

Learned Neural Physics Simulation for Articulated 3D Human Pose Reconstruction

Mykhaylo Andriluka¹, Baruch Tabanpour¹, C. Daniel Freeman^{2*}, and Cristian Sminchisescu¹

¹Google DeepMind ²Anthropic

Abstract. We propose a novel neural network approach to model the dynamics of articulated human motion with contact. Our goal is to develop a faster and more convenient alternative to traditional physics simulators for use in computer vision tasks such as human motion reconstruction from video. To that end we introduce a training procedure and model components that support the construction of a recurrent neural architecture to accurately learn to simulate articulated rigid body dynamics. Our neural architecture (LARP) supports features typically found in traditional physics simulators, such as modeling of joint motors, variable dimensions of body parts, contact between body parts and objects, yet it is differentiable, and an order of magnitude faster than traditional systems when multiple simulations are run in parallel. To demonstrate the value of our approach we use it as a drop-in replacement for a state-of-the-art classical non-differentiable simulator in an existing video-based 3D human pose reconstruction framework [13] and show comparable or better accuracy.

1 Introduction

We introduce a neural network approach to modeling rigid body dynamics [10] often required for physics simulation of articulated human motion with the objective to lower the bar for using physics-based reasoning in human reconstruction and synthesis. Towards that goal, we propose a physically grounded articulated motion model that is comparable in accuracy to state of the art classical physics simulators (*e.g.* [8, 37]) but is significantly faster. Being developed in terms of standard deep learning building blocks, it enables easy integration with other modern optimization and learning components.

Our approach can be interpreted as neural simulator for a subset of rigid body dynamics and is closely related to an established line of work on neural simulation [1, 12, 30]. We refer to our approach as *LARP* (**L**earned **A**rticulated **R**igid body **P**hysics). As common in the literature, we train LARP on examples provided by a standard off-the-shelf physics simulator, essentially learning to approximate the rigid body dynamics by means of a neural network. This neural network approximation has several desirable properties. (1) It can be significantly

* Work done while the author was with Google DeepMind.




Fig. 1: Left: Examples of articulated 3d human pose reconstructions obtained with LARP on public benchmarks [22, 38] and real world video. Right: Comparison to common physics simulators [8, 11, 29, 37] in terms of simulation speed. The x-axis shows the number of parallel simulations, whereas the y-axis shows the total time taken to advance all the simulations to the next step.

faster to execute on parallel hardware as it is composed of simple building blocks well suited for parallelization. Traditional physics simulators typically solve an optimization problem at each simulation step to compute velocities that satisfy collision and joint constraints¹. Neural simulation can learn to approximate such calculations. In doing so it amortizes the solution of such optimization problems at training time, and allows for faster computation at test time. Our work is also related to how neural networks learn to directly predict the solution of optimization problems, for *e.g.* articulated pose estimation [23, 35]. (2) Another advantage of neural simulators is their construction in terms of standard end-to-end differentiable deep learning components. This is in contrast to some of the established physics simulators that are non-differentiable and not natively amenable to parallel simulation [8, 37]. Finally, neural physics engines can rely on parameters estimated directly from real-world data [1, 36], which enables realistic simulation even when analytic formulations cannot be obtained.

We model the motion of a set of articulated bodies using a recurrent neural network (RNN) with an explicit state aggregating the physical parameters of each body part. For each articulated object type we define a neural network that updates its state given the previous state, external forces and internal joint torques applied at the current step. Such networks, implemented as multi-layer perceptrons (MLP), are applied recursively over time. We use different MLPs for each object, *e.g.* person and ball in fig. 2. To model interactions between objects we define a collision sub-network that computes additional inputs for the MLPs associated to each object, via sum-pooling over other objects in the scene, similarly to collision handling in graph neural networks [30].

Contributions. Our main contribution is a neural network architecture and training procedure that results in a model of rigid body dynamics (LARP) that can compute accurate human motion trajectories up to an order of magnitude faster than traditional physics simulators (fig. 1). We measure accuracy both directly as well as in the context of physics-based 3D pose reconstruction from video by integrating LARP into the framework of [13].

¹ See Chapter 11 in [10].




Fig. 2: Left: Overview of our approach (LARP). At time t the input to the neural simulator is given by a state of the scene \mathbf{S}_t and joint control targets \mathbf{Q}_t^p . Here state is composed of the state of the person \mathbf{S}_t^p and ball \mathbf{S}_t^b . LARP propagates the state through time by recurrently applying contact and dynamics networks. We visualize the state of each rigid component of the articulated body using rectangles with a color matching the scene structures. Right: Scenarios used to evaluate our approach: chain of linked capsules (a), two colliding capsule chains (b), articulated pose reconstruction from video (c), human-ball and human-capsule collision handling (d). See Supp. Mat. for videos.

Even though the architecture of LARP is seemingly simple, the results significantly depend on training details such as data augmentation, gradient clipping, input features and length of sequences used for training (§4.1). As an additional contribution, we perform a detailed analysis of the model components and the impact of training parameters, identifying those ingredients that make the model perform well. We plan to make our implementation and pre-trained models publicly available upon publication.

2 Related Work

Our approach can be viewed as a special type of neural physics simulator focused on articulated human motion. Neural physics is an established research area with early work going at least back to [16]. Neural simulation has been applied to phenomena as diverse as weather forecasting [24], simulation of liquids, clothing and deformable objects [6, 26, 31], as well as groups of rigid bodies [2, 4, 17].

