

QuestSim: Human Motion Tracking from Sparse Sensors with Simulated Avatars

ALEXANDER WINKLER, Reality Labs Research, Meta, USA

JUNGDMAM WON, Meta AI Research, USA

YUTING YE, Reality Labs Research, Meta, USA




Fig. 1. User pose reconstructed from the position and orientation of a headset and two hand controllers (left), or from the headset only (right). The same policy can track users of different sizes (left: 167 cm, right: 181 cm). The avatar motion is simulated in a physics engine, which gives access to dynamic quantities such as normal contact forces (red line on the ground).

Real-time tracking of human body motion is crucial for interactive and immersive experiences in AR/VR. However, very limited sensor data about the body is available from standalone wearable devices such as HMDs (Head Mounted Devices) or AR glasses. In this work, we present a reinforcement learning framework that takes in sparse signals from an HMD and two controllers, and simulates plausible and physically valid full body motions. Using high quality full body motion as dense supervision during training, a simple policy network can learn to output appropriate torques for the character to balance, walk, and jog, while closely following the input signals. Our results demonstrate surprisingly similar leg motions to ground truth without any observations of the lower body, even when the input is only the 6D transformations of the HMD. We also show that a single policy can be robust to diverse locomotion styles, different body sizes, and novel environments.

CCS Concepts: • **Computing methodologies** → **Motion capture**.

Additional Key Words and Phrases: Motion Tracking, Character Animation, Reinforcement Learning, Wearable Devices

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SA '22 Conference Papers, December 6–9, 2022, Daegu, Republic of Korea

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9470-3/22/12.

<https://doi.org/10.1145/3550469.3555411>

ACM Reference Format:

Alexander Winkler, Jungdam Won, and Yuting Ye. 2022. QuestSim: Human Motion Tracking from Sparse Sensors with Simulated Avatars. In *SIGGRAPH Asia 2022 Conference Papers (SA '22 Conference Papers)*, December 6–9, 2022, Daegu, Republic of Korea. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3550469.3555411>

1 INTRODUCTION

A promise of AR/VR (Augmented, Virtual Reality) is that it will enable richer forms of self-expression and social experience compared to 2D video. This could be achieved through avatars that accurately capture a user's movement and body language. To enable this, we need sensors and methods to faithfully reproduce the full body motion of a user in real-time.

Optical *marker-based* solutions [Vicon 2022] are often used in industries or research labs that require high precision. However, the setup is complex: It requires multiple cameras placed throughout the room, as well as attaching and calibrating markers on the users to be captured. A solution with less friction is *markerless* motion capture, which does not require any markers attached to the user. However, the sensor still needs to observe the user at all times, so moving between rooms or large scale motion capture is difficult. This motivates motion capture from *wearable* sensors, which rely only on sensors attached to the user with no other external sensing modality. One type of wearable sensor are Inertial Measurement Units (IMUs)

that can capture both linear and angular motion. Since IMUs are prone to drifting, current HMDs often fuse this acceleration data with camera information (SLAM) to estimate their location [Durrant-Whyte and Bailey 2006]. This results in reasonable estimates of the global position and orientation of the headset and the controllers. And since the sensors are wearable they can be used across rooms and even outside.

However, sensor signals that are accessible from AR/VR devices are sparse, with no information about the lower body. To reconstruct full body poses from this data, part of the human pose must be synthesized. Purely kinematic approaches have difficulty synthesizing this missing information in a believable way, especially from sparse inputs, as the space of all possible human poses is vast. This can lead to unnatural artifacts such as jitter, foot skating, and unstable contacts. In this work, we incorporate an off-the-shelf physics simulator into the tracking pipeline in order to constrain the solution space to physically valid poses to mitigate some of these artifacts.

Contributions. We show that sparse upper-body sensors carry enough signal, when combined with physics, to predict the lower-body pose, even from the HMD alone. We demonstrate this by tracking users of different heights from real-world sensor data with a single policy, trained end-to-end with deep reinforcement learning. This creates motions with less artifacts such as foot skating compared to kinematic approaches. The simulated environment can also be used to adapt motions (e.g. to rough terrain), to better fit into the virtual environment.

2 RELATED WORK

We categorize approaches to human motion tracking based on the type of sensor used as input. As the input signal becomes sparser, more of the pose reconstruction must be synthesized. Physics is a sensible prior when reconstructing those aspects of the pose which are not observed by any sensors. Therefore we end this section by reviewing physics-based approaches using Reinforcement Learning (RL), and contrast them to kinematic-based approaches.

2.1 Vision Sensors

Generating 3D full body poses based on camera images can be used for motion tracking. This image can be used to predict parameters for a parameterized and differentiable statistical human body model such as SMPL [Kanazawa et al. 2019; Loper et al. 2015; Rong et al. 2021; Xu et al. 2019]. Oftentimes, camera pixels are preprocessed to extract body keypoints [Cao et al. 2019] or body correspondences from the image [Güler et al. 2018]. One of the difficulties when using monocular cameras is the depth ambiguity, e.g. a short person closer to the camera looks identical to a taller person further away. This depth ambiguity can cause reconstructed 3D poses to accurately match the camera image when viewed from the same angle, whereas different views reveal unnatural leaning, scale and pose. Furthermore, if every frame is reconstructed individually, continuity between frames can be difficult and can cause poses to jitter. Constraining the poses through physics-based priors can help to mitigate these issues [Rempe et al. 2021].

