

Physics-based object-aware ego-centric human pose estimation

Zhengyi Luo

CMU-RI-TR-21-59

August 19, 2021



The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Thesis Committee:

Prof. Kris Kitani, *chair*
Prof. Fernando De La Torre Frade
Prof. David Held
Ye Yuan

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Robotics.*

Copyright © 2021 Zhengyi Luo. All rights reserved.

To all my advisors, mentors, and anyone who has offered me help and advice.

Abstract

We investigate the roles of body kinematics, dynamics, and objects for 3D human pose estimation using a head-mounted camera. Human kinematics models play a key role in encoding the natural range of human motion, while dynamics models can react to the spatial arrangement between humans and objects in the scene. Thus, we propose a method for object-aware 3D egocentric pose estimation that tightly integrates kinematics modeling, dynamics modeling, and scene object information. Unlike prior kinematics or dynamics-based approaches where the two components are used disjointly, we synergize the two approaches via *dynamics-regulated training*. At each timestep, a kinematic model is used to provide a target pose using video evidence and simulation state. Then, a prelearned dynamics model attempts to mimic the kinematic pose in a physics simulator that computes the new simulation state. By comparing the pose instructed by the kinematic model with the pose generated by the dynamics model, we can use their misalignment to further improve the kinematic model. By factoring in the 6DoF pose of objects (e.g., chairs, boxes) in the scene, we demonstrate for the first time, the ability to estimate physically-plausible 3D human-object interactions using a single wearable camera. We evaluate our egocentric pose estimation method in both controlled laboratory settings and real-world scenarios.

Acknowledgments

I would first like to thank my wonderful advisor, Professor Kris Kitani, for being an inspiring, encouraging, and helpful mentor during my time at CMU. Professor Kitani has been incredibly kind, patient, and candid to me throughout my research journey, and helped me grow as a student, researcher, and person. Without his guidance, insight, and support, I would not have been able to push this work to its current form.

I am also extremely grateful to the rest of my committee members, Professor Fernando De La Torre Frade, Professor David Held, and Ye Yuan, for their constant advice and suggestions for this work and academic research in general. Professor De La Torre Frade has been very generous with his time and advice, guiding me to think broadly without losing focus on the problem at hand. I thank Professor Held for agreeing to serve on my committee on a short notice and provide me with valuable feedback about my work and presentation. Ye has been my go-to person for research questions both big and small. He took me under his wing and generously shared his ideas, thoughts, and experiences, without which I would be totally at loss. I can not express my thankfulness enough, and would treasure the learnings and advice I received from them as I move on to my next journey!

Finally, I would like to thank all my friends and family who have been my emotional and mental support as I go through my master's studies during this difficult time. It is your company, love, and encouragement that help me through the stress, pressure, and ups & downs inherent through research. I would not have been able to reach where I am without it.

Contents

1	Introduction	1
2	Background	5
2.1	Third-person human pose estimation.	5
2.2	Egocentric human pose estimation.	6
2.3	Humanoid control inside physics simulation.	6
3	Methods	9
3.1	Dynamics Model - Universal Humanoid Controller (UHC)	10
3.1.1	Formulation.	10
3.1.2	Training procedure.	13
3.1.3	Implementation Details.	14
3.2	Kinematic Model – Object-aware Kinematic Policy.	15
3.2.1	Formulation.	15
3.2.2	Scene context modelling and initialization.	15
3.2.3	Training kinematic policy via supervised learning.	16
3.3	Dynamics-Regulated Training.	18
3.3.1	Test-time evaluation.	20
3.3.2	Implementation details.	21
4	Experiments	23
4.1	Datasets.	23
4.1.1	MoCap dataset.	23
4.1.2	Real-world dataset.	24
4.2	Evaluation metrics.	25
4.3	Comparison with state-of-the-art	27
4.3.1	Baseline methods.	27
4.3.2	MoCap dataset results.	28
4.3.3	Real-world dataset results.	29
4.4	Ablation Study	30
4.4.1	Evaluation of the Universal Humanoid Controller	30
4.4.2	Dynamics-regulated Training	31
4.4.3	Stochasticity	32

5	Conclusions	33
5.1	Failure Cases and Limitations	33
5.2	Conclusion and Future Work	33
5.3	Broader social impact.	34
	Bibliography	35

When this dissertation is viewed as a PDF, the page header is a link to this Table of Contents.

List of Figures

1.1	From egocentric videos, we infer physically-plausible 3D human pose and human-object interaction.	2
3.1	Overview of our dynamics-regulated kinematic policy. Given an egocentric video $\mathbf{I}_{1:T}$, our initialization module $\pi_{\text{KIN}}^{\text{init}}(\mathbf{I}_{1:T})$ computes the object states $\tilde{\mathbf{o}}_{1:T}$, camera poses $\tilde{\mathbf{h}}_{1:T}$, and image features $\phi_{1:T}$. We then roll out our per-step kinematic policy $\pi_{\text{KIN}}^{\text{step}}$ together with the Universal Humanoid Controller to output physically-plausible pose \mathbf{q}_t inside a physics simulator.	10
3.2	Overview of our Universal Dynamics Controller. Given a frame of target pose and current simulation state, our UHC π_{UHC} can dirve the the humanoid to match the target pose.	11
4.1	Our real-world dataset capturing equipment.	24
4.2	Trajectory analysis on our MoCap and real-world datasets. Here we recenter each trajectory using the object position and plot the camera trajectory. The starting point is marked as a red dot.	25
4.3	Results of egocentric pose and human-object interaction estimation from the MoCap dataset.	29

List of Tables

3.1	Hyperparameters used for training the Universal Humanoid Controller.	15
3.2	Hyperparameters used for training the kinematic policy.	20
4.1	Speed analysis of our MoCap dataset and real-world dataset. Unit: (meters/second)	24
4.2	Speed analysis of our MoCap dataset and real-world dataset. Unit: (meters/second)	26
4.3	Quantitative results on pose and physics based metrics on the MoCap and real-world Dataset.	28
4.4	Evaluation of motion imitation for our UHC using target motion from the H36M dataset.	30
4.5	Ablation study of different components of our framework.	31
4.6	Results of our dynamics-regulated kinematic policy on the test split of MoCap and real-world datasets using different random seeds. The “loco” motion in the MoCap dataset corresponds to the generic locomotion action, containing all sequences from the EgoPose [55] Dataset.	32

Chapter 1

Introduction

This thesis focuses on the problem of egocentric human pose estimation. From a video captured by a single head-mounted wearable camera (*e.g.*, smartglasses, action camera, body camera), we aim to infer the wearer’s global 3D full-body pose and interaction with objects in the scene, as illustrated in Fig. 1.1. This is important for applications like virtual and augmented reality, sports analysis, and wearable medical monitoring, where third-person views are often unavailable and proprioception algorithms are needed for understanding the actions of the camera wearer. However, this task is challenging since the wearer’s body is often unseen from a first-person view and the body motion needs to be inferred solely based on the videos captured by the *front-facing* camera. Furthermore, egocentric videos usually capture the camera wearer interacting with objects in the scene, which adds additional complexity in recovering a pose sequence that agrees with the scene context. Despite these challenges, we show that it is possible to infer accurate human motion and human-object interaction from a single head-worn front-facing camera.

Egocentric pose estimation can be solved using two different paradigms: (1) a kinematics perspective and (2) a dynamics perspective. *Kinematics-based approaches* study motion without regard to the underlying forces (*e.g.*, gravity, joint torque) and cannot faithfully emulate human-object interaction without modeling proper contact and forces. They can achieve accurate pose estimates by directly outputting joint angles but can also produce results that violate physical constraints (*e.g.* foot skating and ground penetration). *Dynamics-based approaches*, or *physics-based approaches*,

1. Introduction



Figure 1.1: From egocentric videos, we infer physically-plausible 3D human pose and human-object interaction.

study motions that result from forces. They map directly from visual input to control signals of a human proxy (humanoid) inside a physics simulator and recover 3D poses through simulation. These approaches have the crucial advantage that they output physically-plausible human motion and human-object interaction (*i.e.*, pushing an object will move it according to the rules of physics). However, since no joint torques are captured in human motion datasets, physics-based humanoid controllers are hard to learn, generalize poorly, and are actively being researched [36, 50, 51, 56].

However, unique challenges exist in modelling human dynamics in egocentric vision. Firstly, physics-based humanoid control itself (without visual grounding) is an active research area in vision and graphics [5, 36, 50, 51, 56, 57]. Since no joint torque information is typically captured in human motion datasets and modern physics simulations are not end-to-end differentiable, dynamics-based methods are often optimized through reinforcement learning (RL). In turn, popular RL methods are either limited to mimicking a single clip of motion [36, 56], require high-quality motion capture (MoCap) as input [51], or are constrained to a single type of interaction [5]. Secondly, the handful of works that factor in egocentric visual input and human dynamics [54, 55] are limited to recovering simple motions (walking, running, hopping *etc.*), as estimating non-periodic actions such as sitting down and stepping on a box requires modeling complex and diverse human-object interaction which is challenging.

Thirdly, direct mapping from video to control signals induces compounding errors and covariate shift, which causes large body center drift and catastrophic failures such as losing balance or missing the object. To perform these complex actions in a physics simulator that conforms to the egocentric video evidence, the joint modeling of the human kinematics, dynamics, and scene context is needed.

In this work, we argue that a *hybrid* approach merging the kinematics and dynamics perspectives is needed. Leveraging a large human motion database [29], we first learn a task-agnostic dynamics-based humanoid controller to mimic a wide variety of human behaviors, ranging from every day motion to dancing and kickboxing. The controller is *general-purpose* and can be viewed as providing low-level motor skills of a human. After the controller is learned, we train a scene-aware kinematic policy to specify the target poses for the controller to mimic. One approach to specify the target motion is to let the kinematic model produce estimates only based on the visual input [50, 55, 57]. This approach only uses the physics simulation as a separate post-processing step: the kinematic model is never aware of the simulation state and may output unreasonable target poses. We propose to synchronize the two perspectives by designing a kinematic policy that guides the controller and receives timely feedback (rewards) through comparing its target pose and the resulting simulation state. Our model thus serves as a high-level motion planning module that adapts intelligently based on the current simulation state. In addition, since our kinematic policy only outputs poses and does not model joint torques, it can learn directly from kinematic poses acquired by motion capture (MoCap). While poses from MoCap can provide a initial-guess of target motion, our model can search for better solutions through trial and error. This learning process, dubbed *dynamics-regulated training*, jointly optimizes our model via supervised learning and reinforcement learning, and significantly improves its robustness to real-world use cases.

This thesis is organized as follows: in Chapter 2 we will review the background works in this area, spanning from third person pose estimation, first person pose estimation, and physics-based humanoid control. Chapter 3 will introduce our proposed methods on human dynamics modelling, kinematics modelling, and the hybrid approach that synergize the two components. Here we describe our approach to learn a general-purpose humanoid controller from a large MoCap dataset that can perform a broad range of motions inside a physics simulation; our object-aware

1. Introduction

kinematics model that factors in the current scene context; finally, our dynamics-regulated training procedure that synergizes kinematics, dynamics, and scene context for egocentric vision. Experiments on a controlled motion capture laboratory dataset and a real-world dataset are included in Chapter 4. These experiments demonstrate that our model outperforms other state-of-the-art methods on pose-based and physics-based metrics, while generalizing to videos taken in real-world scenarios. Finally, Chapter 5 discusses the limitation, social impact, and future work.