Particularly relevant to our work is the literature focused on rigid body dynamics. The paradigm in this area is to formulate neural simulation as a graph-convolutional neural network (GCN) [2, 4, 17]. Recent work has shown the ability to simulate large number of objects and correctly handle collisions even for objects with complex shapes [2, 4]. It has been demonstrated that in certain cases neural simulation can improve over standard rigid body simulators by directly training on real data [4, 36]. We believe that LARP is complementary to these works. Whereas they primarily focus on simulating large number of rigid objects passively moving through the scene, we instead focus on modeling objects with a large number of components connected by joints actively controlled by motors. To the best of our knowledge the only work that applies GCN to modeling dynamics of complex articulated objects is [33]. LARP is methodologically similar to [33], but addresses a more complex tasks of modeling human motion and interaction with scene objects (see fig. 6).

Closely related to our work is SuperTrack [12]. As in LARP, [12] trains a fully connected neural network that updates the dynamic state of human body parts. LARP can be seen as a generalization of [12], that enables multiple objects, object collisions, improved performance on long sequences, and support for bodies with variable dimensions of the body parts. We demonstrate (see fig. 7) that compared to SuperTrack, LARP can generate long-term motion trajectories without loss of consistency of the simulation state².

LARP can be seen as an autoregressive motion model that can generate physically plausible human motion conditioned on the control parameters of the body joints. The key difference to methods such as [14, 19, 20, 27, 32] is that LARP is trained on motion trajectories generated by physics simulation and does not include a probabilistic model for the joint control parameters. In contrast, models such as [32] are trained on motion capture data and focus on representing statistical dependencies in the articulated motion, but do not directly incorporate physics-related motion features. We see these approaches as complementary and hope to incorporate probabilistic motion control in the future.

3 Methodology

Overview. The architecture of LARP is shown in fig. 2. We represent the scene by a set of articulated objects, each corresponding to a tree of rigid components connected by joints. We refer to such components as “links” using physics simulation terminology [9]. Each joint is optionally equipped with a motor that can generate torque to actuate the object. The dynamic state of each object is given by position, orientation, and (rotational) velocity of its links, given in world coordinates. At each time step t LARP takes a state of the scene S_t and optionally a set of motor control targets \mathbf{Q}_t as input and produces the state of the scene at the next time step. This process is applied recurrently to obtain motion trajectories over longer time horizons.

Notation. Let us assume that the scene is composed of M articulated bodies each consisting of N links. Let $\mathbf{x}_{tb}^i \in \mathbb{R}^3$ denote the position of link $i \in [1 \dots N]$ of the body $b \in [1 \dots M]$ at time t . We use $\mathbf{q}_{tb}^i \in \mathbb{R}^4$ for the quaternion representing body orientation, and $\mathbf{v}_{tb}^i \in \mathbb{R}^3$ and $\omega_{tb}^i \in \mathbb{R}^3$ for linear and angular velocities, respectively. The torque applied at time t is $\tau_{tb}^i \in \mathbb{R}^3$, set to zero if no external torque was provided. The quaternions representing the joint targets are $\mathbf{Q}_{tb}^i \in \mathbb{R}^4$, and are also set to zero if joint targets are not specified. We assume that in our world representation the gravity is aligned with the z-axis, and the ground plane corresponds to $z = 0$.

3.1 Dynamics Network

The main component of LARP is a per-object dynamics network f_d . The network takes a concatenated vector of features encoding the state of the body links, and

² See sec. 7 and fig. 14 in [12] for discussion of limitations.

outputs the linear and rotational velocity of each link in the next time step. The dynamic network is implemented as a densely connected neural network with $L = 12$ layers and ELU nonlinearities [7]. In the following we describe the features used to represent state of each link. We drop the timestep and the articulated body index to avoid clutter.

Dynamic state. The features encoding the dynamic state of the link correspond to the root-relative position of each link \mathbf{x}_{rel}^i , the z-position of each link \mathbf{x}_z^i , world orientation of each link represented as a flattened rotation matrix (\mathbf{q}_{9d}^i), and linear and angular velocities. We experimented with other encodings of orientation such as quaternions or 6d representation and found that directly passing the rotation matrix works best. This representation is invariant to shifts parallel to the ground plane which we assume orthogonal to the z-axis. Note that we do not encode rotation invariance along the z-axis explicitly and instead induce it via data augmentation (§3.4). We experimentally found that taking shift and rotation invariances into account is essential in making the approach work.

Geometric information. For each link we include its length l^i , radius r^i as well as the displacement error \mathbf{d}^i of the joint to the parent link. Our state representation independently encodes the positions of all links in world coordinates. This is often referred to as a *maximal coordinates* representation in the physics simulation literature [11, 21]. Since positions of all links are updated independently, they might float apart and disagree on the position of the mutual joint after a number of update steps. To mitigate this effect we explicitly provide the difference between positions of the joint computed based on the state of the child and parent link, and add a loss function to penalize such displacement during training (§3.4). We experimentally show in that these components are essential for good performance (fig. 4 and fig. 3).

Control. We include the target orientation \mathbf{Q}^i of the link relative to the parent. For traditional simulators such as [9, 37] target orientation is used by the control algorithm (*e.g.* PD-control) to compute the torque applied by the joint motor, whereas LARP is supposed to learn from data how to transform control targets into updates of the link state. In addition we also provide an external torque τ^i which is used in the experiments in §4.1 to diversify the motion trajectories.

Contact. We include a feature vector $\hat{\mathbf{p}}^i \in \mathbb{R}^6$ computed by the contact network f_c that encodes the information about the interaction with links of other objects. The computation of the contact network f_c is described in §3.2.