2.2 IMU Sensors

Another popular approach is to use sensors mounted on a user's body, for example inertial measurement units (IMUs). They are small, lightweight and don't exhibit occlusion, since they are based on acceleration signal, versus vision. This allows them to be used across rooms, as well as outdoors and makes them agnostic to light and weather conditions. An offline method using IMUs was proposed by Marcard et al. [2017], where the system optimizes the pose parameters of the SMPL model so that it matches the sensor's signal. The system performs best when it can access the entire sensor signal trajectory. In order to overcome this offline constraint, deep learning based approaches using mocap data paired with IMU signals were proposed. These models produce a *local* joint angle pose in realtime [Huang et al. 2018; Nagaraj et al. 2020]. Methods that estimate the full pose, including global 6D root from IMU signal were explored by [Jiang et al. 2022b; Yi et al. 2021]. Since kinematic pose prediction can be jittery or drift, these methods use predicted foot contacts, *Stationary Boundary Points*, or terrain prediction to improve the synthesized poses. Apart from IMUs, sensors based on electromagnetic (EM) fields have also been used to reconstruct full-body poses [Kaufmann et al. 2021]. However, most of the above approaches require sensors attached to the lower body.

2.3 HMD Sensors

As Head Mounted Devices (HMDs) for AR/VR are becoming more widely available, methods that generate full-body poses only from HMD and controllers are being explored. Dittadi et al. [2021] proposed a framework based on a variational autoencoder (VAE), where a VAE is first learned with all poses in the dataset, then the decoder is combined with another encoder which uses IMU signals as input. Aliakbarian et al. [2022] utilized a flow-based model instead of VAEs, where the invertible nature of flow-based models enables learning a shared latent space in a single learning phase. Jiang et al. [2022a] uses a Transformer architecture to predict the global skeleton state from only the HMD and controller poses. However, since these approaches are kinematic and don't enforce physical constraints, they can suffer from artifacts such as foot-skating and jitter.

2.4 Physics-based Approaches

Many of the previous approaches are *kinematic*, meaning that the model has no notion of masses, inertias or forces as it synthesizes the poses. However, obeying physical laws makes synthesized motions more believable. This physics prior is especially important when the problem is under-constrained, where many solutions fulfill the sensor constraints but only a few are desirable. This is the case when using sparse sensor signal (head and controllers) to reconstruct physically accurate full-body poses.

Shimada et al. [2020] optimized the output kinematic motion predicted from a monocular video where physics laws are used as soft constraints. Yi et al. [2022] added a physics-based optimizer as a refinement process to further improve motions generated by kinematic methods such as *TransPose* [Yi et al. 2021], showing impressive results for IMU-based pose reconstruction. We show that we can generate motions of comparable quality without sensors on the lower body.

However, many methods still allow physics to be slightly violated, e.g. they are not enforced as hard-constraints. This can cause the posture of the characters to be abnormally tilted or for generated motions to appear to be floating in the air. We build much of our work on RL-based imitation learning, which adds a physics simulator as the final step before the pose generation, thereby enforcing physics as hard-constraints that cannot be violated [Peng et al. 2018a]. The controller outputs either target joint angles or torques which are used by the physics simulator to generate the next pose of the simulated character. This stream of research traditionally focused on imitating motions from *dense, full-body* user observations. Here, a variety of methods have explored how a single policy can imitate diverse motions, such as walking, jogging, break-dancing etc [Bergamin et al. 2019; Chentanez et al. 2018; Fussell et al. 2021; Park et al. 2019; Won et al. 2020]. This becomes especially relevant for motion tracking, where we cannot know in advance what type of motion the user will perform, but require the policy to be able to track it. Another line of research is training on large mocap datasets to generate fundamental "building blocks" of motions, that can later be reused to achieve other downstream tasks [Merel et al. 2020; Peng et al. 2019, 2021; Won et al. 2021]. In the context of motion tracking these insights could be used to track the motion using a specific style.

Reinforcement Learning has also been used to imitate kinematic motions generated from sparser observations, such as monocular video [Peng et al. 2018b; Yu et al. 2021]. It has been successfully combined with egocentric video to synthesize poses and object interaction [Luo et al. 2021] and predict future motion [Yuan and Kitani 2019]. Luo et al. [2021] learned a universal controller and a kinematic policy and used it to generate simulated motions from a single-view egocentric video. In contrast to these video-based approaches, we show motion synthesis using only the 6 DoF state of the headset and controllers. We use a simple architecture, consisting of a single MLP trained end-to-end. Since our physics simulation doesn't use additional non-physical forces, our approach generates high-quality and believable motions.

3 METHOD

Traditional RL-imitation frameworks often assume that full-body, noise-free observations of the reference are available, which is not the case when tracking a user from only a HMD. In the following we detail our architecture based on Reinforcement learning, with important modifications for motion-tracking of users from sparse and real-world sensors.

