

This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# Physically Plausible Animation of Human Upper Body from a Single Image

Ziyuan Huang\* National Taiwan University Zhengping Zhou\* Stanford University

Yung-Yu Chuang National Taiwan University

Jiajun Wu Stanford University C. Karen Liu Stanford University

#### Abstract

We present a new method for generating controllable, dynamically responsive, and photorealistic human animations. Given an image of a person, our system allows the user to generate Physically plausible Upper Body Animation (PUBA) using interaction in the image space, such as dragging their hand to various locations. We formulate a reinforcement learning problem to train a dynamic model that predicts the person's next 2D state (i.e., keypoints on the image) conditioned on a 3D action (i.e., joint torque), and a policy that outputs optimal actions to control the person to achieve desired goals. The dynamic model leverages the expressiveness of 3D simulation and the visual realism of 2D videos. PUBA generates 2D keypoint sequences that achieve task goals while being responsive to forceful perturbation. The sequences of keypoints are then translated by a poseto-image generator to produce the final photorealistic video.

## 1. Introduction

Physics-based 3D character animation provides powerful tools to create controllable and interactable human agents using physical forces. These techniques can produce entirely novel movements grounded by laws of physics without any training data, but the visual appearance of the agent is still not on par with photorealistic images. In contrast, datadriven image-based motion synthesis methods are effective in synthesizing photorealistic videos, but they are not capable of creating novel movements based on unpremeditated physical interaction with the agent.

Can we combine the expressive power that comes with physics-based 3D animation with the rich appearance offered by 2D videos, such that the agent is physically controllable, interactable, and photorealistic? We propose a new approach that splits the difference between the 2D and 3D representations—our dynamic model operates on a 2D state and a 3D action. We represent states as 2D keypoints to avoid conversion between 3D and 2D representations, because our final products are 2D images. On the other hand, we choose to represent actions as 3D joint torques, because they contain more information to respect human kinematic and dynamic constraints. Using this hybrid representation, we train policies to control the person's upper body in the photo to achieve desired tasks. Combining the policy with the dynamic model, our system can animate the person by dragging her hands to arbitrary 2D locations or pushing the person with an arbitrary force. The person in the photo will follow the user's commands or react to physical forces in a human-like and physically plausible way (Figure 1).

Our system, named Physically plausible Upper Body Animation (PUBA), consists of two modules: motion synthesis and photorealistic rendering. We cast motion synthesis into a reinforcement learning problem with the hybrid representation of 2D states and 3D actions. Specifically, we define a 2D state as a set of keypoints on the person in the image and a 3D action as a torque vector applied to the person's joints. We first train a transition function that predicts the next 2D keypoints from the current 2D keypoints under the effect of 3D torque vectors. Subsequently, we train a policy that takes as input 2D keypoints and outputs a 3D torque vector to maximize the long-term reward that defines the task. Once the policy and the transition function are trained, we can generate sequences of 2D keypoints that achieve the task while being responsive to forceful perturbation applied by the user. The keypoint sequences are then translated into photorealistic images via a pose-to-image generator [5].

PUBA has two key advantages: First, the motion synthesis module and the photorealistic rendering module are decoupled and can be trained separately, allowing it to learn from unpaired 3D mocap data and 2D images; Second, the policy and dynamics models operate fully on 2D keypoints, bypassing the need of monocular 3D pose reconstruction. Our system design avoids the failure case when the predicted 3D pose cannot overlay with the input image and removes the dependency on a complex physics-based animation pipeline.

## 2. Related Works

**Image animation.** A body of research aims to create an animation for a class of or a specific object in a source image

<sup>\*</sup> indicates equal contribution.



Figure 1: Given an image with a person (first column), our system allows the user to animate the person using intuitive 2D interaction in the image space. Top: The user specifies two target locations for the wrists (shown as red dots). Bottom: The user specifies desired trajectories for hands to track (shown as red curves). At t = 20, the user creates a force vector that pushes the person to the left (shown as the red arrow).

by following the movements of a driving video, based on cycle consistency [2], estimation of optical flow [28], or affine transformation [29] of unsupervisedly detected 2D keypoints. A special line of work focuses on human body pose transfer, where the central goal is to synthesize a person in novel poses. Representative works include 2D methods that use direct mapping from pose to image [5], estimated body part spatial transformations [1, 25], attention maps [36], modal bases [7], and disentanglement of color and pose [8], as well as 3D methods that take advantages of dense pose [22], reconstruction of parametric body models [17], or personspecific textured 3D character models [16].