3.2 Contact Network

Given two articulated bodies m and b , the contact feature $\hat{\mathbf{p}}_m^i$ is given by

$$\hat{\mathbf{p}}_m^i = \sum_j f_c(\phi_b^j, \phi_m^i, c(\phi_b^j, \phi_m^i), \theta_c). \quad (1)$$

The contact feature for link i is the summation of f_c over all links j , where f_c is a neural network with $K = 6$ fully-connected layers, the vector ϕ_b^i contains

ALGORITHM 1: Algorithm for training the dynamics and contact models given two articulated bodies.

Input: Dynamics network parameters θ_d , contact network parameters θ_c , a batch of sub-sequences S , dynamics network feature normalization μ_d , σ_d , contact network feature normalization μ_c , σ_c , dynamics network output normalization μ_{do} , σ_{do} , dynamics feature function f_x , contact feature function f_y , $dt=0.01$, $M = 2$, $N_h = 20$

Output: θ_d , θ_c

$n = 0$;
repeat
 $\hat{S}_0 = S_0$;
for $t = 1$ **to** N_h **do**
 $Y_{t,1} = (f_y(\text{stopgrad}(\hat{S}_{t,1})) - \mu_c)/\sigma_c$
 $Y_{t,2} = (f_y(\text{stopgrad}(\hat{S}_{t,2})) - \mu_c)/\sigma_c$
 $\hat{p}_{12}^{ij} = f_c(Y_{t,1}^i, Y_{t,2}^j, h(Y_{t,1}^i, Y_{t,2}^j), \theta_c)$
 $\hat{p}_{21}^{ji} = f_c(Y_{t,2}^j, Y_{t,1}^i, h(Y_{t,2}^j, Y_{t,1}^i), \theta_c)$
 $\hat{p}_1^i = \sum_j \hat{p}_{12}^{ij}$
 $\hat{p}_2^j = \sum_i \hat{p}_{21}^{ji}$
for $b = 1$ **to** M **do**
 $X_t = (f_x(\hat{S}_t, \hat{p}_b^1, \dots, \hat{p}_b^N) - \mu_d)/\sigma_d$
 $(\hat{\mathbf{v}}_{t+1}, \hat{\omega}_{t+1}) = f_d(X_t, \theta_d)$
 $(\hat{\mathbf{v}}_{t+1}, \hat{\omega}_{t+1}) = \sigma_{do}(\hat{\mathbf{v}}_{t+1}, \hat{\omega}_{t+1}) + \mu_{do}$
 $\hat{\mathbf{x}}_{t+1}, \hat{\mathbf{q}}_{t+1} = \text{Integrate}(\hat{S}_t, \hat{\mathbf{v}}_{t+1}, \hat{\omega}_{t+1})$
 $\hat{S}_{t+1} = (\hat{\mathbf{x}}_{t+1}, \hat{\mathbf{q}}_{t+1}, \hat{\mathbf{v}}_{t+1}, \hat{\omega}_{t+1})$
end
end
Compute loss L from Equation 5 using \hat{S} and S
 $(\theta_c, \theta_d) \leftarrow \text{Adam}(\theta_c, \theta_d, L)$
until $n = \text{Epochs}$;

position \mathbf{x}_b^i , quaternion orientation \mathbf{q}_b^i , velocity \mathbf{v}_b^i , angular velocity ω_b^i , capsule length l_b^i , capsule radius r_b^i , and mass m_b^i of link i . The function $c(\cdot)$ computes properties of the collision between two links: link-relative contact point, normal, and penetration distance. We use the capsule-capsule contact implementation in Brax [11] which is implemented in JAX and is differentiable [5]. Before feeding in features to the contact network, features are normalized by their mean and standard deviation obtained from a single batch of offline training data. As motivated empirically in §4.1, we apply a stop-gradient on the contact feature before feeding it into the dynamics network to stabilize training. Notice that the summation of f_c in (1) is similar in spirit to the summation to obtain net force (e.g. Newton’s Second Law), and is also similar to graph neural networks (GNN) used in neural physics simulators [3]; our approach is a special case of GNNs with the graph between articulated objects being fully connected.

3.3 Integrator

To compute the position and orientation of all links in the next timestep, we integrate the velocity outputs of the dynamics network f_d over a timestep dt . The world position and orientation of each link are given by $\hat{\mathbf{x}}_{t+1} = \mathbf{x}_t + \hat{\mathbf{v}}_t dt$ and $\hat{\mathbf{q}}_{t+1} = \mathbf{q}_t + g(\hat{\omega}_t, \mathbf{q}_t) dt$, where $g(\hat{\omega}_t, \mathbf{q}_t) = 0.5 \text{quat}(\hat{\omega}_t) \otimes \mathbf{q}_t$ computes quaternion derivative from rotational velocity, and \otimes is a quaternion product.

3.4 Training

We estimate the parameters of the dynamics and contact networks on a dataset of training examples corresponding to sequences of scene states comprised of the position, orientation, and 6d velocity of all links N and all articulated bodies M , over a period of T steps. Let's denote the given a ground truth sequence of states as $S = \{S_t | t = 1, \dots, T\}$, and a sequence of states produced by the model starting from the initial state S_0 by \hat{S} . The loss used to estimate parameters consists of a position, rotation, and joint displacement loss over a rollout of length T . The position loss is given by

$$L_p = \frac{1}{NMT} \sum_{i,j,t} (\mathbf{x}_t^{ij} - \hat{\mathbf{x}}_t^{ij})^2. \quad (2)$$

The rotation loss is

$$L_r = \frac{1}{NMT} \sum_{i,j,t} 1 - |\mathbf{q}_t^{ij} \cdot \hat{\mathbf{q}}_t^{ij}|, \quad (3)$$

and the joint displacement loss is the mean-squared error between the predicted child and parent joint position

$$L_d = \frac{1}{NMT} \sum_{i,j,t} (\hat{\mathbf{x}}_{pt}^{ij} - \hat{\mathbf{x}}_{ct}^{ij})^2. \quad (4)$$

The joint displacement loss encourages joint constraints to be respected. Notice that if joint constraints are not violated, $L_d = 0$. The total loss is

$$L = w_p L_p + w_r L_r + w_d L_d, \quad (5)$$

where $w_p = 1$, $w_r = 1$, and $w_d = 0.1$ are hyper-parameters that balance the loss terms to have similar weight during training. We summarize the details of model training for the case of $M = 2$ articulated bodies in alg. 1.

