Embodied Scene-aware Human Pose Estimation

Zhengyi Luo[1] * **Shun Iwase** [1]* **Ye Yuan**[1] **Kris Kitani**[1]
[1] Carnegie Mellon University
https://embodiedscene.github.io/embodiedpose/

Figure 1: From a monocular video input in a known scene, we estimate 3D absolute body pose by controlling a simulated character to match 2D observation in a sequential manner.

Abstract

We propose *embodied scene-aware human pose estimation* where we estimate 3D poses based on a simulated agent's proprioception and scene awareness, along with external third-person observations. Unlike prior methods that often resort to multistage optimization, non-causal inference, and complex contact modeling to estimate human pose and human scene interactions, our method is one-stage, causal, and recovers global 3D human poses in a simulated environment. Since 2D third-person observations are coupled with the camera pose, we propose to disentangle the camera pose and use a multi-step projection gradient defined in the global coordinate frame as the movement cue for our embodied agent. Leveraging a physics simulation and prescanned scenes (*e.g.,* 3D mesh), we simulate our agent in everyday environments (library, office, bedroom, *etc.*) and equip our agent with environmental sensors to intelligently navigate and interact with the geometries of the scene. Our method also relies only on 2D keypoints and can be trained on synthetic datasets derived from popular human motion databases. To evaluate, we use the popular H36M and PROX datasets and achieve high quality pose estimation on the challenging PROX dataset without ever using PROX motion sequences for training. Code and videos are available on the project page.

1 Introduction

We consider a 3D human pose estimation setting where a calibrated camera is situated within a prescanned environment (with known 3D meshes geometries). Such a setting has a wide range of

*Equal Contribution.

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

applications in mixed reality telepresence, animation and gaming, at home robotics, video surveillance, *etc.*, where a single fixed RGB camera is often used to observe human movement. However, even under such constraints, severe occlusion, low-textured garments, and human-scene interactions often cause pose estimation methods to fail catastrophically. Especially when estimating the global 3D human pose, state-of-the-art methods often resort to expensive multistage batch optimization [6], motion prior or infilling [31, 47], and depth information to recover reasonable global translation and orientation. Consequently, these methods cannot be applied in a streaming fashion and have limited use cases in time-critical applications. Multistage optimization can also suffer from failure of convergence, where the optimizer cannot find a good solution due to occlusion, numerical instability, or poor initialization.

The effective use of scene information is also critical for human pose estimation. Unfortunately, the lack of datasets with paired ground-truth 3D body pose and annotated contact labels makes it difficult to incorporate scene information to estimate pose. Thus, popular methods often resort to 1) minimizing the distance between predefined contact points in the body and the scene [6], 2) using vertex-based PROXimity [48] as a proxy scene representation, and 3) semantic labels as affordance cues [17]. These methods, while effective in modeling human-scene interactions, are expensive to compute and heavily rely on a batch optimization procedure to achieve good performance. More complex scenes are also harder to model in a batch optimization framework, as the number of contact constraints increases significantly as the number of geometries increases.

We hypothesize that **embodiment** may be a key element in estimating accurate absolute 3D human poses. Here, we define embodiment as having a tangible existence in a simulated environment in which the agent can act and perceive. In particular, humans move according to their current body joint configuration (proprioception), movement destination (goal), and an understanding of their environment's physical laws (scene awareness and embodiment). Under these guiding principles, we propose **embodied scene-aware human pose estimation** as an alternative to optimization and regression-based methods. Concretely, we estimate poses by controlling a physically simulated agent equipped with environmental sensors to mimic observations from third-person in a **streaming** fashion. At each timestep, our agent receives movement goals from a third-person observation and computes its torque actuation based on (1) current agent state: body pose, velocities, *etc.*, (2) comparison between its current pose and third-person observation, and (3) environmental features from the agent's point of view. Acting in a physics simulator, our agent can receive timely feedback about physical plausibility, contact, and scene geometry to compute actions for the next time frame. Although there exist prior methods that also use simulated characters, we draw the distinction in how the next frame pose is computed. Previous work first regresses the kinematic pose [4, 45] without considering egocentric features such as body and environmental states. As a result, when the kinematic pose estimation method fails due to occlusion, these simulated agents will receive unreasonable target poses. Our method, on the other hand, predicts target poses by comparing the current body pose with third-person observations. When 2D observations are unreliable, the target pose will be extrapolated from environmental cues and current body states. Thus, our method is robust to occlusion and can recover from consecutive frames of missing body observations. Our method is also capable of estimating pose for long-term (several minutes to perpetual), in sub-real-time ($\sim$ 10 fps), while creating physically realistic human scene interactions.

To achieve this, we propose a multi-step projection gradient (MPG) to provide movement cues to our embodied agent from 2D observations. At each time step, MPG performs iterative gradient descent starting from the current pose of the humanoid to exploit the temporal coherence of human movement and better avoid local minima. MPG also directly provides temporal body and root movement signals to our agent in the world coordinate frame. Thus, our method does not need to take the 3D camera pose as input explicitly. To control a simulated character, we extend the Universal Humanoid Controller (UHC) and the kinematic policy proposed in previous work [21] to support the control of humanoids with different body shapes in environments with diverse geometries. Finally, we equip our agent with an occupancy-based environmental sensor to learn about affordances and avoidance. Since only 2D keypoint coordinates are needed by our method to estimate body motion, our model can be learned from purely synthetic data derived from popular motion databases.

To summarize, our contribution is as follows: 1) we propose embodied scene-aware human pose estimation, where we control a simulated scene-aware humanoid to mimic visual observations and recover accurate global 3D human pose and human-scene interactions; 2) we formulate a multi-step projection gradient (MPG) to provide a camera-agnostic movement signal for an embodied agent to

match 2D observations; and 3) we achieve state-of-the-art result on the challenging PROX dataset while our method remains sub-real-time, causal, and trained only on synthetic data.

2 Related Work

Local and Absolute 3D Human Pose Estimation. It is popular to estimate the 3D human pose in the local coordinate system (relative to the root) [1, 11, 12, 13, 14, 15, 16, 20, 24, 25, 27, 30, 37, 39] where the translation is discarded and the rotation is defined in the camera coordinate frame. Among these methods, some focus on lifting 2D keypoints to 3D [1, 27], while others use a template 3D human mesh (the SMPL [19] body model and its extensions [26, 32]) and jointly recover body shape and joint angles. Among these methods, some [15, 37, 45] make use of iterative optimization to better fit 2D keypoints. SPIN [15] devises an optimization-in-the-loop fitting procedure to provide a better starting point for a pose regressor. NeuralGD [37] and SimPOE [45] both utilize a learned gradient update to iteratively refine body pose and shape, where optimization starts from either zero pose (NeuralGD) or estimated kinematic pose (SimPOE). Compared to these methods, our MPG starts the optimization process from the previous time-step pose to leverage the temporal coherence in human motion. While estimating pose in the local coordinate frame is useful in certain applications, it is often desirable to recover the pose in the absolute/global space to fully leverage the estimated pose for animation or telepresence. However, recovering the absolute 3D body pose that includes root translation is challenging due to limb articulation, occlusion, and depth ambiguity. Among the work that attempt to estimate the absolute human pose [24, 27, 30, 33], most take a two-stage approach where the root-relative body pose is first regressed from a single image and then a root translation is estimated from image-level features. Compared to these methods, we leverage physical laws and motion cues to autoregressively recover both root movement and body pose.

Scene-aware Human Pose and Motion Modeling The growing interest in jointly modeling human pose and human-scene interactions has led to a growing popularity in datasets containing both human and scenes [7, 23, 41]. In this work, we focus on the PROX [7] dataset, where a single human is recorded moving in daily indoor environments such as offices, libraries, bedrooms, etc. Due to the diversity of motion and human-scene interaction (sitting on chairs, lying on beds, etc.) and lack of ground-truth training data, the PROX dataset remains highly challenging. Consequently, state-of-the-art methods are mainly optimization-based [7, 31, 46] while leveraging the full set of observations (scene scans, camera pose, depth). HuMoR [31] and LEMO [46] achieves impressive result by using large-scale human motion datasets to learn strong prior of human movement to constrain or infill the missing human motion. However, multistage optimization can often become unstable due to occlusion, failures in 2D keypoint detections, or erroneous depth segmentation.

Physics-based humanoid control for pose estimation Physical laws guard every aspect of human motion and there is a growing interest in enforcing physics prior in human pose estimation [4, 9, 21, 35, 36, 42, 43, 44, 45]. Some methods prefer to use physics as a postprocessing step and refine initial kinematic pose estimates [35, 36, 42], while ours follows the line of work where a simulated character is controlled to estimate pose inside a physics simulation [4, 9, 21, 45]. Although both approaches have their respective strong suites, utilizing a physics simulation simplifies the process of incorporating unknown scene geometries and can empower real-time pose estimation applications. Among the simulated character control-based methods, Neural Mocon [9] utilizes sampling-based control and relies on the learned motion distribution from the training split of the dataset to acquire quality motion samples. SimPOE [45] first extracts the kinematic pose using off-the-shelf method [13] and then uses a simulated agent for imitation. PoseTriplet [4] proposes using a self-supervised framework to co-evolve a pose estimator, imitator, and hallucinator. Compared to ours, SimPoE and PoseTriplet both rely on a kinematic pose estimator as backbones [13, 27] and require well-observed 2D keypoints for reliable pose estimates. They also require lengthy policy training time and do not model any scene geometries. Our work extends Kin_poly [21], where a simulated character is used to perform *first-person* pose estimation. While Kin_poly uses egocentric motion as the only motion cue, we propose a Multistep Projection Gradient to leverage thrid-person observations. Compared to their object-pose based scene feature, we use an occupancy-based scene feature inspired by the graphics community [5, 38] to better model scene geometries.