A key limitation of those methods is that complete reference motion must be provided. In contrast, our method allows the user to generate animation by specifying a concise goal, for instance "track this mouse trajectory" or "reach this target location with your wrists". The synthesized motion can be further edited by physical forces, thus generating physically plausible behaviors such as the person "struggles to prevent falling when suddenly pushed to the side".

**Motion prediction.** Opposite to works that synthesize motions based on a reference motion, another line of research focuses on predicting the motion solely based on past or future observations, where the predicted motion can either be represented as RGB videos [21], or as 3D body mesh sequences [33, 13] based on the prior art of parametric body surface models [19] and monocular 3D body pose and shape reconstruction [3, 12, 14, 27, 32, 31]. For future prediction, an autoregressive model is typically employed that takes the output from the last step as input, such as Struct V-RNN [21] or a casual model [33]. The autoregressive model has also proved to be effective for synthesizing long-term 3D motion from scratch, such as the acLSTM model [15].

Since the synthesized motion is based on past predictions or generated purely from scratch, they are barely controllable and are only applicable to highly specific and predictable motions, such as sports or martial arts. However, we borrow the idea of autoregressive prediction for our dynamics model, which is controlled by a per-step torque input.

**Physical simulation-based character animation.** On the spectrum of "motion controllability", physically simulated character animation stands on "the most controllable" extreme, while motion prediction stands on the other. For synthesizing motions that react realistically to perturbations, controllers are widely used for moderating the character's movements in the 3D environment. To bring a similar level of controllability and physical realism into the 2D pixel space, we borrow the idea from the deep reinforcement learning (DRL) based animation systems [23, 24] that learn policies for user-specified goals. The key difference is that our policy takes 2D keypoint-based observations, while they operate fully in 3D and require joint angle-based observations, which are usually difficult to infer from a single view.

# 3. Method

Given a person in a photo  $I^0$ , our PUBA system enables the user to animate the person using intuitive 2D interaction, such as dragging or pushing at different parts of the person, and create a photorealistic video  $I^{1:T}$ . We view the person in the photo as a reinforcement learning agent operating in a hybrid representation of 2D state space and 3D action space. The task of the agent is to achieve the target locations specified by the user and respond dynamically to perturbation applied by the user.

Our system consists of three components (Figure 2):

- a dynamic model  $s^{t+1} = \phi(s^0, a^{0:t})$ ,
- a goal-conditioned policy  $\boldsymbol{a}^t \sim \pi(\cdot | \boldsymbol{s}^t, \boldsymbol{g})$ , and
- a pose-to-image generator  $I^t = \psi(\boldsymbol{x}^t)$ .

A 2D pose,  $x \in \mathbb{R}^{2n}$  is defined by n landmark points on the person's body, such as shoulders and elbows. A state



Figure 2: Overview of our system.

 $s \in \mathbb{R}^{4n}$  is a concatenation of a 2D pose and its offset from the 2D pose at the previous time step,  $s^t = (x^t, \triangle x^t)$ . An action is defined as a 3D torque vector,  $a \in \mathbb{R}^m$ , where mis the number of actuated degrees of freedom (DOFs) of the proxy 3D agent. Finally, our policy is also conditioned on the 2D target locations, g, for the specified body parts.

#### **3.1. Dynamic Model**

The dynamic model is trained to simulate a sequence of states  $s^{1:t+1}$ , given the initial state  $s^0$  and a control sequence  $a^{0:t}$  (Figure 3). Specifically, at time t, the network predicts a sequence of  $\triangle x^{1:t+1}$  and integrates them to get the 2D pose at t + 1 as  $x^{t+1} = x^0 + \sum_{t'=1}^{t+1} \triangle x^{t'}$ , which is then concatenated with  $\triangle x^{t+1}$  to compose the next state  $s^{t+1} = (x^{t+1}, \triangle x^{t+1})$ .

