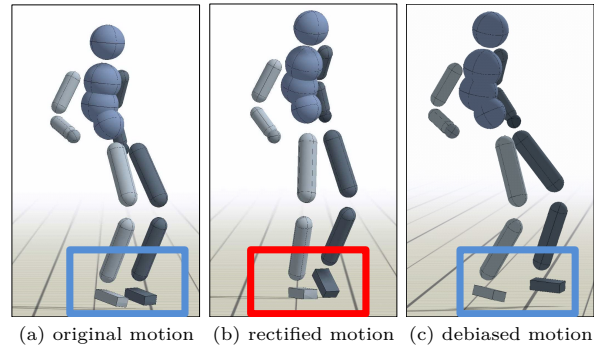


Figure 8. **Examples of our motion debiasing.** The motion bias is introduced by the dynamic mismatch on the *kick* class motion.

though the policy network learns additional external forces from scratch, the learned forces just reduce the positional difference between the character and goal motion. Based on this observation, we propose the PD-residual force that calculates additional external forces with the PD controller based on the positional difference between the character and goal motion. This simple modification allows us to omit the learning of external forces and significantly shorten the training process by reducing the dimensionality of the policy action space to explore.

3.4. Motion Debiasing

Our imitation learning scheme explained in Sec. 3.3 can rectify synthesized motions to be physically-plausible. However, a prediction model trained with these rectified motions cannot entirely reduce the prediction error due to the motion bias introduced by the dynamic mismatch [48] during imitation learning. The dynamic mismatch is the body difference between “real humans with hundreds of bones and muscles and deformable skins” and “the simulated character with torque-actuated joints and rigid body surfaces.” Due to this difference, the simulated character fails to fully imitate the motions even if they are physically-plausible,

especially motions with fine footwork.

To alleviate the motion bias, we propose motion debiasing to offset the biases imposed by this dynamic mismatch. We construct training pairs that are original motions and their modified motions by imitation learning as Sec. 3.3. These pairs only contain the bias imposed by the dynamic mismatch. We propose a simple motion debiasing model of several fully connected layers that map the biased to unbiased motions in frame-wise. We apply this motion debiasing to the rectified motions of synthesized motions by our VAE- and IK-based syntheses. As a result, we obtained the debiased various physically-plausible motions and further improved prediction accuracy, as shown in Fig. 8.

4. Experiments

Our experiments consist of four parts: Ablation study on our VAE-based method. The effects of several components of our proposed method are validated by the physical and contextual closeness between synthesized and test motions (Sec. 4.1). Convergence time comparison on imitation learning with our PD-residual force (Sec. 4.2). Performance evaluation on motion prediction with different augmentations (Sec. 4.3). Augmentation comparison to the previous method DeepMimic [36] (Sec. 4.4).

Dataset: Our experiments were conducted on HDM05 Motion Database [34]. HDM05 is a relatively small and challenging dataset with dynamic motions compared to other standard benchmarks such as Human3.6M [20]. We tested our method with five-fold cross-validation where our models were trained on the motion sets of four actors and tested on one of the last actor. Motion sets of *punch*, *kick*, and *walk* action classes were resampled to 30Hz and used for the experiments. The number of synthesized motions from our VAE and IK is ten times larger than the train set.

4.1. Motion Synthesis by VAE

The effectiveness of our VAE-based motion synthesis is validated by ablation. For comparison, a GAN-based method is also evaluated.

Implementation Details: All encoders, decoders, and discriminators consist of 256-D LSTM cells and one fully connected layer to output poses. The dimension of a latent space is 128 for VAE. The noise dimension for GAN is also 128. We used the SGD optimizer to train models for 20,000 epochs. The number of samples to take mean $n_s = 2$ and clusters $n_c = 3$ is used for sampling-near-samples.

Metrics: The quality of synthesized motions is evaluated with two metrics: the minimum Dynamic Time Warping (DTW) [37] distance and the Maximum Mean

Table 1. Quantitative evaluation of augmented motions.

