
Residual Force Control for Agile Human Behavior Imitation and Extended Motion Synthesis

Ye Yuan Kris M. Kitani
Robotics Institute
Carnegie Mellon University
{yyuan2, kkitani}@cs.cmu.edu

Abstract

Reinforcement learning has shown great promise for synthesizing realistic human behaviors by learning humanoid control policies from motion capture data. However, it is still very challenging to reproduce sophisticated human skills like ballet dance, or to stably imitate long-term human behaviors with complex transitions. The main difficulty lies in the *dynamics mismatch* between the humanoid model and real humans. That is, motions of real humans may not be physically possible for the humanoid model. To overcome the dynamics mismatch, we propose a novel approach, *residual force control (RFC)*, that augments a humanoid control policy by adding external residual forces into the action space. During training, the RFC-based policy learns to apply residual forces to the humanoid to compensate for the dynamics mismatch and better imitate the reference motion. Experiments on a wide range of dynamic motions demonstrate that our approach outperforms state-of-the-art methods in terms of convergence speed and the quality of learned motions. Notably, we showcase a physics-based virtual character empowered by RFC that can perform highly agile ballet dance moves such as pirouette, arabesque and jeté. Furthermore, we propose a dual-policy control framework, where a kinematic policy and an RFC-based policy work in tandem to synthesize multi-modal infinite-horizon human motions without any task guidance or user input. Our approach is the first humanoid control method that successfully learns from a large-scale human motion dataset (Human3.6M) and generates diverse long-term motions. Code and videos are available at <https://www.ye-yuan.com/rfc>.

1 Introduction

Understanding human behaviors and creating virtual humans that act like real people has been a mesmerizing yet elusive goal in computer vision and graphics. One important step to achieve this goal is human motion synthesis which has broad applications in animation, gaming and virtual reality. With advances in deep learning, data-driven approaches [12, 13, 41, 38, 3] have achieved remarkable progress in producing realistic motions learned from motion capture data. Among them are physics-based methods [41, 38, 3] empowered by reinforcement learning (RL), where a humanoid agent in simulation is trained to imitate reference motions. Physics-based methods have many advantages over their kinematics-based counterparts. For instance, the motions generated with physics are typically free from jitter, foot skidding or geometry penetration as they respect physical constraints. Moreover, the humanoid agent inside simulation can interact with the physical environment and adapt to various terrains and perturbations, generating diverse motion variations.

However, physics-based methods have their own challenges. In many cases, the humanoid agent fails to imitate highly agile motions like ballet dance or long-term motions that involve swift transitions between various locomotions. We attribute such difficulty to the *dynamics mismatch* between the humanoid model and real humans. Humans are very difficult to model because they are very complex




Figure 1: **Top:** A ballet dancer performing highly agile moves like jeté, arabesque and pirouette. **Bottom:** A humanoid agent controlled by a policy augmented with the proposed residual forces (blue arrows) is able to dance like the performer. The motion is best viewed in our supplementary [video](#).

creatures with hundreds of bones and muscles. Although prior work has tried to improve the fidelity of the humanoid model [21, 52], it is nonetheless safe to say that these models are not exact replicas of real humans and the dynamics mismatch still exists. The problem is further complicated when motion capture data comprises a variety of individuals with diverse body types. Due to the dynamics mismatch, motions produced by real humans may not be admissible by the humanoid model, which means no control policy of the humanoid is able to generate those motions.

To overcome the dynamics mismatch, we propose an approach termed *Residual Force Control (RFC)* which can be seamlessly integrated into existing RL-based humanoid control frameworks. Specifically, RFC augments a control policy by introducing external residual forces into the action space. During RL training, the RFC-based policy learns to apply residual forces onto the humanoid to compensate for the dynamics mismatch and achieve better motion imitation. Intuitively, the residual forces can be interpreted as invisible forces that enhance the humanoid’s abilities to go beyond the physical limits imposed by the humanoid model. RFC generates a more expressive dynamics that admits a wider range of human motions since the residual forces serve as a learnable time-varying correction to the dynamics of the humanoid model. To validate our approach, we perform motion imitation experiments on a wide range of dynamic human motions including ballet dance and acrobatics. The results demonstrate that RFC outperforms state-of-the-art methods with faster convergence and better motion quality. Notably, we are able to showcase humanoid control policies that are capable of highly agile ballet dance moves like pirouette, arabesque and jeté (Fig. 1).

Another challenge facing physics-based methods is synthesizing multi-modal long-term human motions. Previous work has elicited long-term human motions with hierarchical RL [33, 32, 42] or user interactive control [3, 38]. However, these approaches still need to define high-level tasks of the agent or require human interaction. We argue that removing these requirements could be critical for applications like automated motion generation and large-scale character animation. Thus, we take a different approach to long-term human motion synthesis by leveraging the temporal dependence of human motion. In particular, we propose a *dual-policy control* framework where a kinematic policy learns to predict multi-modal future motions based on the past motion and a latent variable used to model human intent, while an RFC-based control policy learns to imitate the output motions of the kinematic policy to produce physically-plausible motions. Experiments on a large-scale human motion dataset, Human3.6M [14], show that our approach with RFC and dual policy control can synthesize stable long-term human motions without any task guidance or user input.

The main contributions of this work are as follows: (1) We address the dynamics mismatch in motion imitation by introducing the idea of RFC which can be readily integrated into RL-based humanoid control frameworks. (2) We propose a dual-policy control framework to synthesize multi-modal long-term human motions without the need for task guidance or user input. (3) Extensive experiments show that our approach outperforms state-of-the-art methods in terms of learning speed and motion quality. It also enables imitating highly agile motions like ballet dance that evade prior work. With RFC and dual-policy control, we present the first humanoid control method that successfully learns from a large-scale human motion dataset (Human3.6M) and generates diverse long-term motions.

2 Related Work

Kinematics-based models for human motion synthesis have been extensively studied by the computer graphics community. Early approaches construct motion graphs from large motion datasets and

design controllers to navigate through the graph to generate novel motions [20, 43]. Alternatively, prior work has explored learning a low-dimensional embedding space to synthesize motions continuously [45, 23, 24]. Advances in deep learning have enabled methods that use deep neural networks to design generative models of human motions [12, 13, 47]. While the graphics community focuses on user control, computer vision researchers have been increasingly interested in predicting future human motions. A vast body of work has used recurrent neural networks to predict a deterministic future motion from the past motion [9, 16, 25, 30, 50, 39, 1, 10]. To address the uncertainty of future, stochastic approaches develop deep generative models to predict multi-modal future motions [53, 2, 19, 56, 58]. The major drawback of kinematics-based approaches is that they are prone to generating physically-invalid motions with artifacts like jitter, foot skidding and geometry (e.g., body, ground) penetration.