**Training data.** Training the dynamic model  $\phi$  requires a large set of training data that covers the joint space of state and action. We use motion capture dataset  $\mathcal{D}$  from Human 3.6M [11] and augment it with a set of synthetic human motions  $\mathcal{D}'$ . We build a 3D human agent represented as an articulated rigid body system to process the mocap data (Figure 4 Left). We first apply inverse kinematics to obtain joint configurations  $\{q^{1:T}\}$  from the raw mocap data, where  $q \in \mathbb{R}^m$  contains the actuated degrees of freedom of the human agent. To derive the joint torques  $\{a^{0:T}\}$ , we compute inverse dynamics from trajectories of q, as well as the trajectories of  $\dot{q}$  and  $\ddot{q}$  approximated by finite differencing. To obtain the 2D keypoints  $\{\bar{x}^{0:T}\}$ , we simply compute the 3D landmark points via forward kinematics and project them to the 2D image space.

We use two different approaches to create augmented human motion dataset  $\mathcal{D}'$  (Figure 4). First, since the policy training process usually involves adding exploration noise to actions for improving the robustness of the policy, we augment the training data by perturbing the per-step torque  $\{a^{0:T}\}$ , computed by inverse dynamics, with noise sampled from a zero-centered Gaussian distribution, and then forward simulate to get a new state. The standard deviation of the Gaussian matches half of the standard deviation of



Figure 3: The dynamics model. Starting from initial 2D keypoint coordinates  $x^0$  and offsets  $\Delta x^0$ , the GRU model takes one torque  $a^t$  at each step t as input, and predicts the per-step 2D keypoint offsets  $\Delta x^t$ .  $FC_h$ ,  $FC_a$ ,  $FC_{dx}$  are fully connected layers with tanh activation.



Figure 4: Left: The location and DOFs for each joint on our 3D human agent. Right: The training data for the dynamic model. A sample motion sequence from the mocap dataset is shown in black boxes. The augmented poses are shown in red boxes.

the joint torques computed from the mocap dataset. Second, we randomly generate 3D points within the human range of motion as the target points for the specified body parts (e.g., the wrist), and use inverse kinematics to generate additional reaching motion for training. Using these two methods for data augmentation, we can generate as much training data as we need. We augment 330,000 frames into the original dataset, which contains 141,021 frames. The dynamic model trained with such an augmented dataset performs more reliably when being queried by the policy during training, as shown quantitatively in various scenarios in Section 4.

**Model architecture.** We use a gated recurrent unit (GRU) [6] network to model the temporal dependency in human motion. The hidden state, which implicitly models the evolving kinematic state, is initialized through a fully connected layer that takes  $s^0 = (x^0, \Delta x^0)$  as input, and is updated at each step according to the input torque  $a^t$ . We denote the matrix multiplication operator as  $\circ$ , the vector concatenation operator as  $[\cdot]$ , and the GRU cell update rule as GRU( $\cdot$ ):

$$\boldsymbol{h}^{t} = \begin{cases} \tanh(W_{h} \circ [\boldsymbol{x}^{0}, \bigtriangleup \boldsymbol{x}^{0}] + \boldsymbol{b}_{h}), & t = 0; \\ \operatorname{GRU}(\boldsymbol{h}^{t-1}, \tanh(W_{a} \circ \boldsymbol{a}^{t} + \boldsymbol{b}_{a})), & t > 0. \end{cases}$$
(1)

The per-step keypoint offsets output  $\Delta x^t$  (normalized into [-1,1]) is then predicted based on the hidden state  $h^t$  as

 $\Delta \boldsymbol{x}^t = \tanh(W_{\Delta \boldsymbol{x}} \circ \boldsymbol{h}^t + \boldsymbol{b}_{\Delta \boldsymbol{x}}).$ 

**Learning.** We sample all possible length-*T* trajectories from the training dataset. Each trajectory  $\tau$  is annotated with  $\{\bar{x}^{0:T}\}$  and  $\{a^{1:T}\}$ . The loss function is defined as the sum of the mean squared error (MSE) of both the 2D keypoint coordinates and the 2D keypoint offsets:

$$\ell = \sum_{\tau \in \mathcal{D}} \sum_{t=1}^{T} (\|\boldsymbol{x}^t - \bar{\boldsymbol{x}}^t\|^2 + \|\Delta \boldsymbol{x}^t - \Delta \bar{\boldsymbol{x}}^t\|^2).$$
(2)

#### **3.2. Goal-Conditioned Policy**

We can now learn a policy that achieves the given task under the dynamics enforced by the learned dynamic model  $\phi$ . We formulate a Markov Decision Process and solve for a goal-conditioned policy,  $\pi(a^t|s^t, g)$ . In addition to the current state  $s^t$ , the policy also takes as input a goal vector g defined by the task. For example, the reaching task would define the goal vector as  $g^t = [x_R^t - \bar{x}_R, x_L^t - \bar{x}_L]$ , where  $\bar{x}_R$  and  $\bar{x}_L$  are the 2D target positions for the right wrist and the left wrist to reach, respectively.