	Min DTW	MMD
Original Motions	2.92	0.00
GAN	2.90	4.28
VAE	2.73	1.94
VAE + adv training	2.71	1.37
VAE + sampl. near samples	2.72	0.85
VAE + both (proposed)	2.70	0.20

Discrepancy (MMD) [16]. The minimum DTW distance is a DTW distance between a test motion and the synthesized motion (training set for original motions) closest to it. For DTW, frame-wise distances are calculated based on the Euclidean distance in the Euler angle in the radian scale. The sum over all joints except a root joint is evaluated. MMD measures the distribution distance between the test and synthesized motions. Minimum DTW distance and MMD measure how the synthesized motions are close to the test motions physically and contextually. The lower score is better in both metrics.

Results: Table 1 shows that the proposed VAE-based method with adversarial training and sampling-near-samples performs best in both metrics. Meanwhile, a GAN-based method fails to decrease the minimum DTW distance and gets the highest MMD because the training dataset is too small for GAN to learn various patterns and falls mode collapse problem.

4.2. Convergence Time Comparison on Imitation Learning

The convergence times of RFC [48] and our imitation learning with PD-residual force are evaluated.

Implementation details: We use Bullet Physics [9] as the physics engine. We build the humanoid model from the skeleton of the MDM05 Motion Database, which has 52 DoF and 16 rigid bodies. We use the same reward function r_t as RFC [48]. We train both methods on one kick motion for 100 hours with five threads of Intel® Xeon® Gold 6248 CPU.

Metrics: We evaluate the training time vs. normalized reward. The normalized reward is calculated with the obtained reward over the maximum reward on one episode. The normalized reward is calculated based on the motion similarity r_t^{im} and its max value $r_t^{\text{im,max}}$.

$$R_{\text{norm}} = \frac{1}{T} \sum_{t=0}^T \frac{r_t^{\text{im}}}{r_t^{\text{im,max}}}$$

Results: The results are shown in Fig. 9. While RFC requires 30 hours for convergence, our method converges around 9 hours thanks to the dimensionality

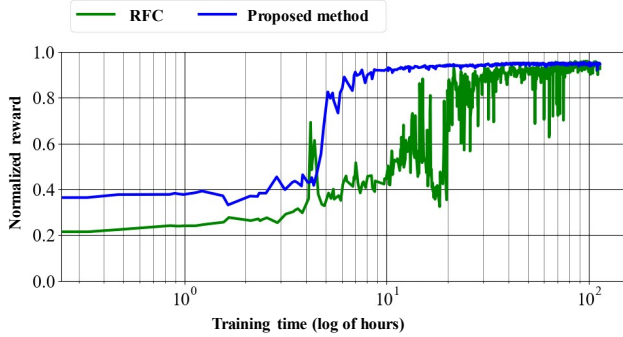


Figure 9. Normalized rewards vs. training time in logarithmic scale.

reduction by our PD-residual force. Furthermore, our method is stabler than RFC in terms of the reward curve because the physical stability is kept throughout the training process.

4.3. Motion Prediction with DA

Implementation details: We use the same parameters for VAE as Sec. 4.1. FABRIK [2] is used as the IK algorithm. The keyframes for our IK synthesis are set as the frame when a foot joint reaches the furthest position from the root joint for all action classes. We chose the foot joint to modify the motion for *punch* class because more motion variation is observed on the foot joint than a hand joint. The IK target sampling spaces are set as fan-shapes shown in Fig. 5 for all action classes. The parameters for fan-shapes are determined based on the position $\mathbf{p}_{\text{key}} = (r, h, \theta)$ of the foot joint on the keyframe in cylindrical coordinates. For *punch*, *kick*, *walk* classes, the IK target positions are sampled from $([0.5, 2.0]r, [1.0, 1.0]h, [-1.7, 1.7] + \theta)$, $([0.8, 1.2]r, [0.8, 1.2]h, [-0.785, 0.785] + \theta)$, and $([0.5, 2.0]r, [1.0, 1.0]h, [-0.3, 0.3] + \theta)$ respectively. Our motion debiasing network is four 512-dim fully-connected layers with the ReLU activation to offset the bias framewise. We also temporally expand and shrink motion sequences in the range of 10% shorter and 10% longer as temporal data augmentation.

Prediction Model and Metrics: We use the heavily benchmarked RNN baseline [32] and the SOTA GCN-based model [31] to evaluate the effectiveness of our motion DA method on the human motion prediction task. We follow the standard evaluation protocol used in [12, 32], and report the Euclidean distance between the predicted and ground-truth joint angles in Euler representation. The reported errors in the radian scale are summed over all joints except a root joint and temporally averaged.