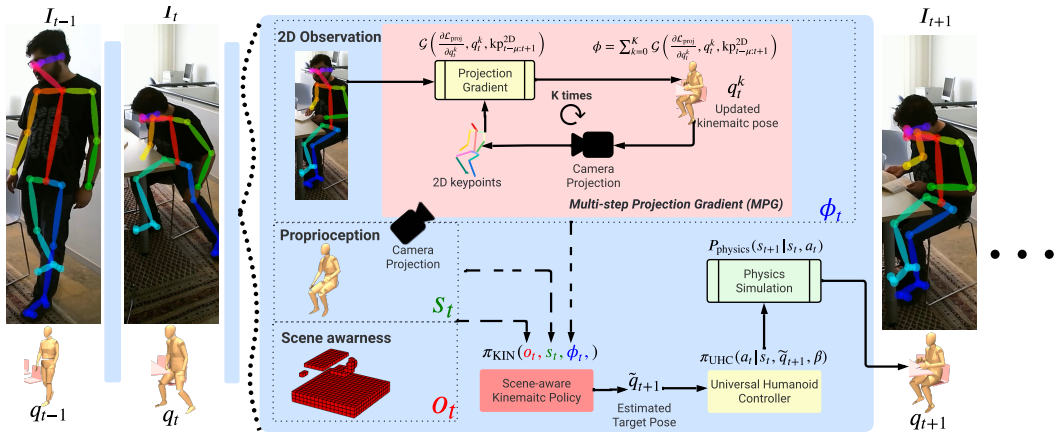

Figure 2: Our embodied human pose estimation framework: given 2D keypoint features, our model computes the next-time frame pose based on third-person 2D observation (multi-step projection gradient $\boldsymbol{\phi}_t$), proprioception features (body state $\boldsymbol{s}_t$), and scene (occupancy $\boldsymbol{o}_t$).

3 Method

Given a monocular video $I_{1:T}$ taken using a calibrated camera in a known environment, our aim is to estimate, track, and simulate 3D human motion in metric space. We denote the 3D human pose as $\boldsymbol{q}_t \triangleq (\boldsymbol{r}_t^{\mathrm{R}}, \boldsymbol{r}_t^{\mathrm{T}}, \boldsymbol{\theta}_t)$, consisting of the 3D root position $\boldsymbol{r}_t^{\mathrm{T}}$, the orientation $\boldsymbol{r}_t^{\mathrm{R}}$ and the angles of the body joints $\boldsymbol{\theta}_t$. Combining 3D pose and velocity, we have the character's state $\boldsymbol{s}_t \triangleq (\boldsymbol{q}_t, \dot{\boldsymbol{q}}_t)$. As a notation convention, we use $\widetilde{\cdot}$ to represent kinematic poses (without physics simulation), $\widehat{\cdot}$ to denote ground-truth quantities from Motion Capture (MoCap), and normal symbols without accents for values from the physics simulation. In Sec 3.1, we will first summarize our extension to the Universal Humanoid Controller [21] to support controlling people with different body proportions. Then, in Sec.3.2, we will describe the state space of our character's control policy: the multi-step projection gradient (MPG) as a feature and occupancy based scene feature. Finally, we will detail the training procedure to learn our embodied pose estimator using synthetic datasets in Sec 3.3. Our pose estimation pipeline is visualized at Fig.2.

3.1 Preliminaries: Physics-based Character and Universal Humanoid Controller (UHC)

To simulate people with varying body shapes, we first use an automatic simulated character creation process similar to SimPoE [45] to generate human bodies with realistic joint limits, body proportions, and other physical properties based on the popular human mesh model SMPL [19]. In SMPL, body joints are recovered from mesh vertices and a joint regressor, so the length of each limb may vary depending on body pose. This is undesirable as varying body limb length is physically inaccurate. Recovering the body joints by first computing a large number of vertices (6890) followed by a joint regressor is also time-consuming. Therefore, we fix the lengths of the limbs using those calculated from the T pose (setting the joint angles to zero) and use forward kinematics to recover the 3D location of the body joints $\boldsymbol{J}_t$. Given the shape of the body β and the pose $\boldsymbol{q}_t$, we define the forward kinematics function as $\mathrm{FK}_\beta(\boldsymbol{q}_t) \rightarrow \boldsymbol{J}_t$.

To learn a universal humanoid controller that can perform everyday human motion with different body types, we follow the standard formulation in physics-based character and model controlling a humanoid for motion imitation as a Markov Decision Process (MDP) defined as a tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{T}, R, \gamma \rangle$ of states, actions, transition dynamics, reward function, and discount factor. The state $\boldsymbol{s} \in \mathcal{S}$, reward $r \in \mathcal{R}$, and the transition dynamics $\mathcal{T}$ are determined by the physics simulator, while the action $\boldsymbol{a} \in \mathcal{A}$ is given by the control policy $\pi_{\mathrm{UHC}}(\boldsymbol{a}_t | \boldsymbol{s}_t, \widehat{\boldsymbol{q}}_{t+1}, \boldsymbol{\beta})$. Here we extend the Universal Humanoid Controller proposed in Kin_poly [21] and condition the control policy on the current simulation state $\boldsymbol{s}_t$, the reference motion $\widehat{\boldsymbol{q}}_t$, and the body shape $\boldsymbol{\beta}$. We employ Proximal Policy Optimization (PPO) [34] to find the optimal control policy π_{UHC} to maximize the expected

discounted reward $\mathbb{E}\left[\sum_{t=1}^{T}\gamma^{t-1}r_t\right]$. For more information on the UHC network architecture and training procedure, please refer to Kin_poly [21] and our Appendix.

3.2 Embodied Scene-Aware Human Pose Estimation

To estimate the global body pose from the video observations, we control a simulated character to follow the 2D keypoint observations in a sequential decision manner. Assuming a known camera pose and scene geometry, we first use an off-the-shelf 2D keypoint pose estimator, DEKR [3], to extract 2D keypoints kp_t^{2D} at a per-frame level. Then, we use HuMoR [31] to extract the first-frame 3D pose $\widetilde{\boldsymbol{q}}_0$ and body shape $\boldsymbol{\beta}$ from the RGB observations as initialization. Given $\widetilde{\boldsymbol{q}}_0$ and $\boldsymbol{\beta}$, we initialize our simulated character inside the pre-scanned scene, and then we attempt to estimate poses which match observed 2D keypoints in a streaming fashion.

Multi-step Projection Gradient (MPG) as a Motion Feature To match the motion of the 3D simulated humanoids to 2D keypoint observations, it is essential to provide 3D information about how poses should be modified in the global coordinate frame (coordinate frame of the scene geometry and the 3D humanoid). As 2D keypoints are coupled with camera pose (the same body pose corresponds to different 2D keypoints based on the camera pose), it is problematic to directly input 2D observations without a proper coordinate transform. To avoid this issue, prior arts that use simulated character control [4, 45] often estimate the kinematic human pose in the camera coordinate frame using a kinematic pose estimator. However, since the up-vector of the camera is not always aligned with the direction of gravity, corrections to the character's z-direction are often needed for a successful simulation [4].

To decouple the pose of the camera from the pose estimation pipeline, we propose using a multi-step projection gradient (MPG) feature as an alternative to estimated kinematic pose as an input to the policy. Intuitively, the MPG is a modified (learned) sum of 2D projection loss gradients that encodes how much the 3D pose should be adjusted for to fit the next frame 2D keypoints. Given the current global state of the humanoid $\boldsymbol{q}_t$ and the camera projection function π_{cam}, we define MPG $\boldsymbol{\phi}$ as:

$$\boldsymbol{\phi}_t = \boldsymbol{q}_t^K - \boldsymbol{q}_t^0 = \sum_{k=0}^{K}\mathcal{G}(\frac{\partial\mathcal{L}_{\text{proj}}}{\partial\boldsymbol{q}_t^k}, \boldsymbol{q}_t^k, \text{kp}_{t-\mu:t+1}^{\text{2D}}),\ \text{where}\ \boldsymbol{q}_t^{k+1} = \boldsymbol{q}_t^k - s\cdot\mathcal{G}(\frac{\partial\mathcal{L}_{\text{proj}}}{\partial\boldsymbol{q}_t^k}, \boldsymbol{q}_t^k, \text{kp}_{t-\mu:t+1}^{\text{2D}}) \tag{1}$$

and $\mathcal{L}_{\text{proj}} = \|\pi_{\text{cam}}(\text{FK}_\beta(\boldsymbol{q}_t)) - \text{kp}_{t+1}^{\text{2D}}\|_2$ is the L_2 reprojection loss with respect to the observed 2D keypoints $\text{kp}_{t+1}^{\text{2D}}$. s is the step size of the gradient and k indexes into the total number of gradient descent steps K. The 2D project loss gradient $\frac{\partial\mathcal{L}}{\partial\boldsymbol{q}_t}$ represents how the root pose and body joints should change to better match the 2D key points of the next frame and contains rich geometric information independent of the camera projection function π_{cam}.

