

# Generalizing Skill Embeddings Across Body Shapes for Physically Simulated Characters

Sammy Christen\*

Nina Schmid\*

Otmar Hilliges

Department of Computer Science, ETH Zurich

{sammy.christen, otmar.hilliges}@inf.ethz.ch

## Abstract

Recent progress in physics-based character animation has enabled learning diverse skills from large motion capture datasets. However, most often only a single character shape is considered. On the other hand, work on controlling various body shapes with one policy is limited to few motions. In this paper, we first evaluate the generalization capabilities of latent skill embeddings on physics-based character control for varying body shapes. We then propose two strategies to learn a single policy that can generalize across different body shapes. In our experiments, we show that these simple but effective strategies significantly improve the performance over state-of-the-art, without having to retrain the skill embeddings from scratch.

## 1. Introduction

Animating simulated characters in a physically plausible manner has many applications in AR/VR, robotics, and graphics. Particularly in Embodied AI, a key goal is to make virtual characters move around and interact with their environment in natural and plausible ways. Recently, a lot of progress has been achieved by leveraging motion capture data and reinforcement learning (RL) to synthesize character motions in physics simulations [1–8, 11]. Approaches mostly rely on tuned reward functions that compare the character motions to ground truth [3, 6, 11] or use a discriminator to indicate if a motion appears realistic [7, 8]. Despite significant progress in training a single policy across a wide variety of skills using latent embeddings [7], a relatively under-researched problem is the variation of body shapes [9, 10]. [10] parameterize the body shape variations and condition a policy explicitly on the body shape. However, they build upon [6] and therefore focus on imitating single motions only. Meanwhile, [9] considers multiple motions, but does not leverage large-scale skill embeddings.

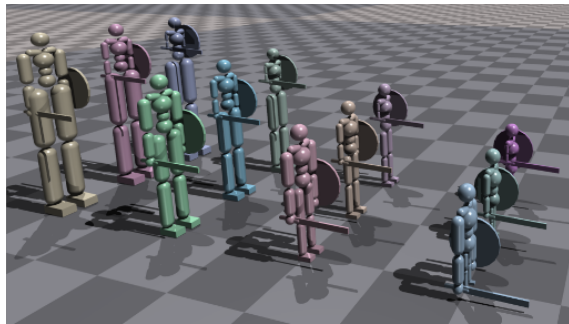


Figure 1. Visual example of a subset of different body shapes.

In this work, we explore the setting of general skill embeddings under varying body shapes. We first evaluate the zero-shot performance of state-of-the-art [7] when the body shape is changed at inference. We find that the policy transfers poorly to unseen body shapes. Thus, we propose two strategies that improve the performance of pre-trained skill embeddings under varying body shapes.

## 2. Approach

Both our approaches finetune the pretrained skill embeddings from [7], using mocap data and a discriminator for training. To generate the different body shapes, we vary the lengths of the arms and legs (see Fig. 1 for a subset of the body shapes). The physical properties of the body parts, i.e., the masses, inertia, and collision geometries, are scaled accordingly. We use 35 body shapes in total.

**Domain Randomization** In this approach, we randomize the body shapes when finetuning the low-level skill embeddings. To balance the sampling and focus on improving low-performing body shapes, we weigh each shape according to the zero-shot performance of the pre-trained low-level policy (see Section 3). During training, we sample the body shapes according to this pre-computed weight distribution. Note that the policy does not have a notion of the body shape in this approach and therefore has to find a general policy that works across all body shapes.

\*Equal Contribution



Figure 2. **Evaluation.** Average success rates across varying body shapes, displayed as heat maps for the baseline [7] and our two approaches. Our method can better adapt to varying body shapes, which is indicated by the higher success rates (brighter colors).




Figure 3. The modifications (dashed) to the original architecture (left) of the actor and critic. The modified actor outputs the corrected action and the modified critic outputs the corrected value.

**Body Shape Aware Policy** In the second approach, we modify the network architecture of the actor and critic from [7]. To this end, we add correction networks conditioned on the output of the low-level networks and the body shape parameters, i.e., in our case the scaling factors of the body parts (see Fig. 3). They adapt the action and value from the low-level networks according to the new body shape. The correction networks’ outputs are added to the original low-level networks’ outputs as a residual for stability reasons. We train the correction networks while freezing the weights of the pre-trained low-level models.