Chapter 2

Background

2.1 Third-person human pose estimation.

The task of estimating the 3D human pose (and sometimes shape) from *third-person* video is a popular research area in the vision community [2, 10, 11, 12, 14, 21, 23, 24, 28, 32, 34, 40, 53, 57, 58], with methods aiming to recover 3D joint positions [14, 25, 34, 45], 3D joint angles with respect to a parametric human model [2, 21, 23, 28], and dense body parts [11]. Notice that all these methods are purely kinematic and disregard physical reasoning. They also do not recover the global 3D root position and are evaluated by zeroing out the body center (root-relative). A smaller number of works factor in human dynamics [4, 38, 42, 48] through postprocessing and physics-based trajectory optimization. These approaches can produce physically-plausible human motion, but since they do not utilize a physics simulator, they can not faithfully model human-object interaction. SimPoE [57], a recent work on third-person pose estimation using simulated character control, is most related to ours, but 1) trains a single and limited humanoid controller per dataset; 2) designs the kinematic model to be independent from simulation states.

2.2 Egocentric human pose estimation.

Compared to third-person human pose estimation, there are only a handful of attempts at estimating 3D full body poses from egocentric videos due to the ill-posed nature of this task. Most existing methods still assume partial visibility of body parts in the image [39, 47, 52], often through a downward-facing camera. Among works where the human body is mostly not observable [20, 33, 54, 55], Jiang *et al.* [20] use a kinematics-based approach where they construct a motion graph from the training motions and recover the pose sequence by solving the optimal pose path. Ng *et al.* [33] focus on modeling person-to-person interactions from egocentric videos and inferring the wearer’s pose conditioning on the other person’s pose. The works most related to ours are [19, 54, 55] which use dynamics-based approaches and map visual inputs to control signals to perform physically-plausible human motions inside a physics simulation. They show impressive results on a set of noninteractive locomotion tasks, but also observe large errors in absolute 3D position tracking—mapping directly from the visual inputs to control signals is a noisy process and prone to error accumulation. In comparison, our work jointly models kinematics and dynamics, and estimates a wider range of human motion and human-object interactions while improving absolute 3D position tracking. To the best of our knowledge, we are the first approach to estimate the 3D human poses from egocentric video while factoring in human-object interactions.

2.3 Humanoid control inside physics simulation.

Our work is also connected to controlling humanoids to mimic reference motion [5, 6, 15, 36, 37, 51, 56] and interact with objects [5, 30] inside a physics simulator. The core motivation of these works is to learn the necessary dynamics to imitate or generate human motion in a physics simulation. Deep RL has been the predominant approach in this line of work since physics simulators are typically not end-to-end differentiable. Goal-oriented methods [1, 5, 30] does not involve motion imitation and are evaluated on task completion (moving an object, sitting on a chair, moving based on user-input *etc.*). Chao *et al.* [5] propose a hierarchical reinforcement learning approach to generate realistic sitting motion. Their framework is goal-oriented and

evaluated based on task completion, so their RL agent only needs to master a subset of possible sitting motions from different initialization positions to successfully complete the action. In our case, humans We also do not limit our human-object interaction to sitting, and show that our framework generalizes (but is not limited to) actions such as stepping on a box, pushing a box, and avoiding obstacles. Mere *et al.* [30] design an approach for enabling humanoid full-body manipulation and locomotion in simulation. They use a *phased task* in which the task policy is trained to solve different stages of the task and show impressive results in human object manipulation. Similarly, their controller is goal-oriented and not grounded on video evidence. Soonhan *et al.* [15] use an action representation in which the target pose is the sum of the kinematic pose and the output of the policy network. Inspired by their work, we also employ this residual action representation to accelerate training and improve stability. Our task is also related to DeepMimic [36, 37], the Residual Force Controller (RFC) [56], and physics-based controllers for graphics [1, 51]. While DeepMimic and RFC have shown remarkable results in imitating human motion, they are limited to performing a single clip of human motion. Bergamin *et al.* [1] mainly focus on controlling a humanoid through a gamepad and can only perform limited locomotion motion. Won *et al.* [51] develop an impressive dynamics controller that can perform thousands of human motion clips or motions generated by pretrained target pose generators. However, it is limited to perform high-quality MoCap data and requires generator-specific fine-tuning to generalize to unseen target motion generators. It also does not model human-object interaction. In contrast, our dynamics controller is general and can be used to perform everyday motion and human-object interactions estimated by a kinematic motion estimator without task-specific fine-tuning.

2. Background

Chapter 3

Methods

The problem of egocentric pose estimation can be formulated as follows: from a wearable camera footage $\mathbf{I}_{1:T}$, we want to recover the wearer’s ground truth global 3D poses $\hat{\mathbf{q}}_{1:T}$. Each pose $\hat{\mathbf{q}}_t \triangleq (\hat{\mathbf{r}}_t^{\text{pos}}, \hat{\mathbf{r}}_t^{\text{rot}}, \hat{\mathbf{j}}_t^{\text{rot}})$ consists of the root position $\hat{\mathbf{r}}_t^{\text{pos}}$, root orientation $\hat{\mathbf{r}}_t^{\text{rot}}$, and body joint angles $\hat{\mathbf{j}}_t^{\text{rot}}$ of the human model. Here we adopt the popular SMPL [27] human model and the humanoid we use in physics simulation is created from the kinematic structure and mean body shape defined by the SMPL. Our framework first learns a Universal Humanoid Controller (UHC) from a large MoCap dataset (Sec. 3.1). The learned UHC can be viewed as providing the lower level muscle skills of a real human, trained by mimicking thousands of human motion sequences. Using the trained UHC, we learn our kinematic policy (Sec. 3.2) through dynamics-regulated training (Sec. 3.3). At the test time, the kinematic policy provides per-step target motion to the UHC, forming a closed-loop system that operates inside the physics simulation to control a humanoid. The result of the UHC and physics simulation is then used as input to the kinematic policy to produce the next-frame target motion, as depicted in Fig. 3.1. As a notation convention, we use $\tilde{\cdot}$ to denote kinematic quantities (obtained without using physics simulation), $\hat{\cdot}$ to denote ground truth quantities, and normal symbols without accents to denote quantities from the physics simulation.

3. Methods

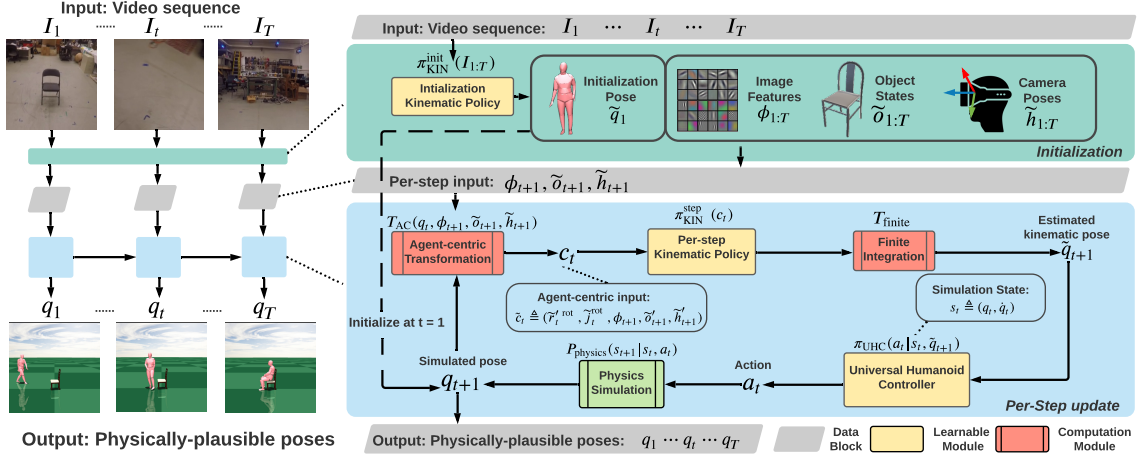


Figure 3.1: Overview of our dynamics-regulated kinematic policy. Given an egocentric video $I_{1:T}$, our initialization module $\pi_{\text{KIN}}^{\text{init}}(I_{1:T})$ computes the object states $\tilde{o}_{1:T}$, camera poses $\tilde{h}_{1:T}$, and image features $\phi_{1:T}$. We then roll out our per-step kinematic policy $\pi_{\text{KIN}}^{\text{step}}$ together with the Universal Humanoid Controller to output physically-plausible pose q_t inside a physics simulator.

3.1 Dynamics Model - Universal Humanoid Controller (UHC)

3.1.1 Formulation.

To learn a task-agnostic dynamics model that can be tightly integrated with a kinematic model, we design our controller’s state space to only rely on the current simulated pose q_t and target pose \hat{q}_{t+1} and remove all phase or sequence-level information found in prior arts [36, 37, 56]. This design allows us to train on an order of magnitude larger dataset of human motion [29] with only pose information and significantly improve our models’ ability to mimic diverse and unseen motions. Formally, we model controlling a humanoid to follow a reference motion $\hat{q}_{1:T}$ as a Markov Decision Process (MDP) defined as a tuple $\mathcal{M} = \langle S, A, P_{\text{physics}}, R, \gamma \rangle$ of states, actions, transition dynamics, reward function, and a discount factor. The state S , reward R , and transition dynamics P_{physics} are provided by the physics simulator, while action A is computed by the policy π_{UHC} . At each timestep t , the agent in state s_t takes an action sampled from the policy while the environment generates the

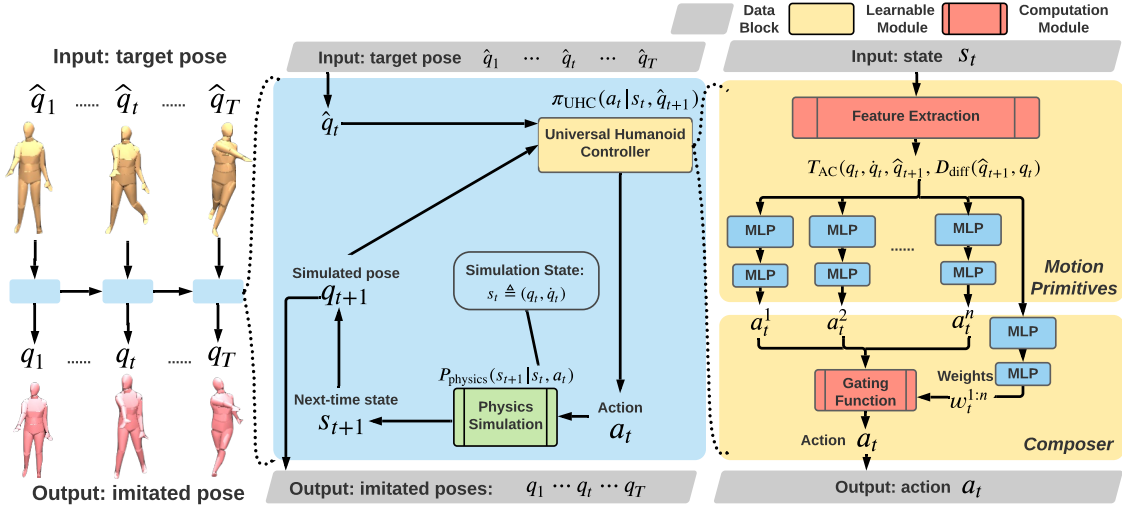


Figure 3.2: Overview of our Universal Dynamics Controller. Given a frame of target pose and current simulation state, our UHC π_{UHC} can drive the humanoid to match the target pose.

next state s_{t+1} and reward r_t .

State. The state $s_t \triangleq (q_t, \dot{q}_t)$ of the humanoid contains the character’s current pose q_t and joint velocity \dot{q}_t . Here, the state s_t encapsulates the humanoid’s full physical state at time step t . It only includes information about the current frame (q_t, \dot{q}_t) and does not include any extra information, enabling our learned controller to be guided by a target pose only.