3.1 RL background

We use reinforcement learning to train the character. The goal is to find a policy π_θ to maximise the expected discounted return

$$J(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_0^T \gamma^t r_t \right], \quad (1)$$

where θ represents the weights of a neural network, $\gamma \in [0, 1]$ is the discount factor, and $\tau = (s_0, a_0, \dots, s_{T+1})$ represents trajectories collected by the policy interacting with the environment. The probability for a specific trajectory τ when using the current policy

π_θ is given by $P(\tau|\theta) = P(s_0) \prod_{t=0}^T P(s_{t+1}|s_t, a_t) \pi_\theta(a_t|s_t)$, where $P(s_0)$ is the initial state distribution and $P(s_{t+1}|s_t, a_t)$ represents the dynamics of the physics simulator. This probability is used to calculate the expectation $\mathbb{E}_{\tau \sim \pi_\theta}$ which quantifies the expected return. To learn the policy weights θ , we use the proximal policy optimization (PPO) algorithm [Schulman et al. 2017].

3.2 Overview

The goal is to reconstruct the full body pose of a user from sparse user observations o_{user} as shown in Figure 2. This output pose s_t is generated by a physics simulator. The simulator is driven by joint torques, which are produced by a neural network termed "Policy". We use Reinforcement Learning, together with an imitation objective [Peng et al. 2018a], to train the policy to produce torques that track a user. During training we use a scalar reward r_t that captures the goal of producing a pose s_t that is as close as possible to the corresponding ground-truth pose $s_{t,\text{gt}}$ from the motion database.

We hypothesize this architecture works well with sparse input data because of the high quality full body supervision signal during training, as well as because of the feedback of the simulated state back into the policy. Since the user observations are sparse and don't provide much information about the user, considering the current state of the physically simulated character significantly reduces ambiguity on possible next poses. This allows the policy to use this *internal* state to decide on optimal actions. For kinematic tracking approaches, this feedback loop is conceptually similar to Recurrent Neural Networks (RNNs) or "Pose Priors". However, complex architectures such as RNNs or Transformers are not required in our case and a 3-layer MLP policy is sufficient.

The policy requires observations as input in order to determine the torque at every simulation step. The observations can be split into three sections discussed in the following: the observations coming from the simulated character, the sparse observations coming from the sensors worn by the user, and the scale of the user.

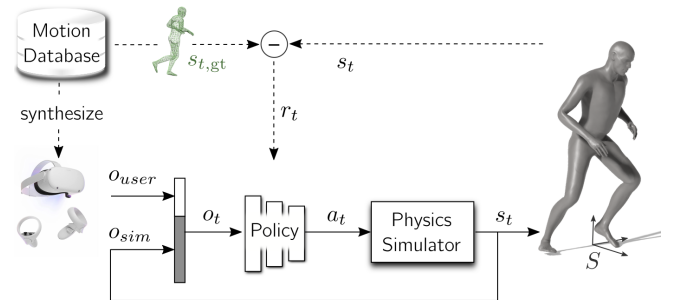


Fig. 2. Overview of our Imitation Learning architecture to track users from sparse sensor data. The dotted paths are only needed to train the policy network (MLP). During inference, this network produces torques for a physics simulator that generates the pose.

3.3 Simulated Character Observations

The simulated avatar has 33 degrees-of-freedom (DoF). It is fully observable, allowing the policy access to any values that are helpful

for determining the torques. We use the joint angles $o_{\text{sim},q} \in \mathbb{R}^{33}$ and joint angle velocities $o_{\text{sim},\dot{q}} \in \mathbb{R}^{33}$. Even though redundant, we also give the policy access to the Cartesian positions and orientations of each link, which speeds up training. All the Cartesian positions $o_{\text{sim},x} \in \mathbb{R}^{16 \times 3}$ are expressed with respect to frame S , which is located on the floor below the avatar and rotates according to its heading direction (see Figure 2). This allows the policy to learn torque mappings independent of the heading direction. The link orientations $o_{\text{sim},R} \in \mathbb{R}^{16 \times 6}$, also in frame S , are encoded by the first two columns of their rotation matrices. This has been shown to be a favorable orientation representation for neural networks compared to quaternions or angle axis, which can introduce discontinuities. We also observe dynamic quantities such as contact forces of each foot $o_f \in \mathbb{R}^{2 \times 3}$, which allows the policy to reason about contact states when determining the torques. The policy also observes the simulated avatar’s linear and angular velocity of each of the links, which is necessary in order to take the inertia of the character into account when producing torques. In total the simulated avatar is observed by $o_{\text{sim}} \in \mathbb{R}^{312}$.

3.4 Synthetic Training Data

Training the policy requires sensor data of the user paired with ground-truth poses $s_{t,\text{gt}}$ from which the reward is computed. Instead of recording a new dataset capturing the full-body pose of a user (e.g. with a marker-based setup) *while* wearing a headset, we synthetically generate this paired data. We offset the ground-truth head and wrist joints to emulate the position and orientation of a headset and left and right controllers as if the subjects were equipped with the devices. Since the real sensor signal is sufficiently clean, noise was not added to the offsets.

Our in-house Mocap data consists of 8 hours of motion clips of 172 subjects. Specifically, the dataset contains 130 minutes of walking, 110 minutes of jogging, 80 minutes of casual conversations with gestures, 90 minutes of whiteboard discussion and 70 minutes of balancing. Existing datasets lacked either subject diversity or motion diversity.