**Reward function.** The goal of reinforcement learning is to solve for a policy that maximizes the long-term reward. Our reward function consists of the following four terms:

$$r = w_{\text{task}} r_{\text{task}} + w_{\text{upright}} r_{\text{upright}} + w_{\text{ctrl}} r_{\text{ctrl}} + w_{\text{alive}} r_{\text{alive}}, \quad (3)$$

where w are weights. Among the four terms,  $r_{task}$  encourages the policy to minimize the deviation from the goal of the task. For example, a reaching task would be defined as

$$r_{\text{goal}}^{t} = -\left(\|\boldsymbol{x}_{R}^{t+1} - \boldsymbol{g}_{R}^{t+1}\| + \|\boldsymbol{x}_{L}^{t+1} - \boldsymbol{g}_{L}^{t+1}\|\right), \quad (4)$$

 $r_{upright}$  encourages the agent to maintain an upright torso by penalizing the positions of the neck and the head being lower than those in the default pose:

$$r_{\text{upright}}^{t} = -(\max(0, \boldsymbol{x}_{\text{neck-y}}^{t+1} - h_{\text{neck}}) + \max(0, \boldsymbol{x}_{\text{head-y}}^{t+1} - h_{\text{head}})).$$
(5)

where  $h_{\text{neck}}$  and  $h_{\text{head}}$  are the height of the neck and of the head in the default pose.  $r_{\text{ctrl}}$  encourages minimal action which leads to smoother motion:  $r_{\text{ctrl}}^{(t)} = ||\boldsymbol{a}^t||$ . Finally,  $r_{\text{alive}}$  penalizes the policy from stepping out of the image boundary or falling down:

$$r_{\text{alive}} = \begin{cases} x^{(t)} \text{ in image boundary and the} \\ 1, & \text{height of neck and head are above} \\ a fixed threshold; \\ 0, & \text{otherwise.} \end{cases}$$
(6)

**Training policy.** We use Proximal Policy Optimization (PPO) [26] to train an MLP policy  $\pi_{\theta}$  with a fixed horizon T = 50. At the beginning of each episode, the initial 2D pose and the targets are independently randomly sampled from the training dataset we used to train the dynamic model. At each time step t, we clip the current 2D pose  $x^t$  into a valid range. This range is determined by the boundary of 2D poses in  $\mathcal{D}$ . The episode is terminated whenever the policy reaches an invalid state that makes  $r_{\text{alive}} = 0$ .

**Testing policy.** Once the policy is trained, we can use it to generate 2D pose sequences that reach or track arbitrary targets/target trajectories while responding to the forceful perturbation applied by the user. The user-specified perturbation is interpreted as a torque applied on specific joints in a time interval. The torque will then be added to the output of the policy model.

#### 3.3. Photorealistic Rendering

Starting from the input image  $I^0$ , we first extract the initial 2D pose  $x^0$  with the state-of-the-art 2D pose detector [4]. We then use the policy  $\pi$  and the dynamic model  $\phi$  to produce a sequence of 2D poses  $\{x^{1:T}\}$  for a given control task. A photorealistic video can then be synthesized frame-by-frame from the 2D poses:  $I^t = \psi(x^t)$ for t = 1 : T, where  $\psi$  is a generator network that maps the 2D pose to the corresponding image of the source person. For best photorealism, we adopt the method proposed in Everybody Dance Now [5] to construct  $\psi$ , and train it with an additional video where the source person performs random poses<sup>1</sup>. Our motion synthesis module is applicable to other pose-to-image generators, such as ADGAN [20]. Although Everybody Dance Now [5] produces more temporally coherent videos. ADGAN is capable of generating results with various human subjects from the DeepFashion Dataset [18]. More results can be found in Section 4. Note that the video training data can be completely independent of the 3D poses in the mocap training set. In this paper, we directly use mocap data from the public mocap database, Human 3.6M dataset [11], and record our own video data.