Action. The action a_t specifies the target joint angles for the proportional derivative (PD) controller [44] at each degree of freedom (DoF) of the humanoid joints except for the root (pelvis). We use the residual action representation [15]: $q_t^d = \hat{q}_{t+1} + a_t$, where q_t^d is the final PD target, a_t is the output of the control policy π_{UHC} , and \hat{q}_{t+1} is the target pose. The torque to be applied at joint i is: $\tau^i = k^p \circ (q_t^d - q_t) - k^d \circ \dot{q}_t$ where k^p and k^d are manually specified gains and \circ is the element-wise multiplication. As observed in prior work [56, 57], allowing the policy to apply external residual forces η_t to the root helps stabilizing the humanoid, so our final action is $a_t \triangleq (\Delta \tilde{q}_t^d, \eta_t)$.

Policy. The policy $\pi_{\text{UHC}}(a_t | s_t, \hat{q}_{t+1})$ takes in the current humanoid state s_t and the target pose \hat{q}_{t+1} to compute the action a_t for the time step t . Its workflow and architecture can be seen in Fig. 3.2. The network $\pi_{\text{UHC}}(a_t | s_t, \hat{q}_{t+1})$ is implemented as a multiplicative compositional policy (MCP) [35] with eight motion primitives,

3. Methods

each being an Multi-layer Perceptron (MLP) with two hidden layers (512, 256). The composer is another MLP with two hidden layers (300, 200) and outputs the multiplicative weights $\mathbf{w}_t^{1:n}$ for the n motion primitives. As studied in MCP [35], this hierarchical control policy increases the model’s capacity to learn multiple skills simultaneously. The output action \mathbf{a}_t is represented by a Gaussian distribution with a fixed diagonal covariance matrix Σ following the formulation of the Policy Gradient algorithm. $\mathbf{a}_t \in \mathbb{R}^{75}$ is a vector concatenation of the target angles of the PD controller mounted on the 23 no-root joints (each has 3 DoF), plus the residual force [56]: $\boldsymbol{\eta}_t \in \mathbb{R}^6$. Recall that each target pose $\hat{\mathbf{q}}_t \in \mathbb{R}^{76}$, $\hat{\mathbf{q}}_t \triangleq (\hat{\mathbf{r}}_t^{\text{pos}}, \hat{\mathbf{r}}_t^{\text{rot}}, \hat{\mathbf{j}}_t^{\text{rot}})$ consists of the root position $\hat{\mathbf{r}}_t^{\text{pos}} \in \mathbb{R}^3$, root orientation in quaternions $\hat{\mathbf{r}}_t^{\text{rot}} \in \mathbb{R}^4$, and body joint angles in Euler angles $\hat{\mathbf{j}}_t^{\text{rot}} \in \mathbb{R}^{69}$ of the human model. The use of quaternions and Euler angles follows the specification of Mujoco [46]. Given the input \mathbf{s}_t and $\hat{\mathbf{q}}_{t+1}$, our UHC first transforms the inputs to a feature vector using $\mathbf{T}_{\text{AC}}(\mathbf{q}_t, \dot{\mathbf{q}}_t, \hat{\mathbf{q}}_{t+1}, D_{\text{diff}}(\hat{\mathbf{q}}_{t+1}, \mathbf{q}_t))$ to output a 640 dimensional vector that is a concatenation of the following values:

$$\mathbf{T}_{\text{AC}} : (\mathbf{q}_t, \dot{\mathbf{q}}_t, \hat{\mathbf{q}}_{t+1}, D_{\text{diff}}(\hat{\mathbf{q}}_{t+1}, \mathbf{q}_t)) \rightarrow (\mathbf{h}_t^q, \mathbf{q}'_t, \hat{\mathbf{q}}'_t, (\mathbf{q}_t - \hat{\mathbf{q}}_t), \dot{\mathbf{q}}_t, (\boldsymbol{\psi}_t - \hat{\boldsymbol{\psi}}_t), (\mathbf{j}_t^{\text{pos}} - \hat{\mathbf{j}}_t^{\text{pos}}), (\mathbf{j}_t^{\prime\text{pos}} - \hat{\mathbf{j}}_t^{\prime\text{pos}}), (\mathbf{j}_t^{\text{rot}} - \hat{\mathbf{j}}_t^{\text{rot}}), (\mathbf{j}_t^{\prime\text{rot}} - \hat{\mathbf{j}}_t^{\prime\text{rot}})) \quad (3.1)$$

It consists of: root orientation $\mathbf{h}_t^q \in \mathbb{R}^4$ in agent-centric coordinates; simulated pose $\mathbf{q}'_t \in \mathbb{R}^{74}$ ($\mathbf{q}'_t \triangleq (\mathbf{r}'_t, \mathbf{r}'_t, \mathbf{j}'_t)$, root height $\mathbf{r}'_t \in \mathbb{R}^1$, root orientation $\mathbf{r}'_t \in \mathbb{R}^4$, and body pose $\mathbf{j}'_t \in \mathbb{R}^{69}$ expressed in Euler angles) in agent-centric coordinates; target pose $\hat{\mathbf{q}}'_t \in \mathbb{R}^{74}$ in agent-centric coordinates; $(\mathbf{q}_t - \hat{\mathbf{q}}_t) \in \mathbb{R}^{76}$ is the difference between the simulated and target pose (in world coordinate), calculated as $(\mathbf{q}_t - \hat{\mathbf{q}}_t) \triangleq (\hat{\mathbf{r}}_t^{\text{pos}} - \mathbf{r}_t^{\text{pos}}, \hat{\mathbf{r}}_t^{\text{rot}} \ominus \mathbf{r}_t^{\text{rot}}, \hat{\mathbf{j}}_t^{\text{rot}} \ominus \mathbf{j}_t^{\text{rot}})$, where \ominus calculates the quaternion difference; $\dot{\mathbf{q}}_t \in \mathbb{R}^{75}$ is the joint velocity computed by Mujoco; $(\boldsymbol{\psi}_t - \hat{\boldsymbol{\psi}}_t) \in \mathbb{R}^1$ is the difference between the current heading (yaw) of the target and simulated root orientation; $(\mathbf{j}_t^{\text{pos}} - \hat{\mathbf{j}}_t^{\text{pos}}) \in \mathbb{R}^{72}$ and $(\mathbf{j}_t^{\prime\text{pos}} - \hat{\mathbf{j}}_t^{\prime\text{pos}}) \in \mathbb{R}^{72}$ are joint position differences, calculated in the global and agent-centric space, respectively; $(\mathbf{j}_t^{\text{rot}} - \hat{\mathbf{j}}_t^{\text{rot}}) \in \mathbb{R}^{96}$ and $(\mathbf{j}_t^{\prime\text{rot}} - \hat{\mathbf{j}}_t^{\prime\text{rot}}) \in \mathbb{R}^{96}$ are joint rotation (quaternion) differences, calculated in the global and agent-centric space, respectively. Note that our feature vector contains redundant information and can be further simplified.

Reward function. For UHC, the reward function is designed to encourage the simulated pose \mathbf{q}_t to better match the target pose $\widehat{\mathbf{q}}_{t+1}$. The imitation reward function per timestep, similar to the reward defined in Yuan *et al.* [56] is as follows:

$$r_t^{\text{UHC}} = w_{\text{jr}}r_{\text{jr}} + w_{\text{jp}}r_{\text{jp}} + w_{\text{jv}}r_{\text{jv}} + w_{\text{res}}r_{\text{res}}, \quad (3.2)$$

where $w_{\text{jr}}, w_{\text{jp}}, w_{\text{jv}}, w_{\text{res}}$ are the weights of each reward. The joint rotation reward r_{jr} measures the difference between the simulated joint rotation $\mathbf{j}_t^{\text{rot}}$ and the target $\widehat{\mathbf{j}}_t^{\text{rot}}$ in quaternion for each joint on the humanoid. The joint position reward r_{jp} computes the distance between each joint’s position $\mathbf{j}_t^{\text{pos}}$ and the target joint position $\widehat{\mathbf{j}}_t^{\text{pos}}$. The joint velocity reward r_{jv} penalizes the deviation of the estimated joint angular velocity $\dot{\mathbf{j}}_t^{\text{rot}}$ from the target $\widehat{\dot{\mathbf{j}}}_t^{\text{rot}}$. The target velocity is computed from the data via finite difference. All above rewards include every joint on the humanoid model (including the root joint), and are calculated in the world coordinate frame. Finally, the residual force reward r_{res} encourages the policy to rely less on the external force and penalize for a large $\boldsymbol{\eta}_t$:

$$\begin{aligned} r_{\text{jr}} &= \exp \left[-2.0 \left(\left\| \mathbf{j}_t^{\text{rot}} \ominus \widehat{\mathbf{j}}_t^{\text{rot}} \right\|^2 \right) \right], & r_{\text{jp}} &= \exp \left[-5 \left(\left\| \mathbf{j}_t^{\text{pos}} - \widehat{\mathbf{j}}_t^{\text{pos}} \right\|^2 \right) \right], \\ r_{\text{jv}} &= \exp \left[-0.005 \left\| \dot{\mathbf{j}}_t^{\text{rot}} \ominus \widehat{\dot{\mathbf{j}}}_t^{\text{rot}} \right\|^2 \right], & r_{\text{res}} &= \exp \left[- \left(\left\| \boldsymbol{\eta}_t \right\|^2 \right) \right]. \end{aligned} \quad (3.3)$$

3.1.2 Training procedure.

Data processing. We train our controller on the AMASS [29] dataset, which contains 11505 high-quality MoCap sequences with 4000k frame of poses (after removing sequences involving human-object interaction like running on a treadmill). The original AMASS dataset contains 13944 motion sequences, and around 2600 of them contain human-object interactions such as sitting on a chair, walking on a treadmill, and walking on a bench. Since AMASS does not contain object information, we can not faithfully recreate and simulate the human-object interactions. Thus, we use a combination of heuristics and visual inspection to remove these sequences. For instance, we detect sitting sequences through finding combinations of the humanoid’s

root, leg, and torso angles that correspond to the sitting posture; we find walking-on-a-bench sequences through detecting a prolonged airborne period; for sequences that are difficult to detect automatically, we conduct manual visual inspection. After the data cleaning process, we obtain 11299 motion sequences that do not contain human-object interaction for our UHC to learn from.

Data sampling and optimization. At the beginning of each episode, a random fixed length sequence (300 frames) is sampled from the dataset for training. While prior works [50, 51] uses more complex motion clustering techniques to sample motions, we devise a simple yet empirically effective sampling technique by inducing a probability distribution based on the value function. For each pose frame $\hat{\mathbf{q}}_j$ in the dataset, we first compute an initialization state \mathbf{s}_1^j : $\mathbf{s}_1^j \triangleq (\hat{\mathbf{q}}_j, \mathbf{0}, \hat{\mathbf{q}}_j)$, and then score it using the value function to access how well the policy can mimic the sequence $\hat{\mathbf{q}}_{j:T}$ starting from this pose: $V(\mathbf{s}_1^j) = v_j$. Intuitively, the higher v_j is, the more confident our policy is in mimicking this sequence, and the less likely we should pick this frame. The probability of choosing frame j , comparing against all frames J in the AMASS dataset, is then $P(\hat{\mathbf{q}}_j) = \frac{\exp(-v_j/\tau)}{\sum_i \exp(-v_i/\tau)}$ where τ is the sampling temperature. We employ Proximal Policy Optimization (PPO) [41] to find the optimal policy π_{UHC}^* that maximizes the expected discounted return $\mathbb{E}[\sum_{t=1}^T \gamma^{t-1} r_t]$.