3.5 User Pose Observations

The sensor data, either coming from the real headset or synthetically generated for training, is given by the position and orientation of the headset h , the left controller l and the right controller r . Like the Cartesian observations of the simulated character, all user positions and orientations are relative to frame S as shown in Figure 2. The orientations $sR_i \in \mathbb{R}^6$ denote the first two columns of the rotation matrix of the head, left and right controller relative to frame S (see Sec. 3.3 for details).

$$o_{\text{user}, t} = [h_S, sR_h, l_S, sR_l, r_S, sR_r]. \quad (2)$$

The policy also has access to 6 future user observations. This gives the simulated avatar the ability to better anticipate and thereby track a user’s motion by knowing what they will do next. The trade-off is that relying on future poses introduces a 160ms latency, which makes real-time tracking less responsive. In total the pose observations coming from the user are $o_{\text{user}} \in \mathbb{R}^{6 \times 3 \times (3+6)} \in \mathbb{R}^{162}$.

3.6 User Scale observations

Many imitation learning approaches generate policies that are specific to one character of a particular scale. In order to deliver a general motion tracking solution, we want to be able to track users of any scale (tall or short). One approach is to generate individual policies per user scale and then blend them during inference. The downside is that this requires a unique policy for every user scale and there might not be sufficient mocap training data available for each user. So instead, we learn a single policy which generalizes to various user scales. To achieve this we give the policy access to the user scale in form of a scalar value, which specifies the user’s height in meters.

$$o_{\text{user}, \text{scale}} \in \mathbb{R}. \quad (3)$$

This allows the policy to learn to adjust torques based on the user, for instance apply larger torques when tracking taller users with larger masses and inertia.

During training, when imitating a clip recorded by one of the 172 individual subjects, we use the first pose in the clip, which is an A-pose, to extract the height of the subject. This single scale value is used to initialize a simulated character with approximately the same scale. During inference, when a user wears the headset, we require the user to initially stand upright and use the height estimated by the headset to initialize their avatar.

3.7 Reward

The reward r_t is used to generate a desired behavior of the character. The goal for the simulated character is to imitate a user’s motion as close as possible. We build on the imitation reward introduced by Peng et al. [2018a]. Since during training the sparse observations are synthetically generated from the full-body mocap pose, the corresponding full-body pose $s_{t,\text{gt}}$ is known and can be used to formulate this reward. This way the policy has dense supervision during training while requiring only sparse data during inference. Our reward function is given by the dot product

$$r_t = \mathbf{w}[r(q), r(\dot{q}), r(x), r(\dot{x}), r_f]^T, \quad (4)$$

where the terms $r(q)$ and $r(\dot{q})$ quantify the difference of joint angles and joint angle velocities between the simulated avatar and the ground truth mocap pose and $r(x)$ and $r(\dot{x})$ quantify the difference in Cartesian positions and velocities of each joint. Each term is expressed using a Gaussian kernel $r(s) = \exp(-k_s \sum_j \|s_{\text{sim}} - s_{\text{gt}}\|_2^2)$, where s_{sim} and s_{gt} are the particular representation of the simulated character and the ground truth pose respectively and k_s is the sensitivity of the kernel for each representation. Weights \mathbf{w} and kernel sizes \mathbf{k} can be found in the Appendix.

Using only the above terms when no lower-body user observations are available results in the simulated avatar taking short, high-frequency steps, which look unnatural. The additional reward term $r_f = \exp(k_f \sum_{i=L,R} \max(0, f_{y,i,\text{prev}} - f_{y,i}))$ reduces this, by penalizing an abrupt decrease in vertical contact force $f_y \in \mathbb{R}$, where L and R denote the left and right foot, and k_f the sensitivity of the Gaussian kernel. This encourages the avatar to first naturally unload a leg before lifting it, but still allows to forcefully step down to make contact.

3.8 Training

The policy directly outputs torque values instead of PD target angles, which tracked stably and reduced code complexity compared to e.g. *Stable PD* [Tan et al. 2011] or other more involved controller formulations. We run the policy and simulation at 36 frames per second (fps), which we found was the largest stable timestep. This also facilitated downsampling of the HMD sensor data during inference, which is provided at 72 fps. As a physics simulator we use Nvidia’s *PhysX*, which has been wrapped by the RL training framework *IsaacGym* [Makoviychuk et al. 2021]. This allows us to simulate physics on the GPU and gives access to the simulation data through PyTorch tensors [Paszke et al. 2019]. We combine *IsaacGym* with the open-source, GPU-ready PPO implementation *rs_l_rl* [Rudin 2021; Rudin et al. 2021]. This allows us to train 4000 characters in parallel on an NVIDIA GTX 3080, which trains to 14 billion policy steps, equivalent to 13 years of human experience, in 48 hours. This is roughly 2 orders of magnitude faster than CPU-based approaches. For a detailed description of RL algorithm parameters see Appendix.

