ShapeVerse: Physics-based Characters with Varied Body Shapes

Bharat Vyas
Trinity College Dublin
Dublin, Ireland
vyasb@tcd.ie

Carol O'Sullivan
Trinity College Dublin
Dublin, Ireland
carol.osullivan@tcd.ie

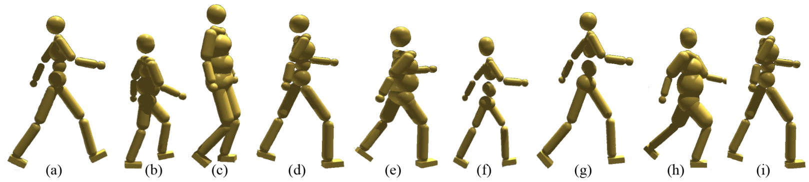

Figure 1: Human body representations with motion variations based on individual body shape parameters

KEYWORDS

Physics-based Animation, Body Shape, Motion Capture

ACM Reference Format:
Bharat Vyas and Carol O'Sullivan. 2023. ShapeVerse: Physics-based Characters with Varied Body Shapes. In *Proceedings of ** (MIG '23)*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

To animate realistic virtual characters, the motions of human actors are often captured and retargeted to the character's body model. In this work, we use deep reinforcement learning (DRL) to generate physics-based characters with varying body shapes that closely resemble real humans. Our objective is 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 serves as a sequence of poses, providing a target for the generated characters to mimic. We employ Proximal Policy Optimization (PPO) [Schulman et al. 2017], 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 is the variation of body shape parameters (β parameters), based on the SMPL body model [Loper et al. 2015], to create a diverse population of characters. The length and width of capsules fitted to the body represent the β parameters, allowing for flexibility and customization of the characters' physical appearance. We propose a reward system similar to Won and Lee [2019] that combines imitation rewards and regularization (energy) rewards. We further control the balance between these rewards using a parameter, thereby allowing for flexibility in achieving the desired motion characteristics.

MIG '23,
,

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

2 FRAMEWORK FOR MOTION VARIATION

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

The state $\boldsymbol{s}$ captures the configuration of the character in the environment. It encompasses the joint angles ($\boldsymbol{q}$) that define the posture, as well as their corresponding velocities ($\dot{\boldsymbol{q}}$). Additionally, we incorporate the body shape state ($\boldsymbol{s_b}$), which includes the length and width of the rigid bodies used to represent the character's body. All features are 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 $\boldsymbol{a}$ generated by the policy denotes the deviation from the reference motion's posture ($\Delta\boldsymbol{q}$). We utilize a Proportional-Derivative (PD) controller to drive the character's joints by applying torque. The action space serves 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 are represented using scalar values denoting joint angles.

For training our policy, we utilize common reward terms from imitation learning [Peng et al. 2018; Won and Lee 2019]. These rewards consist of an imitation reward (R_i) and a regularization (energy) reward.

The imitation reward (R_i) aims to make the generated animations closely resemble the desired reference motions and is composed of several components:

$$R_i = w^p r_p + w^v r_v + w^e r_e + w^c r_c$$

where r_p encourages alignment of joint orientations with the reference motion at each time step, r_v promotes matching joint velocities, r_e enforces correspondence between the character's end-effectors (hands and feet) and their positions in the reference motion, while r_c accounts for the difference in center-of-mass deviation and encourages the character to follow the same trajectory as the reference.