Physics-based methods for motion synthesis address the limitation of kinematics-based models by enforcing physical constraints. Early work has adopted model-based methods for tracking reference motions [54, 36, 22, 28, 29]. Recently, deep RL has achieved great success in imitating human motions with manually-designed rewards [26, 27, 41]. GAIL [11] based approaches have been proposed to eliminate the need for reward engineering [31, 51]. RL-based humanoid control has also been applied to estimating physically-plausible human poses from videos [55, 57, 15]. To synthesize long-term human motions, prior work has resorted to hierarchical RL with predefined high-level task objectives [33, 32, 42]. Alternatively, recent works use deep RL to learn controllable policies to generate long-term motions with user input [3, 38]. Different from previous work, our dual-policy control framework exploits the temporal dependence of human motion and synthesizes multi-modal long-term motions by forecasting diverse futures, which can be used to replace manual task guidance or user input. Furthermore, our proposed residual force control addresses the dynamics mismatch in humanoid control and enables imitating agile motions like ballet dance that evade prior work.

Inferring external forces from human motion has been an active research area in biomechanics. Researchers have developed models that regress ground reaction forces from human motion using supervised learning [4, 37, 17]. These approaches require expensive force data collected in laboratory settings to train the models. On the other hand, machine learning researchers have proposed differentiable physics engines that enable learning forces to control simple simulated systems [6, 7]. Trajectory optimization based approaches [35, 34] have also been used to optimize external contact forces to synthesize human motions. A recent work [8] predicts forces acting on rigid objects in simulation to match image evidence with contact point supervision. Unlike prior work, we use deep RL to learn residual forces that complement contact forces to improve motion imitation without the supervision of forces or contact points.

3 Preliminaries

The task of humanoid control-based motion imitation can be formulated as a Markov decision process (MDP), which is defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, R, \gamma)$ of states, actions, transition dynamics, a reward function, and a discount factor. A humanoid agent interacts with a physically-simulated environment according to a policy $\pi(\mathbf{a}|s)$, which models the conditional distribution of choosing an action $\mathbf{a} \in \mathcal{A}$ given the current state $s \in \mathcal{S}$. Starting from some initial state s_0 , the agent iteratively samples an action \mathbf{a}_t from the policy π and the simulation environment with transition dynamics $\mathcal{T}(s_{t+1}|s_t, \mathbf{a}_t)$ generates the next state s_{t+1} and gives the agent a reward r_t . The reward is assigned based on how the agent’s motion aligns with a given reference motion. The agent’s goal is to learn an optimal policy π^* that maximizes its expected return $J(\pi) = \mathbb{E}_\pi[\sum_t \gamma^t r_t]$. To solve for the optimal policy, one can apply one’s favorite reinforcement learning algorithm (e.g., PPO [44]). In the following, we will give a more detailed description of the states, actions, policy and rewards to show how motion imitation fits in the standard reinforcement learning (RL) framework.

States. The state s is formed by the humanoid state $\mathbf{x} = (\mathbf{q}, \dot{\mathbf{q}})$ which includes all degrees of freedom (DoFs) \mathbf{q} of the humanoid and their corresponding velocities $\dot{\mathbf{q}}$. Specifically, the DoFs $\mathbf{q} = (\mathbf{q}_r, \mathbf{q}_{nr})$ include 6 root DoFs \mathbf{q}_r (global position and orientation) as well as the angles of other joints \mathbf{q}_{nr} . We transform \mathbf{q}_r to the root’s local coordinate to remove dependency on global states.

Actions. As noticed in previous work [40, 57], using proportional derivative (PD) controllers at each joint yields more robust policies than directly outputting joint torques. Thus, the action \mathbf{a} consists of the target angles \mathbf{u} of the PD controllers mounted at non-root joint DoFs \mathbf{q}_{nr} (root DoFs \mathbf{q}_r are not

actuated). The joint torques $\boldsymbol{\tau}$ can then be computed as

$$\boldsymbol{\tau} = \mathbf{k}_p \circ (\mathbf{u} - \mathbf{q}_{nr}) - \mathbf{k}_d \circ \dot{\mathbf{q}}_{nr}, \quad (1)$$

where \mathbf{k}_p and \mathbf{k}_d are manually-specified gains and \circ denotes element-wise multiplication.

Policy. As the action \mathbf{a} is continuous, we use a parametrized Gaussian policy $\pi_\theta(\mathbf{a}|\mathbf{s}) = \mathcal{N}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma})$ where the mean $\boldsymbol{\mu}_\theta$ is output by a neural network with parameters θ and $\boldsymbol{\Sigma}$ is a fixed diagonal covariance matrix. At test time, instead of sampling we use the mean action to achieve best performance.

Rewards. Given a reference motion $\hat{\mathbf{x}}_{0:T} = (\hat{\mathbf{x}}_0, \dots, \hat{\mathbf{x}}_{T-1})$, we need to design a reward function to incentivize the humanoid agent to imitate $\hat{\mathbf{x}}_{0:T}$. To this end, the reward $r_t = r_t^{\text{im}}$ is defined by an imitation reward r_t^{im} that encourages the state \mathbf{x}_t of the humanoid agent to match the reference state $\hat{\mathbf{x}}_t$. The detailed definition of the imitation reward r_t^{im} can be found in the supplementary materials.

During RL training, the agent’s initial state is initialized to a random frame from the reference motion $\hat{\mathbf{x}}_{0:T}$. The episode ends when the agent falls to the ground or the episode horizon H is reached.

4 Residual Force Control (RFC)

As demonstrated in prior work [41, 57], we can apply the motion imitation framework described in Sec. 3 to successfully learn control policies that imitate human locomotions (e.g., walking, running, crouching) or acrobatics (e.g, backflips, cartwheels, jump kicks). However, the motion imitation framework has its limit on the range of motions that the agent is able to imitate. In our experiments, we often find the framework unable to learn more complex motions that require sophisticated foot interaction with the ground (e.g., ballet dance) or long-term motions that involve swift transitions between different modes of locomotion. We posit that the difficulty in learning such highly agile motions can be attributed to the *dynamics mismatch* between the humanoid model and real humans, i.e., the humanoid transition dynamics $\mathcal{T}(s_{t+1}|s_t, \mathbf{a}_t)$ is different from the real human dynamics. Thus, due to the dynamics mismatch, a reference motion $\hat{\mathbf{x}}_{0:T}$ generated by a real human may not be admissible by the transition dynamics \mathcal{T} , which means no policy under \mathcal{T} can generate $\hat{\mathbf{x}}_{0:T}$.