4 RESULTS

We show that the positions and orientations of Meta’s Quest headset and controllers contain enough signal to reasonably estimate the full-body pose of a user. The framework is able to distinguish between various locomotion modes, turning, and their transitions. It is also able to track motions in which upper and lower body are less correlated like writing on a whiteboard or boxing. The lower body pose matches the user surprisingly accurately, so the correct foot is often in contact at the right time and position. For a qualitative evaluation, we encourage the readers to compare the reconstructed simulated poses with the image references in Figure 3 and in the video. A quantitative evaluation is shown in Table 1 and described in the following.

Table 1. Pose reconstruction from synthesized and real Meta Quest HMD (H) and controllers (C). We show that despite not having sensors on the lower body, our metrics match state-of-the-art methods like PIP [Yi et al. 2022] that use IMUs attached to the legs. Limitations are discussed in Sec. 5.4.

Method	Ours	Ours	Ours	PIP
Sensors	H+2C	H+2C	H	6 IMUs
Test data	Lafan	Real	Real	TC
MPJRE [deg]	5.7	-	-	-
MPJPE [cm]	3.7	-	-	-
RootE [cm]	1.8	-	-	-
SIP [deg]	12.3	-	-	12.9
Jitter [km/s ³]	0.3	0.1	0.2	0.2
MHPE [cm]	3.7	6.3	6.2	-
MHRE [deg]	8.4	14.3	14.7	-

4.1 Quantitative Evaluation

The tracking accuracy for a variety of motions is summarized in Table 1. We evaluate our trained model on synthetic input generated from the Lafan dataset [Harvey et al. 2020] as well as on real

data recorded while wearing the headset. We evaluate poses reconstructed using the HMD and controllers (H+2C), and using only the 6 DoF headset pose (H). We compare our metrics to a state-of-the-art solution [Yi et al. 2022], which uses 6 IMUs, two attached to the legs and evaluated on the TotalCapture (TC) dataset [Trumble et al. 2017].

As metrics we use MPJRE (Mean Per Joint Rotation Error) in degrees and MPJPE (Mean Per Joint Position Error) in centimeters. We also compute the global root error in centimeters. To compare to [Yi et al. 2022], we calculate the “SIP” error, which represents the mean orientation error of the upper arms and legs in the global space in degrees. Jitter represents the smoothness between poses, where smaller is more smooth. This is quantified by the *jerk*, the derivative of the acceleration, and calculated as done in [Yi et al. 2022]. Finally, we compute the MHPE/MHRE (Mean Headset Position/Rotation Error), the error between the the simulated avatars head/hands (+offset) and the global position and orientation of the Quest devices.

5 DISCUSSION

We discuss three applications of our framework. The main application is using sparse sensor information to reconstruct full-body poses (Sec. 5.1). Then, we demonstrate how the same policy can track users of different scale (Sec. 5.2). Finally, we demonstrate how physics simulation can modify tracked motions to generate believable interactions with the environment (Sec. 5.3). We end with limitations of this approach (Sec. 5.4).

5.1 Headset and Controller Tracking

When comparing reconstructed poses in Figure 3, we observe that the framework can distinguish different types of motion. To better understand how the model is achieving this, we plot the real sensor input driving these animations in Figure 4. This specific clip contains three different motion types: writing, walking and jogging. Starting with the writing section, we notice that the vertical head position dips, hinting that the user is crouching and must be bending the knees. Other than that there is little correlation between controllers and head. During the walking section, the controllers oscillate in the horizontal directions (red and blue) relative to the head. This might be a pattern the model uses to identify a walking gait. In contrast to walking, jogging motions have no oscillation in the x direction (red), but displays a very unique and high frequency oscillation of the vertical HMD (green).

5.2 Tracking users of different scale

We show that a single policy can track users of different scale. The user in Figure 3 is 180 cm tall, whereas the user in Figure 5 is 167 cm. The pose reconstruction is done by the same policy, without retraining. The user scale during inference is determined based on the initial height of the HMD, which causes the appropriately scaled avatar to be initialized (see Sec. 3.6 for details).

Apart from tracking users of different scale, we can also use avatars of *different* scale to track the *same* user. Figure 5 shows the same user being tracked by a larger avatar, which is a naive form of retargeting. Since the avatar is larger, but tries to track the position of e.g. the head, it naturally results in a more crouched

Input: Headset and Controllers (H+2C)



Input: Headset only (H)




Fig. 3. Qualitative evaluation of pose reconstruction on motion sequences not used during training, corrected for latency. Most of the reconstructed poses match the video reference footage, despite having no lower-body sensor signal available for reconstruction. The four sequences demonstrate writing on a whiteboard, walking, jogging, and backwards walking with turns. Readers are encouraged to view the video for more examples.

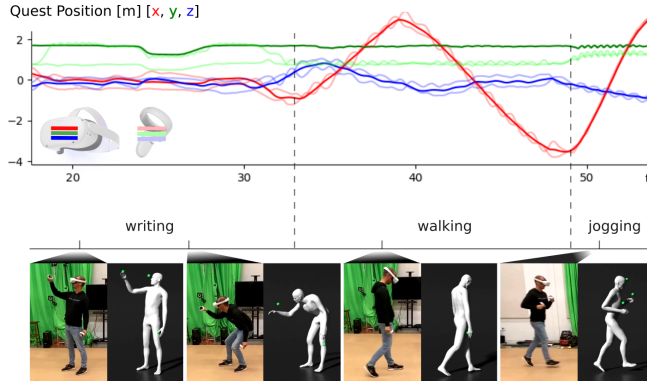


Fig. 4. Sparse input used to generate the full-body pose: the position of the headset (saturated) and left and right controller (less saturated). Y is the vertical dimension. Orientation information is also used by the model but not visualized here.