Implementation details. Depending on the dataset and setting we generate state trajectories with a length between 80 and 200 steps. At training time we then subdivide them into batches of shorter subsequences of length T . Note that $T = 1$ amounts to doing “teacher forcing” on every step, meaning that the model always gets one of the ground-truth states as input. We experimentally observe that models trained on short subsequences ($T < 5$) perform poorly (see §4.1). This is likely because there are subtle differences between ground-truth states and states generated by the model, and the model needs to be

exposed to sufficiently diverse generated states at training time. To stabilize the training for larger T we add gradient norm clipping and drop training batches with the gradient norm is above the threshold of 0.3. Note that the dynamics network outputs link velocities, but the loss in eq. (5) compares link positions and orientations. We experimented with a loss that compares velocities directly, but found that models trained with such loss performed poorly. Finally, models did not perform well unless we augmented training data by randomly rotating each batch around an axis orthogonal to the ground plane (z-axis in our scenes).

4 Results

In §4.1 we first analyze the importance of various components and the effect of training parameters using simple articulated objects corresponding to capsule chains (fig. 2 (a, b)). Equipped with the best settings from this analysis, in §4.2 we evaluate LARP on a more complex humanoid model with joint motors (fig. 2 (c, d)). Specifically, we compare LARP to off-the-shelf physics simulator on the task of reconstruction of human motion from monocular video. To that end we use the physics-based articulated pose estimation approach from [13] that employs the Bullet simulator [8], and replace it with LARP. We further evaluate the accuracy of simulating collisions between a human and an external object represented by a ball (or capsule), quantitatively compare LARP with related work and benchmark LARP simulation speed against several established physics simulation engines.




Fig. 3: Example of a generated sequence obtained with a model that includes displacement features and displacement loss (bottom row), and without either of these elements (middle row). We highlight inconsistencies of the joint positions with the red circles.

4.1 Analysis of the model components

Datasets and training. We begin our evaluation using the setting shown in fig. 2 (a) in which four capsules linked by joints are falling on the ground subject to random initial state and randomly applied initial torque. The capsules can freely rotate at the joint, and can slide and roll on the ground following impact. This setting is simple enough to enable quick experimentation yet is generally non-trivial since the motion of each capsule needs to obey gravity, ground collision, and joint constraints. We generate 500k sequences of length 200 for training, and 25k sequences for testing and validation. Each sequence starts with a unique random position, orientation, and velocity. We apply a random torque to the top capsule of the chain for the first 10 steps, and randomize capsule lengths and

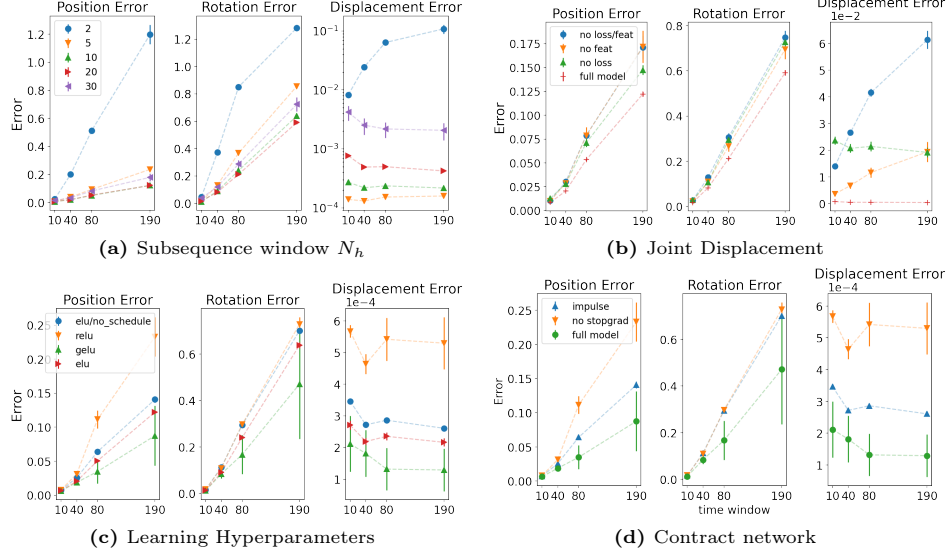


Fig. 4: Experiments with the dynamics network: training subsequence length N_h (a), ablations of the joint displacement feature and loss (b), ablations of non-linearity type and learning rate schedule (c), and evaluation of variants for the contact network (d). Error bars show one standard deviation calculated over 5 runs. The x-axis sweeps the time window over which metrics are computed.

radii for every link. We show an example of a generated sequence in fig. 3 (top row). We train the models for 200 epochs using batch size of 30k and learning rate of $1e-3$. To reduce the effect of random initialization of neural network parameters in the following experiments we train each network 5 times and use the validation set to pick the best model.