Results: In Tables 2 and 3, we show quantitative results for human motion prediction with data augmen-




Figure 10. Augmentation comparison to the base method, DeepMimic [36]. (a) shows an original motion. (b) and (c) show 100 augmented motions by DeepMimic with two rewards weighting $\{\alpha = 0.7, \beta = 0.3\}$ and $\{\alpha = 0.3, \beta = 0.7\}$ respectively. (d) shows 100 augmented motions by our method.

tation combinations. The prediction errors are shown on three timesteps (100, 200, 400ms) for three action classes (*punch*, *kick*, *walk*). The motion syntheses themselves (rows with no checkmark) often fail to decrease the prediction errors compared to “No Aug” because the motion prediction model learns unrealistic motions that are far from test motion data recorded in the real world. The motion syntheses with physical correction (rows with one checkmark) also fail because the prediction model learns biased motion data different from test motions. Our proposed motion data augmentation (rows with two checkmarks) achieved the lowest prediction error in all cases by a large margin compared to the previous method “Noise”.

4.4. Augmentation Comparison to Previous Method

We compared the augmentation capability of our method and the additional tasks of DeepMimic [36].

Experimental set up: We choose one *kick* class motion and independently augment it with DeepMimic, and our IK-based motion synthesis with motion correction. DeepMimic can also augment motions by training characters to solve additional tasks besides the original imitation. The used additional task is defined as DeepMimic’s *Strike* reward r_t^{strike} that rewards the character when the foot strikes randomly placed targets. The targets are randomly placed within the same target sampling space used in our IK-based motion synthesis. DeepMimic is tested in two rewards weighting $\{\alpha = 0.7, \beta = 0.3\}$ and $\{\alpha = 0.3, \beta = 0.7\}$ in the following equation:

$$r = \alpha r_t + \beta r_t^{\text{strike}} \quad (1)$$

Results: We show the results in Fig. 10. Augmented motions from DeepMimic have limited diversity in both reward weightings because the policy suffers from the

Table 2. Quantitative results of motion data augmentation on the RNN-based human motion prediction [32].

Methods			Prediction errors↓[rad] on each action class & timesteps [ms]								
	physical correction	motion debiasing	punch			kick			walk		
			100	200	400	100	200	400	100	200	400
No aug	-	-	1.58	2.2	2.59	1.28	1.89	2.45	0.74	1.15	1.49
Noise	-	-	1.57	2.19	2.57	1.26	1.84	2.35	0.72	1.12	1.45
VAE			1.48	2.17	2.71	1.28	1.90	2.46	0.68	1.11	1.55
IK			1.58	2.41	3.14	1.44	2.20	2.95	0.71	1.11	1.47
VAE & IK			1.56	2.38	3.03	1.21	1.78	2.25	0.67	1.07	1.43
VAE	✓		1.57	2.19	2.75	1.27	1.92	2.57	0.71	1.14	1.56
IK	✓		1.52	2.26	3.02	1.26	1.94	2.59	0.74	1.22	1.78
VAE & IK	✓		1.53	2.20	2.97	1.24	1.91	2.60	0.71	1.19	1.71
VAE	✓	✓	1.50	2.10	2.63	1.11	1.64	2.06	0.66	1.08	1.48
IK	✓	✓	1.39	1.93	2.45	1.05	1.52	1.85	0.57	0.91	1.22
VAE & IK	✓	✓	1.37	1.91	2.49	1.01	1.48	1.80	0.58	0.92	1.24

Table 3. Quantitative results of motion data augmentation on the SOTA GCN-based human motion prediction [31].