To overcome the dynamics mismatch, our goal is to come up with a new transition dynamics \mathcal{T}' that admits a wider range of motions. The new transition dynamics \mathcal{T}' should ideally satisfy two properties: (1) \mathcal{T}' needs to be expressive and overcome the limitations of the current dynamics \mathcal{T} ; (2) \mathcal{T}' needs to be physically-valid and respect physical constraints (e.g., contacts), which implies that kinematics-based approaches such as directly manipulating the resulting state s_{t+1} by adding some residual δs are not viable as they may violate physical constraints.

Based on the above considerations, we propose *residual force control (RFC)*, that considers a more general form of dynamics $\tilde{\mathcal{T}}(s_{t+1}|s_t, \mathbf{a}_t, \tilde{\mathbf{a}}_t)$ where we introduce a corrective control action $\tilde{\mathbf{a}}_t$ (i.e., external residual forces acting on the humanoid) alongside the original humanoid control action \mathbf{a}_t . We also introduce a corresponding RFC-based composite policy $\tilde{\pi}_\theta(\mathbf{a}_t, \tilde{\mathbf{a}}_t|s_t)$ which can be decomposed into two policies: (1) the original policy $\tilde{\pi}_{\theta_1}(\mathbf{a}_t|s_t)$ with parameters θ_1 for humanoid control and (2) a residual force policy $\tilde{\pi}_{\theta_2}(\tilde{\mathbf{a}}_t|s_t)$ with parameters θ_2 for corrective control. The RFC-based dynamics and policy are more general as the original policy $\tilde{\pi}_{\theta_1}(\mathbf{a}_t|s_t) \equiv \tilde{\pi}_\theta(\mathbf{a}_t, \mathbf{0}|s_t)$ corresponds to a policy $\tilde{\pi}_\theta$ that always outputs zero residual forces. Similarly, the original dynamics $\mathcal{T}(s_{t+1}|s_t, \mathbf{a}_t) \equiv \tilde{\mathcal{T}}(s_{t+1}|s_t, \mathbf{a}_t, \mathbf{0})$ corresponds to the dynamics $\tilde{\mathcal{T}}$ with zero residual forces. During RL training, the RFC-based policy $\tilde{\pi}_\theta(\mathbf{a}_t, \tilde{\mathbf{a}}_t|s_t)$ learns to apply proper residual forces $\tilde{\mathbf{a}}_t$ to the humanoid to compensate for the dynamics mismatch and better imitate the reference motion. Since $\tilde{\mathbf{a}}_t$ is sampled from $\tilde{\pi}_{\theta_2}(\tilde{\mathbf{a}}_t|s_t)$, the dynamics of the original policy $\tilde{\pi}_{\theta_1}(\mathbf{a}_t|s_t)$ is parametrized by θ_2 as $\mathcal{T}'_{\theta_2}(s_{t+1}|s_t, \mathbf{a}_t) \equiv \tilde{\mathcal{T}}(s_{t+1}|s_t, \mathbf{a}_t, \tilde{\mathbf{a}}_t)$. From this perspective, $\tilde{\mathbf{a}}_t$ are learnable time-varying dynamics correction forces governed by $\tilde{\pi}_{\theta_2}$. Thus, by optimizing the composite policy $\tilde{\pi}_\theta(\mathbf{a}_t, \tilde{\mathbf{a}}_t|s_t)$, we are in fact jointly optimizing the original humanoid control action \mathbf{a}_t and the dynamics correction (residual forces) $\tilde{\mathbf{a}}_t$. In the following, we propose two types of RFC, each with its own advantages.

4.1 RFC-Explicit

One way to implement RFC is to explicitly model the corrective action $\tilde{\mathbf{a}}_t$ as a set of residual force vectors $\{\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_M\}$ and their respective contact points $\{\mathbf{e}_1, \dots, \mathbf{e}_M\}$. As the humanoid model is formed by a set of rigid bodies, the residual forces are applied to M bodies of the humanoid, where

ξ_j and e_j are represented in the local body frame. To reduce the size of the corrective action space, one can apply residual forces to a limited number of bodies such as the hip or feet. In RFC-Explicit, the corrective action of the policy $\tilde{\pi}_\theta(\mathbf{a}, \tilde{\mathbf{a}}|s)$ is defined as $\tilde{\mathbf{a}} = (\xi_1, \dots, \xi_M, e_1, \dots, e_M)$ and the humanoid control action is $\mathbf{a} = \mathbf{u}$ as before (Sec. 3). We can describe the humanoid motion using the equation of motion for multibody systems [46] augmented with the proposed residual forces:

$$\mathbf{B}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}) = \begin{bmatrix} \mathbf{0} \\ \boldsymbol{\tau} \end{bmatrix} + \underbrace{\sum_i \mathbf{J}_{v_i}^T \mathbf{h}_i}_{\text{Contact Forces}} + \underbrace{\sum_{j=1}^M \mathbf{J}_{e_j}^T \xi_j}_{\text{Residual Forces}}, \quad (2)$$

where we have made the residual forces term explicit. Eq. (2) is an ordinary differential equation (ODE), and by solving it with an ODE solver we obtain the aforementioned RFC-based dynamics $\tilde{\mathcal{T}}(s_{t+1}|s_t, \mathbf{a}_t, \tilde{\mathbf{a}}_t)$. On the left hand side $\ddot{\mathbf{q}}, \mathbf{B}, \mathbf{C}, \mathbf{g}$ are the joint accelerations, the inertial matrix, the matrix of Coriolis and centrifugal terms, and the gravity vector, respectively. On the right hand side, the first term contains the torques $\boldsymbol{\tau}$ computed from \mathbf{a} (Sec. 3) applied to the non-root joint DoFs \mathbf{q}_{nr} and $\mathbf{0}$ corresponds to the 6 non-actuated root DoFs \mathbf{q}_r . The second term involves existing contact forces \mathbf{h}_i on the humanoid (usually exerted by the ground plane) and the contact points v_i of \mathbf{h}_i , which are determined by the simulation environment. Here, $\mathbf{J}_{v_i} = dv_i/d\mathbf{q}$ is the Jacobian matrix that describes how the contact point v_i changes with the joint DoFs \mathbf{q} . By multiplying $\mathbf{J}_{v_i}^T$, the contact force \mathbf{h}_i is transformed from the world space to the joint space, which can be understood using the principle of virtual work, i.e., the virtual work in the joint space equals that in the world space or $(\mathbf{J}_{v_i}^T \mathbf{h}_i)^T d\mathbf{q} = \mathbf{h}_i^T dv_i$. Unlike the contact forces \mathbf{h}_i which are determined by the environment, the policy can control the corrective action $\tilde{\mathbf{a}}$ which includes the residual forces ξ_j and their contact points e_j in the proposed third term. The Jacobian matrix $\mathbf{J}_{e_j} = de_j/d\mathbf{q}$ is similarly defined as \mathbf{J}_{v_i} . During RL training, the policy will learn to adjust ξ_j and e_j to better imitate the reference motion. Most popular physics engines (e.g., MuJoCo [49], Bullet [5]) use a similar equation of motion to Eq. (2) (without residual forces), which makes our approach easy to integrate.