Due to noise in the detection, occlusion, and depth ambiguity of 2D key points, a single-step instantaneous projection gradient $\frac{\partial\mathcal{L}_{\text{proj}}}{\partial\boldsymbol{q}_t^k}$ can be erroneous and noisy, especially the gradient of root translation and orientation. Thus, we augment $\frac{\partial\mathcal{L}_{\text{proj}}}{\partial\boldsymbol{q}_t^k}$ with a geometric and learning-based transform $\mathcal{G}$ that is intended to remove noise. In this way, the MPG feature is an aggregation of a collection of modified gradients to compute a more stable estimate of the 2D projection loss gradient.

To better recover root rotation, we use a temporal convolution model (TCN) similar to VideoPose3D [27] and PoseTriplet [4] to regress camera-space root orientation from 2D keypoints. We notice that while TCN may can produce an unreasonable pose estimate during server occlusion, its root orientation estimation remains stable. We compute the camera space root orientation using a causal TCN $\mathcal{F}_{\text{TCN}}$ and transform it into the world coordinate system $r_t'^{\text{R}}$ using the inverse camera rotation: $r_t'^{\text{R}} = \pi_{\text{cam}}^{-1}(\mathcal{F}_{\text{TCN}}(\text{kp}_{t-\mu:t+1}^{\text{2D}}))$. We then calculate the difference between the current root rotation r_t^{R} and $r_t'^{\text{R}}$ as an additional movement gradient signal. Given the current 3D pose $\boldsymbol{q}_t^k$ and the projection function π_{cam}, we can obtain a better root translation $\boldsymbol{r}_t^T$ by treating the 3D keypoints obtained as fixed on a rigid body and solving for translation using linear least squares. Combining the $\mathcal{F}_{\text{TCN}}$ model and the geometric transformation step, we obtain the function $\mathcal{G}$ to obtain a better root translation and orientation gradient. By taking multiple steps, we avoid getting stuck in local minima and obtain better gradient estimates. For a more complete formulation of MPG, see our Appendix.

Scene Representation To provide our embodied agent with an understanding of its environment, we provide our policy with two levels of scene-awareness: 1) we simulate the geometry of the scene inside a physics simulation, which implicitly provides contact and physics constraints for our policy. As the scanned meshes provided by the PROX dataset often have incorrect volumes, we use a semi-automatic geometric fitting procedure to approximate the scene by simple shapes such as rectangles and cylinders. The modified scene can be visualized together in Eq.2. For more information about our scene simplification process, please refer to Appendix. 2) We use a voxel-based occupancy map to indicate the location and size of the scene geometries, similar to [5, 38]. Occupancy provides basic-level scene and object awareness to our agent and is robust to changing scene geometries. Concretely, we create a cube of 1.8 meters edge length as the sensing range of our agent's surroundings and create a 16 by 16 by 16 evenly spaced grid. At the center of each point on the grid, we sample the signed distance function of the scene $\mathcal{F}_s$ to obtain how far the sampling point is from the geometries of the scene. Then, we threshold the gird feature to provide a finer scene information. In time step t, for threshold $\sigma = 0.05$, the value of point p_i in the grid can be computed as o_t^i, and $\boldsymbol{o}_t = \{o_0^i, o_1^i, \dots, \}$:

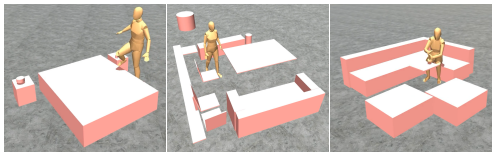

Figure 3: Our modified scene for simulation.

$$o_t^i = \begin{cases} 1 & \mathcal{F}_s(p_i) \leq 0 \\ 0 & \mathcal{F}_s(p_i) > \sigma \\ 1 - \mathcal{F}_s(p_i)/\sigma & \text{otherwise} \end{cases} \quad (2)$$

Scene-aware Kinematic Policy To train and control our embodied agent in the physics simulator, we extend a policy network similar to Kin_poly [21], with the important distinction that our networks performs third-person pose estimation in a real-world scene. Under the guiding principle of embodied human pose estimation, our kinematic policy $\boldsymbol{\pi}_{\text{KIN}}$ computes a residual 3D movement $\tilde{\dot{\boldsymbol{q}}}_t$ based on proprioception $\boldsymbol{s}_t$, movement goal $\boldsymbol{\phi}_t$, and scene awareness $\boldsymbol{o}_t$:

$$\tilde{\dot{\boldsymbol{q}}}_t = \boldsymbol{\pi}_{\text{KIN}}(\boldsymbol{s}_t, \boldsymbol{\phi}_t, \boldsymbol{o}_t), \quad \tilde{\boldsymbol{q}}_{t+1} = \tilde{\dot{\boldsymbol{q}}}_t + \tilde{\boldsymbol{q}}_t. \quad (3)$$

Specifically, the input $\boldsymbol{q}_t$ contains the current body joint configurations, $\boldsymbol{\phi}_t$ provides the movement signal per step, while $\boldsymbol{o}_t$ provides awareness to the agent's surroundings. Within $\boldsymbol{\pi}_{\text{KIN}}$, an agent-centric transform T_{AC} is used to transform $\boldsymbol{q}_t$, $\boldsymbol{\phi}_t$ and $\boldsymbol{s}_t$ into the agent-centirc coordinate system. The output $\tilde{\dot{\boldsymbol{q}}}_t$, is defined as $\tilde{\dot{\boldsymbol{q}}}_t \triangleq (\tilde{\dot{\boldsymbol{r}}}_t^{\text{R}}, \tilde{\dot{\boldsymbol{r}}}_t^{\text{T}}, \tilde{\dot{\boldsymbol{\theta}}}_t)$, which computes the velocity in root translation $\tilde{\dot{\boldsymbol{r}}}_t^{\text{R}}$ and orientation $\tilde{\dot{\boldsymbol{r}}}_t^{\text{T}}$, as well as the joint angles $\tilde{\dot{\boldsymbol{\theta}}}_t$. Notice that $\tilde{\dot{\boldsymbol{q}}}_t$ computes the residual body movement instead of the full kinematic body pose $\tilde{\boldsymbol{q}}_{t+1}$ and establishes a better connection with the movement feature $\boldsymbol{\phi}_t$. During inference, $\boldsymbol{\pi}_{\text{KIN}}$ works in tandem with $\boldsymbol{\pi}_{\text{UHC}}$ and computes the next-step target pose $\tilde{\boldsymbol{q}}_{t+1}$ for the humanoid controller to mimic. In turn, $\boldsymbol{\pi}_{\text{KIN}}$ receives timely feedback on physical realism and contact information based on the simulated agent state.

3.3 Learning the embodied human pose estimator

Since $\boldsymbol{\pi}_{\text{KIN}}$ relies only on 2D keypoints, we can leverage large-scale human motion databases and train our pose estimation network using synthetic data. We use the AMASS [22], H36M [10], and Kin_poly [21] datasets to dynamically generate 2D keypoints and 3D pose pairs for training. At the beginning of each episode, we sample a motion sequence and randomly change its direction and starting position in a valid range, similar to the procedure defined in NeuralGD [37]. Then we use a randomly selected camera projection function π_{cam} to project 3D keypoints onto 2D images to obtain paired 2D keypoints and human motion sequences. During training, we initialize the agent with the ground-truth 3D location and pose. The H36M and Kin_poly dataset contains simple human-scene interactions, such as sitting on chairs and stepping on boxes. To better learn object affordance and collision avoidance, we also randomly sample scenes from the PROX dataset and pair them with motions from the AMASS dataset. Although these motion sequences do not correspond to the sampled scene, the embodied agent will hit the scene geometries while trying to follow the 2D observations. We stop the simulation when the agent deviates too far away from the ground-truth motion to learn scene affordance and avoidance.

ALGORITHM 1: Learning embodied pose estimator via dynamics-regulated training.

Input: Pre-trained controller $\boldsymbol{\pi}_{\text{UHC}}$, ground truth motion dataset $\widehat{\boldsymbol{Q}}$;
while *not converged* **do**
  $\boldsymbol{M}_{\text{dyna}} \leftarrow \emptyset$ initialize sampling memory ;
  while $\boldsymbol{M}_{dyna}$ *not full* **do**
    $\boldsymbol{q}_{1:T} \leftarrow$ sample random sequence of human pose ;
    $\text{kp}^{2D}_{1:T} \leftarrow \pi_{\text{cam}}(\boldsymbol{q}_{1:T})$ `// compute paired 2D projection.` ;
    for $t \leftarrow 1...T$ **do**
      $\boldsymbol{s}_t \leftarrow (\boldsymbol{q}_t, \dot{\boldsymbol{q}}_t)$
      $\boldsymbol{\phi}_t \leftarrow \sum_{k=0}^{K} \mathcal{G}(\frac{\partial \mathcal{L}_{\text{proj}}}{\partial \boldsymbol{q}_t^k}, \boldsymbol{q}_t^k, \text{kp}^{\text{2D}}_{t-\mu:t+1})$ `// movement cues from 2D observations.` ;
      $\widetilde{\boldsymbol{q}}_{t+1} \leftarrow \boldsymbol{\pi}_{\text{KIN}}(\boldsymbol{s}_t, \boldsymbol{\phi}_t, \boldsymbol{o}_t) + \boldsymbol{q}_t$
      $\boldsymbol{s}_{t+1} \leftarrow \mathcal{T}(\boldsymbol{s}_{t+1}|\boldsymbol{s}_t, \boldsymbol{a}_t)$, $\boldsymbol{a}_t \leftarrow \boldsymbol{\pi}_{\text{UHC}}(\boldsymbol{a}_t|\boldsymbol{s}_t, \widetilde{\boldsymbol{q}}_{t+1})$ `// physics simulation and agent control using` $\boldsymbol{\pi}_{\text{UHC}}$
      store $(\boldsymbol{s}_t, \boldsymbol{a}_t, \boldsymbol{r}_t, \boldsymbol{s}_{t+1}, \widehat{\boldsymbol{q}}_t, \widetilde{\boldsymbol{q}}_{t+1})$ into memory $\boldsymbol{M}_{\text{dyna}}$
    end
  end
end
$\boldsymbol{\pi}_{\text{KIN}} \leftarrow$ Supervised learning update using experiences collected in $\boldsymbol{M}_{\text{dyna}}$ for 10 epoches.

