ShapeVerse: Physics-based Characters with Varied Body Shapes

Bharat Vyas[1] and Carol O'Sullivan[1]

[1]Trinity College Dublin, Ireland

Figure 1: *Human body representations with motion variations based on individual body shape parameters*

Abstract

Computer animation of realistic human characters remains a significant challenge. This work used deep reinforcement learning to generate physics-based characters with diverse body shapes. We aimed to replicate reference motions like walking or jogging while considering individual variations in body shape and mass. Reference motions served as training targets, accounting for differences in shape parameters to accommodate mass variations. This method produced animations that accurately capture human motion details, leading to diverse and lifelike character performances.

CCS Concepts

• ***Computing methodologies*** → *Physical simulation;* ***Procedural animation;*** *Motion capture;*

1. Introduction

To animate varied virtual characters, the motion of multiple human actors of different body shapes is redundantly motion-captured. In this work, we used deep reinforcement learning (DRL) to generate physics-based characters with varying body shapes that closely resemble real humans. Our objective was to create a diverse population of characters capable of mimicking reference motions, such as walking or jogging, where the effects of individual body shape parameters and mass are simulated. The reference motion served as a sequence of poses, providing a target for the generated characters to mimic. We employed Proximal Policy Optimization (PPO) [SWD*17], a popular DRL algorithm, to optimize the characters' motions and ensure a close match between their shape parameters and those of the reference motion actor.

A key feature of our approach was the variation of body shape parameters (β parameters), based on the SMPL body model [LMR*15], to create a diverse population of characters. We used shape parameters for the characters' bodies and proposed a reward system that is dependent on these parameters. We further controlled these rewards using a parameter, thereby allowing for flexibility in achieving the desired motion characteristics.

2. Framework for Motion Variation

To formulate our problem as a Deep Reinforcement Learning (DRL) task, we imitated a reference motion represented by a sequence of target poses (q_t), where the objective of our policy was to replicate this desired motion using physics-based simulation.

The state s captured the configuration of the character in the environment. It encompassed the joint angles (q) that define the posture, as well as their corresponding velocities ($\dot{q}$). Additionally, we incorporated the body shape state (s_b), which includes the length and width of the rigid bodies used to represent the character's body. All features were computed relative to the character's local coordinate frame, with the root at the origin and the x-axis aligned with the root link's facing direction.

The action a generated by the policy deviates the generated motion from the reference motion's posture (Δq). We utilized a Proportional-Derivative (PD) controller to drive the character's joints by applying torque. The action space served as the target input for this PD controller. Joints with three degrees of freedom (DOF) or spherical joints are represented using axis-angle notation, while joints with one DOF or revolute joints were represented using scalar values denoting joint angles.

DOI: 10.2312/egp.20241042

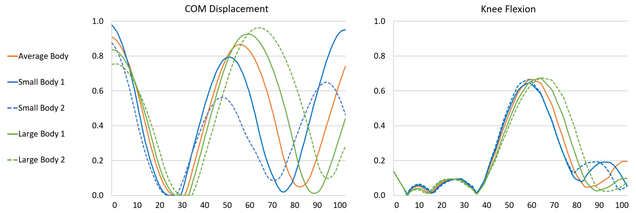

Figure 2: *Biomechanical metrics for the average, small, and large bodies, for one complete gait cycle: normalized root COM (Center Of Mass) displacement (top); and normalized values of knee flexion angle (bottom).*

We trained our policy using common reward terms from imitation learning [PALvdP18, WL19], including an imitation reward (R_i) and a regularization (energy) reward.

The imitation reward (R_i) consists of several components:

$$R_i = w^p r_p + w^v r_v + w^e r_e + w^c r_c \tag{1}$$

where r_p encourages joint orientations alignment with the reference motion, r_v promotes matching joint velocities, r_e enforces correspondence between the character's end-effectors and their positions in the reference motion, and r_c accounts for the difference in center-of-mass deviation.