As the residual forces are designed to be a correction mechanism to the original humanoid dynamics \mathcal{T} , we need to regularize the residual forces so that the policy only invokes the residual forces when necessary. Consequently, the regularization keeps the new dynamics $\tilde{\mathcal{T}}$ close to the original dynamics \mathcal{T} . Formally, we change the RL reward function by adding a regularizing reward r_t^{reg} :

$$r_t = r_t^{\text{im}} + w_{\text{reg}} r_t^{\text{reg}}, \quad r_t^{\text{reg}} = \exp\left(-\sum_{j=1}^M (k_{\text{f}} \|\xi_j\|^2 + k_{\text{cp}} \|e_j\|^2)\right), \quad (3)$$

where $w_{\text{reg}}, k_{\text{f}}$ and k_{cp} are weighting factors. The regularization constrains the residual force ξ_j to be as small as possible and pushes the contact point e_j to be close to the local body origin.

4.2 RFC-Implicit

One drawback of RFC-explicit is that one must specify the number of residual forces and the contact points. To address this issue, we also propose an implicit version of RFC where we directly model the total joint torques $\boldsymbol{\eta} = \sum \mathbf{J}_{e_j}^T \xi_j$ of the residual forces. In this way, we do not need to specify the number of residual forces or the contact points. We can decompose $\boldsymbol{\eta}$ into two parts ($\boldsymbol{\eta}_r, \boldsymbol{\eta}_{\text{nr}}$) that correspond to the root and non-root DoFs respectively. We can merge $\boldsymbol{\eta}$ with the first term on the right of Eq. (2) as they are both controlled by the policy, which yields the new equation of motion:

$$\mathbf{B}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}) = \begin{bmatrix} \boldsymbol{\eta}_r \\ \boldsymbol{\tau} \pm \boldsymbol{\eta}_{\text{nr}} \end{bmatrix} + \sum_i \mathbf{J}_{v_i}^T \mathbf{h}_i, \quad (4)$$

where we further remove $\boldsymbol{\eta}_{\text{nr}}$ (crossed out) because the torques applied at non-root DoFs are already modeled by the policy $\tilde{\pi}_\theta(\mathbf{a}, \tilde{\mathbf{a}}|s)$ through $\boldsymbol{\tau}$ which can absorb $\boldsymbol{\eta}_{\text{nr}}$. In RFC-Implicit, the corrective action of the policy is defined as $\tilde{\mathbf{a}} = \boldsymbol{\eta}_r$. To regularize $\boldsymbol{\eta}_r$, we use a similar reward to Eq. (3):

$$r_t = r_t^{\text{im}} + w_{\text{reg}} r_t^{\text{reg}}, \quad r_t^{\text{reg}} = \exp\left(-k_r \|\boldsymbol{\eta}_r\|^2\right), \quad (5)$$

where k_r is a weighting factor. While RFC-Explicit provides more interpretable results by exposing the residual forces and their contact points, RFC-Implicit is computationally more efficient as it only

increases the action dimensions by 6 which is far less than that of RFC-Explicit and it does not require Jacobian computation. Furthermore, RFC-Implicit does not make any underlying assumptions about the number of residual forces or their contact points.

5 Dual-Policy Control for Extended Motion Synthesis

So far our focus has been on imitating a given reference motion, which in practice is typically a short and segmented motion capture sequence (e.g., within 10 seconds). In some applications (e.g., behavior simulation, large-scale animation), we want the humanoid agent to autonomously exhibit long-term behaviors that consist of a sequence of diverse agile motions. Instead of guiding the humanoid using manually-designed tasks or direct user input, our goal is to let the humanoid learn long-term behaviors directly from data. To achieve this, we need to develop an approach that (i) infers future motions from the past and (ii) captures the multi-modal distribution of the future.

As multi-modal behaviors are usually difficult to model in the control space due to non-differentiable dynamics, we first model human behaviors in the kinematic space. We propose a *dual-policy control* framework that consists of a kinematic policy κ_ψ and an RFC-based control policy $\tilde{\pi}_\theta$. The ψ -parametrized kinematic policy $\kappa_\psi(\mathbf{x}_{t:t+f}|\mathbf{x}_{t-p:t}, \mathbf{z})$ models the conditional distribution over a f -step future motion $\mathbf{x}_{t:t+f}$, given a p -step past motion $\mathbf{x}_{t-p:t}$ and a latent variable \mathbf{z} used to model human intent. We learn the kinematic policy κ_ψ with a conditional variational autoencoder (CVAE [18]), where we optimize the evidence lower bound (ELBO):

$$\mathcal{L} = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}_{t-p:t}, \mathbf{x}_{t:t+f})} [\log \kappa_\psi(\mathbf{x}_{t:t+f}|\mathbf{x}_{t-p:t}, \mathbf{z})] - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x}_{t-p:t}, \mathbf{x}_{t:t+f})||p(\mathbf{z})), \quad (6)$$

where $q_\phi(\mathbf{z}|\mathbf{x}_{t-p:t}, \mathbf{x}_{t:t+f})$ is a ϕ -parametrized approximate posterior (encoder) distribution and $p(\mathbf{z})$ is a Gaussian prior. The kinematic policy κ_ψ and encoder q_ϕ are instantiated as Gaussian distributions whose parameters are generated by two recurrent neural networks (RNNs) respectively. The detailed architectures for κ_ψ and q_ϕ are given in the supplementary materials.

