Enabling Physical VR Interaction with Deep RL Agents

Anonymous Author(s)

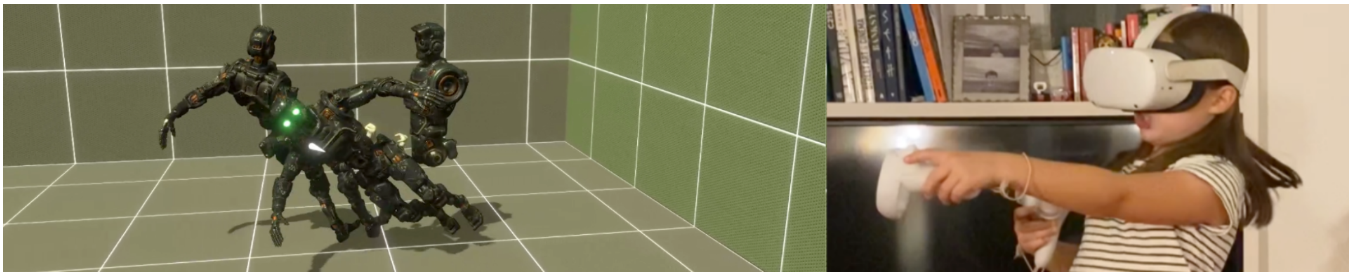

Figure 1: An 8-year-old user interacting with two RL agents in Virtual Reality (VR).

ABSTRACT

Natural interactions between avatars and autonomous agents are essential for enhancing VR applications and games. Our solution consists in immersing the user into a fully-physically-based head/hands avatar, able to interact with Deep Reinforcement Learning agents. The avatar includes novel PD-joints from the virtual hands to the head, retargeting both to maintain consistency between virtual and real hand's placement relative to the user's viewpoint. We enhance the training procedure of Deep Reinforcement Learning agents to make them robust to user interaction and provide a method to adjust the relative strengths of the avatar and the RL agent.

ACM Reference Format:
Anonymous Author(s). 2023. Enabling Physical VR Interaction with Deep RL Agents. In *Proceedings of ACM SIGGRAPH Conference on Motion, Interaction and Games (MIG '23).* ACM, New York, NY, USA, 2 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION AND RELATED WORK

Real-time physically-based animation and interaction have shown much promise in delivering more natural experiences in both video games and VR. The physics simulation drives natural interaction between all physics objects, but controlling them has to be done with forces, which is a challenging task, because of the indirect control on animation.

As shown by a previous user study [8], physically-based VR body retargeting, used so far without viewpoint retargeting, already improves embodiment. Physically simulated VR hands were used to interact with static and moving objects in a virtual environment [2, 5]. Several VR games, such as Boneworks[9], implemented physics-based VR interaction systems with humanoid characters. However, the characters were not fully physically simulated and were artificially kept upright. In contrast, we rely on Deep Reinforcement Learning to train interaction-ready physically-based autonomous agents, inspired by Peng's AMP method [6].

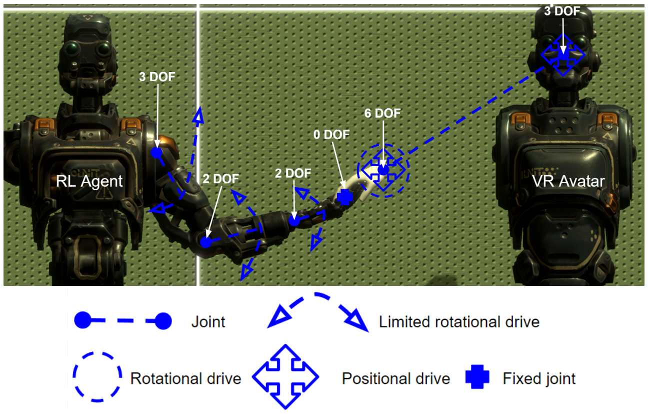

Figure 2: Physical models for the RL agents (left) and the VR avatar (Right), in a hand grabbing interaction.

The goal of this work is to create a fully physically integrated framework where the user, immersed in VR, can interact with autonomous agents in an intuitive and coherent manner, i.e. physically, through opposite interaction forces. Our contributions are:

- A fully physically-based head/hands avatar for the user. it uses a novel PD-controlled joint between virtual hands and head to align them with their real counterparts, either by displacing the hand or the head if necessary.
- A training procedure for Deep RL walking agents, making them robust to arbitrary physical interactions.
- A method to calibrate the strength of the VR avatar in order to match the strength of the agent.

2 USER-DRIVEN PHYSICS-BASED AVATAR

The VR character is composed of dynamic rigid bodies for the virtual head-and-body and each hand (see fig 2). The virtual hands are connected to the head with PD controllers that apply forces and torques in order to attain the target positions and rotations given

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MIG '23, November 15–17, 2023, Rennes, France
© 2023 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/XXXXXXX.XXXXXXX

by VR device outputs (hand-device and headset).

$$P^{target}_{hand} - P_{head} = P_{hand-device} - P_{headset} \tag{1}$$

The target position P^{target}_{hand} of the virtual hand is set to overlap with the VR hand-device from the user's point of view, which makes it consistent between the user's kinesthetic and visual awareness of the position of their own hand. The PD controller applies a force on the hand and an opposite force on the head to satisfy this constraint. Since the hands are much lighter, they generally move first with no head movement. If the hand cannot move to satisfy the constraint, the force becomes strong enough to move the head (and body), and therefore, change the user's camera viewpoint in the virtual world.