3.1.3 Implementation Details.

Proxy humanoid. The proxy humanoid we use for simulation is created automatically using the mesh, bone and kinematic tree defined in the popular SMPL [27] human model. Similar to the procedure in [57], given the SMPL body vertices $V = 6890$ and bones $B = 25$, we generate our humanoid based on the skinning weight matrix $\mathbf{W} \in \mathbb{R}^{V \times B}$ that defines the association between each vertex and bone. The geometry of each bone’s mesh is defined by the convex hull of all vertices assigned to the bone. The mass of each bone is in turn defined by the volume of the mesh. To simplify the simulation process, we discard all body shape information from the AMASS [29] dataset, and use the mean body shape of the SMPL model. Since AMASS and our MoCap dataset are recorded by people with different height, we manually adjust the starting height of the MoCap pose to make sure each of the humanoid’s feet are touching the ground at the starting point of the episode.

Hyperparameters. We train our Universal Humanoid Controller for 10000 epoches, which takes about 5 days. Hyperparameters for training the UHC can be found in Table 3.1:

Table 3.1: Hyperparameters used for training the Universal Humanoid Controller.

	γ	Batch Size	Value Learning Rate	Policy Learning Rate	PPO clip ϵ	Covariance Std
Value	0.95	50000	3×10^{-4}	5×10^{-5}	0.2	0.1
	w_{jr}	w_{jp}	w_{jv}	w_{res}	Sampling Temperature	
Value	0.3	0.55	0.1	0.05	2	

3.2 Kinematic Model – Object-aware Kinematic Policy.

3.2.1 Formulation.

To leverage the power of our learned UHC, we design an auto-regressive and scene-aware kinematic policy to generate per-frame target motion from egocentric inputs. We synchronize the state space of our kinematic policy and UHC such that the policy can be learned with or without physics simulation. When trained without physics simulation, the model is purely kinematic and can be optimized via supervised learning; when trained with a physics simulation, the model can be optimized through a combination of supervised learning and reinforcement learning. The latter procedure, coined *dynamics-regulated training*, enables our model to distill human dynamics information inside UHC learned from large-scale MoCap data into the kinematic model and learns a policy more robust to covariate shifts. In this section, we will describe the architecture of the policy itself and the training through supervised learning (without physics simulation).

3.2.2 Scene context modelling and initialization.

To serve as a high-level target motion estimator for egocentric videos with potential human-object interaction, our kinematic policy needs to be object-aware and grounded with visual input. To this end, given an input image sequence $\mathbf{I}_{1:T}$, we compute

3. Methods

the object states $\tilde{\mathbf{o}}_{1:T}$, image features $\phi_{1:T}$, and (optional) camera trajectory $\tilde{\mathbf{h}}_{1:T}$ as inputs to the kinematic policy. Notice that only our kinematic model factors in the scene context, enabling us to train our kinematic and dynamics model independently. The object states, $\tilde{\mathbf{o}}_t \triangleq (\tilde{\mathbf{o}}_t^{cls}, \tilde{\mathbf{o}}_t^{\text{pos}}, \tilde{\mathbf{o}}_t^{\text{rot}})$, is modeled as a vector concatenation of the main object-of-interest’s class $\tilde{\mathbf{o}}_t^{cls}$, 3D position $\tilde{\mathbf{o}}_t^{\text{pos}}$, and rotation $\tilde{\mathbf{o}}_t^{\text{rot}}$. $\tilde{\mathbf{o}}_t$ is computed using an off-the-shelf object detector and pose estimator [17]. When there are no objects in the current scene (for walking and running *etc.*), the object states vector is set to zero. The image features $\phi_{1:T}$ contains crucial information about the wearer’s movement and is computed using an optical flow extractor [43] and ResNet [13]. Since visual input can be noisy and modern smartglasses and bodycams are often equipped with built-in SLAM or Visual Inertial Odometry (VIO) [9, 49] capabilities, we utilize this additional data modalitiy and compute the 6DoF camera pose from input Images. Using an off-the-shelf VIO method [16], we extracts camera trajectory as: $\tilde{\mathbf{h}}_t \triangleq (\tilde{\mathbf{h}}_t^{\text{pos}}, \tilde{\mathbf{h}}_t^{\text{rot}})$ (position $\tilde{\mathbf{h}}_t^{\text{pos}}$ and orientation $\tilde{\mathbf{h}}_t^{\text{rot}}$). Notice that the camera trajectory $\tilde{\mathbf{h}}_t$ information is optional but can further improve the performance of our framework as shown in our ablation studies (Sec. 4.4).

To provide our UHC with a plausible initial state for simulation, we estimate $\tilde{\mathbf{q}}_1$ from the scene context features $\phi_{1:T}$, $\tilde{\mathbf{o}}_{1:T}$, and $\tilde{\mathbf{h}}_{1:T}$. We use an Gated Recurrent Unit (GRU) [7] based network to regress the initial agent pose $\tilde{\mathbf{q}}_1$. Combining the above procedures, we obtain the context modelling and initialization model $\pi_{\text{KIN}}^{\text{init}}: \mathbf{I}_{1:T} \rightarrow \tilde{\mathbf{q}}_1, \phi_{1:T}, \tilde{\mathbf{o}}_{1:T}, \tilde{\mathbf{h}}_{1:T}$.

3.2.3 Training kinematic policy via supervised learning.

After initialization, a per-timestep model $\pi_{\text{KIN}}^{\text{step}}$ is used to compute the next frame pose. At each timestep t we aim to estimate the next frame pose base on the next frame observations: we obtain an egocentric input vector $\tilde{\mathbf{c}}_t$ through the agent-centric transformation function $\tilde{\mathbf{c}}_t = \mathbf{T}_{\text{AC}}(\tilde{\mathbf{q}}_t, \phi_{t+1}, \tilde{\mathbf{o}}_{t+1}, \tilde{\mathbf{h}}_{t+1})$ where $\tilde{\mathbf{c}}_t \triangleq (\tilde{\mathbf{r}}_t^{\text{rot}}, \tilde{\mathbf{j}}_t^{\text{rot}}, \phi_{t+1}, \tilde{\mathbf{o}}_{t+1}, \tilde{\mathbf{h}}_{t+1}')$ contains the current agent-centric root orientation $\tilde{\mathbf{r}}_t^{\text{rot}}$, joint angles $\tilde{\mathbf{j}}_t^{\text{rot}}$, image feature for next frame ϕ_{t+1} , object state $\tilde{\mathbf{o}}_{t+1}'$, and camera pose $\tilde{\mathbf{h}}_{t+1}'$. From $\tilde{\mathbf{c}}_t$, the kinematic policy $\pi_{\text{KIN}}^{\text{step}}$ computes the root angular velocity $\tilde{\mathbf{w}}_t$, linear velocity $\tilde{\mathbf{v}}_t$, and next frame joint rotation $\tilde{\mathbf{j}}_{t+1}^{\text{rot}}$. The next frame pose is computed through a finite integration module $\mathbf{T}_{\text{finite}}$ with time difference $\delta t = 1/30\text{s}$:

Algorithm 1 Learning kinematic policy via supervised learning.

```

1: Input: Egocentric videos  $\mathbf{I}$  and paired ground truth motion dataset  $\widehat{\mathbf{Q}}$ 
2: while not converged do
3:    $\mathbf{M}_{\text{SL}} \leftarrow \emptyset$  ▷ initialize sampling memory
4:   while M not full do
5:      $\mathbf{I}_{1:t} \leftarrow$  random sequence of images  $\mathbf{I}_{1:T}$  from the dataset  $\mathbf{I}$ 
6:      $\tilde{\mathbf{q}}_1, \phi_{1:T}, \tilde{\mathbf{o}}_{1:T}, \tilde{\mathbf{h}}_{1:T} = \pi_{\text{KIN}}^{\text{init}}(\mathbf{I}_{1:T})$  ▷ compute scene context and initial pose
7:     for  $i \leftarrow 1 \dots T$  do
8:        $\tilde{\mathbf{c}}_t \leftarrow \mathbf{T}_{\text{AC}}(\tilde{\mathbf{q}}_t, \phi_{t+1}, \tilde{\mathbf{o}}_{t+1}, \tilde{\mathbf{h}}_{t+1})$  ▷ compute agent-centric input features
9:        $\tilde{\mathbf{q}}_{t+1} \leftarrow \mathbf{T}_{\text{finite}}(\pi_{\text{KIN}}^{\text{step}}(\tilde{\mathbf{c}}_t), \tilde{\mathbf{q}}_t)$ 
10:      store  $(\tilde{\mathbf{q}}_t, \hat{\mathbf{q}}_t)$  into memory  $\mathbf{M}_{\text{SL}}$ 
11:    end for
12:  end while
13:   $\pi_{\text{KIN}}^{\text{step}}, \pi_{\text{KIN}}^{\text{init}} \leftarrow$  supervised learning update using data collected in  $\mathbf{M}_{\text{SL}}$  for 10
    epoches.
14: end while

```

$$\tilde{\omega}_t, \tilde{\mathbf{v}}_t, \tilde{\mathbf{j}}_{t+1}^{\text{rot}} = \pi_{\text{KIN}}^{\text{step}}(\tilde{\mathbf{c}}_t), \quad \tilde{\mathbf{q}}_{t+1} = \mathbf{T}_{\text{finite}}(\tilde{\omega}_t, \tilde{\mathbf{v}}_t, \tilde{\mathbf{j}}_{t+1}^{\text{rot}}, \tilde{\mathbf{q}}_t). \quad (3.4)$$

When trained without physics simulation, we auto-regressively apply the kinematic policy: the input $\tilde{\mathbf{c}}_{t+1}$ for the timestep $t + 1$ is computed using the output of the timestep t . This procedure is outlined at Alg. 1. At the beginning of each training epoch, we sample random sequences of input images and compute the estimated kinematic pose. After sampling sufficient sequences from the dataset for the epoch, we optimize the kinematic policy through mini-batched gradient descent using the Adam optimizer. Since all mentioned calculations are end-to-end differentiable, we can directly optimize our $\pi_{\text{KIN}}^{\text{init}}$ and $\pi_{\text{KIN}}^{\text{step}}$ through supervised learning by comparing the estimated kinematic pose $\tilde{\mathbf{q}}_{1:T}$ with ground truth $\widehat{\mathbf{q}}_{1:T}$. Our loss is computed as the difference between the desired and ground truth values of the following quantities: agent root position ($\widehat{\mathbf{r}}_t^{\text{pos}}$ vs $\tilde{\mathbf{r}}_t^{\text{pos}}$) and orientation ($\widehat{\mathbf{r}}_t^{\text{rot}}$ vs $\tilde{\mathbf{r}}_t^{\text{rot}}$), agent-centric object position ($\widehat{\mathbf{o}}_t^{\text{pos}}$ vs $\tilde{\mathbf{o}}_t^{\text{pos}}$) and orientation ($\widehat{\mathbf{o}}_t^{\text{rot}}$ vs $\tilde{\mathbf{o}}_t^{\text{rot}}$), and agent joint orientation ($\widehat{\mathbf{j}}_t^{\text{rot}}$ vs $\tilde{\mathbf{j}}_t^{\text{rot}}$) and position ($\widehat{\mathbf{j}}_t^{\text{pos}}$ vs $\tilde{\mathbf{j}}_t^{\text{pos}}$, computed using forward kinematics). Given a sequence of kinematic poses of T time steps:

$$\mathcal{L}_{\text{SL}} = \sum_{i=1}^T \|\tilde{\mathbf{r}}_t^{\text{rot}} \ominus \hat{\mathbf{r}}_t^{\text{rot}}\|^2 + \|\tilde{\mathbf{r}}_t^{\text{pos}} - \hat{\mathbf{r}}_t^{\text{pos}}\|^2 + \|\tilde{\mathbf{o}}_t^{\text{rot}} \ominus \hat{\mathbf{o}}_t^{\text{rot}}\|^2 + \|\tilde{\mathbf{o}}_t^{\text{pos}} - \hat{\mathbf{o}}_t^{\text{pos}}\|^2 + \|\tilde{\mathbf{j}}_t^{\text{rot}} \ominus \hat{\mathbf{j}}_t^{\text{rot}}\|^2 + \|\tilde{\mathbf{j}}_t^{\text{pos}} - \hat{\mathbf{j}}_t^{\text{pos}}\|^2. \quad (3.5)$$