### 3. Experiments

We now evaluate the performance of the ASE method and compare it to our proposed approaches. For each method, we train high-level policies based on the respective low-level embeddings according to [7]. We focus on the three distinct tasks "GetUp", "Location", and "Reach". For the baseline, we use the task-specific pretrained models from [7]. We evaluate each of the 35 character shapes individually and illustrate the results as heat maps in Fig. 2. The heat maps indicate the average success rates over all three tasks, where each task is rolled out 100 times with random




Figure 4. Reach task with our domain randomization approach (left) and the body shape aware policy (right).

initial and goal conditions. The x and y axis represent the scaling factors of the lower and upper body, respectively.

The ASE model performs well for the base character which was used during training (i.e., 0.98 for scaling factors 1.0). We observe a gradual decrease in success rate with an increase in body shape variation as indicated by the darker colors. On the contrary, our domain randomization approach retains a success rate over 0.76 for all body shapes. This indicates that a simple change to the training procedure can already produce policies that are robust to varying body shapes. Similarly, our body shape aware approach improves upon the baseline. However, it does not perform equally well across all characters as the domain randomization approach. In particular, it struggles with large scaling factors of the lower body, likely because larger legs require more adjustments for stable control than shorter legs due to the higher center of mass. Qualitatively, we find that the domain randomization approach learns a strategy that often results in a crouching position, which leads to a stable humanoid across body shapes (see Fig. 4). While the body shape aware policy does not perform as well for all characters, we find that it results in better qualitative results because it can adapt to the specific body shape.

In conclusion, we have evaluated the zero-shot performance of large skill embeddings for physics-based characters across varying body shapes and found that it struggles to generalize to unseen body shapes. Hence, we have introduced two simple strategies to alleviate these issues. Our results show that our proposed approaches can effectively generalize to different body shapes. In the future, this can be further investigated by varying body shapes beyond scaling limbs and retargeting the reference data to the body shapes.

## References

- [1] Kevin Bergamin, Simon Clavet, Daniel Holden, and James Richard Forbes. Drecon: Data-driven responsive control of physics-based characters. *ACM Trans. Graph.*, 38(6), nov 2019. [1](#)
- [2] Sammy Christen, Muhammed Kocabas, Emre Aksan, Jemin Hwangbo, Jie Song, and Otmar Hilliges. D-grasp: Physically plausible dynamic grasp synthesis for hand-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#)
- [3] Sammy Christen, Stefan Stevsic, and Otmar Hilliges. Demonstration-guided deep reinforcement learning of control policies for dexterous human-robot interaction. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2019. [1](#)
- [4] Zhengyi Luo, Ryo Hachiuma, Ye Yuan, and Kris Kitani. Dynamics-regulated kinematic policy for egocentric pose estimation. In *Advances in Neural Information Processing Systems*, 2021. [1](#)
- [5] Soohwan Park, Hoseok Ryu, Seyoung Lee, Sunmin Lee, and Jehee Lee. Learning predict-and-simulate policies from unorganized human motion data. *ACM Trans. Graph.*, 38(6), 2019. [1](#)
- [6] 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, July 2018. [1](#)
- [7] Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. Ase: Large-scale reusable adversarial skill embeddings for physically simulated characters. *ACM Trans. Graph.*, 41(4), July 2022. [1](#), [2](#)
- [8] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Trans. Graph.*, 40(4), July 2021. [1](#)
- [9] Alexander Winkler, Jungdam Won, and Yuting Ye. Questsim: Human motion tracking from sparse sensors with simulated avatars. In *SIGGRAPH Asia 2022 Conference Papers*, SA '22, New York, NY, USA, 2022. Association for Computing Machinery. [1](#)
- [10] Jungdam Won and Jehee Lee. Learning body shape variation in physics-based characters. *ACM Transactions on Graphics (TOG)*, 38(6):1–12, 2019. [1](#)
- [11] Ye Yuan, Shih-En Wei, Tomas Simon, Kris Kitani, and Jason Saragih. Simpoe: Simulated character control for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [1](#)