We train the kinematic policy as a per-step model together with physics simulation, following the dynamics-regulated training procedure defined in [21]. Combining our UHC and our kinematic policy, we can leverage the pre-learned motor skills from the UHC and let the kinematic policy act directly in a simulated physical space to obtain feedback on physical plausibility. The procedure for dynamics-regulated training is described in Alg. 1. In each episode, we connect the UHC and the kinematic policy using the result of the UHC and the physics simulation $\boldsymbol{q}_{t+1}$ to compute the input features to the kinematic policy. Using $\boldsymbol{q}_{t+1}$ verifies that $\widetilde{\boldsymbol{q}}_t$ can be followed by our controller and informs our $\boldsymbol{\pi}_{\text{KIN}}$ of the current humanoid state to encourage the policy to adjust its predictions to improve stability. Here, we only train our $\boldsymbol{\pi}_{\text{KIN}}$ using supervised learning, as third-person pose estimation is a more well-posed problem and requires less exploration than estimating pose from only egocentric videos [21]. For each motion sequence, our loss function is defined as $L1$ distance between the predicted kinematic pose and the ground truth plus a prior loss using a learned motion transition prior p_θ and the encoder q_ϕ from HuMoR [31]. The prior term encourages the generated motion to be smoother:

$$\mathcal{L} = \sum_{t=0}^{T} w_{\text{L1}} \|\hat{\boldsymbol{q}}_t - \tilde{\boldsymbol{q}}_t\|_1 + w_{\text{prior}} D_{\text{KL}}(q_\phi(\boldsymbol{z}_t|\tilde{\boldsymbol{q}}_t, \tilde{\boldsymbol{q}}_{t-1})||p_\theta(\boldsymbol{z}_t|\tilde{\boldsymbol{q}}_{t-1})), \tag{4}$$

4 Experiments

Implementation details. We use the MuJoCo [40] physics simulator and use modified scene geometries as environments. The training process or our embodied pose estimator takes around 1 day on a RTX 3090 with 30 CPU threads. During inference, our network is causal and runs at 10 FPS on an Intel desktop CPU (not counting the time for 2D keypoint pose estimation, which runs at around 30 FPS [3]). For more details on the implementation and network architecture, please refer to the Appendix.

Datasets. Our Universal Humanoid Controller is trained on the AMASS dataset training split that contains high-quality SMPL parameters of 11402 sequences curated from the various MoCap datasets. For training our pose estimator, we use motions from the AMASS training split, the Kin_poly dataset, and H36M dataset. The Kin_poly dataset contains 266 motion sequences of type sitting, stepping, pushing, and avoiding, and the H36M dataset consists of 210 indoor videos with 51 different unique actions. Note that we only use the motion sequences from these datasets and not the paired videos. For evaluation, we use videos from the PROX [7] dataset and the test split of H36M.

Metrics and Baselines. We compare with a series of state-of-the-art methods HyperIK [16], METRAbs [33], VIBE [13], and SimPOE [45] in the H36M dataset. On the PROX dataset, we compare

with the HuMoR [31], LEMO [46], and HyberIK with RootNet [24]. We use the official implementation of the above methods. We use pose-based and plausibility-based metrics for evaluation. To evaluate pose estimation when there is ground-truth MoCap pose, we report the absolute/global and root-relative mean per joint position error: A-MPJPE and MPJPE. On the PROX dataset, since there is no ground truth 3D pose annotation, we report: success rate SR (whether the optimization procedure converges and whether the simulated agents lose balance), one-way average chamfer distance from the point cloud (ACD, acceleration (ACC), and scene penetration frequency (Freq.) and distance (Pen.). For a full definition of our evaluation metrics, please refer to the Appendix.

4.1 Results

Results on H36M We report the pose estimation on the H36M dataset in Table 1. Here, we can see that our method achieves a reasonable pose estimation result on both root-relative and global pose metrics. Notice that our experiment with the H36M dataset is meant to be a sanity check on our pose estimation framework and is not a fair comparison. Our method only uses a sparse set of detected 2D keypoints (12 keypoints) as input and does not directly use any image feature. We also evaluate a different subset of 3D keypoints, as our use of forward kinematics FK_β causes our joint definition to be different from prior works. For more clarification on the H36M evaluation, please refer to the Appendix.

Table 1: Our pose estimation result on the H36M dataset. *Since our method does not directly use any image-level feature and uses a different set of 2D keypoints for observation and 3D keypoints for evaluation, it is not a completely fair comparison against SOTA methods.

	Image-features	A-MPJPE ↓	MPJPE ↓	PA-MPJPE ↓	Accel.-D ↓
MeTRAbs [33]	✓	223.9	49.3	34.7	-
VIBE [16]	✓	-	61.3	43.1	15.2
SimPoE [45]	✓	-	56.7	41.6	6.7
HybrIK [16] w/ RootNet [24]	✓	135.9	65.2	43.4	-
Ours*	✗	282.5	95.5	71.6	9.3

Table 2: Evaluation of plausibility on the PROX dataset. Since PROX has no ground truth annotation, we employ ACD (one-way Average Chamfer Distance) between the segmented point cloud and predicted body mesh as a proxy metric for pose estimation plausibility, as formulated in HuMoR[31]. Notice that the ACD values correspond to a different number of sequences as some sequences fail during optimization or simulation (which corresponds to the success rate). There are in total 60 sequences from the PROX dataset. Since all of the metrics can only approximate plausibility, we encourage our readers to refer to our project website for visualization of **all** 60 sequences on PROX of our method. Time(s) is the per-frame average runtime evaluated on RTX 2080 Super. Note that only HybrIK [16] w/ RootNet [24] and ours are not multi-stage optimization based methods (indicated by the coloum "Opti.") and offers a relatively fair comparison.

								Ground		Scene		
		Phys.	Scene	Opti.	SR ↑	ACD ↓	Accel. ↓	Freq. ↓	Pen. ↓	Freq. ↓	Pen. ↓	Time(s) ↓
RGB-D	PROX [7]	✗	✓	✓	98.3%	181.1	157.4	10.6 %	109.1	0.7	46.8 %	73.1
	HuMoR [31]	✗	✗	✓	100%	325.5	3.4	2.4%	113.6	0.4 %	47.5	13.72
	LEMO [46]	✗	✓	✓	71.7%	173.2	3.0	0 %	14.5	0.2%	22.8	75.19
RGB	PROX [7]	✗	✓	✓	90.0%	>1000	381.6	4.6 %	> 1000	0.39 %	36.7	47.64
	HuMoR [31]	✗	✗	✓	100%	311.5	3.1	0.1 %	38.1	0.3%	36.4	11.73
	HybrIK [16] w/ RootNet [24]	✗	✗	✗	-	> 1000	310.6	4.4%	384.7	0.6 %	48.9	**0.08**
	Ours	✓	✓	✗	98.3%	**171.7**	**9.4**	**0%**	**24.9**	**0 %**	**16.6**	0.12

Results on PROX Table 2 shows our result in the PROX dataset where we compare against SOTA multi-stage batch optimization based methods: LEMO [46] and HuMoR [31]. We also report the result using HyperIK with RootNet [16] to show that severe occlusion and human scene interaction on the PROX dataset is also challenging to SOTA regression methods. From Table 2, we can see that our method achieves the best ACD (chamfer distance) *without* using any depth information and is competitive in all other metrics against both RGB and RGB-D based methods. Note that LEMO and HuMoR directly optimize the chamfer distance and ground penetration during optimization, while our method relies only on 2D keypoints and laws of physics to achieve plausible pose estimates. Methods with exceedingly high ACD values (> 1000) can be attributed to unstable optimization or unreliable depth estimation, which causes the humanoid to fly away for some sequences.

From Fig.4 we can see that our method can estimate accurate 3D absolute poses that align with the 3D scene. Compared to batch multistage optimization methods that utilize RGB-D input, PROX [7] and LEMO [46], our method can achieve comparable results while running in near-real-time. Compared

Figure 4: Results of HybrIK [16], PROX [7], LEMO [46], and HuMoR [31]. Last two rows are failure cases.

to the regression-based method, HybrIK [16] with RootNet [24], our method can estimate much more accurate root translation and body pose. In terms of physical plausibility, optimization-based methods often perform well but can still exhibit floating and penetration. While the optimization methods can take over a minute per-frame to estimate poses, catastrophic failure can still occur (as shown in the fourth row). Since motion is best seen in videos, please refer to our project website for a quantitative video comparison.