Fig. 5. In this example a single user controls avatars of two different sizes using the same policy.

position than the user. But generally the avatar imitates the users motion fairly well. More importantly, due to the physics simulator, the motion doesn't violate any physical laws, such as foot skating or floor penetration that often appear when retargeting motions. With this naive baseline, future work could investigate how to further decouple users from avatar representation, allowing them to be embodied by whatever avatar they choose.

5.3 Environment Interaction and Adaptation




Fig. 6. Top: The avatar pose is influenced by the environment and it can interact with virtual objects. Bottom: If a user controls an avatar in an environment different than their play space, the motion is adapted to match the virtual world (e.g. rough ground).

Since the characters are physically simulated (compared to kinematically tracked), their pose is influenced by collisions with the simulated environment (floor, objects). This adaptation of the tracked pose by physics can compensate for missing sensor data and convincingly integrate characters with their virtual world. Two scenarios are shown in Figure 6.

In the top scenario a ball and a rigid object are loaded into the physics simulator. As a user is tracked, we observe a two-way interaction: The avatar influences the external objects (ball kick), and the objects can influence the pose of the avatar (trip hazard). This poses new research questions, for instance, how much should this semi-autonomous avatar be controlled by physics, and when should physics be violated to closely track a user?

The bottom scenario demonstrates an avatar tracking a motion on rough terrain, while the user and sensor signal is recorded on flat

ground. This example demonstrates how the reconstructed motion adapts to unseen virtual worlds to create believable animations, for instance the foot adapts to different slopes in the terrain. A similar application would be a user sitting on a chair, in which case the avatar’s mesh should not penetrate the virtual chair. It might be difficult to measure the user to this accuracy with sensors. However, the simulated collisions and gravity can correct the tracked avatars pose to avoid penetration or hovering.

5.4 Limitations

While this method can produce high quality tracking results for some motions, it can also fail to track others entirely. This is different than kinematic-leaning approaches [Yi et al. 2022], which, on difficult motions, might perform worse (e.g more jitter, less accuracy), but nonetheless still track the motion. But since our motions are simulated without additional non-physical forces, moving the root of the character to a desired position requires a precise sequence of joint torques. Furthermore, physics simulation does not allow teleportation, so as the character drifts further from the user, it can become increasingly difficult to catch up. For these reasons the simulated character can fall when attempting to imitate a dynamic out-of-distribution motion for which it hasn’t yet learned the torque controls (e.g. break-dance, jumping).

Another difficulty comes from uncorrelated upper-lower body motions, resulting in different motions being represented by the same upper body sensor data. In this case, the policy will synthesize a natural and physically-valid lower body pose, but this might not match the user’s pose.

6 CONCLUSION AND FUTURE WORK

We presented a method to track users from sparse sensor data building on approaches from imitation learning. We show that physics simulation compensates for missing sensor information by synthesizing poses in a physically plausible way.

There are a variety of avenues for future work. In terms of motion quality, the simulated avatar can still look stiff and unnatural. Different reward strategies or GAN-based approaches [Peng et al. 2021] might improve the style with the goal of achieving VFX-quality animations from sparse input data. Secondly, due to the reliance on future observations, our real-time system currently has a latency of 160 ms. One way to reduce this and make the tracking more responsive could be to predict future poses while tracking [Yuan and Kitani 2019], maybe using Motion VAEs [Ling et al. 2020]. Next, the scale value (Sec. 3.6) is only a very coarse approximation of the user, since we linearly scale all elements (link length, mass, collision geometry, inertia) of a default skeleton equally. In reality, users of the same height can still have different proportions, resulting in different dynamics [Won and Lee 2019]. In the future we want to supply the policy with more detailed skeleton and body shape information. Finally, we want to increase the diversity of motions the avatars can imitate. This could be achieved using mixture-of-expert policies [Won et al. 2020; Xie et al. 2022], pre-trained low-level controllers that facilitate learning high-level tasks [Peng et al. 2022; Won et al. 2022], or more informative observation representations [Starke et al. 2022].