This mechanism enables consistent viewpoint retargeting during interactions: Grabbing an object creates a rigid joint between the avatar's hand and the object. Lowering the VR hand-device then results in lowering the object, or, if this does not happen fast enough, the virtual head will start to lift. This enables users to climb by simply grabbing a wall and lowering their hand, or to jump by rapidly pushing towards the wall with enough force and friction.

We use a new type of PD controller for the Head-Hand joint, making the target velocity proportional to the error:

$$\begin{gathered} F = k_P \cdot error + k_D(targetVelocity - velocity) \\ error = P^{target}_{hand} - P^{actual}_{hand} \text{ and } targetVelocity = (k_T \cdot error) \end{gathered} \tag{2}$$

When getting closer to the target, the bodies therefore slow down which is reduces overshooting and oscillation. Tuning is also made easier: We can discard k_P, use k_D to control the acceleration/deceleration strength and k_T to control the latency.

3 INTERACTION-READY RL AGENTS

Inspired by [6], we use Deep Reinforcement Learning to train the motion controller of humanoid physics-based agents, using the PPO[7] algorithm paired with GAIL[1], available in the Unity ML-Agents library[4]. The neural network outputs target rotation angles mapped between articulation limits and spring damper parameters to pass to the joint drives for torque.

To prevent the agent from immediately falling when touched by the avatar, we train it with random forces applied on its different body parts. We also add new inputs to the neural network : the sums of the applied interaction forces and torques (excepting the reaction force from the floor and gravity forces), in order to enable the agent to learn how to recuperate.

4 AVATAR-AGENT CALIBRATION

Thanks to the physics engine and force-based control of both the VR avatar and the RL agents, they can interact and move each other in a physically consistent manner. However, some tuning is necessary to make such two-way interaction visually interesting, since the forces they apply must be of the same order of magnitude.

To achieve this, we first lower the RL agents energy usage during training, by adding an energy-related penalty to the reward. This is only done when the agent already succeeds in its navigation task, so that the learning capacity is not impacted. In order to maintain the same maximum force capacity for the avatar and the agents, we calculate or set the maximum force applicable on the hands of either character as follows: Fixing its hand in space, we train the RL agent to pull as hard as possible on its hand, and get the maximum force and torque it can apply. In VR, we then simply set these values as maximum forces and torques on the avatar's joint drives, capping potential hand strength.

5 CONCLUSION

We achieved natural interaction in VR between a user's avatar and physically based autonomous agents. So far, at least 20 people have tried this demo informally, with great success. Despite viewpoint retargeting, they did not complain of motion sickness, a problem reported in many VR games [3]. Conducting a formal user study on our system will nevertheless be essential to fully validate it. Furthermore, both the avatar and agents still possess too much strength, and the agent is too stiff which limits interactability. Further work is therefore needed to better calibrate forces.

REFERENCES

[1] Jonathan Ho and Stefano Ermon. 2016. Generative Adversarial Imitation Learning. *CoRR* abs/1606.03476 (2016). arXiv:1606.03476 http://arxiv.org/abs/1606.03476

[2] Markus Höll, Markus Oberweger, Clemens Arth, and Vincent Lepetit. 2018. Efficient Physics-Based Implementation for Realistic Hand-Object Interaction in Virtual Reality. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 175–182. https://doi.org/10.1109/VR.2018.8448284

[3] Quinate Ihemedu-Steinke, Stanislava Thull, Michael Weber, Rainer Erbach, Gerrit Meixner, and Nicola Marsden. 2017. Simulation Sickness Related to Virtual Reality Driving Simulation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 521–532. https://doi.org/10.1007/978-3-319-57987-0_42

[4] Arthur Juliani, Vincent-Pierre Berges, Ervin Teng, Andrew Cohen, Jonathan Harper, Chris Elion, Chris Goy, Yuan Gao, Hunter Henry, Marwan Mattar, and Danny Lange. 2020. Unity: A general platform for intelligent agents. *arXiv preprint arXiv:1809.02627* (2020). https://arxiv.org/pdf/1809.02627.pdf

[5] Jenny Lin, Xingwen Guo, Jingyu Shao, Chenfanfu Jiang, Yixin Zhu, and Song-Chun Zhu. 2016. A Virtual Reality Platform for Dynamic Human-Scene Interaction. In *SIGGRAPH ASIA 2016 Virtual Reality Meets Physical Reality: Modelling and Simulating Virtual Humans and Environments* (Macau) *(SA '16)*. Association for Computing Machinery, New York, NY, USA, Article 11, 4 pages. https://doi.org/10.1145/2992138.2992144

[6] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. 2021. AMP. *ACM Transactions on Graphics* 40, 4 (jul 2021), 1–20. https://doi.org/10.1145/3450626.3459670

[7] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *CoRR* abs/1707.06347 (2017). arXiv:1707.06347 http://arxiv.org/abs/1707.06347

[8] Yujie Tao, Cheng Yao Wang, Andrew D Wilson, Eyal Ofek, and Mar Gonzalez-Franco. 2023. Embodying Physics-Aware Avatars in Virtual Reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 254, 15 pages. https://doi.org/10.1145/3544548.3580979

[9] Stress Level Zero. 2019. Boneworks. https://en.wikipedia.org/wiki/Boneworks.

Received 22 September 2023