#### 4. Experiments

This section presents the experiment setup, compares our method to baselines via quantitative metrics and user studies, and provides results for ablation studies and applications.

#### 4.1. Experimental Setup

**Motion synthesis.** For training the dynamic model and the control policy, we employ a 3D mocap dataset from the Human 3.6M dataset [11], which is publicly available on requests. We currently focus on upper body motions; therefore, we only include action classes with minimal lower body movements: *Directions, Discussion, Greeting, Posing, TakingPhoto.* For all experiments in the paper, the motion is synthesized for T = 50 steps, at 50 FPS. After the motion is synthesized, we apply an average smoothing with window size 11 to improve the visual quality. We use a perspective camera with predefined intrinsics to obtain the 2D keypoint coordinates in the front view.

<sup>&</sup>lt;sup>1</sup>Note that it is possible to omit this additional video input, if we do not care too much about the quality of the generated video. In that case, a pose transfer method (e.g. [1]) can be used. This would require fewer inputs but also produce more artifacts in the generated video.

Methods	realism	precision
PUBA (ours) vs. Ground truth PUBA (ours) vs. Hao et al. [9] PUBA (ours) vs. 2D simulation	28% 88%	36% 79%
PUBA (ours) vs. 3D simulation	60%	61%

Methods	realism	precision	perturbation
PUBA (ours) vs. Hao et al. [9]	85%	74%	75%
PUBA (ours) vs. 3D simulation	65%	74%	64%

(b) User assessment on tracking task with perturbation

(a) User assessment on tracking and reaching task

Table 1: The results of the user study. The matrices include the percentage (%) of users that prefer our method over the methods indicated in the first column. The following columns are the user assessment of video realism, tracking precision, and the ability to recover from a perturbed force, respectively. (a) User assessment of tracking and reaching videos. (b) User assessment of tracking videos that have an external force added on the pelvis during the movements.

**Photorealistic rendering.** For training the generator of Everybody Dance Now [5], the target person is required to perform random upper body movements in front of a camera that is roughly oriented in the front view. In practice, we found that 6–8 minutes of recording is enough for a 30 FPS camera. At the beginning of the recording, the person is expected to start with a T-pose, so that a 4-DoF linear transformation could be computed to normalize the 2D poses in the photo view into the same pose space as the motion synthesis module. The videos are synthesized at resolution  $256 \times 256$ , 50 FPS. For comparisons, we use [5] as the generator, while using ADGAN [20] pre-trained on DeepFashion for some applications.

**Baseline.** We consider two baselines. First, *Hao et al.* [9] propose a video generation model that allows detailed control over the motion of the generated video. Given an image and sparse flow vectors specified by users, the model can generate a dense flow map that warps the input frame along flow vectors. We use the ground truth trajectories as the input flow vectors to their model, and compare their synthesized videos with ours. We also consider a 3D simulation baseline. Given an input image of a person and user-specified 2D target points, we first predict the 3D COCO joints of the person in the input image using HMR [12]. We then run inverse kinematics (IK) to have a 3D humanoid mimic the initial pose according to the predicted body joints. In the meantime, we map the user-specified 2D target points into the 3D space. We fix the depth of the 3D target points to the initial depth of human wrists. With the above settings, we use IK to have the 3D humanoid track the 3D target points. Finally, we reproject the computed 3D joints back to 2D and use the photorealistic render to synthesize a video.

**Evaluation metrics.** We conduct a human preference study on Amazon Mechanical Turk to compare the perceptual quality of our results and the baseline methods. Also, we compare the quality of the synthesized videos frame by frame using two metrics: PSNR and LPIPS [34]. The two metrics help estimate the similarity between the ground truth and the generated images. However, we find PSNR could deviate from human preference because of the diversity of plausible human motions.

# 4.2. Results

For quantitative comparisons, we consider three tasks:

- *Tracking*: Given an RGB image of a person and the desired trajectories of the left and right wrists, generate a video of the person tracking the trajectories;
- *Reaching*: Similar to tracking, but fixed target locations for wrists are provided instead of full trajectories;
- *Perturbation recovery*: In this scenario, the person is perturbed suddenly by an impulse torque (e.g., pushing the person to the left) while following the policy to achieve the goal. The expected output is a video of the person trying to recover from the perturbation and complete the original task. For Hao et al. [9], we put an extra flow vector on the person's pelvis to represent an impulse torque. For 3D simulation, we use inverse dynamics to calculate the internal force of the 3D humanoid and add an external force on the upper torso within a time interval.