REFERENCES

- Sadeh Aliakbarian, Pashmina Cameron, Federica Bogo, Andrew Fitzgibbon, and Tom Cashman. 2022. FLAG: Flow-based 3D Avatar Generation from Sparse Observations. In *2022 Computer Vision and Pattern Recognition*. <https://www.microsoft.com/en-us/research/publication/flag-flow-based-3d-avatar-generation-from-sparse-observations/>
- Kevin Bergamin, Simon Clavet, Daniel Holden, and James Richard Forbes. 2019. DReCon: Data-driven Responsive Control of Physics-based Characters. *ACM Trans. Graph.* 38, 6, Article 206 (2019). <http://doi.acm.org/10.1145/3355089.3356536>
- Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. 2019. OpenPose: Real-time Multi-Person 2D Pose Estimation using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019).
- Nuttapong Chentanez, Matthias Müller, Miles Macklin, Viktor Makovychuk, and Stefan Jeschke. 2018. Physics-based motion capture imitation with deep reinforcement learning. In *Motion, Interaction and Games, MIG 2018*. ACM, 1:1–1:10. <https://doi.org/10.1145/3274247.3274506>
- Andrea Dittadi, Sebastian Dziadzio, Darren Cosker, Ben Lundell, Tom Cashman, and Jamie Shotton. 2021. Full-Body Motion From a Single Head-Mounted Device: Generating SMPL Poses From Partial Observations. In *International Conference on Computer Vision 2021*.
- H. Durrant-Whyte and T. Bailey. 2006. Simultaneous localization and mapping: part I. *IEEE Robotics Automation Magazine* 13, 2 (2006), 99–110. <https://doi.org/10.1109/MRA.2006.1638022>
- Levi Fussell, Kevin Bergamin, and Daniel Holden. 2021. SuperTrack: Motion Tracking for Physically Simulated Characters using Supervised Learning. *ACM Trans. Graph.* 40, 6, Article 197 (2021). <https://dl.acm.org/doi/10.1145/3478513.3480527>
- Ruza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. 2018. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7297–7306.
- Félix G. Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. 2020. Robust Motion In-Betweening. 39, 4 (2020).
- Yinghao Huang, Manuel Kaufmann, Emre Aksan, Michael J. Black, Otmar Hilliges, and Gerard Pons-Moll. 2018. Deep Inertial Poser: Learning to Reconstruct Human Pose from Sparse Inertial Measurements in Real Time. *ACM TOG* 37, 6 (12 2018).
- Jiaxi Jiang, Paul Streli, Huajian Qiu, Andreas Fender, Larissa Laich, Patrick Snape, and Christian Holz. 2022a. AvatarPoser: Articulated Full-Body Pose Tracking from Sparse Motion Sensing. <https://doi.org/10.48550/ARXIV.2207.13784>
- Yifeng Jiang, Yuting Ye, Deepak Gopinath, Jungdam Won, Alexander W Winkler, and C Karen Liu. 2022b. Transformer Inertial Poser: Real-time Human Motion Reconstruction from Sparse IMUs with Simultaneous Terrain Generation. *journal = ACM Trans. Graph.* (2022).
- Angjoo Kanazawa, Jason Y. Zhang, Panna Felsen, and Jitendra Malik. 2019. Learning 3D Human Dynamics from Video. In *Computer Vision and Pattern Recognition (CVPR)*.
- Manuel Kaufmann, Yi Zhao, Chengcheng Tang, Lingling Tao, Christopher Twigg, Jie Song, Robert Wang, and Otmar Hilliges. 2021. EM-POSE: 3D Human Pose Estimation from Sparse Electromagnetic Trackers. In *International Conference on Computer Vision (ICCV)*.
- Huvi Yu Ling, Fabio Zinno, George Cheng, and Michiel Van De Panne. 2020. Character controllers using motion vaes. *ACM Transactions on Graphics (TOG)* (2020).
- Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. 2015. SMPL: A Skinned Multi-Person Linear Model. *ACM TOG* 34, 6 (Oct. 2015), 248:1–248:16.
- Zhengyi Luo, Ryo Hachiuma, Ye Yuan, and Kris Kitani. 2021. Dynamics-regulated kinematic policy for egocentric pose estimation. *Advances in Neural Information Processing Systems* 34 (2021).
- Viktor Makovychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. 2021. Isaac Gym: High Performance GPU-Based Physics Simulation For Robot Learning. <https://doi.org/10.48550/ARXIV.2108.10470>
- Josh Merel, Saran Tunyasuvunakool, Arun Ahuja, Yuval Tassa, Leonard Hasenclever, Vu Pham, Tom Erez, Greg Wayne, and Nicolas Heess. 2020. Catch and Carry: Reusable Neural Controllers for Vision-Guided Whole-Body Tasks. *ACM Trans. Graph.* 39, 4, Article 39 (2020). <https://doi.org/10.1145/3386569.3392474>
- Deepak Nagraj, Erik Schake, Patrick Leiner, and Dirk Werth. 2020. An RNN-Ensemble Approach for Real Time Human Pose Estimation from Sparse IMUs. In *Proceedings of the 3rd International Conference on Applications of Intelligent Systems (Las Palmas de Gran Canaria, Spain) (APPIS 2020)*. Article 32, 6 pages.
- Soohwan Park, Hoseok Ryu, Seyoung Lee, Sunmin Lee, and Jehee Lee. 2019. Learning Predict-and-simulate Policies from Unorganized Human Motion Data. *ACM Trans. Graph.* 38, 6, Article 205 (2019). <http://doi.acm.org/10.1145/3355089.3356501>
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019), 8026–8037.
- Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. 2018a. DeepMimic: Example-guided Deep Reinforcement Learning of Physics-based Character