3.3 Dynamics-Regulated Training.

So far, our kinematics and dynamics model has been largely disjoint; the kinematic model trained through auto-regressively applying itself has no way of gaining information about physical plausibility and can exhibit typical issues residing in purely kinematic pose estimation methods (foot skating, penetration, unrealistic joint bending etc). To tightly integrate our kinematic and dynamics models, we design a *dynamics-regulated training* procedure, where the kinematic policy learns from explicit physics simulation. In the procedure described in the previous section, the next-frame pose fed into the network is computed through finite integration and is not checked by physical laws: whether a real human can perform the computed pose is never verified. Intuitively, this amounts to mentally think about moving in a physical space *without actually moving*. Combining our UHC and our kinematic policy, we can leverage the prelearned motor skills from UHC and let the kinematic policy act directly in a simulated physical space to obtain feedback about physical plausibility. The procedure for dynamics-regulated training is outlined in Alg. 2. In each episode, we use $\pi_{\text{KIN}}^{\text{init}}$ and $\pi_{\text{KIN}}^{\text{step}}$ as in Alg. 1, with the key distinction being: at the next timestep $t + 1$, the input to the kinematic policy is the result of UHC and physics simulation \mathbf{q}_{t+1} instead of $\tilde{\mathbf{q}}_{t+1}$. \mathbf{q}_{t+1} explicitly verify that the $\tilde{\mathbf{q}}_{t+1}$ produced by the kinematic policy can be successfully followed by a motion controller. Using \mathbf{q}_{t+1} also informs our $\pi_{\text{KIN}}^{\text{step}}$ of the current humanoid state and encourages the policy to adjust its predictions to improve humanoid stability.

Dynamics-regulated optimization. Since the physics simulation is not differentiable, we cannot directly optimize the simulated pose \mathbf{q}_t ; however, we can optimize \mathbf{q}_t through reinforcement learning and $\tilde{\mathbf{q}}_t$ through supervised learning. As we know that $\hat{\mathbf{q}}_t$ is a *good guess* reference motion for UHC, we can directly optimize $\tilde{\mathbf{q}}_t$ via supervised learning as done in Sec. 3.2 using the supervised loss \mathcal{L}_{RL} defined in Eq.

Algorithm 2 Learning kinematic policy via dynamics-regulated training.

```

1: Input: Pre-trained controller  $\pi_{\text{UHC}}$ , egocentric videos  $\mathbf{I}$ , and paired ground truth motion
   dataset  $\tilde{\mathbf{Q}}$ 
2: Train  $\pi_{\text{KIN}}^{\text{init}}$ ,  $\pi_{\text{KIN}}^{\text{step}}$  using Alg. 1 for 20 epoches (optional).
3: while not converged do
4:    $\mathbf{M}_{\text{dyna}} \leftarrow \emptyset$  ▷ initialize sampling memory
5:   while  $\mathbf{M}_{\text{dyna}}$  not full do
6:      $\mathbf{I}_{1:T} \leftarrow$  random sequence of images  $\mathbf{I}_{1:T}$ 
7:      $\mathbf{q}_1 \leftarrow \tilde{\mathbf{q}}_1, \phi_{1:T}, \tilde{\mathbf{o}}_{1:T}, \tilde{\mathbf{h}}_{1:T} \leftarrow \pi_{\text{KIN}}^{\text{init}}(\mathbf{I}_{1:T})$  ▷ compute scene context and initial pose
8:      $\mathbf{s}_1 \leftarrow (\mathbf{q}_1, \dot{\mathbf{q}}_1)$  ▷ compute initial state for simulation
9:     for  $i \leftarrow 1 \dots T$  do
10:       $\mathbf{c}_t \leftarrow \mathbf{T}_{\text{AC}}(\mathbf{q}_t, \phi_{t+1}, \tilde{\mathbf{o}}_{t+1}, \tilde{\mathbf{h}}_{t+1})$  ▷ compute agent-centric features using simulated
pose  $\mathbf{q}_t$ 
11:       $\tilde{\mathbf{q}}_{t+1} \sim \mathbf{T}_{\text{finite}}(\pi_{\text{KIN}}^{\text{step}}(\mathbf{c}_t), \mathbf{q}_t)$  ▷ sample from  $\pi_{\text{KIN}}^{\text{step}}$  as a gaussian policy
12:       $\mathbf{s}_t \leftarrow (\mathbf{q}_t, \dot{\mathbf{q}}_t)$ 
13:       $\mathbf{s}_{t+1} \leftarrow P_{\text{physics}}(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t), \mathbf{a}_t \leftarrow \pi_{\text{UHC}}(\mathbf{a}_t|\mathbf{s}_t, \tilde{\mathbf{q}}_{t+1})$  ▷ physics simulation using  $\pi_{\text{UHC}}$ 
14:       $\mathbf{q}_{t+1} \leftarrow \mathbf{s}_{t+1}, \mathbf{r}_t^{\text{KIN}} \leftarrow$  reward from Eq. 3.6 ▷ extract reward and  $\mathbf{q}_{t+1}$  from
simulation
15:      store  $(\mathbf{s}_t, \mathbf{a}_t, \mathbf{r}_t, \mathbf{s}_{t+1}, \hat{\mathbf{q}}_t, \tilde{\mathbf{q}}_{t+1})$  into memory  $\mathbf{M}_{\text{dyna}}$ 
16:    end for
17:  end while
18:   $\pi_{\text{KIN}}^{\text{step}} \leftarrow$  Reinforcement learning updates using experiences collected in  $\mathbf{M}_{\text{dyna}}$  for 10 epoches.
19:   $\pi_{\text{KIN}}^{\text{init}}, \pi_{\text{KIN}}^{\text{step}} \leftarrow$  Supervised learning update using experiences collected in  $\mathbf{M}_{\text{dyna}}$  for 10
   epoches.
20: end while

```

3.5. Since the data samples are collected through physics simulation, the input \mathbf{q}_t is physically-plausible and more diverse than those collected purely through auto-regressively applying $\pi_{\text{KIN}}^{\text{step}}$ in Alg. 1. This way, our dynamics-regulated training procedure performs a powerful data augmentation step, exposing $\pi_{\text{KIN}}^{\text{step}}$ with diverse states collected from simulation.

However, MoCap pose $\hat{\mathbf{q}}_t$ is imperfect and can contain physical violations itself (foot-skating, penetration *etc.*), so asking the policy $\pi_{\text{KIN}}^{\text{step}}$ to produce $\hat{\mathbf{q}}_t$ as reference motion *regardless of the current humanoid state* can lead to instability and cause the humanoid to fall. The kinematic policy should adapt to the current simulation state and provide reference motion $\tilde{\mathbf{q}}_t$ that can lead to poses similar to $\hat{\mathbf{q}}_t$ yet still physically-plausible. Such behavior will not emerge through supervised learning and require *trial and error*. Thus, we optimize $\pi_{\text{KIN}}^{\text{step}}$ through reinforcement learning and reward maximization. We design our RL reward to have two components: motion imitation and dynamics self-supervision. The motion imitation reward encourages

3. Methods

the policy to match the computed camera trajectory $\tilde{\mathbf{h}}_t$ and MoCap pose $\hat{\mathbf{q}}_t$, and serves as a regularization on motion imitation quality. The dynamics self-supervision reward is based on the insight that the disagreement between $\tilde{\mathbf{q}}_t$ and \mathbf{q}_t contains important information about the quality and physical plausibility of $\tilde{\mathbf{q}}_t$: the better $\tilde{\mathbf{q}}_t$ is, the easier it should be for UHC to mimic it. Formally, we define the reward/loss for $\boldsymbol{\pi}_{\text{KIN}}^{\text{step}}$ as:

$$r_t^{\text{KIN}} = -\mathcal{L}_{\text{RL}} = w_{\text{hp}} e^{-45.0(\|\mathbf{h}^{\text{pos}}_t - \tilde{\mathbf{h}}_t^{\text{pos}}\|^2)} + w_{\text{hq}} e^{-45.0(\|\mathbf{h}^{\text{rot}}_t \ominus \tilde{\mathbf{h}}_t^{\text{rot}}\|^2)} + w_{\text{jr}}^{\text{gt}} e^{-0.005(\|\dot{\mathbf{j}}_t^{\text{rot}} \ominus \widehat{\dot{\mathbf{j}}}_t^{\text{rot}}\|^2)} + w_{\text{jr}}^{\text{gt}} e^{-50.0(\|\mathbf{j}_t^{\text{rot}} \ominus \widehat{\mathbf{j}}_t^{\text{rot}}\|^2)} + w_{\text{jr}}^{\text{dyna}} e^{-50.0(\|\mathbf{j}_t^{\text{rot}} \ominus \tilde{\mathbf{j}}_t^{\text{rot}}\|^2)} + w_{\text{jp}}^{\text{dyna}} e^{-50.0(\|\mathbf{j}_t^{\text{pos}} - \tilde{\mathbf{j}}_t^{\text{pos}}\|^2)}, \quad (3.6)$$

w_{hp} , w_{hq} are weights for matching the extracted camera position $\tilde{\mathbf{h}}_t^{\text{pos}}$ and orientation $\tilde{\mathbf{h}}_t^{\text{rot}}$; $w_{\text{jr}}^{\text{gt}}$, $w_{\text{jp}}^{\text{gt}}$ are for matching ground truth joint angles $\widehat{\mathbf{j}}_t^{\text{rot}}$ and angular velocities $\widehat{\dot{\mathbf{j}}}_t^{\text{rot}}$. $w_{\text{jr}}^{\text{dyna}}$, $w_{\text{jp}}^{\text{dyna}}$ are weights for the dynamics self-supervision rewards, encouraging the policy to match the target kinematic joint angles $\tilde{\mathbf{j}}_t^{\text{rot}}$ and positions $\tilde{\mathbf{j}}_t^{\text{pos}}$ to the simulated joint angles $\mathbf{j}_t^{\text{rot}}$ and positions $\mathbf{j}_t^{\text{pos}}$. As demonstrated in Sec. 4.4, the RL loss is particularly helpful in adapting to challenging real-world sequences, which requires the model to adjust to domain shifts and unseen motion.

3.3.1 Test-time evaluation.

At the test time, we follow the same per-episode procedure outlined in Alg 2 and Fig.3.1 to roll out our policy to obtain simulated pose $\mathbf{q}_{1:T}$ given a sequence of images $\mathbf{I}_{1:T}$. The only difference being instead of sampling from $\boldsymbol{\pi}_{\text{KIN}}^{\text{step}}(\mathbf{c}_t)$ as a Gaussian policy, we output the mean action directly.

Table 3.2: Hyperparameters used for training the kinematic policy.

	γ	Batch Size	Value Learning Rate	Policy Learning Rate	PPO clip ϵ	Covariance Std
Value	0.95	10000	3×10^{-4}	5×10^{-4}	0.2	0.04
	w_{hp}	w_{hq}	$w_{\text{jr}}^{\text{gt}}$	$w_{\text{jp}}^{\text{gt}}$	$w_{\text{jr}}^{\text{dyna}}$	$w_{\text{jp}}^{\text{dyna}}$
Value	0.15	0.15	0.2	0.1	0.2	0.2

3.3.2 Implementation details.

The kinematic policy is implemented as a Gated Recurrent Unit (GRU) [7] based network with 1024 hidden units, followed by a three-layer MLP (1024, 512, 256) with ReLU activation. The value function for training the kinematic policy through reinforcement learning is a two-layer MLP (512, 256) with ReLU activation. We use a fixed diagonal covariance matrix and train for 1000 epoches using the Adam [22] optimizer. Hyperparameters for training can be found in Table. 3.2.