Methods			Prediction errors↓[rad] on each action class & timesteps [ms]								
	physical correction	motion debiasing	punch			kick			walk		
			100	200	400	100	200	400	100	200	400
No aug	-	-	1.31	1.87	2.33	1.08	1.68	2.26	0.52	0.88	1.24
Noise	-	-	1.31	1.90	2.35	1.06	1.65	2.25	0.52	0.87	1.21
VAE			1.28	1.88	2.34	1.06	1.63	2.17	0.52	0.91	1.28
IK			1.21	1.73	2.25	0.96	1.39	1.73	0.50	0.85	1.18
VAE & IK			1.22	1.81	2.29	0.95	1.38	1.71	0.49	0.85	1.20
VAE	✓		1.31	1.89	2.36	1.03	1.60	2.14	0.52	0.89	1.25
IK	✓		1.31	1.89	2.49	1.06	1.66	2.20	0.53	0.94	1.35
VAE & IK	✓		1.28	1.84	2.35	1.03	1.65	2.17	0.54	0.94	1.37
VAE	✓	✓	1.22	1.74	2.07	1.00	1.52	1.89	0.52	0.89	1.25
IK	✓	✓	1.27	1.79	2.24	0.92	1.35	1.65	0.48	0.81	1.11
VAE & IK	✓	✓	1.22	1.70	2.06	0.90	1.32	1.60	0.47	0.80	1.11

tradeoff between the imitation and the additional strike tasks. Although more weight on the strike reward slightly improves the diversity, the resulting motions lose the original motion details. On the other hand, our method produces diverse motions by dividing the augmentation to the synthesis and physical correction where the policy focuses only on the imitation task.

5. Limitations

Our motion augmentation has two limitations.

First, motion correction still takes several hours to rectify one motion, even with our proposal to accelerate the training. This cost makes it hard to apply our motion augmentation to more extensive motion prediction benchmarks such as Human3.6M [20] due to computational cost. However, the training time could be shortened by using meta-learning [11] for better policy initialization or the fast physics simulation environment accelerated with GPU rather than CPU.

Second, our motion augmentation is not immediately applicable to the partially-observed motion sequences, such as only observed upper body or 2D motion sequences, because our motion correction only accepts 3D motion sequences for a whole-body 3D character. We need to estimate missing joints or 2D-3D pose lifting [40] to apply our method to these situations.

6. Conclusion

This work presented a new human motion augmentation approach using VAE- and IK-based motion syntheses and motion correction with physics simulation. Experiments demonstrated that our augmentation outperformed previous methods because our VAE- and IK-based motion syntheses improve the diversity of training motion data, and our motion correction rectifies the unrealistic artifacts without motion biases. Our future work includes a new motion synthesis approach and faster motion correction based on meta-learning and GPU acceleration for larger-scale datasets.