Once the kinematic policy κ_ψ is learned, we can generate multi-modal future motions $\hat{\mathbf{x}}_{t:t+f}$ from the past motion $\mathbf{x}_{t-p:t}$ by sampling $\mathbf{z} \sim p(\mathbf{z})$ and decoding \mathbf{z} with κ_ψ . To produce physically-plausible motions, we use an RFC-based control policy $\tilde{\pi}_\theta(\mathbf{a}, \tilde{\mathbf{a}}|\mathbf{x}, \hat{\mathbf{x}}, \mathbf{z})$ to imitate the output motion $\hat{\mathbf{x}}_{t:t+f}$ of κ_ψ by treating $\hat{\mathbf{x}}_{t:t+f}$ as the reference motion in the motion imitation framework (Sec. 3 and 4). The state \mathbf{s} of the policy now includes the state \mathbf{x} of the humanoid, the reference state $\hat{\mathbf{x}}$ from κ_ψ , and the latent code \mathbf{z} . To fully leverage the reference state $\hat{\mathbf{x}}$, we use the non-root joint angles $\hat{\mathbf{q}}_{\text{nr}}$ inside $\hat{\mathbf{x}}$ to serve as bases for the target joint angles \mathbf{u} of the PD controllers. For this purpose, we change the humanoid control action \mathbf{a}_t from \mathbf{u} to residual angles $\delta\mathbf{u}$, and \mathbf{u} can be computed as $\mathbf{u} = \hat{\mathbf{q}}_{\text{nr}} + \delta\mathbf{u}$. This additive action will improve policy learning because $\hat{\mathbf{q}}_{\text{nr}}$ provides a good guess for \mathbf{u} .

Algorithm 1 Learning RFC-based policy $\tilde{\pi}_\theta$ in dual-policy control

- 1: **Input:** motion data \mathcal{X} , pretrained kinematic policy κ_ψ
 - 2: $\theta \leftarrow$ random weights
 - 3: **while** not converged **do**
 - 4: $\mathcal{D} \leftarrow \emptyset$ ▷ initialize sample memory
 - 5: **while** \mathcal{D} is not full **do**
 - 6: $\hat{\mathbf{x}}_{0:p} \leftarrow$ random motion from \mathcal{X}
 - 7: $\mathbf{x}_{p-1} \leftarrow \hat{\mathbf{x}}_{p-1}$ ▷ initialize humanoid state
 - 8: **for** $t \leftarrow p, \dots, p + nf - 1$ **do**
 - 9: **if** $(t - p) \bmod f = 0$ **then** ▷ if reaching end of reference motion segment
 - 10: $\mathbf{z} \sim p(\mathbf{z})$
 - 11: $\hat{\mathbf{x}}_{t:t+f} \leftarrow \kappa_\psi(\hat{\mathbf{x}}_{t:t+f}|\hat{\mathbf{x}}_{t-p:t}, \mathbf{z})$ ▷ generate next reference motion segment
 - 12: **end if**
 - 13: $\mathbf{s}_t \leftarrow (\mathbf{x}_{t-1}, \hat{\mathbf{x}}_{t-1}, \mathbf{z}); \mathbf{a}_t, \tilde{\mathbf{a}}_t \leftarrow \tilde{\pi}_\theta(\mathbf{a}_t, \tilde{\mathbf{a}}_t|\mathbf{s}_t)$
 - 14: $\mathbf{x}_t \leftarrow$ next state from simulation with \mathbf{a}_t and $\tilde{\mathbf{a}}_t$
 - 15: $r_t \leftarrow$ reward from Eq. (3) or (5)
 - 16: $\mathbf{s}_{t+1} \leftarrow (\mathbf{x}_t, \hat{\mathbf{x}}_t, \mathbf{z})$
 - 17: store $(\mathbf{s}_t, \mathbf{a}_t, \tilde{\mathbf{a}}_t, r_t, \mathbf{s}_{t+1})$ into memory \mathcal{D}
 - 18: **end for**
 - 19: **end while**
 - 20: $\theta \leftarrow$ PPO [44] update using trajectory samples in \mathcal{D} ▷ update control policy $\tilde{\pi}_\theta$
 - 21: **end while**
-

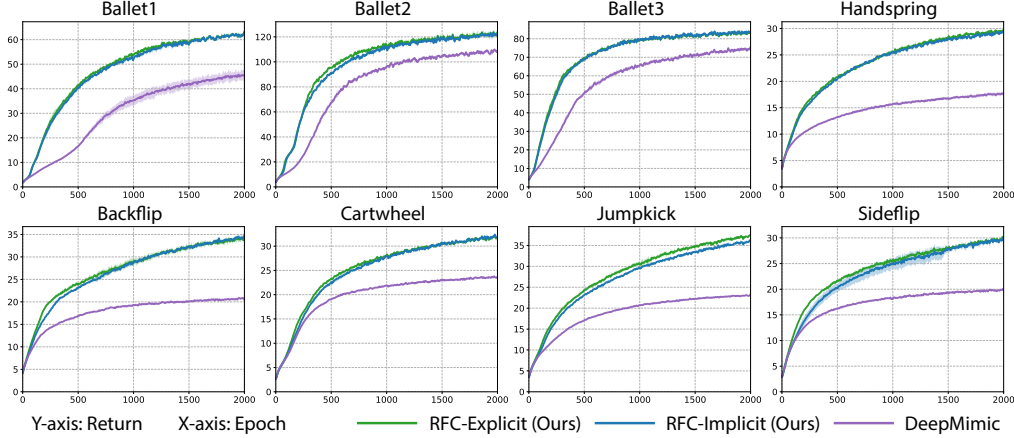


Figure 2: Learning curves of our RFC models and DeepMimic for imitating various agile motions.

The learning procedure for the control policy $\tilde{\pi}_\theta$ is outlined in Alg. 1. In each RL episode, we autoregressively apply the kinematic policy n times to generate reference motions $\hat{\mathbf{x}}_{p:p+nf}$ of nf steps, and the agent with policy $\tilde{\pi}_\theta$ is rewarded for imitating $\hat{\mathbf{x}}_{p:p+nf}$. The reason for autoregressively generating n segments of future motions is to let the policy $\tilde{\pi}_\theta$ learn stable transitions through adjacent motion segments (e.g., $\hat{\mathbf{x}}_{p:p+f}$ and $\hat{\mathbf{x}}_{p+f:p+2f}$). At test time, we use the kinematic policy κ_ψ and control policy $\tilde{\pi}_\theta$ jointly to synthesize infinite-horizon human motions by continuously forecasting futures with κ_ψ and physically tracking the forecasted motions with $\tilde{\pi}_\theta$.