- Skills. *ACM Trans. Graph.* 37, 4, Article 143 (July 2018), 143:1–143:14 pages.
- Xue Bin Peng, Michael Chang, Grace Zhang, Pieter Abbeel, and Sergey Levine. 2019. MCP: Learning Composable Hierarchical Control with Multiplicative Compositional Policies. In *Advances in Neural Information Processing Systems* 32. 3681–3692.
- Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. 2022. ASE: Large-scale Reusable Adversarial Skill Embeddings for Physically Simulated Characters. *ACM Trans. Graph.* 41, 4, Article 94 (July 2022).
- Xue Bin Peng, Angjoo Kanazawa, Jitendra Malik, Pieter Abbeel, and Sergey Levine. 2018b. SFV: Reinforcement Learning of Physical Skills from Videos. *ACM Trans. Graph.* 37, 6, Article 178 (Nov. 2018), 14 pages.
- Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. 2021. AMP: Adversarial Motion Priors for Stylized Physics-Based Character Control. *ACM Trans. Graph.* 40, 4, Article 1 (July 2021), 15 pages. <https://doi.org/10.1145/3450626.3459670>
- Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J Guibas. 2021. Humor: 3d human motion model for robust pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 11488–11499.
- Yu Rong, Takaaki Shiratori, and Hanbyul Joo. 2021. FrankMocap: A Monocular 3D Whole-Body Pose Estimation System via Regression and Integration. In *IEEE International Conference on Computer Vision Workshops*.
- Nikita Rudin. 2021. Github repository: github.com/leggedrobotics/rsl_rl.
- Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. 2021. Learning to Walk in Minutes Using Massively Parallel Deep Reinforcement Learning. <https://doi.org/10.48550/ARXIV.2109.11978>
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. <https://doi.org/10.48550/ARXIV.1707.06347>
- Soshi Shimada, Vladislav Golyanik, Weipeng Xu, and Christian Theobalt. 2020. PhysCap: Physically Plausible Monocular 3D Motion Capture in Real Time. *ACM TOG* 39, 6 (12 2020).
- Sebastian Starke, Ian Mason, and Taku Komura. 2022. DeepPhase: Periodic Autoencoders for Learning Motion Phase Manifolds. *ACM Trans. Graph.* 41, 4, Article 136 (jul 2022), 13 pages. <https://doi.org/10.1145/3528223.3530178>
- Jie Tan, Karen Liu, and Greg Turk. 2011. Stable Proportional-Derivative Controllers. *IEEE Computer Graphics and Applications* 31, 4 (2011), 34–44. <https://doi.org/10.1109/MCG.2011.30>
- Matt Trumble, Andrew Gilbert, Charles Malleson, Adrian Hilton, and John Collomosse. 2017. Total Capture: 3D Human Pose Estimation Fusing Video and Inertial Sensors. In *2017 British Machine Vision Conference (BMVC)*.
- Systems Vicon. 2022. Vicon Motion Systems <https://www.vicon.com/>. Last visited: 01/26/2022.
- Timo von Marcard, Bodo Rosenhahn, Michael Black, and Gerard Pons-Moll. 2017. Sparse Inertial Poser: Automatic 3D Human Pose Estimation from Sparse IMUs. *Computer Graphics Forum* 36(2), *Proceedings of the 38th Annual Conference of the European Association for Computer Graphics (Eurographics)* (2017), 349–360.
- Jungdam Won, Deepak Gopinath, and Jessica Hodgins. 2020. A scalable approach to control diverse behaviors for physically simulated characters. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 33–1.
- Jungdam Won, Deepak Gopinath, and Jessica Hodgins. 2021. Control Strategies for Physically Simulated Characters Performing Two-Player Competitive Sports. *ACM Trans. Graph.* 40, 4, Article 146 (2021). <https://doi.org/10.1145/3450626.3459761>
- Jungdam Won, Deepak Gopinath, and Jessica Hodgins. 2022. Physics-Based Character Controllers Using Conditional VAEs. *ACM Trans. Graph.* 41, 4, Article 96 (jul 2022), 12 pages. <https://doi.org/10.1145/3528223.3530067>
- Jungdam Won and Jehee Lee. 2019. Learning body shape variation in physics-based characters. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–12.
- Zhaoming Xie, Sebastian Starke, Hung Yu Ling, and Michiel van de Panne. 2022. Learning Soccer Juggling Skills with Layer-wise Mixture-of-Experts. (2022).
- Yuanlu Xu, Song-Chun Zhu, and Tony Tung. 2019. Denserac: Joint 3d pose and shape estimation by dense render-and-compare. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7760–7770.
- Xinyu Yi, Yuxiao Zhou, Marc Habermann, Soshi Shimada, Vladislav Golyanik, Christian Theobalt, and Feng Xu. 2022. Physical Inertial Poser (PIP): Physics-aware Real-time Human Motion Tracking from Sparse Inertial Sensors. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xinyu Yi, Yuxiao Zhou, and Feng Xu. 2021. TransPose: Real-time 3D Human Translation and Pose Estimation with Six Inertial Sensors. *ACM TOG* 40, 4 (8 2021).
- Ri Yu, Hwangpil Park, and Jehee Lee. 2021. Human Dynamics from Monocular Video with Dynamic Camera Movements. *ACM Trans. Graph.* 40, 6, Article 208 (2021), 14 pages. <https://doi.org/10.1145/3478513.3480504>
- Ye Yuan and Kris Kitani. 2019. Ego-Pose Estimation and Forecasting as Real-Time PD Control. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 10082–10092.