References

- [1] Mohammad Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Lars Petersson, and Stephen Gould. A stochastic conditioning scheme for diverse human motion prediction. In *CVPR*, pages 5222–5231, 2020. [2](#)
- [2] Andreas Aristidou and Joan Lasenby. FABRIK: A fast, iterative solver for the inverse kinematics problem. *Graph. Models*, 73(5):243–260, Sept. 2011. [7](#)
- [3] Emad Barsoum, John Kender, and Zicheng Liu. HP-GAN: probabilistic 3d human motion prediction via GAN. In *CVPR Workshops*, pages 1418–1427, 2018. [2](#)
- [4] Kevin Bergamin, Simon Clavet, Daniel Holden, and James Richard Forbes. Drecon: data-driven responsive control of physics-based characters. *ACM Trans. Graph.*, 38(6):206:1–206:11, 2019. [2](#)
- [5] Schubert R. Carvalho, Ronan Boulic, and Daniel Thalmann. Interactive low-dimensional human motion synthesis by combining motion models and PIK. *Comput. Animat. Virtual Worlds*, 18(4-5):493–503, 2007. [2](#)
- [6] Hsu-Kuang Chiu, Ehsan Adeli, Borui Wang, De-An Huang, and Juan Carlos Niebles. Action-agnostic human pose forecasting. In *WACV*, pages 1423–1432, 2019. [2](#)
- [7] Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In *ICCV*, pages 6829–6839, 2019. [2](#)
- [8] Enric Corona, Albert Pumarola, Guillem Alenyà, and Francesc Moreno-Noguer. Context-aware human motion prediction. In *CVPR*, pages 6990–6999, 2020. [1](#)
- [9] Erwin Coumans. Bullet physics simulation. In *SIG-GRAPH*, page 7:1, 2015. [6](#)
- [10] Fernando De la Torre, Jessica Hodgins, Adam Bargteil, Xavier Martin, Justin Macey, Alex Collado, and Pep Beltran. Guide to the carnegie mellon university multimodal activity (cmu-mmact) database. 2009. [2](#)
- [11] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pages 1126–1135, 2017. [8](#)
- [12] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *ICCV*, pages 4346–4354, 2015. [2](#), [7](#)
- [13] Partha Ghosh, Jie Song, Emre Aksan, and Otmar Hilliges. Learning human motion models for long-term predictions. In *3DV*, pages 458–466, 2017. [2](#)
- [14] Michael Gleicher. Motion path editing. In John F. Hughes and Carlo H. Séquin, editors, *SI3D*, pages 195–202, 2001. [2](#)
- [15] Anand Gopalakrishnan, Ankur Arjun Mali, Dan Kifer, C. Lee Giles, and Alexander G. Ororbia II. A neural temporal model for human motion prediction. In *CVPR*, pages 12116–12125, 2019. [2](#)
- [16] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander J. Smola. A kernel two-sample test. *J. Mach. Learn. Res.*, 13:723–773, 2012. [6](#)
- [17] Edmond S. L. Ho and Taku Komura. Character motion synthesis by topology coordinates. *Comput. Graph. Forum*, 28(2):299–308, 2009. [2](#)
- [18] Xueshi Hou and Sujit Dey. Motion prediction and pre-rendering at the edge to enable ultra-low latency mobile 6dof experiences. *IEEE Open J. Commun. Soc.*, 1:1674–1690, 2020. [1](#)
- [19] Sheng-Wei Huang, Che-Tsung Lin, Shu-Ping Chen, Yen-Yi Wu, Po-Hao Hsu, and Shang-Hong Lai. Auggan: Cross domain adaptation with gan-based data augmentation. In *ECCV*, pages 731–744, 2018. [2](#)
- [20] 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 Trans. Pattern Anal. Mach. Intell.*, 36(7):1325–1339, 2014. [2](#), [6](#), [8](#)
- [21] Ashesh Jain, Amir Roshan Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *CVPR*, pages 5308–5317, 2016. [2](#)
- [22] Wansoo Kim, Jinoh Lee, Luka Peternel, Nikolaos G. Tsagarakis, and Arash Ajoudani. Anticipatory robot assistance for the prevention of human static joint overloading in human-robot collaboration. *IEEE Robotics Autom. Lett.*, 3(1):68–75, 2018. [1](#)
- [23] Wansoo Kim, Marta Lorenzini, Pietro Balatti, Yuqiang Wu, and Arash Ajoudani. Towards ergonomic control of collaborative effort in multi-human mobile-robot teams. In *IROS*, pages 3005–3011, 2019. [1](#)
- [24] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014. [1](#)
- [25] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, pages 1106–1114, 2012. [1](#)
- [26] Jogendra Nath Kundu, Maharshi Gor, and R. Venkatesh Babu. Bihmp-gan: Bidirectional 3d human motion prediction GAN. In *AAAI*, pages 8553–8560, 2019. [2](#)
- [27] Seunghwan Lee, Moon Seok Park, Kyoung-Min Lee, and Jehee Lee. Scalable muscle-actuated human simulation and control. *ACM Trans. Graph.*, 38(4):73:1–73:13, 2019. [2](#)
- [28] Hongyi Liu and Lihui Wang. Human motion prediction for human-robot collaboration. *Journal of Manufacturing Systems*, 44:287–294, 2017. [1](#)
- [29] Marta Lorenzini, Wansoo Kim, Elena De Momi, and Arash Ajoudani. A synergistic approach to the real-time estimation of the feet ground reaction forces and centers of pressure in humans with application to human-robot collaboration. *IEEE Robotics Autom. Lett.*, 3(4):3654–3661, 2018. [1](#)
- [30] Yuexin Ma, Xinge Zhu, Sibozhang, Ruigang Yang, Wenping Wang, and Dinesh Manocha. Trafficpredict: Trajectory prediction for heterogeneous traffic-agents. In *AAAI*, pages 6120–6127, 2019. [1](#)