6 Experiments

Our experiments consist of two parts: (1) Motion imitation, where we examine whether the proposed RFC can help overcome the dynamics mismatch and enable the humanoid to learn more agile behaviors from reference motions; (2) Extended motion synthesis, where we evaluate the effectiveness of the proposed dual-policy control along with RFC in synthesizing long-term human motions.

6.1 Motion Imitation

Reference Motions. We use the CMU motion capture (MoCap) database ([link](#)) to provide reference motions for imitation. Specifically, we deliberately select eight clips of highly agile motions to increase the difficulty. We use clips of ballet dance with signature moves like pirouette, arabesque and jeté, which have sophisticated foot-ground interaction. We also include clips of acrobatics such as handsprings, backflips, cartwheels, jump kicks and side flips, which involve dynamic body rotations.

Implementation Details. We use MuJoCo [49] as the physics engine. We construct the humanoid model from the skeleton of subject 8 in the CMU Mocap database while the reference motions we use are from various subjects. The humanoid model has 38 DoFs and 20 rigid bodies with properly assigned geometries. Following prior work [41], we add the motion phase to the state of the humanoid agent. We also use the stable PD controller [48] to compute joint torques. The simulation runs at 450Hz and the policy operates at 30Hz. We use PPO [44] to train the policy for 2000 epochs, each with 50,000 policy steps. Each policy takes about 1 day to train on a 20-core machine with an NVIDIA RTX 2080 Ti. More implementation details can be found in the supplementary materials.

Comparisons. We compare the two variants – RFC-Explicit and RFC-Implicit – of our approach against the state-of-the-art method for motion imitation, DeepMimic [41]. For fair comparison, the only differences between our RFC models and the DeepMimic baseline are the residual forces and the regularizing reward. Fig. 2 shows the learning curves of our models and DeepMimic, where we plot the average return per episode against the training epochs for all eight reference motions. We train three models with different initial seeds for each method and each reference motion. The return is computed using only the motion imitation reward and excludes the regularizing reward. We can see that both variants of RFC converge faster than DeepMimic consistently. Moreover, our RFC models always converge to better motion policies as indicated by the higher final returns. One can also observe that RFC-Explicit and RFC-Implicit perform similarly, suggesting that they are equally

Table 1: Quantitative results for human motions synthesis.

Method	Physics-based	Human3.6M (Mix)		Human3.6M (Cross)		EgoMocap	
		MAE ↓	FAE ↓	MAE ↓	FAE ↓	MAE ↓	FAE ↓
RFC-Explicit (Ours)	✓	2.498	2.893	2.379	2.802	0.557	0.710
RFC-Implicit (Ours)	✓	2.498	2.905	2.377	2.802	0.556	0.701
EgoPose [57]	✓	2.784	3.732	2.804	3.893	0.922	1.164
ERD [9]	✗	2.770	3.223	3.066	3.578	0.682	1.092
acLSTM [25]	✗	2.909	3.315	3.386	3.860	0.715	1.130

capable of imitating agile motions. Since the motion quality of learned policies is best seen in videos, we encourage the reader to refer to the supplementary [video](#)¹ for qualitative comparisons. One will observe that RFC can successfully imitate the sophisticated ballet dance skills while DeepMimic fails to reproduce them. We believe the failure of DeepMimic is due to the dynamics mismatch between the humanoid model and real humans, which results in the humanoid unable to generate the external forces needed to produce the motions. On the other hand, RFC overcomes the dynamics mismatch by augmenting the original humanoid dynamics with learnable residual forces, which enables a more flexible new dynamics that admits a wider range of agile motions. We note that the comparisons presented here are only for simulation domains (e.g., animation and motion synthesis) since external residual forces are not directly applicable to real robots. However, we do believe that RFC could be extended to a warm-up technique to accelerate the learning of complex policies for real robots, and the residual forces needed to overcome the dynamics mismatch could be used to guide agent design.

6.2 Extended Motion Synthesis

Datasets. Our experiments are performed with two motion capture datasets: Human3.6M [14] and EgoMocap [57]. Human3.6M is a large-scale dataset with 11 subjects (7 labeled) and 3.6 million total video frames. Each subject performs 15 actions in 30 takes where each take lasts from 1 to 5 minutes. We consider two evaluation protocols: (1) Mix, where we train and test on all 7 labeled subjects but using different takes; (2) Cross, where we train on 5 subjects (S1, S5, S6, S7, S8) and test on 2 subjects (S9 and S11). We train a model for each action for all methods. The other dataset, EgoMocap, is a relatively small dataset including 5 subjects and around 1 hour of motions. We train the models using the default train/test split in the mixed subject setting. Both datasets are resampled to 30Hz to conform to the policy.

Implementation Details. The simulation setup is the same as the motion imitation task. We build two humanoids, one with 52 DoFs and 18 rigid bodies for Human3.6M and the other one with 59 DoFs and 20 rigid bodies for EgoMocap. For both datasets, the kinematic policy κ_{ψ} observes motions of $p = 30$ steps (1s) to forecast motions of $f = 60$ steps (2s). When training the control policy $\tilde{\pi}_{\theta}$, we generate $n = 5$ segments of future motions with κ_{ψ} . Please refer to the supplementary materials for additional implementation details.

Baselines and Metrics. We compare our approach against two well-known kinematics-based motion synthesis methods, ERD [9] and acLSTM [25], as well as a physics-based motion synthesis method that does not require task guidance or user input, EgoPose [57]. We use two metrics, mean angle error (MAE) and final angle error (FAE). MAE computes the average Euclidean distance between predicted poses and ground truth poses in angle space, while FAE computes the distance for the final frame. Both metrics are computed with a forecasting horizon of 2s. For stochastic methods, we generate 10 future motion samples to compute the mean of the metrics.

Results. In Table 1, we show quantitative results of all models for motion forecasting over the 2s horizon, which evaluates the methods’ ability to infer future motions from the past. For all datasets and evaluation protocols, our RFC models with dual-policy control outperform the baselines consistently in both metrics. We hypothesize that the performance gain over the other physics-based method, EgoPose, can be attributed to the use of kine-

Table 2: Ablation Study.