For each pair of the RGB image and trajectories, we generate three videos using our method and the two baseline methods and compare them using the user study and evaluation metrics.

Human preference study. We conduct the user study using MTurk. We communicate with participants and obtain their consent to use their responses. For each example, the study presents to the participants a pair of videos, generated by our model and a baseline, superimposed with the target points/trajectories. Participants are asked to answer the following two questions: (1) Which video looks more realistic? (2) Which video follows the target points/trajectories better? There is no time limit for answering questions. For tracking and reaching, we take 40 videos of a person doing some upper body movement as ground truth, and extract the 2D wrist keypoint coordinates as the target trajectories input to both our and baseline methods. Furthermore, we also perform pairwise comparisons between the ground truth and our results. Table 1(a) shows the results in preference matrices among 100-200 answers for each question. In terms of realism, our PUBA system outperforms the baselines, and 28% of the participants even think our results are more realistic



Figure 5: The qualitative comparison with other methods on the perturbation task. Our method can recover from the perturbation better and keep tracking the trajectories more precisely.  ${}^{\ddagger}SfV$  [24] requires the whole ground truth video as the input for training. See Section 4.2 for more details.

than the ground truth. Note that a perfect method is expected to confuse people 50% of the time, so a 28% preference rate is encouraging.

For the perturbation recovery task, there is an additional question: In which video, does the person better recover from the sudden impulse torque? Note that there is no ground truth for the perturbation task. Table 1(b) reports the user preference of our method against the baselines. Our method outperforms the baselines in all aspects. It is not surprising because our method leverages knowledge about human dynamics while [9] attempts to model generic videos. Although 3D simulation utilizes the physics engine to generate reasonable results, the reconstructed initial pose contains inevitable errors. With an erroneous initial pose, the 3D simulation eventually generates a series of poses with more errors. Visually, we can observe misalignments of initial poses and wrists trajectories in Figure 5. Besides, [5] is susceptible to body misalignment, so the generated person by the 3D baseline is not as sharp as our method.

**Frame by frame evaluation.** In order to conduct a quantitative evaluation, we use a set of ten videos as the ground truth and evaluate the fidelity of the generated videos. Table 2 reports the average LPIPS [35], FID [10], and FVD [30] scores of all methods. LPIPS and FID generally capture the similarity between two images, and FVD takes the temporal coherence of the generative videos into consideration. Our method outperforms the baselines in all metrics, and this is in line with the results of the user study.

	$\text{LPIPS}\downarrow$	$FVD \downarrow$	$FID\downarrow$
Hao et al. [9]	0.134	2197.59	330.96
<b>3D Simulation</b>	0.164	1092.60	148.51
PUBA (ours)	0.121	947.23	130.07

Table 2: Evaluation of our method and the baselines.

Methods	w/o perturbation		w/ perturbation	
	$\mathcal{L}_1$	$\mathcal{L}_2$	$\mathcal{L}_1$	$\mathcal{L}_2$
SfV [24]	26.41	20.66	31.73	25.10
3D simulation	21.35	19.08	27.93	22.13
PUBA (ours)	18.79	15.05	23.47	18.91

Table 3: Comparisons with the physical simulation methods. We compare the accuracy of the keypoints generated by different methods for the scenarios without and with perturbation.



Figure 6: Evaluating the dynamic model against two ablations: (1) no data augmentation and (2) single-step prediction. *x*-axis: step, *y*-axis: per-step average keypoint  $\mathcal{L}_2$  error. *Best viewed in color.* 

**Other comparisons.** In addition to 3D simulation, we also compare our model with a more sophisticated 3D method, SfV [24]. Given a recorded human video, SfV enables a physically simulated 3D character to imitate the motions in the video. Note that we do not consider SfV as a baseline because SfV takes a whole ground truth video as input, and needs to train an RL model for each video to get the best performance, while our method only requires the first frame and wrist movements as input, and uses the same learned model for all videos. After inferring the motions, 3D simulation, SfV, and our method generate desired videos using the same photorealistic generator. Thus, we can compare the fidelity of the generated motions using the accuracy of the generated 2D keypoints. Table 3 reports the  $\mathcal{L}_1$  and  $\mathcal{L}_2$ distances between the 2D keypoints projected from Human 3.6M dataset [11] and the ones generated by different methods. Our method outperforms the compared 3D methods in both metrics. Although SfV can mimic the target motion well, the tracking precision is worse than our method.