4.2 Ablation

In Table 3, we ablate the main components of our framework. Row 1 (R1) corresponds to not using our modified scene geometry and uses the mesh provided by the PROX dataset; row 2 corresponds to not using any scene awareness (voxel features); R2 uses MPG with only one time-step, which results in nosier gradients; R3 trains a model without using the geometric translation step in $\mathcal{G}$; R4 does not use the TCN model $\mathcal{F}_{\text{TCN}}$ in $\mathcal{G}$, R5 does not leverage modified 3D scenes and uses the original scene scans. R1 shows that our modified geometry leads to better overall performance, as the scanned mesh does not correspond well to the actual scene. R2 demonstrates that scene awareness further improves our result, where our agent learns important affordance information about chairs and sofas to improve the success rate. R3 highlights the importance of taking the gradient step multiple times to obtain better gradients. When comparing R4 and R5 with R6, we can see that the learned function $\mathcal{G}$ is essential to obtain a better root translation and orientation movement signal, without which the agent drifts significantly.

Table 3: Ablation study of different components of our framework.

								Ground		Scene	
w/ Modified Scene	w/ Scene Awareness	MGP w/o multi-step	MPG w/o Geometric	MPG w/o TCN	SR ↑	ACD ↓	Accel. ↓	Freq. ↓	Pen. ↓	Freq. ↓	Pen. ↓
✗	✓	✓	✓	✓	95.0%	252.1	9.8	0.0%	26.4	0.0%	18.4
✓	✗	✓	✓	✓	95%	265.0	7.8	0.0%	25.9	0.0%	15.4
✓	✓	✗	✓	✓	96.7%	326.9	8.6	0.0%	24.8	0.0%	**11.4**
✓	✓	✓	✗	✓	91.7%	604.6	10.8	0.0%	25.3	0.0%	11.9
✓	✓	✓	✓	✗	95.0%	223.2	**7.8**	0.0%	26.6	0.0%	18.5
✓	✓	✓	✓	✓	98.3%	**171.7**	9.4	**0.0%**	**24.9**	**0 %**	16.6

5 Discussions

Failure Cases We show the failure cases of our proposed method in the last two rows of Fig. 4. Although our method can estimate human pose and human-scene interaction from the PROX dataset with a high success rate, there is still much room for improvement. We simplify all humanoids and scene geometries as rigid bodies and do not model any soft materials such as human skin, cushions, or beds. As a result, our sitting and lying motion can often result in jittery motion, since contact occurs only in a very small surface area on the simulated body and hard surfaces. Due to the nature of physics simulation, our simulated humanoid sometimes gets stuck on a narrow gap between chairs and tables, or when lying on beds. In these scenarios, our method can often recover quickly by jolting out of the current state and trying to follow the target kinematic pose. Nevertheless, if the drift of global translation is severe or many 2D keypoints are occluded, our agent may lose balance. Our causal formulation may also lead to jittery motion when 2D keypoint observations jitter. Currently, our agent has no freedom to deviate from the 2D observations and always tries to match the 2D motion even when it is stuck. More intelligent control modeling will be needed to improve these scenarios.

Conclusion and Future Work We introduce embodied scene-aware human pose estimation, where we formulate pose estimation in a known environment by controlling a simulated agent to follow 2D visual observations. Based on the 2D observation, its current body pose, and environmental features, our embodied agent acts in a simulated environment to estimate pose and interact with the environment. A multi-step projection gradient (MPG) is proposed to provide global movement cues to the agent by comparing the current body pose and target 2D observations. The experimental result shows that our method, trained only on synthetic datasets derived from popular motion datasets, can estimate accurate global 3D body pose on the challenging PROX dataset while using only RGB input. Future work includes providing a movement prior for the embodied agent to intelligently sample body motion based on scene features to better deal with occlusion.

Acknowledgements: This project is sponsored in part by the JST AIP Acceleration Research Grant (JPMJCR20U1).

References

[1] Y.-F. Cheng, B. Yang, B. Wang, Y. Wending, and R. T. Tan. Occlusion-aware networks for 3d human pose estimation in video. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 723–732, 2019.

[2] K. Cho, B. V. Merrienboer, Çaglar Gülçehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *ArXiv*, abs/1406.1078, 2014.

[3] Z. Geng, K. Sun, B. Xiao, Z. Zhang, and J. Wang. Bottom-up human pose estimation via disentangled keypoint regression. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14671–14681, 2021.

[4] K. Gong, B. Li, J. Zhang, T. Wang, J. Huang, M. B. Mi, J. Feng, and X. Wang. Posetriplet: Co-evolving 3d human pose estimation, imitation, and hallucination under self-supervision. *ArXiv*, abs/2203.15625, 2022.

[5] M. Hassan, D. Ceylan, R. Villegas, J. Saito, J. Yang, Y. Zhou, and M. J. Black. Stochastic scene-aware motion prediction. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11354–11364, 2021.

[6] M. Hassan, V. Choutas, D. Tzionas, and M. J. Black. Resolving 3d human pose ambiguities with 3d scene constraints. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2282–2292, 2019.

[7] M. Hassan, V. Choutas, D. Tzionas, and M. J. Black. Resolving 3D human pose ambiguities with 3D scene constraints. In *International Conference on Computer Vision*, Oct. 2019.

[8] E. Hedlin, H. Rhodin, and K. M. Yi. A simple method to boost human pose estimation accuracy by correcting the joint regressor for the human3.6m dataset. *ArXiv*, abs/2205.00076, 2022.

[9] B. Huang, L. Pan, Y. Yang, J. Ju, and Y. Wang. Neural mocon: Neural motion control for physically plausible human motion capture. *ArXiv*, abs/2203.14065, 2022.

[10] C. Ionescu, D. Papava, V. Olaru, and C. 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.

[11] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik. End-to-end recovery of human shape and pose. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7122–7131, 2018.

[12] A. Kanazawa, J. Y. Zhang, P. Felsen, and J. Malik. Learning 3d human dynamics from video. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5607–5616, 2019.

[13] M. Kocabas, N. Athanasiou, and M. 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.

[14] M. Kocabas, C.-H. P. Huang, J. Tesch, L. Muller, O. Hilliges, and M. J. Black. Spec: Seeing people in the wild with an estimated camera. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11015–11025, 2021.

[15] N. Kolotouros, G. Pavlakos, M. J. Black, and K. 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.

[16] J. Li, C. Xu, Z. Chen, S. Bian, L. Yang, and C. Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3383–3393, 2021.

[17] X. Li, S. Liu, K. Kim, X. Wang, M.-H. Yang, and J. Kautz. Putting humans in a scene: Learning affordance in 3d indoor environments. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12360–12368, 2019.

[18] T.-Y. Lin, M. Maire, S. J. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.

[19] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. Smpl: a skinned multi-person linear model. *ACM Trans. Graph.*, 34:248:1–248:16, 2015.

[20] Z. Luo, S. A. Golestaneh, and K. M. Kitani. 3d human motion estimation via motion compression and refinement. In *ACCV*, 2020.

[21] Z. Luo, R. Hachiuma, Y. Yuan, and K. Kitani. Dynamics-regulated kinematic policy for egocentric pose estimation. In *Advances in Neural Information Processing Systems*, 2021.

[22] N. Mahmood, N. Ghorbani, N. Troje, G. Pons-Moll, and M. J. Black. Amass: Archive of motion capture as surface shapes. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5441–5450, 2019.

[23] A. Monszpart, P. Guerrero, D. Ceylan, E. Yumer, and N. J. Mitra. imapper: Interaction-guided joint scene and human motion mapping from monocular videos. *ArXiv*, abs/1806.07889, 2019.

[24] G. Moon, J. Y. Chang, and K. M. Lee. Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10132–10141, 2019.