- [31] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *ICCV*, pages 9488–9496, 2019. [2](#), [7](#), [8](#)
- [32] Julieta Martinez, Michael J. Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *CVPR*, pages 4674–4683, 2017. [1](#), [2](#), [7](#), [8](#)
- [33] Takuya Matsumoto, Kodai Shimosato, Takahiro Maeda, Tatsuya Murakami, Koji Murakoso, Kazuhiko Mino, and Norimichi Ukita. Automatic human pose annotation for loose-fitting clothes. In *MVA*, 2019. [1](#)
- [34] M. Müller, T. Röder, M. Clausen, B. Eberhardt, B. Krüger, and A. Weber. Documentation mocap database hdm05. Technical Report CG-2007-2, Universität Bonn, June 2007. [6](#)
- [35] Dario Pavllo, David Grangier, and Michael Auli. Quaternet: A quaternion-based recurrent model for human motion. In *BMVC*, page 299, 2018. [2](#)
- [36] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph.*, 37(4):143:1–143:14, 2018. [2](#), [4](#), [6](#), [7](#)
- [37] H. Sakoe and S. Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):43–49, 1978. [6](#)
- [38] Veit Sandfort, Ke Yan, Perry J Pickhardt, and Ronald M Summers. Data augmentation using generative adversarial networks (cycleGAN) to improve generalizability in ct segmentation tasks. *Scientific reports*, 9(1):1–9, 2019. [2](#)
- [39] Martin A Tanner and Wing Hung Wong. The calculation of posterior distributions by data augmentation. *Journal of the American statistical Association*, 82(398):528–540, 1987. [1](#)
- [40] Denis Tomè, Chris Russell, and Lourdes Agapito. Lifting from the deep: Convolutional 3d pose estimation from a single image. In *CVPR*, pages 5689–5698, 2017. [8](#)
- [41] Stefano Tortora, Stefano Michieletto, Francesca Stival, and Emanuele Menegatti. Fast human motion prediction for human-robot collaboration with wearable interface. In *IEEE CIS*, pages 457–462, 2019. [1](#)
- [42] Toan Tran, Trung Pham, Gustavo Carneiro, Lyle J. Palmer, and Ian D. Reid. A bayesian data augmentation approach for learning deep models. In *NeurIPS*, pages 2797–2806, 2017. [2](#)
- [43] Norimichi Ukita, Michiro Hirai, and Masatsugu Kidode. Complex volume and pose tracking with probabilistic dynamical models and visual hull constraints. In *ICCV*, 2009. [1](#)
- [44] Norimichi Ukita and Takeo Kanade. Gaussian process motion graph models for smooth transitions among multiple actions. *Comput. Vis. Image Underst.*, 116(4):500–509, 2012. [1](#)
- [45] Borui Wang, Ehsan Adeli, Hsu-Kuang Chiu, De-An Huang, and Juan Carlos Nieves. Imitation learning for human pose prediction. In *ICCV*, pages 7123–7132, 2019. [2](#)
- [46] Xinchen Yan, Akash Rastogi, Ruben Villegas, Kalyan Sunkavalli, Eli Shechtman, Sunil Hadap, Ersin Yumer, and Honglak Lee. MT-VAE: learning motion transformations to generate multimodal human dynamics. In *ECCV*, pages 276–293, 2018. [2](#)
- [47] Ye Yuan and Kris Kitani. Dlow: Diversifying latent flows for diverse human motion prediction. In *ECCV*, pages 346–364, 2020. [2](#)
- [48] Ye Yuan and Kris Kitani. Residual force control for agile human behavior imitation and extended motion synthesis. In *NeurIPS*, 2020. [2](#), [5](#), [6](#)
- [49] Ye Yuan, Shih-En Wei, Tomas Simon, Kris Kitani, and Jason M. Saragih. Simpoe: Simulated character control for 3d human pose estimation. In *CVPR*, pages 7159–7169, 2021. [2](#)
- [50] Chengyuan Zhang, Lei Zhu, Shichao Shang, and Weiren Yu. PAC-GAN: an effective pose augmentation scheme for unsupervised cross-view person re-identification. *Neurocomputing*, 387:22–39, 2020. [2](#)
- [51] Yi Zhou, Zimo Li, Shuangjiu Xiao, Chong He, Zeng Huang, and Hao Li. Auto-conditioned recurrent networks for extended complex human motion synthesis. In *ICLR*, 2018. [2](#)