We also add an external force on the humanoids and see how the 3D methods resist the perturbation. We conduct HMMR [13] to predict the 3D motion of the ground truth video, and use a PD controller to track the ground truth motion with forward simulation. During perturbation, we add a force along the humanoid's pelvis. The PD controller



Figure 7: Results of our approach tracking two different userdrawn trajectories. Leftmost: input images and the two input wrist trajectories (green); The rest columns show the trajectories of the synthesized wrist positions (red) overlaid on top of the synthesized frames. Top: video generated by Chan et al. [5]; Bottom: video generated by Men et al. [20]. *Best viewed in color*.



Figure 8: Results of our approach reaching two different target wrist positions specified by the user. In the input images, the target wrist positions are marked with green circles; We show the synthesized frames at t = 10 and t = 20, demonstrating how the subjects reach targets. Top: video generated by the Chan et al. [5]; Bottom: video generated by Men et al. [20]. *Best viewed in color*.

will force the agent back to the trajectory after perturbation. Table 3 shows the results between the ground truth keypoints and the ones generated by different methods with perturbation. Our method generates more accurate poses because 3D simulation cannot recover from the perturbation well, and SfV fails to track the trajectory precisely after perturbation. Figure 5 compares them visually.

#### 4.3. Ablation Study

We compare the dynamic model against two ablations. (1) *no data augmentation*: the dynamic model is only trained on forces presented in the mocap dataset, without data aug-



Figure 9: The person recovers from a push while tracking userspecified trajectories. The target trajectories are shown in (green), and the trajectories of the synthesized wrist positions are shown in (red). At t = 20, an impulse torque is applied to the waist, which causes the person to rotate. *Best viewed in color*.



Figure 10: Force control. Arms dropping passively in response to gravity from t = 10 to t = 20.

mentation; (2) *single-step prediction*: the RNN is replaced with an MLP that takes in the 2D pose and 2D offset from the last step and predicts the 2D offset of the next step.

We evaluate the model on three setups. (1) *synthetic*: The augmented data as described in Section 4.1. This tests how well the dynamic model generalizes when faced with highly perturbed torques, similar to the situation of policy training; (2) *mocap*: The original mocap data, testing the performance on torques that lead to regular movements; (3) *idle*: Starts from poses of the mocap data, but a zero torque is applied at each step. This tests the behavior when the person is only driven by gravity. The evaluation trajectories are drawn every 50 steps.

Figure 6 shows the results. Our dynamic model achieves the lowest prediction error on all three setups. It indicates that both expanding the coverage of force and including history information improve the robustness of the model.

# 4.4. Applications

**Interactive photo.** The user can draw trajectories or specify points on the image to be the targets for the wrists of the person. Although only the targets of the wrists are provided, our tracking policy is able to generate natural movements for the entire upper body. Figure 7 shows the result for the tracking task, and Figure 8 demonstrates the reaching task.

The user can also push the person in the image and see the person respond to the push and recover from the perturbation. Figure 9 shows examples of a person perturbed by a sudden leftward push around the waist area, while following the



Figure 11: Retargeting. Top: An input image of the source subject and the results of retargeting. Bottom: Input images that specify the pose for retargeting.



Figure 12: Results of the full-body tracking task. The leftmost image shows the input keypoints and four user-specified trajectories for wrists and ankles (green). Other images show the trajectories of the synthesized wrist and ankle positions (red) overlaid on top of the synthesized frames. *Best viewed in color*.

tracking policy. The impulse exerts from t = 20 to t = 30.

**Force control.** No goal is specified in this scenario, and thus the person is passively reacting to gravity and external forces. Figure 10 shows that the person gradually drops her arms due to gravity. The generated movements are consistent with our physics intuition. More results can be found in the supplemental materials.

**Retargeting.** We can also retarget the motion of one subject to other subjects with different appearances by aligning the source 2D pose sequence with the target generator's scale and pelvis position (Figure 11).

**Full body control.** We have tested our method on a small full-body dataset by fixing the pelvis and setting other joints free. The preliminary results are qualitatively similar to the upper-only model (Figure 12); however, further tuning of reward function weight is required to achieve similar quantitative results. Further evaluation and optimization will remain for future research.