[25] G. Moon and K. M. 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.
[26] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. A. Osman, D. Tzionas, and M. J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019.
[27] D. Pavllo, C. Feichtenhofer, D. Grangier, and M. Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7745–7754, 2019.
[28] X. Peng, M. Chang, G. Zhang, P. Abbeel, and S. Levine. Mcp: Learning composable hierarchical control with multiplicative compositional policies. In *NeurIPS*, 2019.
[29] X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph.*, 37:143:1–143:14, 2018.
[30] N. Reddy, L. Guigues, L. Pischulini, J. Eledath, and S. G. Narasimhan. Tessetrack: End-to-end learnable multi-person articulated 3d pose tracking. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15185–15195, 2021.
[31] D. Rempe, T. Birdal, A. Hertzmann, J. Yang, S. Sridhar, and L. J. Guibas. Humor: 3d human motion model for robust pose estimation. In *International Conference on Computer Vision (ICCV)*, 2021.
[32] J. Romero, D. Tzionas, and M. J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ArXiv*, abs/2201.02610, 2022.
[33] I. Sárándi, T. Linder, K. O. Arras, and B. Leibe. MeTRAbs: metric-scale truncation-robust heatmaps for absolute 3D human pose estimation. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(1):16–30, 2021.
[34] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *ArXiv*, abs/1707.06347, 2017.
[35] S. Shimada, V. Golyanik, W. Xu, P. Pérez, and C. Theobalt. Neural monocular 3d human motion capture. *ACM Transactions on Graphics*, 40(4), aug 2021.
[36] S. Shimada, V. Golyanik, W. Xu, and C. Theobalt. Physcap: Physically plausible monocular 3d motion capture in real time. *ACM Trans. Graph.*, 39:235:1–235:16, 2020.
[37] J. Song, X. Chen, and O. Hilliges. Human body model fitting by learned gradient descent. *ArXiv*, abs/2008.08474, 2020.
[38] S. Starke, H. Zhang, T. Komura, and J. Saito. Neural state machine for character-scene interactions. *ACM Transactions on Graphics (TOG)*, 38:1 – 14, 2019.
[39] Y. Tian, H. Zhang, Y. Liu, and L. Wang. Recovering 3d human mesh from monocular images: A survey. *ArXiv*, abs/2203.01923, 2022.
[40] 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.
[41] Z. Wang, L. Chen, S. Rathore, D. Shin, and C. C. Fowlkes. Geometric pose affordance: 3d human pose with scene constraints. *ArXiv*, abs/1905.07718, 2019.
[42] K. Xie, T. Wang, U. Iqbal, Y. Guo, S. Fidler, and F. Shkurti. Physics-based human motion estimation and synthesis from videos. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11512–11521, 2021.
[43] Y. Yuan and K. Kitani. 3d ego-pose estimation via imitation learning. In *The European Conference on Computer Vision (ECCV)*, Sep. 2018.
[44] Y. Yuan and K. Kitani. Ego-pose estimation and forecasting as real-time pd control. In *IEEE International Conference on Computer Vision (ICCV)*, pages 10082–10092, Oct. 2019.
[45] Y. Yuan, S.-E. Wei, T. Simon, K. Kitani, and J. Saragih. Simpoe: Simulated character control for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
[46] S. Zhang, Y. Zhang, F. Bogo, P. Marc, and S. Tang. Learning motion priors for 4d human body capture in 3d scenes. In *International Conference on Computer Vision (ICCV)*, Oct. 2021.
[47] S. Zhang, Y. Zhang, F. Bogo, M. Pollefeys, and S. Tang. Learning motion priors for 4d human body capture in 3d scenes. *ArXiv*, abs/2108.10399, 2021.
[48] S. Zhang, Y. Zhang, Q. Ma, M. J. Black, and S. Tang. Place: Proximity learning of articulation and contact in 3d environments. *2020 International Conference on 3D Vision (3DV)*, pages 642–651, 2020.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Sec.5.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] In the supplementary material
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Code included in supplemental material
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Included in supplemental material
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Appendix

A Qualitative Results **14**

A.1 Supplementary Video 14

B Multi-step Projection Gradient (MPG) **14**

B.1 Forward Kinematics and Keypoint Definition 14

B.2 MPG and Geometric and Learning-based Transform $\mathcal{G}$ 15

C Details about Evaluation **16**

C.1 Evaluation Metrics 16

C.2 Simplified Scenes 17

C.3 Clarification about evaluation on H36M 17

D Implementation Details **18**

D.1 Shape-based Universal Humanoid Controller 18

D.2 Embodied Scene-aware Kinematic Policy 18

E Broader Social Impact **19**

In this Appendix, we provide extended qualitative results, more details about our Multi-step Projection Gradient as feature, evaluation, and implementation. **All code will be released for research purposes**. For supplementary videos, please refer to our project page: https://embodiedscene.github.io/embodiedpose/.

A Qualitative Results

A.1 Supplementary Video

As motion is best seen in videos, we provide a supplemental video at our project site to showcase our method. Specifically, we provide qualitative results on the PROX and H36M datasets and also compare with state-of-the-art RGB-based methods (HuMoR-RGB and HyberIK). From the visual result, we can see that our method can estimate physically valid human motion from 2D keypoint inputs in a causal fashion and control a simulated character to interact with the scenes. Compared to non-physics-based method, ours demonstrates less penetration and remain relatively stable under occlusion. For future work, incorporating more physical prior and simulating soft objects such as cushions can further improve motion quality.

B Multi-step Projection Gradient (MPG)

In this section, we provide a detailed formulation for our multi-step projection gradient as a feature formulation. First, we provide our keypoint definition based on forward kinematics in Sec. B.1. Then, we will describe the geometric and learning-based transformation function $\mathcal{G}$ for augmenting MPG.

B.1 Forward Kinematics and Keypoint Definition

The SMPL body model defines 24 body joints (green dots visualized in Fig. 5), which correspond to the 24 angles of the joints. For data sets such as H36M [10] and COCO [18], some joints of the SMPL human body do not have a one-to-one correspondence. Thus, prior SMPL-based human pose estimation methods often make use of a predefined joint regressor to recover the corresponding joints from the body vertices. As described in Sec.3.1, for the same body shape, the lengths of the limbs can change according to body pose, and changing the joint regressor may even increase performance

[8]. However, it is computationally expensive to regress 3D joint locations with linear blend skinning at every MPG step. Thus, we opt to use a skeleton with fixed limb length and forward kinematics for our embodied agent toward real-time humanoid control. To obtain limb length, we use the rest pose (T-pose) of the SMPL body and recover the joint positions using the pre-defined joint regressor. From the joint positions, we compute the fixed limb lengths following the parent-child relationship defined in the kinematics tree.

For both the COCO 2D keypoints [18] and H36M 3D joint positions of H36M [10], we use a subset of these points to obtain the corresponding points on the SMPL kinematics tree. Specifically, we use 12 joints, a similar joint definition to HuMoR [31], as visualized in Fig. 5. To obtain correspondences, we choose the following joint numbers from the 24 keypoints based on SMPL: #15, #16, #17, #20, #21, #22, #2, #3, #6, #7 without any processing. Then we average the position of joint #15 and #20 to obtain the mid-chest point. The pelvis point is obtained by averaging the 3 joints #0, #1, and #5. Following the process mentioned above, we obtain a sparse set of points (12) that we can directly obtain from the SMPL body pose parameters using forward kinematics to match visual observations and evaluation. Since not all 2D keypoints are confidently estimated, we use a confidence threshold of 0.1. For detected 2D keypoints of confidence < 0.1, we do not consider them and omit them during MPG calculation.

Figure 5: Visualization of the 12 joints we use for pose estimation and the 2D keytpoins obtained from SMPL using forward kinematics. Red is the DEKR 2D keypoints detection result [3], white is our pose estimation result, and green (numbered) is the original 2D keypoints obtained from the projection joints of SMPL.

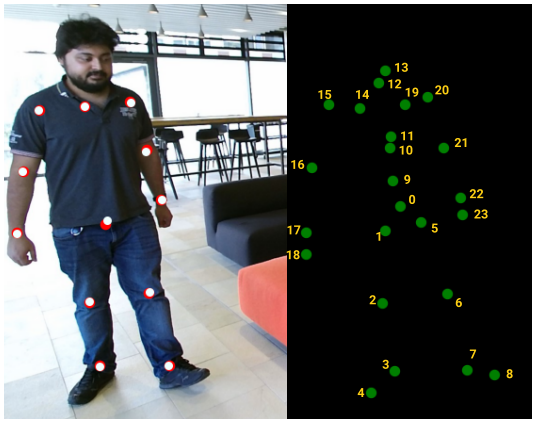

B.2 MPG and Geometric and Learning-based Transform $\mathcal{G}$

We define the 3D body pose for time step t as $\boldsymbol{q}_t \triangleq (\boldsymbol{r}_t^{\mathrm{R}}, \boldsymbol{r}_t^{\mathrm{T}}, \boldsymbol{\theta}_t)$ where $\boldsymbol{q}_t \in \mathcal{R}^{75}$. $\boldsymbol{r}_t^{\mathrm{R}} \in \mathcal{R}^3$ and $\boldsymbol{\theta}_t \in \mathcal{R}^{69}$ are root and body joint orientations in axis-angles, while $\boldsymbol{r}_t^{\mathrm{T}} \in \mathcal{R}^3$ is the global translation.

Geometric Transform of $\mathcal{G}$ for $\partial \boldsymbol{r}_t^{\mathbf{R}}$ Due to depth ambiguity and severe occlusion, recovering global translation accurately from 2D observations alone is an ill-posed problem. Previous attempts that use image features, RootNet [24] and 2D keypoints [27], can often result in unreliable translation estimate (as can be seen in Tab.2 and our quantitative video). Our pose gradient, on the other hand, provides *temporal* movement information instead of estimating global translation in one-shot. Using the temporal smoothness of human movement and the laws of physics, our pose gradient is more robust to occlusion and will not produce sudden jumps in pose estimates. To provide better root translation, we also take advantage of a geometric transformation step to further improve $\partial \boldsymbol{r}_t^{\mathrm{T}}$. Given 2D keypoint observation $\mathrm{kp}_{t+1}^{\mathrm{2D}}$ and body pose $\boldsymbol{q}_t^{\boldsymbol{k}}$ in this gradient descent iteration, we can treat a human skeleton as a rigid body to obtain a better translation. Specifically, we first compute the world-space 3D keypoints using forward kinematics $J_t = \mathrm{FK}_\beta(\boldsymbol{q}_t^{\boldsymbol{k}})$. Then, we optimize the following reprojection error;

$$\underset{\Delta \boldsymbol{r}_t^{\mathrm{T}-k}}{\arg\min} ||\pi_{\mathrm{cam}}(\boldsymbol{J}_t + \Delta \boldsymbol{r}_t^{\mathrm{T}-k}) - \mathrm{kp}_{t+1}^{\mathrm{2D}}||_2^2 \tag{5}$$

We solve this equation using SVD to obtain $\Delta \boldsymbol{r}_t^{\text{T-k}}$ and $\partial \boldsymbol{r}_t^{\text{T-k}\prime}$ is updated by Eq. 6.

$$\partial \boldsymbol{r}_t^{\text{T-k}\prime} = \partial \boldsymbol{r}_t^{\text{T-k}} + \Delta \boldsymbol{r}_t^{\text{T-k}} \tag{6}$$

Notice that $\partial \boldsymbol{r}_t^{\text{T-k}\prime}$ is guaranteed to lead to a 2D reprojection error no bigger than the original $\partial \boldsymbol{r}_t^{\mathrm{T}}$, so we use a direct additive transform to obtain a better global root translation gradient.