Analysis of the dynamics model. Our major finding is that a combination of training on longer sequences while using position, rotation, and joint displacement-based features and losses (eq. 4) with applying gradient clipping, leads to significantly improved results than reported in related work (e.g. [12]). We observe that results improve considerably when dynamics is unrolled for larger number of steps during training (see fig. 4 (a)). Position error is reduced from 1.1 meter to 11 cm when increasing the training sequence length from 2 to 20 steps. Training on longer sequences leads to somewhat larger joint displacement error, but consistently reduces the global position and rotation error. Given the setting of $N_h = 20$ steps for the training sequence length we evaluate the importance of displacement loss and displacement features in fig. 4 (b). We observe that removing these components has a significant negative impact: position error increases from 11 to 15cm, and joint displacement error increases from nearly 0 to 6.1cm. Gradient clipping was essential for training on longer sequences with $N_h > 10$. This is consistent with observations from training recurrent neural networks where gradient clipping is common in stabilizing training [15]. We also

	$N_h=10$	$N_h=40$	$N_h=80$	$N_h=190$
1 link	0.008	0.021	0.048	0.173
2 link	0.007	0.023	0.047	0.137

(a) Average difference between Bullet and LARP trajectories (in meters) for two colliding 1-link and 2-link capsule chains.

	$N_h=10$	$N_h=38$	$N_h=60$
Human	0.010	0.031	0.048
Ball	0.006	0.033	0.061
Human	0.010	0.031	0.050
Capsule	0.010	0.042	0.069

(b) Average difference between Bullet and LARP trajectories (in meters) for a human kicking a ball/capsule for sequences of 10, 38 and 60 simulation steps.

Fig. 5: Evaluation of LARP on datasets with colliding objects.

found that supervising our neural simulator for positions instead of velocities is essential for good performance. Note that our simulator outputs velocities that are used to compute positions via integration. Replacing the position loss (eq. 2) with a corresponding velocity loss degrades position accuracy from 11 cm. to 95 cm. Our explanation is that the network must occasionally output small deviations from target velocities to correct the drift in the body joints and using the velocity loss hinders such desirable behavior. In fig. 4 (c) we verify that, as previously reported in [12] and [36], using soft activation function such as ELU or GeLU is favorable compared to more standard ReLU in the context of neural dynamics simulation. Finally, in fig. 4 (c) we show that learning rate schedule helps to improve accuracy.

Analysis of the contact model. We analyze the variants of the contact network using a diagnostic environment with two chains composed of two connected capsules shown in fig. 2 (b). As in experiments with the dynamics network we report position, rotation, and joint displacement error for each variant of the contact model. We consider the following variants: (1) our full “contact feature” model shown in Alg. 1 where the output of the contact network is fed as a feature to the dynamic network, (2) an alternative “contact impulse” variant where the output of the contact network is used directly in the integrator as an additive factor to linear and rotational velocities \mathbf{v} and ω . For the primary variant (1) we evaluate a version with and without stop-gradient. Results in fig. 4 (d) indicate that all variants are able to meaningfully handle the contact, and that the “contact impulse” formulation performs somewhat worse than “contact feature” variant. Note that adding stop-gradient is important for good performance and that without stop-gradient the “contact feature” model exhibits high variance across retraining runs. We show results for the best variant on the held-out test dataset in tab. 5a.

4.2 Simulation of articulated human motion

The goal of this paper is to propose a component for modeling articulated motion dynamics suitable for variety of tasks (*e.g.* prior for physics-based reconstruction or motion synthesis). We demonstrate one such application by using LARP as a replacement for Bullet in a physics-based human motion reconstruction pipeline [13], which reconstructs human motion via trajectory optimization. Given an input video, [13] infers a sequence of control parameters that results in physically

Dataset	Model	MPJPE-G	MPJPE	MPJPE-PA	MPJPE-2d	Velocity	Foot skate
H3.6M	DiffPhy [18]	139	82	56	13	-	7.4
	[13] + Bullet	143	84	56	13	0.24	4
	[13] + LARP	143	85	56	13	0.25	5.4
AIST-easy	DiffPhy [18]	150	106	66	12	-	19.6
	[13] + Bullet	154	113	69	13	0.41	4
	[13] + LARP	150	113	70	13	0.37	5
AIST-hard	[13] + Bullet	654	437	83	16	0.17	8.5
	[13] + LARP	643	442	73	14	0.13	7.2

Table 1: Quantitative results obtained with LARP on Human3.6M [22] and AIST datasets [38] and comparison to related work.

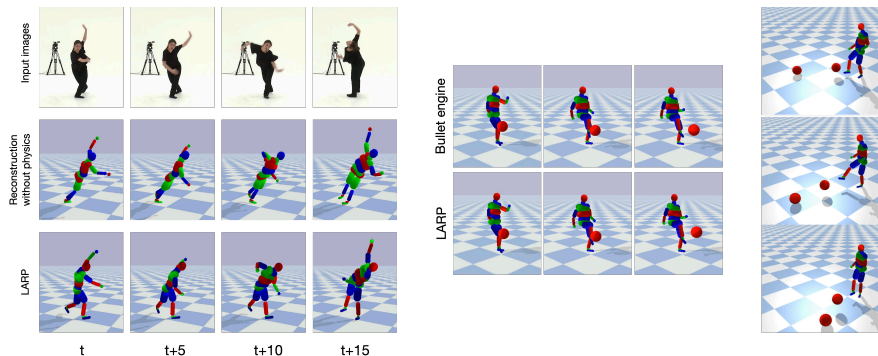


Fig. 6: Left: Reconstructed 3d poses on four consecutive video frames on AIST-hard dataset. Middle row shows results obtained with the kinematic pipeline from [13, 18] that LARP uses for initialization. Bottom row show results obtained with LARP integrated into [13]. Middle: Motion sequence with person-ball collision simulated with LARP (bottom) and comparison to Bullet engine [8] (top). Right: Examples of generated human motion sequences of a person kicking a ball for three different ball targets. In each image we show position of the ball right after the kick and at the end of the sequence. Note that the person pose differs considerably depending on the ball target.

simulated human motion that agrees well with observations (*e.g.* locations of 2d body joints in the frames of the video input). We refer to [13] for the details and focus on our results below.