Component		Metric	
AddAct	ResForce	MAE ↓	FAE ↓
✓	✓	2.498	2.893
✓	✗	2.610	3.150
✗	✓	3.099	3.634

¹Video: <https://youtu.be/XuzH1u78o1Y>.

matic policy and the residual forces. To verify this hypothesis, we conduct an ablation study in the Human3.6M (Mix) setting. We train model variants of RFC-Implicit by removing the residual forces (ResForce) or the additive action (AddAct) that uses the kinematic policy’s output. Table 2 demonstrates that, in either case, the performance decreases for both metrics, which supports our previous hypothesis. Unlike prior physics-based methods, our approach also enables synthesizing sitting motions even when the chair is not modeled in the physics environment, because the learned residual forces can provide the contact forces need to support the humanoid. Furthermore, our model allows infinite-horizon stable motion synthesis by autoregressively applying dual policy control. As motions are best seen in videos, please refer to the supplementary [video](#) for qualitative results.

7 Conclusion

In this work, we proposed residual force control (RFC), a novel and simple method to address the dynamics mismatch between the humanoid model and real humans. RFC uses external residual forces to provide a learnable time-varying correction to the dynamics of the humanoid model, which results in a more flexible new dynamics that admits a wider range of agile motions. Experiments showed that RFC outperforms state-of-the-art motion imitation methods in terms of convergence speed and motion quality. RFC also enabled the humanoid to learn sophisticated skills like ballet dance which have eluded prior work. Furthermore, we proposed a dual-policy control framework to synthesize multi-modal infinite-horizon human motions without any task guidance or user input, which opened up new avenues for automated motion generation and large-scale character animation. We hope our exploration of the two aspects of human motion, dynamics and kinematics, can encourage more work to view the two from a unified perspective. One limitation of the RFC framework is that it can only be applied to simulation domains (e.g., animation, motion synthesis, pose estimation) in its current form, as real robots cannot generate external residual forces. However, we do believe that RFC could be applied as a warm-up technique to accelerate the learning of complex policies for real robots. Further, the residual forces needed to overcome the dynamics mismatch could also be used to inform and optimize agent design. These are all interesting avenues for future work.

Broader Impact

The proposed techniques, RFC and dual policy control, enable us to create virtual humans that can imitate a variety of agile human motions and autonomously exhibit long-term human behaviors. This is useful in many applications. In the context of digital entertainment, animators could use our approach to automatically animate numerous background characters to perform various motions. In game production, designers could make high-fidelity physics-based characters that interact with the environment robustly. In virtual reality (VR), using techniques like ours to improve motion fidelity of digital content could be important for applications such as rehabilitation, sports training, dance instruction and physical therapy. The learned motion policies could also be used for the preservation of cultural heritage such as traditional dances, ceremonies and martial arts.

Our research on physics-based human motion synthesis combined with advances of human digitalization in computer graphics could be used to generate highly realistic human action videos which are visually and physically indistinguishable from real videos. Similar to the creation of ‘deepfakes’ using image synthesis technology, the technology developed in this work could enable more advanced forms of fake video generation, which could lead to the propagation of false information. To mitigate this issue, it is important that future research should continue to investigate the detection of synthesized videos of human motions.

Acknowledgment

This project was sponsored in part by IARPA (D17PC00340).

References

- [1] E. Aksan, M. Kaufmann, and O. Hilliges. Structured prediction helps 3d human motion modelling. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7144–7153, 2019.

- [2] E. Barsoum, J. Kender, and Z. Liu. Hp-gan: Probabilistic 3d human motion prediction via gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1418–1427, 2018.
- [3] K. Bergamin, S. Clavet, D. Holden, and J. R. Forbes. Drecon: data-driven responsive control of physics-based characters. *ACM Transactions on Graphics (TOG)*, 38(6):1–11, 2019.
- [4] A. Choi, J.-M. Lee, and J. H. Mun. Ground reaction forces predicted by using artificial neural network during asymmetric movements. *International Journal of Precision Engineering and Manufacturing*, 14(3): 475–483, 2013.
- [5] E. Coumans. Bullet physics engine. *Open Source Software: <http://bulletphysics.org>*, 1(3):84, 2010.
- [6] F. de Avila Belbute-Peres, K. Smith, K. Allen, J. Tenenbaum, and J. Z. Kolter. End-to-end differentiable physics for learning and control. In *Advances in Neural Information Processing Systems*, pages 7178–7189, 2018.
- [7] J. Degraeve, M. Hermans, J. Dambre, and F. Wyffels. A differentiable physics engine for deep learning in robotics. *Frontiers in neurorobotics*, 13:6, 2019.
- [8] K. Ehsani, S. Tulsiani, S. Gupta, A. Farhadi, and A. Gupta. Use the force, luke! learning to predict physical forces by simulating effects. *arXiv preprint arXiv:2003.12045*, 2020.
- [9] K. Fragkiadaki, S. Levine, P. Felsen, and J. Malik. Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4346–4354, 2015.
- [10] A. Gopalakrishnan, A. Mali, D. Kifer, L. Giles, and A. G. Ororbia. A neural temporal model for human motion prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12116–12125, 2019.
- [11] J. Ho and S. Ermon. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, pages 4565–4573, 2016.
- [12] D. Holden, J. Saito, and T. Komura. A deep learning framework for character motion synthesis and editing. *ACM Transactions on Graphics (TOG)*, 35(4):1–11, 2016.
- [13] D. Holden, T. Komura, and J. Saito. Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017.
- [14] C. Ionescu, D. Papava, V. Olaru, and C. 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, 2013.
- [15] M. Isogawa, Y. Yuan, M. O’Toole, and K. M. Kitani. Optical non-line-of-sight physics-based 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7013–7022, 2020.
- [16] A. Jain, A. R. Zamir, S. Savarese, and A. Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5308–5317, 2016.
- [17] W. R. Johnson, J. Alderson, D. Lloyd, and A. Mian. Predicting athlete ground reaction forces and moments from spatio-temporal driven cnn models. *IEEE Transactions on Biomedical Engineering*, 66(3):689–694, 2018.
- [18] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [19] J. N. Kundu, M. Gor, and R. V. Babu. Bihmp-gan: bidirectional 3d human motion prediction gan. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8553–8560, 2019.
- [20] J. Lee, J. Chai, P. S. Reitsma, J. K. Hodgins, and N. 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.
- [21] S. Lee, M. Park, K. Lee, and J. Lee. Scalable muscle-actuated human simulation and control. *ACM Transactions on Graphics (TOG)*, 38(4):1–13, 2019.
- [22] Y. Lee, S. Kim, and J. Lee. Data-driven biped control. In *ACM SIGGRAPH 2010 papers*, pages 1–8. 2010.
- [23] Y. Lee, K. Wampler, G. Bernstein, J. Popović, and Z. Popović. Motion fields for interactive character locomotion. In *ACM SIGGRAPH Asia 2010 papers*, pages 1–8. 2010.
- [24] S. Levine, J. M. Wang, A. Haraux, Z. Popović, and V. Koltun. Continuous character control with low-dimensional embeddings. *ACM Transactions on Graphics (TOG)*, 31(4):1–10, 2012.
- [25] Z. Li, Y. Zhou, S. Xiao, C. He, Z. Huang, and H. Li. Auto-conditioned recurrent networks for extended complex human motion synthesis. *arXiv preprint arXiv:1707.05363*, 2017.
- [26] L. Liu and J. Hodgins. Learning to schedule control fragments for physics-based characters using deep q-learning. *ACM Transactions on Graphics (TOG)*, 36(3):29, 2017.
- [27] L. Liu and J. Hodgins. Learning basketball dribbling skills using trajectory optimization and deep reinforcement learning. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018.