3. Methods

Chapter 4

Experiments

4.1 Datasets.

As no public dataset contains synchronized ground-truth full-body pose, object pose, and egocentric videos with human-object interactions, we record two egocentric datasets: one inside a MoCap studio, another in the real-world. In the following sections, we will describe the dataset capture procedure and their statistics in detail.

4.1.1 MoCap dataset.

Our MoCap dataset (202 training sequences, 64 testing sequences, in total 148k frames) is captured in a MoCap studio with three different subjects. Each motion clip contains paired first-person footage of a person performing one of the five tasks: sitting down and (standing up from) a chair, avoiding an obstacle, stepping on a box, pushing a box, and generic locomotion (walking, running, crouching). Each action has around 50 sequences. The locomotion part of our dataset is merged from the egocentric dataset from EgoPose [54] since the two datasets are captured using a compatible system. MoCap markers are attached to the camera wearer and the objects to get the 3D full-body human pose and 6DoF object pose. To diversify the way actions are performed, we instruct the actors to vary their performance for each action (varying starting position and facing gait, speed *etc.*). We followed the Institutional Review Board’s guidelines and obtained approval for the collection of

4. Experiments

Table 4.1: Speed analysis of our MoCap dataset and real-world dataset. Unit: (meters/second)

MoCap dataset					Real-world dataset				
Action	Mean	Min	Max	Std	Action	Mean	Min	Max	Std
Sit	0.646	0.442	0.837	0.098	Sit	0.556	0.227	0.891	0.171
Push	0.576	0.320	0.823	0.119	Push	0.526	0.234	0.762	0.127
Avoid	0.851	0.567	1.084	0.139	Avoid	0.668	0.283	0.994	0.219
Step	0.844	0.576	1.029	0.118	Step	0.729	0.395	1.092	0.196

this dataset. To study the diversity of our MoCap dataset, we plot the trajectory taken by the actors in Fig. 4.2. We can see that our trajectories are diverse and are spread out around a circle with varying distance from the objects. Table 4.2 shows the speed statistics for our MoCap dataset.

4.1.2 Real-world dataset.

Our real world dataset (183 testing sequences, in total 55k frames) is captured in everyday settings (living room and hallway) with an additional subject.

It contains the same four types of interactions as our MoCap dataset and is captured from a head-mounted iPhone using a VR headset (demonstrated in Fig.4.1). Each action has around 40 sequences. As can be seen in the camera trajectory in Fig. 4.2, the real-world dataset is more heterogeneous than the MoCap dataset, and has more curves and banks overall. Speed analysis in Table 4.2 also shows that our real-world dataset has a larger standard deviation in terms of walking velocity and has a larger overall spread than the MoCap dataset. In all, our real-world dataset has more diverse trajectories and motion patterns than our MoCap dataset, and our dynamics-regulated kinematic policy can still estimate the sequences recorded in this dataset.



Figure 4.1: Our real-world dataset capturing equipment.

Notice that our framework is starting *position and orientation invariant*, since all

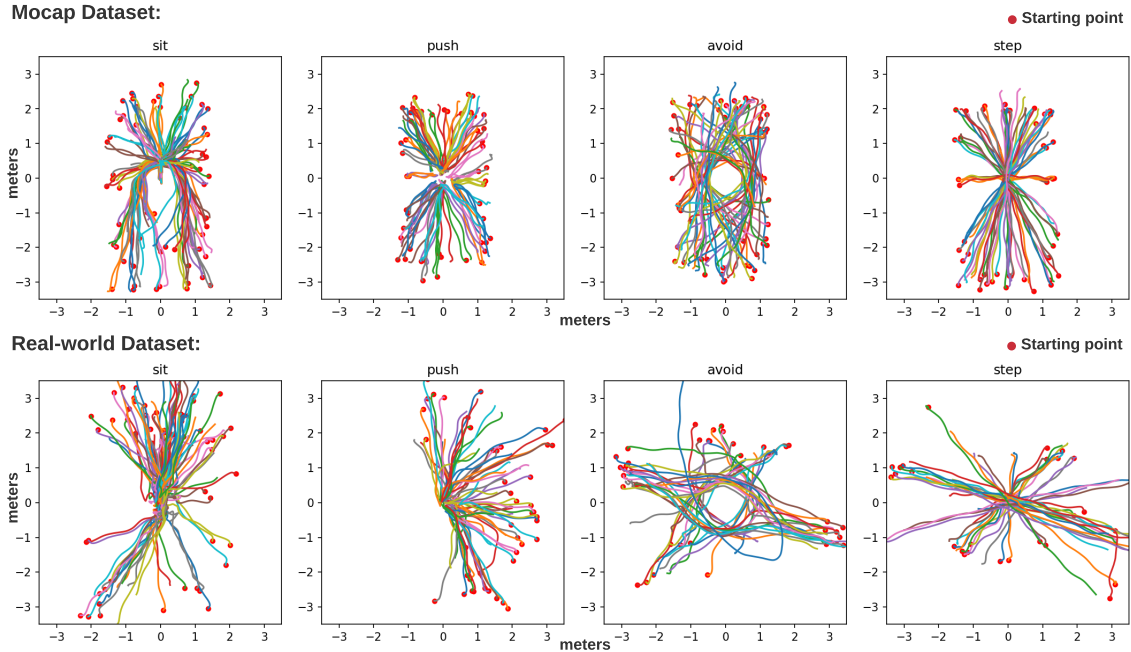


Figure 4.2: Trajectory analysis on our MoCap and real-world datasets. Here we recenter each trajectory using the object position and plot the camera trajectory. The starting point is marked as a red dot.

of our input features are transformed into the agent-centric coordinate system using the transformation function T_{AC} .

4.2 Evaluation metrics.

We use both pose-based and physics-based metrics for evaluation. To evaluate the 3D global pose accuracy, we report the root pose error (E_{root}) and root-relative mean per joint position error [23] (E_{mpjpe}). When ground-truth root/pose information is unavailable (for real-world dataset), we substitute E_{root} with E_{cam} to report camera pose tracking error. We also employ four physics based pose metrics: acceleration error (E_{acc}), foot skating (FS), penetration (PT), and interaction success rate (S_{inter}). E_{acc} (mm/frame^2) compares the ground truth and estimated average joint acceleration; FS (mm) is defined the same as in Ling *et al.* [26]; PT (mm) measures the average penetration distance between our humanoid and the scene (ground floor and objects). Notice that our MoCap dataset has *an penetration of 7.182 mm and foot sliding of*

4. Experiments

Table 4.2: Speed analysis of our MoCap dataset and real-world dataset. Unit: (meters/second)

MoCap dataset					Real-world dataset				
Action	Mean	Min	Max	Std	Action	Mean	Min	Max	Std
Sit	0.646	0.442	0.837	0.098	Sit	0.556	0.227	0.891	0.171
Push	0.576	0.320	0.823	0.119	Push	0.526	0.234	0.762	0.127
Avoid	0.851	0.567	1.084	0.139	Avoid	0.668	0.283	0.994	0.219
Step	0.844	0.576	1.029	0.118	Step	0.729	0.395	1.092	0.196

2.035 mm per frame, demonstrating that the MoCap data is imperfect and may not serve as the best target motion. S_{inter} is defined as whether the objects of interest has been moved enough (pushing and avoiding) or if desired motion is completed (stepping and sitting). If the humanoid falls down at any point, $S_{\text{inter}} = 0$.

- Root error: E_{root} compares the estimated and ground truth root rotation and orientation, measuring the difference in the respective 4×4 transformation matrix (\mathbf{M}_t): $\frac{1}{T} \sum_{t=1}^T \|I - (\mathbf{M}_t \widehat{\mathbf{M}}_t^{-1})\|_F$. This metric reflects both the position and orientation tracking quality.
- Mean per joint position error: E_{mpjpe} (mm) is the popular 3D human pose metric [21, 23, 24] and is defined as $\frac{1}{J} \|\mathbf{j}^{\text{pos}} - \widehat{\mathbf{j}}^{\text{pos}}\|_2$ for J number of joints. This value is root-relative and is computed after setting the root translation to zero.
- Acceleration error: E_{acc} (mm/frame²) measures the difference between the ground truth and estimated joint position acceleration: $\frac{1}{J} \|\ddot{\mathbf{j}}^{\text{pos}} - \widehat{\ddot{\mathbf{j}}}^{\text{pos}}\|_2$.
- Foot sliding: FS (mm) is computed similarly as in [26], *i.e.* $\text{FS} = d(2 - 2^{h/H})$ where d is the foot displacement and h is the foot height of two consecutive poses. We use a height threshold of $H = 33$ mm, the same as in [26].
- Penetration: PT (mm) is provided by the physics simulation. It measures the per-frame average penetration distance between our simulated humanoid and the scene (ground and objects). Notice that Mujoco uses a soft contact model where a larger penetration will result in a larger repulsion force, so a small amount of penetration is expected.
- Camera trajectory error: E_{cam} is defined the same as the root error, and measures the camera trajectory tracking instead of the root. To extract the

camera trajectory from the estimated pose \mathbf{q}_t , we use the head pose of the humanoid and apply a delta transformation based on the camera mount’s vertical and horizontal displacement from the head.

- Human-object interaction success rate: S_{inter} measures whether the desired human-object interaction is successful. If the humanoid falls down at any point during the sequence, the sequence is deemed unsuccessful. The success rate is measured automatically by querying the position, contact, and simulation states of the objects and humanoid. For each action:
 - Sitting down: successful if the humanoid’s pelvis or the roots of both legs come in contact with the chair at any point in time.
 - Pushing a box: successful if the box is moved more than 10 cm during the sequence.
 - Stepping on a box: successful if the humanoid’s root is raised at least 10 cm off the ground and either foot of the humanoid has come in contact with the box.
 - Avoiding an obstacle: successful if the humanoid has not come in contact with the obstacle and the ending position of the root/camera is less than 50 cm away from the desired position (to make sure the humanoid does not drift far away from the obstacle).

4.3 Comparison with state-of-the-art

4.3.1 Baseline methods.

To show the effectiveness of our framework, we compare with the previous state-of-the-art egocentric pose estimation methods: (1) the best dynamics-based approach *EgoPose* [55] and (2) the best kinematics-based approach, also proposed in [55], *PoseReg*. We use the official implementation and augment their input with additional information ($\tilde{\mathbf{o}}_t$ and $\tilde{\mathbf{h}}_t$) for a fair comparison.

4. Experiments

Table 4.3: Quantitative results on pose and physics based metrics on the MoCap and real-world Dataset.