Datasets. Following [13] we evaluate on Human3.6M [22] and AIST [38]. Human3.6M is a large dataset of videos of people performing common everyday activities with ground-truth 3D poses acquired using marker-based motion capture. We follow the protocol proposed in [34] that excludes “Seating” and “Eating” activities. This leaves a test set of 20 videos with 19,690 frames in total. AIST include diverse videos of dancing people with motions that are arguably more complex and dynamic compared to Human3.6M. The evaluation on AIST uses pseudo-ground truth 3d joint positions computed by triangulation from multiple camera views as defined by [25]. We use the subset of 15 AIST videos as used for evaluation in [13, 18]. In addition, we evaluate on 15 AIST videos with




Fig. 7: Left and middle: average mean per-joint position error and joint displacement error computed over all test sequences of the Human3.6M dataset. Right: displacement loss for the “S9-WalkDog” sequence. The units for y-axis are meters.



Fig. 8: Example results obtained with LARP on real-world videos.

more challenging motions that involve fast turns and rotations. We refer to these subsets of AIST as “AIST-easy” and “AIST-hard” in tab. 1.

Model training. We employ the physics-based human model introduced in [13] which is based on the GHUM model of human shape and pose [39]. The model is composed of 26 rigid components represented as capsules. We generate training data for LARP by running the sampling-based optimization from [13] on the training set and recording a subset of human motion samples and corresponding joint torques. This produces a significantly more diverse set of examples compared to original motion capture sequences, including examples of people falling on the ground and various self-collisions. We use the best settings for LARP as identified in the experiments in §4.1.

Metrics. We use standard mean per joint position (MPJPE) metrics for evaluation. The MPJPE-G metric measures mean joint error in the world coordinates after aligning estimated and ground-truth 3d pose sequences with respect to pelvis position in the first frame. MPJPE does pelvis alignment independently per-frame, whereas MPJPE-PA relies on Procrustes to align both position and orientation for each frame. Finally MPJPE-2D measures 2d joint localization accuracy. In addition we use two metrics to measure physical motion plausibility. “Velocity” compares the velocity error between the estimated motion and the ground-truth and is high when estimated motion is “jittery”. “Footskate” is implemented as in [18] and corresponds to the percentage of frames where the foot moves between adjacent frames by more than 2 cm, while in contact with the ground. “Footskate” measures the presence of a common artifact in video-based pose estimation where the positions of a person’s feet unreasonably shift between nearby frames.

Results. Our results are presented in tab. 1 and indicate that reconstructions using LARP are similar or better compared to Bullet. For example, LARP achieves

150 mm. MPJPE-G on AIST-easy and 643 mm. on AIST-hard, compared to 154 mm. and 654 mm. for Bullet. Note that [13] performs in general slightly worse compared to DiffPhy [18] that makes use of a differentiable physics simulator and gradient-based optimization. In principle our approach should work in combination with [18], and we plan to explore this in the future work.

In fig. 7 (left, center) we present plots of the position error (MPJPE) and joint displacement error averaged over all test sequences of the Human3.6m dataset³ and compare to the SuperTrack approach of [12]⁴. We show joint displacement error for one of the longer sequences “S9-WalkDog” in fig. 7 (right). We observe that for LARP the simulation indeed accumulates inconsistencies over time, albeit rather slowly, joint displacement error grows to about 2 cm. after 1500 sim. steps. This appears to have negligible effect on pose estimation accuracy: note that MPJPE does not increase with the simulation steps. Our observation is that instability (e.g. inconsistent simulation state with high joint displacement error) is not necessarily a function of the number of simulation steps, but rather happens during complex motions that are underrepresented in the training data. For example in fig. 7 (right) the displacement error changes little for most of the 25 sec. long sequence, but jumps up during the abrupt transition from walking into rushing forward (see fig. 9). We believe that training on larger and more diverse dataset might mitigate this.




Fig. 9: Example of estimated pose from “S9-WalkDog” seq. after 11 sec. of input. Left: input frame, middle: result obtained with SuperTrack trained on longer sequences, right: result obtained with LARP.

We show examples of the human pose reconstructions on AIST-easy in fig. 1 (left plot, middle column) and AIST-hard in fig. 6 (left). In fig. 6 (left) we compare output of LARP (bottom row) to 3d pose reconstruction that does not employ physics-based constraints and is used by LARP for initialization (middle row). Note that inference with LARP was able to correct physically implausible leaning of the person and estimate correct pose of the lower body. Finally, in fig. 8 and fig. 1 (left, third column) we include a few examples of 3d pose reconstructions obtained with LARP on real-world videos. These results confirm that LARP generalizes beyond AIST and Human3.6M datasets, is able to handle motions different from the training set (e.g. karate kick in fig. 8) and can handle complex contact with the ground (e.g. safety roll in fig. 1).

Simulation speed. We measure the speed of our approach using humanoid environment in Fig. 2 (c) and compare it to Bullet [9], MuJoCo [37], and Brax [11]. Bullet and MuJoCo were run on an AMD 48-core CPU for best performance. Brax, MuJoCo XLA (MJX) [29], and LARP are implemented to run on SIMD ar-

³ When computing the average we truncate the output to the shortest sequence.