- [28] L. Liu, K. Yin, M. van de Panne, T. Shao, and W. Xu. Sampling-based contact-rich motion control. *ACM Transactions on Graphics (TOG)*, 29(4):128, 2010.
- [29] L. Liu, M. V. D. Panne, and K. Yin. Guided learning of control graphs for physics-based characters. *ACM Transactions on Graphics (TOG)*, 35(3):29, 2016.
- [30] J. Martinez, M. J. Black, and J. 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.
- [31] J. Merel, Y. Tassa, S. Srinivasan, J. Lemmon, Z. Wang, G. Wayne, and N. Heess. Learning human behaviors from motion capture by adversarial imitation. *arXiv preprint arXiv:1707.02201*, 2017.
- [32] J. Merel, A. Ahuja, V. Pham, S. Tunyasuvunakool, S. Liu, D. Tirumala, N. Heess, and G. Wayne. Hierarchical visuomotor control of humanoids. *arXiv preprint arXiv:1811.09656*, 2018.
- [33] J. Merel, L. Hasenclever, A. Galashov, A. Ahuja, V. Pham, G. Wayne, Y. W. Teh, and N. Heess. Neural probabilistic motor primitives for humanoid control. *arXiv preprint arXiv:1811.11711*, 2018.
- [34] I. Mordatch, Z. Popović, and E. Todorov. Contact-invariant optimization for hand manipulation. In *Proceedings of the ACM SIGGRAPH/Eurographics symposium on computer animation*, pages 137–144, 2012.
- [35] I. Mordatch, E. Todorov, and Z. Popović. Discovery of complex behaviors through contact-invariant optimization. *ACM Transactions on Graphics (TOG)*, 31(4):1–8, 2012.
- [36] U. Muico, Y. Lee, J. Popović, and Z. Popović. Contact-aware nonlinear control of dynamic characters. In *ACM SIGGRAPH 2009 papers*, pages 1–9. 2009.
- [37] S. E. Oh, A. Choi, and J. H. Mun. Prediction of ground reaction forces during gait based on kinematics and a neural network model. *Journal of biomechanics*, 46(14):2372–2380, 2013.
- [38] S. Park, H. Ryu, S. Lee, S. Lee, and J. Lee. Learning predict-and-simulate policies from unorganized human motion data. *ACM Transactions on Graphics (TOG)*, 38(6):1–11, 2019.
- [39] D. Pavllo, D. Grangier, and M. Auli. Quaternet: A quaternion-based recurrent model for human motion. *arXiv preprint arXiv:1805.06485*, 2018.
- [40] X. B. Peng and M. van de Panne. Learning locomotion skills using deeprl: Does the choice of action space matter? In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 1–13, 2017.
- [41] X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018.
- [42] X. B. Peng, M. Chang, G. Zhang, P. Abbeel, and S. Levine. Mcp: Learning composable hierarchical control with multiplicative compositional policies. In *Advances in Neural Information Processing Systems*, pages 3681–3692, 2019.
- [43] A. Safonova and J. K. Hodgins. Construction and optimal search of interpolated motion graphs. In *ACM SIGGRAPH 2007 papers*, pages 106–es. 2007.
- [44] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [45] H. J. Shin and J. Lee. Motion synthesis and editing in low-dimensional spaces. *Computer Animation and Virtual Worlds*, 17(3–4):219–227, 2006.
- [46] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo. *Robotics: modelling, planning and control*. Springer Science & Business Media, 2010.
- [47] S. Starke, H. Zhang, T. Komura, and J. Saito. Neural state machine for character-scene interactions. *ACM Trans. Graph.*, 38(6):209–1, 2019.
- [48] J. Tan, K. Liu, and G. Turk. Stable proportional-derivative controllers. *IEEE Computer Graphics and Applications*, 31(4):34–44, 2011.
- [49] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.
- [50] R. Villegas, J. Yang, Y. Zou, S. Sohn, X. Lin, and H. Lee. Learning to generate long-term future via hierarchical prediction. In *International Conference on Machine Learning*, pages 3560–3569, 2017.
- [51] Z. Wang, J. S. Merel, S. E. Reed, N. de Freitas, G. Wayne, and N. Heess. Robust imitation of diverse behaviors. In *Advances in Neural Information Processing Systems*, pages 5320–5329, 2017.
- [52] J. Won and J. Lee. Learning body shape variation in physics-based characters. *ACM Transactions on Graphics (TOG)*, 38(6):1–12, 2019.
- [53] X. Yan, A. Rastogi, R. Villegas, K. Sunkavalli, E. Shechtman, S. Hadap, E. Yumer, and H. Lee. Mt-vae: Learning motion transformations to generate multimodal human dynamics. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 265–281, 2018.
- [54] K. Yin, K. Loken, and M. Van de Panne. Simbicon: Simple biped locomotion control. *ACM Transactions on Graphics (TOG)*, 26(3):105–es, 2007.

- [55] Y. Yuan and K. Kitani. 3d ego-pose estimation via imitation learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 735–750, 2018.
- [56] Y. Yuan and K. Kitani. Diverse trajectory forecasting with determinantal point processes. *arXiv preprint arXiv:1907.04967*, 2019.
- [57] Y. Yuan and K. Kitani. Ego-pose estimation and forecasting as real-time pd control. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10082–10092, 2019.
- [58] Y. Yuan and K. Kitani. Dlow: Diversifying latent flows for diverse human motion prediction. *arXiv preprint arXiv:2003.08386*, 2020.