MoCap dataset									
Method	Physics	$S_{\text{inter}} \uparrow$	$E_{\text{root}} \downarrow$	$E_{\text{mpjpe}} \downarrow$	$E_{\text{acc}} \downarrow$	FS \downarrow	PT \downarrow		
PoseReg	\times	-	0.857	87.680	12.981	8.566	42.153		
Kin_poly: supervised learning (ours)	\times	-	0.176	33.149	6.257	5.579	10.076		
EgoPose	\checkmark	48.4%	1.957	139.312	9.933	2.566	7.102		
Kin_poly: dynamics-regulated (ours)	\checkmark	96.9%	0.205	40.443	7.064	2.474	0.686		

Real-world dataset									
Method	Physics	$S_{\text{inter}} \uparrow$	$E_{\text{cam}} \downarrow$	FS \downarrow	PT \downarrow	Per class success rate $S_{\text{inter}} \uparrow$			
						Sit	Push	Avoid	Step
PoseReg	\times	-	1.260	6.181	50.414				
Kin_poly: supervised learning (ours)	\times	-	0.491	5.051	34.930				
EgoPose	\checkmark	9.3%	1.896	2.700	1.922	7.93%	6.81%	4.87%	0.2%
Kin_poly: dynamics-regulated (ours)	\checkmark	92.3%	0.476	2.742	1.229	98.4%	90.9%	100%	74.2%

4.3.2 MoCap dataset results.

Table 4.3 shows the quantitative comparison of our method with the baselines. All results are averaged across five actions and all models have access to the same inputs. We observe that our method, trained either with supervised learning or dynamics-regulated, outperform the two state-of-the-art methods across all metrics. Not surprisingly, our purely kinematic model performs the best on pose-based metrics, while our dynamics-regulated trained policy excels at the physics-based metrics. Comparing the kinematics-only models we can see that our method has a much lower (79.4% error reduction) root and joint position error (62.1% error reduction) than PoseReg, which shows that our object-aware and autoregressive design of the kinematic model can better utilize the provided visual and scene context and avoid compounding errors. Comparing with the dynamics-based methods, we find that the humanoid controlled by EgoPose has a much larger root drift, often falls down to the ground, and has a much lower success rate in human-object interaction (48.4 % vs 96.9%). Upon visual inspection in Fig. 4.3, we can see that our kinematic policy can faithfully produce human-object interaction on almost every test sequence from our MoCap dataset, while PoseReg and EgoPose often miss the object-of-interest (as can be reflected by the large root tracking error). Both of the dynamics-based methods has smaller acceleration error, foot skating, and penetration; some even smaller than MoCap (which has 2 mm FS and 7mm PT).




Figure 4.3: Results of egocentric pose and human-object interaction estimation from the MoCap dataset.

4.3.3 Real-world dataset results.

The real-world dataset is far more challenging, having similar number of sequences (183 clips) as our training set (202 clips) and recorded using different equipment, environments, and motion patterns. Since no ground-truth 3D poses are available, we report our results on camera tracking and physics-based metrics. As shown in Table 4.3, our method outperforms the baseline methods by a large margin in almost all metrics: although EgoPose has less foot-skating (as it also utilizes a physics simulator), its human-object interaction success rate is extremely low. This can be also be reflected by the large camera trajectory error, indicating that the humanoid is drifting far away from the objects. The large drift can be attributed to the domain shift and challenging locomotion from the real-world dataset, causing EgoPose’s humanoid controller to accumulate error and lose balance easily. On the other hand, our method is able to generalize and perform successful human-object interactions, benefiting from our pretrained UHC and kinematic policy’s ability to adapt to new domains and motion. Table 4.3 also shows a success rate breakdown by action. Here we can see that “stepping on a box” is the most challenging action as it requires the humanoid lifting its feet at a precise moment and pushing itself up. Note that our

UHC has never been trained on any stepping or human-object interaction actions (as AMASS has no annotated object pose) but is able to perform these action. As motion is best seen in videos, we refer readers to our [supplementary video](#).

4.4 Ablation Study

4.4.1 Evaluation of the Universal Humanoid Controller

To evaluate our Universal Humanoid Controller, we use the popular Human 3.6M (H36M) dataset [18]. We first fit the SMPL body model to ground truth 3D keypoints similar to the process in [31] and obtain motion sequences in SMPL parameters. Notice that this fitting process is imperfect and the resulting motion sequence is of less quality than the original MoCap sequence. These sequences

Table 4.4: Evaluation of motion imitation for our UHC using target motion from the H36M dataset.

Method	H36M dataset			
	$S_{\text{inter}} \uparrow$	$E_{\text{root}} \downarrow$	$E_{\text{mpjpe}} \downarrow$	$E_{\text{acc}} \downarrow$
DeepMimic	0.0%	0.609	107.895	28.881
UHC w/o MCP	89.3%	0.200	36.972	4.723
UHC	92.0%	0.194	40.424	3.672

are also never seen by our UHC during training. As observed in Moon *et al.* [31], the fitted SMPL poses have a mean per joint position error of around 10mm. We use the train split of H36M (150 unique motion sequences) as the target pose for our UHC to mimic. We use a subset of metrics from egocentric pose estimation to evaluate the motion imitation results of UHC. Namely, we report S_{inter} , E_{root} , E_{mpjpe} , E_{acc} , where the human-object interaction S_{inter} indicates whether the humanoid has become unstable and falls down during the imitation process. The baseline we compare against is the popular motion imitation method DeepMimic [36]. Since our framework uses a different physics simulation (Bullet [8] vs Mujoco [46]), we use an in-house implementation of DeepMimic. From the results shown in Table 4.4, we can see that our UHC can imitate the unseen motion in H36M with high accuracy and success rate, and outperforms the baseline method significantly. Upon visual inspection, we can see that the failure cases often result from losing balance while the humanoid is crouching down or starts running suddenly. Since our controller does not use any sequence level information, it has no way of knowing the upcoming speedup of the

target motion and can result in instability. This indicates the importance of the kinematic policy adjusting its target pose based on the current simulation state to prevent the humanoid from falling down, and signifies that further investigation is needed to obtain a better controller. For visual inspection of motion imitation quality and failure cases, please refer to our supplementary video.

4.4.2 Dynamics-regulated Training

To evaluate the importance of our components, we train our kinematic policy under different configurations and study its effects on the *real-world dataset*, which is much harder than the MoCap dataset.

The results are summarized in Table 4.5. Row 1 (R1) corresponds to training the kinematic policy only with Alg. 1 only and use UHC to mimic the target kine-

Table 4.5: Ablation study of different components of our framework.

Component				Metric			
SL	Dyna_reg	RL	VIO	$S_{\text{inter}} \uparrow$	$E_{\text{cam}} \downarrow$	FS \downarrow	PT \downarrow
✓	✗	✗	✓	73.2%	0.611	4.234	1.986
✓	✓	✗	✓	80.9%	0.566	3.667	4.490
✓	✓	✓	✗	54.1%	1.129	7.070	5.346
✓	✓	✓	✓	92.3%	0.476	2.742	1.229

matic motion as a post-processing step. Row 2 (R2) are the results of using dynamics-regulated training but only performs the supervised learning part. R3 show a variant trained without the estimated camera pose from VIO. Comparing R1 and R2, the lower interaction success rate (73.2% vs 80.9%) indicates that exposing the kinematic policy to states from the physics simulation serves as a powerful data augmentation step and leads to a model more robust to real-world scenarios. R2 and R4 show the benefit of the RL loss in dynamics-regulated training: allowing the kinematic policy to deviate from the MoCap poses makes the model more adaptive and achieves higher success rate. R3 and R4 demonstrate the importance of *intelligently* incorporating extracted camera pose as input: visual features ϕ_t can be noisy and suffer from domain shifts, and using techniques such as SLAM and VIO to extract camera poses as an additional input modality can largely reduce the root drift. Note that our kinematic policy *without using* extracted camera trajectory outperforms EgoPose that *uses camera pose* in both success rate and camera trajectory tracking. Upon visual inspection, the humanoid in R3 largely does not fall down (compared to in EgoPose) and mainly attributes the failure cases to drifting too far from the object.

4. Experiments

Table 4.6: Results of our dynamics-regulated kinematic policy on the test split of MoCap and real-world datasets using different random seeds. The “loco” motion in the MoCap dataset corresponds to the generic locomotion action, containing all sequences from the EgoPose [55] Dataset.

MoCap dataset										
$S_{\text{inter}} \uparrow$	$E_{\text{root}} \downarrow$	$E_{\text{mpjpe}} \downarrow$	$E_{\text{acc}} \downarrow$	FS \downarrow	PT \downarrow	Per class success rate $S_{\text{inter}} \uparrow$				
						Sit	Push	Avoid	Step	Loco
96.87% \pm 1.27%	0.21 \pm 0.01	39.46 \pm 0.52	6.27 \pm 0.1	3.22 \pm 0.11	0.69 \pm 0.03	100%	97.20% \pm 3.96%	100%	86.70% \pm 4.71%	97.4% \pm 3.63%

Real-world dataset									
$S_{\text{inter}} \uparrow$	$E_{\text{root}} \downarrow$	FS \downarrow	PT \downarrow	Per class success rate $S_{\text{inter}} \uparrow$					
				Sit	Push	Avoid	Step		
92.17% \pm 1.41%	0.49 \pm 0.01	2.72 \pm 0.03	1.03 \pm 0.16	94.7% \pm 4.20%	93.10% \pm 1.84%	100.0%	77.1% \pm 2.37%		

4.4.3 Stochasticity

Our kinematic policy is trained through physics simulation and samples a random sequence from the MoCap dataset for each episode. Here we study the stochasticity that rises from this process. We train our full pipeline with three different random seeds and report its results with error bars on both the MoCap test split and the real-world dataset. As can be seen in Table 4.6, our method has very small stochasticity and maintains high performance on *both* the MoCap test split and the real-world dataset, demonstrating the robustness of our dynamics-regulated kinematic policy. Across different random seeds, we can see that “stepping” is consistently the hardest action and “avoiding” is the easiest. Intuitively, “stepping” requires precise coordination between the kinematic policy and the UHC for lifting the feet and pushing up, while “avoiding” only requires basic locomotion skills.

Chapter 5

Conclusions

5.1 Failure Cases and Limitations

Although our method can produce realistic human pose and human-object interaction estimation from egocentric videos, we are still at the early stage of this challenging task. Our method performs well in the MoCap studio setting and generalizes to real-world settings, but is limited to a predefined set of interactions where we have data to learn from. Object class and pose information is computed by off-the-shelf methods such as Apple’s [17], and we only factor in the 6DoF object pose in our state representation and discard all other object geometric information. The lower success rate on the real-world dataset also indicates that our method still suffers from covariate shifts and can become unstable when the shift becomes too extreme. Our Universal Humanoid Controller can imitate everyday motion with high accuracy, but can still fail at extreme motion. To enable pose and human-object interaction estimation for arbitrary actions and objects, better scene understanding and kinematic motion planning techniques need to be developed.

5.2 Conclusion and Future Work

In this paper, we tackle, for the first time, estimating physically-plausible 3D poses from an egocentric video while the person is interacting with objects. We collect a

motion capture dataset and real-world dataset to develop and evaluate our method, and extensive experiments have shown that our method outperforms all prior arts. We design a dynamics-regulated kinematic policy that can be directly trained and deployed inside a physics simulation, and we propose a general-purpose humanoid controller that can be used in physics-based vision tasks easily. Through our real-world experiments, we show that it is possible to estimate 3D human poses and human-object interactions from just an egocentric view captured by consumer hardware (iPhone). In the future, we would like to support more action classes and further improve the robustness of our method by techniques such as using a learned motion prior. Applying our dynamics-regulated training procedure to other vision tasks such as visual navigation and third-person pose estimation can also be of interest.

5.3 Broader social impact.

Our overall framework can be used in extracting first-person camera wearer’s physically-plausible motion and our humanoid controller can be a plug-and-play model for physics-based humanoid simulation, useful in the animation and gaming industry for creating physically realistic characters. There can be also negative impact from this work. Our humanoid controller can be used as a postprocessing tool to make computer generated human motion physically and visually realistic and be misused to create fake videos using Deepfake-like technology. Improved egocentric pose estimation capability can also mean additional privacy concerns for smart glasses and bodycam users, as the full-body pose can now be inferred from front-facing cameras only. As the realism of motion estimation and generation methods improves, we encourage future research in this direction to investigate more in detecting computer generated motion [3].