⁴ We use a reimplement of SuperTrack in this comparison since the original paper didn’t make the code or pre-trained models public.

chitecture and are thus run on a Tesla V100 Nvidia GPU for best performance⁵. For each simulator we measure the time required to perform a simulation step for different number of simulations run in parallel. We run open-loop rollouts of 100 steps averaged over 5 runs for each number of parallel simulation. The compilation time is excluded. The results are shown in fig. 1. Overall, we observe that LARP exhibits better performance already when performing a single simulation (0.07ms for LARP vs. 0.19ms for MuJoCo and 1.8 ms. for Bullet). When running multiple simulations in parallel we improve over other simulators by an order of magnitude, *e.g.* 1.3ms for 4,096 parallel simulations compared to about 20 ms for MJX and MuJoCo.

Human-object collision handling. We include two types of experiments to evaluate accuracy of human object collisions. In the first experiment a ball or a capsule is added in a random position and orientation in front of the humanoid. We then run simulation in Bullet engine [8] and in LARP and measure the difference between trajectories of the ball and humanoid. We show an example of simulation results in fig. 6 (middle) and present quantitative evaluation in tab. 5b. Overall we observe that LARP fairly accurately approximates the output of a Bullet simulation. For example, the difference in the ball position after 60 simulation steps is about 6 cm. on average. In the second experiment we demonstrate that we can use LARP to control a humanoid in order to kick the ball to a particular target. The control parameters are inferred via model predictive control by minimizing residuals of: (1) the distance between the ball trajectory and the target, and (2) a term that keeps the resulting articulated human motion close to the reference kicking motion. We obtain an average error between the ball and the target of 5.7 cm. The qualitative results are shown in fig. 6 (right).

5 Conclusion

We have presented a methodology (LARP) to train neural networks that can simulate the complex physical motion of an articulated human body. LARP supports features found in classical rigid body dynamics simulators, such as joint motors, variable dimensions of the body component volumes, and contact between body parts or objects. Our experiments demonstrate that neural physics predictions produce results comparable to traditional simulation, while being considerably simpler architecturally, comparable in accuracy, and much more efficient computationally. Our neural modeling replaces the complex, computationally expensive operations in traditional physics simulators with efficient forward state propagation in recurrent neural networks. We discuss and illustrate the capability of LARP in challenging scenarios involving reconstruction of human motion from video and collisions of articulated bodies, with promising results.

⁵ Brax PBD is an implementation of position based dynamics [28] which is unstable for the given mass/inertia in the humanoid model, but is added for speed comparison.

References

1. Allen, K.R., Guevara, T.L., Rubanova, Y., Stachenfeld, K., Sanchez-Gonzalez, A., Battaglia, P., Pfaff, T.: Graph network simulators can learn discontinuous, rigid contact dynamics. In: Liu, K., Kulic, D., Ichnowski, J. (eds.) *Proceedings of The 6th Conference on Robot Learning. Proceedings of Machine Learning Research*, vol. 205, pp. 1157–1167. PMLR (14–18 Dec 2023), <https://proceedings.mlr.press/v205/allen23a.html> **1, 2**
2. Allen, K.R., Guevara, T.L., Rubanova, Y., Stachenfeld, K., Sanchez-Gonzalez, A., Battaglia, P., Pfaff, T.: Graph network simulators can learn discontinuous, rigid contact dynamics. In: Liu, K., Kulic, D., Ichnowski, J. (eds.) *Proceedings of The 6th Conference on Robot Learning. Proceedings of Machine Learning Research*, vol. 205, pp. 1157–1167. PMLR (14–18 Dec 2023) **3**
3. Allen, K.R., Guevara, T.L., Rubanova, Y., Stachenfeld, K., Sanchez-Gonzalez, A., Battaglia, P., Pfaff, T.: Graph network simulators can learn discontinuous, rigid contact dynamics. In: *Conference on Robot Learning*. pp. 1157–1167. PMLR (2023) **6**
4. Allen, K.R., Rubanova, Y., Lopez-Guevara, T., Whitney, W.F., Sanchez-Gonzalez, A., Battaglia, P., Pfaff, T.: Learning rigid dynamics with face interaction graph networks. *International Conference on Learning Representations (ICLR)* (2023) **3**
5. Bradbury, J., Frostig, R., Hawkins, P., Johnson, M.J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., Zhang, Q.: JAX: composable transformations of Python+NumPy programs (2018), <http://github.com/google/jax> **6**
6. Cai, X., Coevoet, E., Jacobson, A., Kry, P.: Active learning neural c-space signed distance fields for reduced deformable self-collision. In: *Graphics Interface 2022* (2022) **3**
7. Clevert, D.A., Unterthiner, T., Hochreiter, S.: Fast and accurate deep network learning by exponential linear units (elus). In: Bengio, Y., LeCun, Y. (eds.) *ICLR* (2016) **5**
8. Coumans, E.: Bullet physics simulation. In: *ACM SIGGRAPH 2015 Courses*, p. 1 (2015) **1, 2, 8, 11, 14**
9. Coumans, E., Bai, Y.: Pybullet, a python module for physics simulation for games, robotics and machine learning (2016) **4, 5, 13**
10. Featherstone, R.: *Rigid Body Dynamics Algorithms*. Springer-Verlag, Berlin, Heidelberg (2007) **1, 2**
11. Freeman, C.D., Frey, E., Raichuk, A., Girgin, S., Mordatch, I., Bachem, O.: Brax—a differentiable physics engine for large scale rigid body simulation. *arXiv preprint arXiv:2106.13281* (2021) **2, 5, 6, 13**
12. Fussell, L., Bergamin, K., Holden, D.: Supertrack: Motion tracking for physically simulated characters using supervised learning. *ACM Transactions on Graphics (TOG)* **40**(6), 1–13 (2021) **1, 4, 9, 10, 13**
13. Gärtner, E., Andriluka, M., Xu, H., Sminchisescu, C.: Trajectory optimization for physics-based reconstruction of 3d human pose from monocular video. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13106–13115 (2022) **1, 2, 8, 10, 11, 12, 13**
14. Ghorbani, S., Wloka, C., Etemad, A., Brubaker, M.A., Troje, N.F.: Probabilistic character motion synthesis using a hierarchical deep latent variable model. *Computer Graphics Forum* **39** (2020) **4**

15. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2016), <http://www.deeplearningbook.org> 9
16. Grzeszczuk, R., Terzopoulos, D., Hinton, G.: Neuroanimator: Fast neural network emulation and control of physics-based models. In: Proceedings of the 25th annual conference on Computer graphics and interactive techniques. pp. 9–20 (1998) 3
17. Guo, M., Jiang, Y., Spielberg, A.E., Wu, J., Liu, K.: Benchmarking rigid body contact models. In: Matni, N., Morari, M., Pappas, G.J. (eds.) Proceedings of The 5th Annual Learning for Dynamics and Control Conference. Proceedings of Machine Learning Research, vol. 211, pp. 1480–1492. PMLR (15–16 Jun 2023) 3
18. Gärtner, E., Andriluka, M., Coumans, E., Sminchisescu, C.: Differentiable dynamics for articulated 3d human motion reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022) 11, 12, 13
19. Habibie, I., Holden, D., Schwarz, J., Yearsley, J., Komura, T.: A recurrent variational autoencoder for human motion synthesis. In: Proceedings of the 28th British Machine Vision Conference (2017) 4
20. Henter, G.E., Alexanderson, S., Beskow, J.: Moglow: Probabilistic and controllable motion synthesis using normalising flows. ACM Transactions on Graphics (TOG) 39(6), 1–14 (2020) 4
21. Howell, T., Le Cleac’h, S., Bruedigam, J., Kolter, Z., Schwager, M., Manchester, Z.: Dojo: A differentiable simulator for robotics. arXiv preprint arXiv:2203.00806 (2022), <https://arxiv.org/abs/2203.00806> 5
22. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: 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) 2, 11
23. Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2019) 2
24. Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., Merose, A., Hoyer, S., Holland, G., Vinyals, O., Stott, J., Pritzel, A., Mohamed, S., Battaglia, P.: Learning skillful medium-range global weather forecasting. Science 0(0), eadi2336 (2023). <https://doi.org/10.1126/science.adi2336>, <https://www.science.org/doi/abs/10.1126/science.adi2336> 3
25. Li, R., Yang, S., Ross, D.A., Kanazawa, A.: Learn to dance with aist++: Music conditioned 3d dance generation (2021) 11
26. Li, Y., Wu, J., Tedrake, R., Tenenbaum, J.B., Torralba, A.: Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. In: International Conference on Learning Representations (ICLR) (2019) 3
27. Ling, H.Y., Zinno, F., Cheng, G., van de Panne, M.: Character controllers using motion vaes. ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH) 39 (2020) 4
28. Macklin, M., Müller, M., Chentanez, N.: Xpbd: position-based simulation of compliant constrained dynamics. In: Proceedings of the 9th International Conference on Motion in Games. pp. 49–54 (2016) 14
29. MuJoCo team: Mujoco xla (2023), <https://mujoco.readthedocs.io/en/stable/mjx.html> 2, 13
30. Mrowca, D., Zhuang, C., Wang, E., Haber, N., Fei-Fei, L.F., Tenenbaum, J., Yamins, D.L.: Flexible neural representation for physics prediction. Advances in neural information processing systems 31 (2018) 1, 2

31. Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., Battaglia, P.: Learning mesh-based simulation with graph networks. In: International Conference on Learning Representations (ICLR) (2021), outstanding Paper [3](#)
32. Rempe, D., Birdal, T., Hertzmann, A., Yang, J., Sridhar, S., Guibas, L.J.: Humor: 3d human motion model for robust pose estimation. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021) [4](#)
33. Sanchez-Gonzalez, A., Heess, N., Springenberg, J.T., Merel, J., Riedmiller, M., Hadsell, R., Battaglia, P.: Graph networks as learnable physics engines for inference and control. In: Proceedings of the 35th International Conference on Machine Learning (ICML) (2018) [3](#)
34. Shimada, S., Golyanik, V., Xu, W., Theobalt, C.: Physcap: Physically plausible monocular 3d motion capture in real time. *ACM Transactions on Graphics* **39**(6) (dec 2020) [11](#)
35. Song, J., Chen, X., Hilliges, O.: Human body model fitting by learned gradient descent. In: European Conference on Computer Vision. pp. 744–760. Springer (2020) [2](#)
36. Sukhija, B., Kohler, N., Zamora, M., Zimmermann, S., Curi, S., Krause, A., Coros, S.: Gradient-based trajectory optimization with learned dynamics. arXiv preprint arXiv:2204.04558 (2022) [2](#), [3](#), [10](#)
37. Todorov, E., Erez, T., Tassa, Y.: Mujoco: A physics engine for model-based control. In: 2012 IEEE/RSJ international conference on intelligent robots and systems. pp. 5026–5033. IEEE (2012) [1](#), [2](#), [5](#), [13](#)
38. Tsuchida, S., Fukayama, S., Hamasaki, M., Goto, M.: Aist dance video database: Multi-genre, multi-dancer, and multi-camera database for dance information processing. In: Proceedings of the 20th International Society for Music Information Retrieval Conference, ISMIR 2019. pp. 501–510. Delft, Netherlands (Nov 2019) [2](#), [11](#)
39. Xu, H., Bazavan, E.G., Zanfır, A., Freeman, W.T., Sukthankar, R., Sminchisescu, C.: GHUM & GHUML: Generative 3d human shape and articulated pose models. pp. 6184–6193 (2020) [12](#)