Learning-based Transform of $\mathcal{G}$ for $\partial \boldsymbol{r}_t^{\mathbf{T}}$. Unlike global translation where depth ambiguity can cause sudden jumps in depth perception, we find that global rotation can be robustly estimated using

a sequence of 2D keypoints only. Inspired by PoseTriplet [4], we use the temporal convolution model (TCN) [27] to recover the camera space 3D body orientation. Unlike PoseTriplet [4], we *only* regress the 3D body orientation and do not regress the 3D body pose: the TCN model is not quite robust to occluded or incorrect 2D keypoints (e.g. missing leg joints, which is prevalent when people are moving behind furniture). However, we notice that the TCN can still recover a reasonable root orientation by reasoning about reliable body joints (such as arms and torsos). Thus, we use the TCN to regress the camera space root orientation $r_t'^{\mathrm{R}} = \pi_{\mathrm{cam}}^{-1}(\mathcal{F}_{\mathrm{TCN}}(\mathrm{kp}^{\mathrm{2D}}_{t-\mu:t+1}))$ based on a window of 2D keypoints. We then calculate the difference between current root rotationr_t^{R} and $r_t'^{\mathrm{R}}$: $\Delta r_t'^{\mathrm{R}} = (r_t^{\mathrm{R}})^{-1} r_t'^{\mathrm{R}}$. We directly add $\Delta r_t'^{\mathrm{R}}$ as an additional entry to MPG instead of modifying $\partial \boldsymbol{r}_t^{\mathrm{R}-k}$ since there is no guarantee that the 2D reprojection error from using $\Delta r_t'^{\mathrm{R}}$ will be better. We compute $\Delta r_t'^{\mathrm{R}}$ only once instead of K times. Combining the above two components, we obtain our complete MPG feature.

$$\phi_t = [\Delta r_t^{\mathrm{R}}, \partial \boldsymbol{r}_t^{\mathrm{R}}, \partial \boldsymbol{r}_t^{\mathrm{T\text{-}k'}}, \partial \boldsymbol{q_t}] = \sum_{k=0}^{K} \mathcal{G}(\frac{\partial \mathcal{L}_{\mathrm{proj}}}{\partial \boldsymbol{q}_t^k}, \boldsymbol{q}_t^k, \mathrm{kp}^{\mathrm{2D}}_{t-\mu:t+1}) \tag{7}$$

C Details about Evaluation

C.1 Evaluation Metrics

Here we provide a detailed list of our metric definition.

- SR: success rate measures whether the pose sequence has been estimated successfully. For optimization-based methods, LEMO [46] and HuMoR [31], it indicates whether the optimization process has successfully converged or not. For physics-based methods (ours), it indicates whether the agent has lost balance or not during pose estimation. We measure losing balance as when the position of the simulated agent $\boldsymbol{q}_t$ and the output of the kinematic policy $\tilde{\boldsymbol{q}}_t$ deviates too much from each other (> 0.15 meters per joint on average). For regression-based methods, *e.g.*, HyberIK [16], we do not use this metric.
- ACD (mm): one-side average chamfer distance measures the chamfer distance between the estimated human mesh and the point cloud captured from the RGB-D camera. This metric is a proxy measurement of how well the policy can estimate global body translation, as used in HuMoR [31]. Notice that RGB-D based methods such as HuMoR-RGBD and Prox-D both directly optimize this value, while our method does not use any depth information. Unstable optimization and degenerate solutions (where the humanoid flies away) can lead to a very large ACD value.
- ACC (mm/frame^2): acceleration measures the joint acceleration and indicates whether the estimated motion is jittery. From the MoCap motion sequences from the AMASS dataset, a value of around $3 \sim 10$ mm/frame^2 indicates human motion that are smooth. When there is no ground-truth motion available, the joint acceleration serves as a proxy measurement for physical realism.
- Freq. ($\%$): frequency of penetration measures how often a scene or ground penetration of more than 100 mm occurs. We obtain the scene penetration measurement by using the mesh vertices obtained from the SMPL body model.
- Pen (mm): penetration measures the mean absolute values of the penetrated vertices in the scene / ground, similar to the penetration metric in SimPOE [45].
- MPJPE (mm): mean per-joint position error is a popular pose-based metric that measures how well the estimated joint matches the ground-truth MoCap joint locations. This metric is calculated in a root-relative fashion where the root joint ($\#0$, as shown in Fig.5 is aligned. In H36M, this is commonly referred to as Protocol 1.
- PA-MPJPE (mm): mean per-joint position error after procrustes alignment. In H36M, this is usually referred to as Protocol 2.
- A-MPJPE (mm): absolute mean per-joint position error measures the absolute (non root-relative) joint position error in the world coordinate frame.
- ACC-D (mm/frame^2): acceleration error measures the difference between the estimated joint position acceleration and the ground truth joint acceleration.

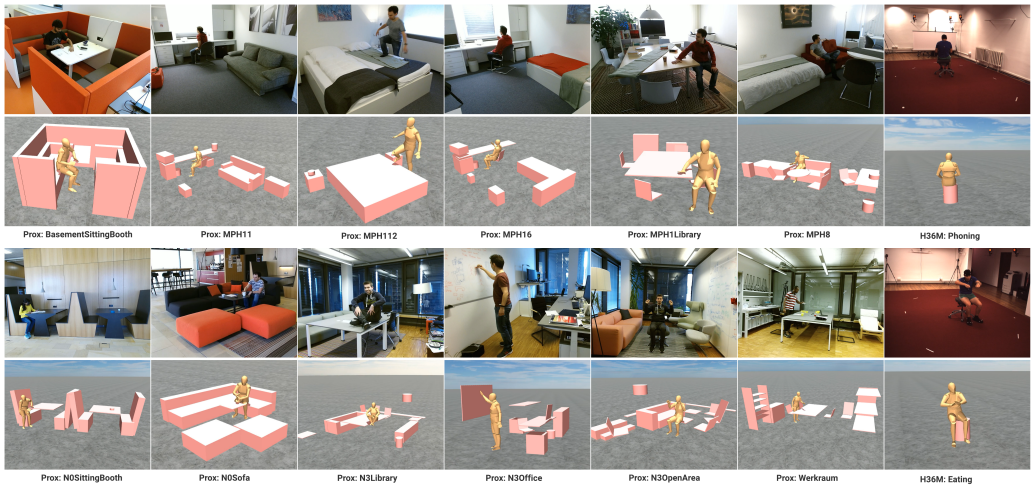

Figure 7: Visualization of the 3D scenes used for simulation for both Prox and H36M dataset. For each scene, we approximate the scene geometries with simple shapes. The humanoids visualized in these scenes are our pose estimation result from 2D keypoints.

C.2 Simplified Scenes

As described in Sec.3.2, we make simplified scenes based on the scanned meshes to better approximate the actual 3D scene. Due to inaccuracy of the scanning procedure and point cloud to mesh algorithms, the scanned meshes are often bulkier than the real environment. As a result, the gaps between chairs and tables are often too small in the mesh form, as visualized in Fig. 6. These inaccuracies can often cause the humanoid to get stuck between chairs and tables when trying to follow the 2D observations. To resolve this issue, we created simplified scenes from the scanned meshes to better match the videos. We create these scenes semi-automatically by selecting keypoints in the mesh vertices and create rectangular cuboids and cylinders. Another advantage of using these simplified scenes is that the signed distance functions for these simplified gems can be computed analytically and efficiently. Simplified scenes are created for all scenes from the Prox and H36M datasets. All scenes are visualized in Fig.7.

Figure 6: Comparison between the scanned scene and our simplified scene based on simple shapes. Notice that the scanned scene has much thicker chairs and tables, which leave smaller gaps for the agent to sit down.