Bibliography

- [1] Kevin Bergamin, Simon Clavet, Daniel Holden, and J. Forbes. Drecon. *ACM Transactions on Graphics (TOG)*, 38:1 – 11, 2019. [2.3](#)
- [2] Federica Bogo, A. Kanazawa, Christoph Lassner, P. Gehler, J. Romero, and Michael J. Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *The European Conference on Computer Vision (ECCV)*, 2016. [2.1](#)
- [3] Nicolò Bonettini, Edoardo Daniele Cannas, Sara Mandelli, Luca Bondi, Paolo Bestagini, and Stefano Tubaro. Video face manipulation detection through ensemble of cnns. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 5012–5019, 2021. doi: 10.1109/ICPR48806.2021.9412711. [5.3](#)
- [4] M. Brubaker, L. Sigal, and David J. Fleet. Estimating contact dynamics. *2009 IEEE 12th International Conference on Computer Vision*, pages 2389–2396, 2009. [2.1](#)
- [5] Yu-Wei Chao, Jimei Yang, Weifeng Chen, and Jia Deng. Learning to sit: Synthesizing human-chair interactions via hierarchical control. *ArXiv*, abs/1908.07423, 2019. [1](#), [2.3](#)
- [6] N. Chentanez, M. Müller, M. Macklin, Viktor Makoviychuk, and S. Jeschke. Physics-based motion capture imitation with deep reinforcement learning. *Proceedings of the 11th Annual International Conference on Motion, Interaction, and Games*, 2018. [2.3](#)
- [7] Kyunghyun Cho, B. V. Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *ArXiv*, abs/1406.1078, 2014. [3.2.2](#), [3.3.2](#)
- [8] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016. [4.4.1](#)
- [9] J. Engel, J. Sturm, and D. Cremers. Semi-dense visual odometry for a monocular camera. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1449–1456, Los Alamitos, CA, USA, Dec. 2013. IEEE Computer Society. [3.2.2](#)

- [10] Georgios V. Georgakis, Ren Li, Srikrishna Karanam, Terrence Chen, Jana Kosecka, and Ziyang Wu. Hierarchical kinematic human mesh recovery. *ArXiv*, abs/2003.04232, 2020. 2.1
- [11] Riza Alp Güler, N. Neverova, and I. Kokkinos. Densepose: Dense human pose estimation in the wild. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7297–7306, 2018. 2.1
- [12] I. Habibie, W. Xu, D. Mehta, G. Pons-Moll, and C. Theobalt. In the wild human pose estimation using explicit 2d features and intermediate 3d representations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10897–10906, Jun. 2019. 2.1
- [13] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 3.2.2
- [14] Mir Rayat Imtiaz Hossain and J. Little. Exploiting temporal information for 3d human pose estimation. In *The European Conference on Computer Vision (ECCV)*, 2018. 2.1
- [15] Soo hwan Park, Hoseok Ryu, Seyoung Lee, Sunmin Lee, and J. Lee. Learning predict-and-simulate policies from unorganized human motion data. *ACM Transactions on Graphics (TOG)*, 38:1 – 11, 2019. 2.3, 3.1.1
- [16] Apple Inc. Estimating camera pose with arkit. <https://developer.apple.com/documentation/arkit/arcamera>, 2021. Accessed: 2021-03-16. 3.2.2
- [17] Apple Inc. Scanning and detecting 3d objects with arkit. https://developer.apple.com/documentation/arkit/content_anchors/scanning_and_detecting_3d_objects, 2021. Accessed: 2021-03-16. 3.2.2, 5.1
- [18] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, jul 2014. 4.4.1
- [19] Mariko Isogawa, Ye Yuan, Matthew O’Toole, and Kris M Kitani. Optical non-line-of-sight physics-based 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7013–7022, 2020. 2.2
- [20] Hao Jiang and Kristen Grauman. Seeing invisible poses: Estimating 3d body pose from egocentric video. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3501–3509, Jun. 2016. 2.2
- [21] A. Kanazawa, Michael J. Black, D. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. *2018 IEEE/CVF Conference on Computer*

- Vision and Pattern Recognition*, pages 7122–7131, 2018. [2.1](#), [4.2](#)
- [22] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. [3.3.2](#)
- [23] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. Vibe: Video inference for human body pose and shape estimation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5252–5262, 2020. [2.1](#), [4.2](#)
- [24] Nikos Kolotouros, Georgios Pavlakos, Michael J. Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2252–2261, 2019. [2.1](#), [4.2](#)
- [25] Sijin Li and Antoni B. Chan. 3d human pose estimation from monocular images with deep convolutional neural network. In *ACCV*, 2014. [2.1](#)
- [26] Hung Yu Ling, Fabio Zinno, George H. Cheng, and M. V. D. Panne. Character controllers using motion vaes. *ACM Transactions on Graphics (TOG)*, 39:40:1 – 40:12, 2020. [4.2](#)
- [27] M. Loper, Naureen Mahmood, J. Romero, Gerard Pons-Moll, and Michael J. Black. Smpl: a skinned multi-person linear model. *ACM Trans. Graph.*, 34: 248:1–248:16, 2015. [3](#), [3.1.3](#)
- [28] Zhengyi Luo, S. Alireza Golestaneh, and Kris M. Kitani. 3d human motion estimation via motion compression and refinement. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, November 2020. [2.1](#)
- [29] Naureen Mahmood, N. Ghorbani, N. Troje, Gerard Pons-Moll, and Michael J. Black. Amass: Archive of motion capture as surface shapes. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5441–5450, 2019. [1](#), [3.1.1](#), [3.1.2](#), [3.1.3](#)
- [30] J. Merel, S. Tunyasuvunakool, Arun Ahuja, Yuval Tassa, Leonard Hasenclever, Vu Pham, T. Erez, Greg Wayne, and N. Heess. Reusable neural skill embeddings for vision-guided whole body movement and object manipulation. *ArXiv*, abs/1911.06636, 2019. [2.3](#)
- [31] Gyeongsik Moon and Kyoung Mu Lee. I2l-meshnet: Image-to-lixel prediction network for accurate 3d human pose and mesh estimation from a single rgb image. *ArXiv*, abs/2008.03713, 2020. [4.4.1](#)
- [32] Gyeongsik Moon, Juyong Chang, and Kyoung Mu Lee. Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In *IEEE Conference on International Conference on Computer Vision (ICCV)*, pages 10113–10142, Oct. 2019. [2.1](#)

- [33] Evonne Ng, Donglai Xiang, Hanbyul Joo, and Kristen Grauman. You2me: Inferring body pose in egocentric video via first and second person interactions. *CoRR*, abs/1904.09882, 2019. URL <http://arxiv.org/abs/1904.09882>. 2.2
- [34] Dario Pavlo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2019. 2.1
- [35] X. Peng, M. Chang, Grace Zhang, P. Abbeel, and Sergey Levine. Mcp: Learning composable hierarchical control with multiplicative compositional policies. In *NeurIPS*, 2019. 3.1.1
- [36] 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, 7 2018. ISSN 0730-0301. 1, 2.3, 3.1.1, 4.4.1
- [37] Xue Bin Peng, Angjoo Kanazawa, Jitendra Malik, Pieter Abbeel, and Sergey Levine. Sfv: Reinforcement learning of physical skills from videos. *ACM Trans. Graph.*, 37(6), November 2018. 2.3, 3.1.1
- [38] Davis Rempe, L. Guibas, Aaron Hertzmann, Bryan C. Russell, R. Villegas, and Jimei Yang. Contact and human dynamics from monocular video. In *SCA*, 2020. 2.1
- [39] Helge Rhodin, Christian Richardt, Dan Casas, Eldar Insafutdinov, Mohammad Shafiei, Hans-Peter Seidel, Bernt Schiele, and Christian Theobalt. Egocap: Egocentric marker-less motion capture with two fisheye cameras. *ACM Trans. Graph.*, 35(6), November 2016. ISSN 0730-0301. 2.2
- [40] Grégory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. LCR-Net++: Multi-person 2D and 3D Pose Detection in Natural Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 2.1
- [41] John Schulman, F. Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *ArXiv*, abs/1707.06347, 2017. 3.1.2
- [42] Soshi Shimada, Vladislav Golyanik, Weipeng Xu, and C. Theobalt. Physcap: Physically plausible monocular 3d motion capture in real time. *ACM Trans. Graph.*, 39:235:1–235:16, 2020. 2.1
- [43] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2018. 3.2.2
- [44] J. Tan, K. Liu, and G. Turk. Stable proportional-derivative controllers. *IEEE Computer Graphics and Applications*, 31(4):34–44, Jul. 2011. ISSN 1558-1756.

3.1.1

- [45] Bugra Tekin, Isinsu Katircioglu, M. Salzmann, Vincent Lepetit, and P. Fua. Structured prediction of 3d human pose with deep neural networks. *ArXiv*, abs/1605.05180, 2016. 2.1
- [46] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5026–5033, Oct. 2012. 3.1.1, 4.4.1
- [47] Denis Tome, Patrick Peluse, Lourdes Agapito, and Hernan Badino. xr-egopose: Egocentric 3d human pose from an hmd camera. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 7728–7738, Oct. 2019. 2.2
- [48] M. Vondrak, L. Sigal, J. Hodgins, and O. C. Jenkins. Video-based 3d motion capture through biped control. *ACM Transactions on Graphics (TOG)*, 31:1–12, 2012. 2.1
- [49] S. Wang, R. Clark, H. Wen, and N. Trigoni. Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2043–2050, May 2017. 3.2.2
- [50] Tingwu Wang, Yunrong Guo, Maria Shugrina, and S. Fidler. Unicon: Universal neural controller for physics-based character motion. *ArXiv*, abs/2011.15119, 2020. 1, 3.1.2
- [51] Jungdam Won, Deepak Gopinath, and Jessica Hodgins. A scalable approach to control diverse behaviors for physically simulated characters. *ACM Trans. Graph.*, 39(4), 2020. URL <https://doi.org/10.1145/3386569.3392381>. 1, 2.3, 3.1.2
- [52] Weipeng Xu, Avishek Chatterjee, Michael Zollhoefer, Helge Rhodin, Pascal Fua, Hans-Peter Seidel, and Christian Theobalt. Mo²Cap² : Real-time mobile 3d motion capture with a cap-mounted fisheye camera. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2019. ISSN 1077-2626. 2.2
- [53] Yuanlu Xu, S. Zhu, and Tony Tung. Denserac: Joint 3d pose and shape estimation by dense render-and-compare. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7759–7769, 2019. 2.1
- [54] Ye Yuan and Kris Kitani. 3d ego-pose estimation via imitation learning. In *The European Conference on Computer Vision (ECCV)*, Sep. 2018. 1, 2.2, 4.1.1
- [55] Ye Yuan and Kris Kitani. Ego-pose estimation and forecasting as real-time pd control. In *IEEE International Conference on Computer Vision (ICCV)*, pages 10082–10092, Oct. 2019. (document), 1, 2.2, 4.3.1, 4.6

Bibliography

- [56] Ye Yuan and Kris Kitani. Residual force control for agile human behavior imitation and extended motion synthesis. In *Advances in Neural Information Processing Systems*, 2020. [1](#), [2.3](#), [3.1.1](#), [3.1.1](#), [3.1.1](#)
- [57] Ye Yuan, Shih-En Wei, Tomas Simon, Kris Kitani, and Jason Saragih. Simpoe: Simulated character control for 3d human pose estimation. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [1](#), [2.1](#), [3.1.1](#), [3.1.3](#)
- [58] Hongwen Zhang, Jie Cao, Guo Lu, Wanli Ouyang, and Z. Sun. Learning 3d human shape and pose from dense body parts. *ArXiv*, abs/1912.13344, 2019. [2.1](#)