The regularization reward (R_e) focuses on energy efficiency by minimizing joint torques:

$$R_e = w^e exp(-\sum \|m_J \tau_i\|^2) \tag{2}$$

where m_J represents the total mass of the rigid bodies connected to the J_{th} joint, and τ_i is the joint torque for the i_{th} joint. The relative weights (w^*) for these reward terms were manually tuned during the policy training process.

Body shape variations were introduced using b parameters to transform the base character (B) via the SMPL body model's first two Principal Components ($\beta(0)$ and $\beta(1)$). This created diverse body shapes through a set of "capsules" representing body surfaces, preserving stability and structure.

The total reward (R) combined imitation (R_i) and regularization (R_e) rewards, controlled by θ. At $\theta = 0$, imitation got prioritized, while increasing θ emphasizes energy-efficient and smoother motions. θ enabled trade-offs between motion fidelity and energy optimization. It was set based on beta parameter deviations between the base character B and the newly generated character B', ranging from 0 to 1.

3. Results & Discussion

We trained our policy (π_θ) using Proximal Policy Optimization (PPO) on the Isaac Gym physics simulation platform [MWG*21]. This policy-controlled character operated at a frequency of 60 Hz through a proportional-derivative (PD) controller.

For our evaluation, we used motions from a motion-captured dataset and extracted corresponding beta parameters for actors' body shapes using MoSh [LMB14]. This dataset included actors with varying Body Mass Index (BMI) values, categorized into "Large" (high BMI) and "Small" (low BMI) groups. For our base character B, we selected an actor with an average BMI and trained a policy for four additional actors' body shapes. Notably, our framework effectively handled variations in body shape within a single trained policy.

To assess the impact of body shape on generated motion, we examined specific lower-body biological parameters, as prior studies have demonstrated the accuracy of lower limb and trunk models in capturing center of mass (CoM) kinematics. We analyzed the same dataset for actors with different BMI values. Our findings revealed distinctive pelvis trajectories for each character, highlighting the role of total body mass in determining CoM displacement. Moreover, we analyzed the normalized knee flexion values ($^\circ$) of the left leg, which revealed variations in knee joint motion between characters with larger and smaller body shapes as shown in figure 2. These results aligned with observations in medical studies by Browning & Kram [BK07] and MacLean et al. [MCM16].

In summary, we presented a deep reinforcement learning (DRL) framework for simulating physics-based characters with diverse body shapes and sizes, based on real human body data. We highlighted the influence of physics parameters, particularly mass, on motion styles and patterns. Our results demonstrated the impact of body dimensions on motion metrics and the effectiveness of our framework in generating realistic character animations.

4. Acknowledgement

This work has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska Curie grant agreement No 860768 (CLIPE project).

References

[BK07] Browning R. C., Kram R.: Effects of obesity on the biomechanics of walking at different speeds. *Medicine and science in sports and exercise 39*, 9 (2007), 1632. 2

[LMB14] Loper M., Mahmood N., Black M. J.: Mosh: motion and shape capture from sparse markers. *ACM Trans. Graph. 33*, 6 (2014), 220–1. 2

[LMR*15] Loper M., Mahmood N., Romero J., Pons-Moll G., Black M. J.: Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG) 34*, 6 (2015), 1–16. 1

[MCM16] MacLean K. F., Callaghan J. P., Maly M. R.: Effect of obesity on knee joint biomechanics during gait in young adults. *Cogent Medicine 3*, 1 (2016), 1173778. 2

[MWG*21] Makoviychuk V., Wawrzyniak L., Guo Y., Lu M., Storey K., Macklin M., Hoeller D., Rudin N., Allshire A., Handa A., et al.: Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470* (2021). 2

[PALvdP18] Peng X. B., Abbeel P., Levine S., van de Panne M.: Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics (TOG) 37*, 4 (2018), 1–14. 2

[SWD*17] Schulman J., Wolski F., Dhariwal P., Radford A., Klimov O.: Proximal policy optimization algorithms, 2017. `arXiv:1707.06347`. 1

[WL19] Won J., Lee J.: Learning body shape variation in physics-based characters. *ACM Transactions on Graphics (TOG) 38*, 6 (2019), 1–12. 2

Proceedings published by Eurographics - The European Association for Computer Graphics.