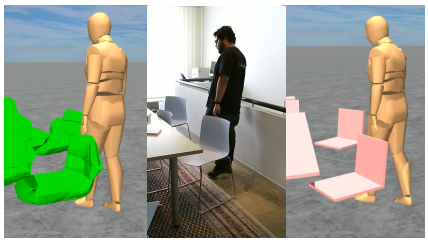

C.3 Clarification about evaluation on H36M

We conduct quantitative evaluations on the H36M dataset, as shown in Table 1. Since we do not make use of a joint regressor, we cannot perform the standard evaluation using the 14 joints commonly used for the evaluation on the H36M dataset. Instead, we use the 12 joints defined in Sec.B.1 to calculate MPJPE and A-MPJPE. As input, we use the same set of 12 2D keypoints. The 2D keypoints are estimated using a pre-trained DEKR [3] without additional fine-tuning on the H36M dataset. Although our method does not obtain the state-of-the-art result on H36M, we want to point out that 1) our method is trained on purely synthetic data and has never seen detected 2D keypoints from H36M videos, 2) we do not use any image-level information and only relies on a small set of 2D keypoints to infer global pose. Under such conditions, it is especially challenging when the person's movement direction is parallel to the forward direction of the camera (the person is moving along the depth direction in the image plane). For quantitative results, please refer to our supplementary video. This experiment also shows our method's ability to estimate pose in a generic setting where there are almost no scene elements (except for the ground and sporadically a chair).

D Implementation Details

D.1 Shape-based Universal Humanoid Controller

We extend the universal humanoid controller (UHC) proposed in Kin_poly [21] to also support humanoids with different body shapes. Specifically, we added the shape parameter β to the state space of UHC to form our body shape conditioned UHC: $\pi_{\text{UHC}}(\boldsymbol{a}_t|\boldsymbol{s}_t, \widehat{\boldsymbol{q}}_t, \beta)$. In the original UHC, the body shape information is discarded during training, and the policy is learned using a humanoid based on the mean SMPL body shape. In ours, we create different humanoids for each body shape during training using the ground-truth body shape annotation from AMASS. In this way, we can learn a humanoid controller that can control humanoids with different body shapes. The action space and policy network architecture, and the reward function remain the same.

Training procedure. We train on the training split of the AMASS dataset [22] and remove motion sequences that involve human-object interactions (such as stepping on stairs or leaning on tables). This results in 11402 high-quality motion sequences. At the beginning of each episode, we sample a fixed-length sequence (maximum length 300 frames) to train our motion controller. Reference state initialization [29] is uesed to randomly choose the starting point for our sampled sequence and we terminate the simulation if the simulated character deviates too far from the reference motion $\widehat{\boldsymbol{q}}_{\boldsymbol{t}}^{\boldsymbol{p}}$. To choose which motion sequence to learn from, we employ a hard-negative mining technique based on the success rate of the sequence during training. For each motion sequence $\widehat{\boldsymbol{Q}^{\boldsymbol{i}}} = \widehat{\boldsymbol{q}}_{\boldsymbol{0}}, \ldots, \widehat{\boldsymbol{q}}_{\boldsymbol{T}}$, we maintain a history queue (of maximum length 50) for the most recent times the sequence is sampled: $h_1^i, \ldots h_{50}^i$ where each h_0^i is a Boolean variable indicating whether the controller has successfully imitated the sequence during training. We compute the expected success rate s^i of the sequence based on the exponentially weighted moving average using the history $s^i = \text{ewma}(h_1^i, \ldots h_{50}^i)$. The probability of sampling sequence i is then $P\left(\widehat{\boldsymbol{Q}}^i\right) = \frac{\exp\left(-s^i/\tau\right)}{\sum_i^j \exp(-s^i/\tau)}$ where τ is a temperature parameter. Intuitively, the more we fail at imitating a sequence, the more likely we will sample it.

D.2 Embodied Scene-aware Kinematic Policy

Network Architecture and Hyperparameters Our kinematic policy is an RNN-based policy with multiplicative compositional control [28]. It consists of a Gated Recurrent Unit (GRU)[2] based feature extractor to provide memory for our policy. Then, 6 primitives $\pi_{\text{KIN}}^{\text{prim}}$ are used to provide a diverse set of pose estimates based on 2D observation $\boldsymbol{\phi}_i$ (MPG) and current body state $\boldsymbol{s}_t$. We process the environmental features separately from the pose features due to its high dimensionality, using an MLP $\mathcal{S}_t$. The output of $\mathcal{S}_t$, concatenated with the RNN output, is fed into a composer $\pi_{\text{KIN}}^{\text{comp}}$ to produce the action weights. Intuitively, the primitive produces potential body movement based on the 2D observation and current body pose. Then, the composer chooses plausible action based on the environmental features. The detailed data flow of our policy is visualized in Fig. 8. The hyperparameters we used for training are shown in Table 4.

Training procedure To train our kinematic policy, we sample motion sequences from AMASS [22], H36M [10], and Kin_poly [20] to generate synthetic paired 2D keypoints. To learn a network that is more robust to 2D keypoints detected from real videos, we randomly sample confidence from a uniform $[0, 1]$ distribution as the detection confidence. Keypoints with confidence < 0.1 will be omitted in our model. For each episode, we randomly sample a sequence, randomly change its heading, and sample a valid starting point in the camera frustum as visualized in Fig.9. We constrain the newly sampled sequence to be always visible in the image plane and re-sample if not. With a

Table 4: Hyperparameters for our scene-aware kinematic policy.

	Batch Size	Learning Rate	Voxel grid	w_{L1}	w_{prior}
Value	5000	5×10^{-4}	$16 \times 16 \times 16$	5	0.0001
	TCN window	MPG - #steps	MLP-hidden	RNN-hidden	# primitives
Value	81	5	[1024, 512, 256]	512	6

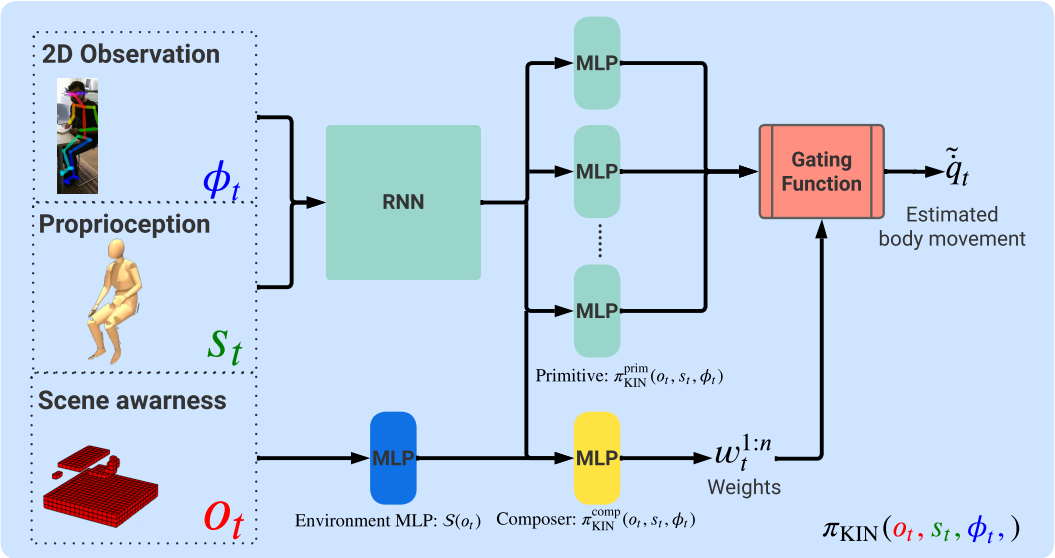

Figure 8: Network architecture for our scene-aware kinematic policy

Figure 9: Visualization of a randomly sampled motion sequence and paired 2D keypoints for training. Here, we sample from the AMASS dataset (sequence name in the caption).

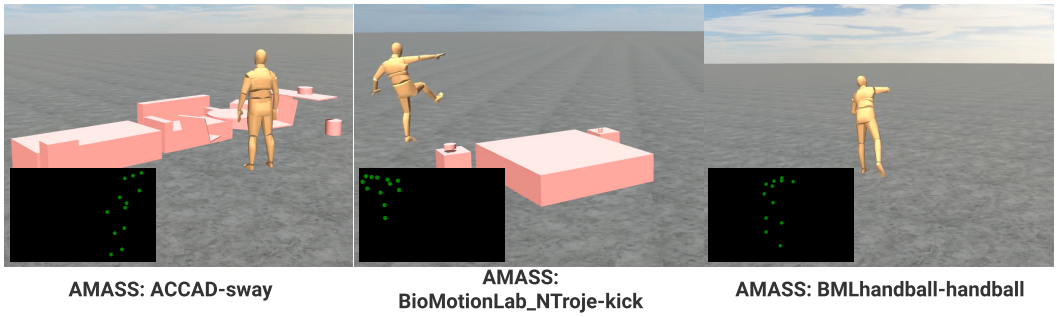

probability of 0.5, we also sample a random scene from the PROX dataset to let our agent learn human-scene interactions. Although the randomly sampled motion sequence does not always yield reasonable human-scene interactions (e.g. the reference motion may walk through a chair), collision and physics simulation will stop our agent automatically. In this way, our agent can learn to use the scene interaction term based on acting inside the scene.

E Broader Social Impact

This research focuses on estimating physically valid human poses from monocular videos. Such a method can be positively used for animation, telepresence, at home robotics, etc. where natural looking human pose can be served as the basis for higher-level methods. It can also lead to malicious use cases, such as illegal surveillance and video synthesis. As our method focuses on fast and causal inference, it is possible to significantly improve its speed and deploy it to real-time scenarios. Thus, it is essential to deploy these algorithms with care and make sure that the extracted human poses are with consent and not misused.