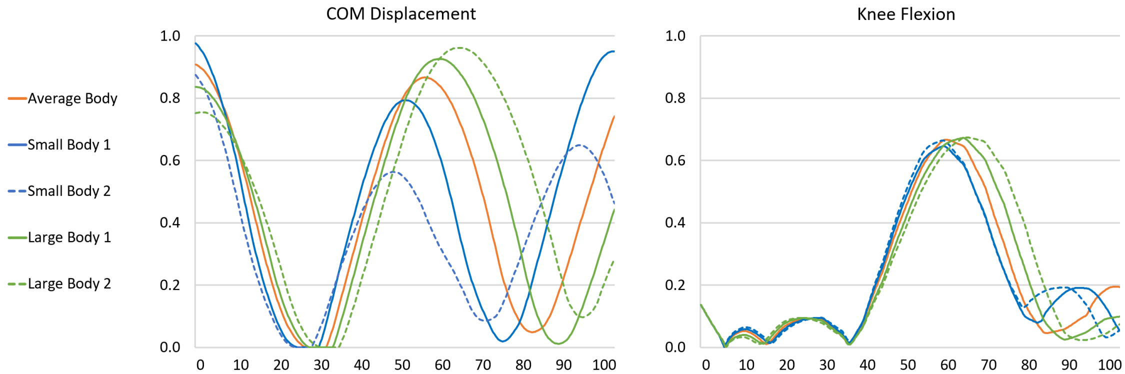

Figure 2: Biomechanical metrics for the average, small and large bodies, for one complete gait cycle: normalized root COM (Center Of Mass) displacement (left); and normalized values of knee flexion angle (right).

The regularization reward (R_e) focuses on promoting energy efficiency by minimizing joint torques. It is computed using the formula:

$$R_e = w^e exp(-\sum \left\|\boldsymbol{m}_J \tau_i\right\|^2)$$

Here, $\boldsymbol{m_J}$ represents the total mass of the rigid bodies connected to the J_{th} joint. The relative weights (w^*) for these reward terms are manually tuned during the policy training process.

In each training episode, we introduce body shape variations using $\boldsymbol{b}$ parameters to transform the base character ($\boldsymbol{B}$) via the SMPL body model's first two Principal Components ($\beta(0)$ and $\beta(1)$). This manipulation creates diverse body shapes through a set of "capsules" representing body surfaces. These capsules, with radius and axis length, align with the character's kinematic chain, forming 21 parts. This approach preserves stability and structure.

The total reward ($\boldsymbol{R}$) combines imitation (R_i) and regularization (R_e) rewards, controlled by θ. At $\theta = 0$, imitation is prioritized, while θ increasing from 0 to 1 emphasizes energy-efficient and smoother motions. θ enables trade-offs between motion fidelity and energy optimization. To set θ, we compare beta parameters between the base character $\boldsymbol{B}$ and newly generated character $\boldsymbol{B'}$. If $\Delta\beta = 0$, then $\theta = 0$. Otherwise, θ varies between 0 and 1, reflecting beta parameter deviations and different regularization levels.

3 RESULTS & DISCUSSION

We trained our policy (π_θ) using Proximal Policy Optimization (PPO) on the Isaac Gym physics simulation platform [Makoviychuk et al. 2021]. This policy controlled character operates at a frequency of 60 Hz through a proportional-derivative (PD) controller.

For our evaluation, we utilized motions from a motion captured dataset and extracted corresponding beta parameters for actors' body shapes using MoSh [Loper et al. 2014]. 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 $\boldsymbol{B}$, 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. Additionally, differences in lower body extremities were observed among individuals with varying body shapes and sizes [MacLean et al. 2016]. 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 (°) of the left leg, which revealed variations in knee joint motion between characters with larger and smaller body shapes. These results align with observations in medical studies by Browning & Kram [2007] and MacLean et al. [2016].

In summary, we have 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 demonstrate the impact of body dimensions on motion metrics and the effectiveness of our framework in generating realistic character animations. Future research directions include the integration of additional physiological factors and the incorporation of human feedback to enhance motion quality and diversity.

REFERENCES

Raymond C Browning and Rodger Kram. 2007. Effects of obesity on the biomechanics of walking at different speeds. *Medicine and science in sports and exercise* 39, 9 (2007), 1632.

Matthew Loper, Naureen Mahmood, and Michael J Black. 2014. MoSh: motion and shape capture from sparse markers. *ACM Trans. Graph.* 33, 6 (2014), 220–1.

Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. 2015. SMPL: A skinned multi-person linear model. *ACM transactions on graphics (TOG)* 34, 6 (2015), 1–16.

Kathleen FE MacLean, Jack P Callaghan, and Monica R Maly. 2016. Effect of obesity on knee joint biomechanics during gait in young adults. *Cogent Medicine* 3, 1 (2016), 1173778.

Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al. 2021. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470* (2021).

Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. 2018. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–14.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347 [cs.LG]

Jungdam Won and Jehee Lee. 2019. Learning body shape variation in physics-based characters. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–12.