## 5. Discussion

We present a novel system, PUBA, that allows the user to generate photorealistic and physically plausible human ani-

mation by controlling a person in the image using intuitive 2D interaction. Experiments demonstrate that our framework is highly flexible and produces more compelling animation compared to the baseline. Importantly, we found that 2D keypoints already provide rich information and can be directly controlled effectively by 3D torques, removing the need to solve a challenging 3D pose reconstruction problem. Our method does not require 3D mocap data to be aligned with 2D images, largely simplifying the data acquisition process and enabling the composition of movements and appearances of different people. We also show that physics-based character animation can be an effective tool to synthesize training data that boost the performance of learning.

Our work has limitations. The dynamic model is not accurate when the input state-action pair is very different from those in the training data. We observe some bone stretching artifacts when the arms move to positions where training data are sparse. The data augmentation indeed helps, but it is somewhat task-dependent. When training a different category of motion, such as dancing, the policy might still visit the areas in state-action space in which the dynamic model performs poorly, resulting in a suboptimal dancing policy. Training a policy that handles balance and physical interaction with the environment in the image space presents many interesting research challenges and applications.

In our research, we collected videos of two human subjects and used data from the DeepFashion dataset [18]. Our research has been deemed IRB exempt by our institution, because we only used the human subject data as annotations, without studying the human subjects themselves. DeepFashion was released in 2016, widely used and cited over 1,000 times; the IRB approval situation of DeepFashion, including whether an IRB is needed, is unclear to us. We choose to use the dataset, because it was used by prior methods for state-of-the-art image generation results, which is critical for a fair comparison in our experiments.

Our research has many potential positive societal impacts, with future applications in the fashion industry, choreography, computer games, and personal and family entertainment with personalized, physically plausible avatars. It may also find uses in assistive technologies in the rehabilitation of disabled persons, and in automatic sign language generation for hearing impaired people. On the other hand, like all other visual content generation methods, our method might be exploited by malicious users to generate disinformation. Thus, we urge users of our models to be aware of ethical and societal concerns and to apply them with good intent. We will use techniques such as watermarking to identify and label visual content generated by our system.

Acknowledgments. This work is in part supported by the Toyota Research Institute (TRI), the Stanford Institute for Human-Centered AI (HAI), Samsung, and Amazon.

# References

- Guha Balakrishnan, Amy Zhao, Adrian V Dalca, Fredo Durand, and John Guttag. Synthesizing images of humans in unseen poses. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8340–8348, 2018.
- [2] Aayush Bansal, Shugao Ma, Deva Ramanan, and Yaser Sheikh. Recycle-gan: Unsupervised video retargeting. In Proceedings of the European Conference on Computer Vision, pages 119–135, 2018.
- [3] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *Proceedings of the European Conference on Computer Vision*, pages 561–578. Springer, 2016.
- [4] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):172–186, 2019.
- [5] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5933–5942, 2019.
- [6] Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Workshop on Deep Learning, December 2014*, 2014.
- [7] Abe Davis, Justin G Chen, and Frédo Durand. Image-space modal bases for plausible manipulation of objects in video. *ACM Transactions on Graphics*, 34(6):1–7, 2015.
- [8] Patrick Esser, Ekaterina Sutter, and Björn Ommer. A variational u-net for conditional appearance and shape generation. In *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition, pages 8857–8866, 2018.
- [9] Zekun Hao, Xun Huang, and Serge Belongie. Controllable video generation with sparse trajectories. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition, 2018.
- [10] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [11] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2013.
- [12] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7122–7131, 2018.
- [13] Angjoo Kanazawa, Jason Y. Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019.

- [14] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2252–2261, 2019.
- [15] Zimo Li, Yi Zhou, Shuangjiu Xiao, Chong He, Zeng Huang, and Hao Li. Auto-conditioned recurrent networks for extended complex human motion synthesis. arXiv preprint arXiv:1707.05363, 2017.
- [16] Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Hyeongwoo Kim, Florian Bernard, Marc Habermann, Wenping Wang, and Christian Theobalt. Neural rendering and reenactment of human actor videos. ACM Transactions on Graphics, 38(5):1– 14, 2019.
- [17] Wen Liu, Zhixin Piao, Jie Min, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5904–5913, 2019.
- [18] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1096–1104, 2016.
- [19] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multiperson linear model. ACM transactions on graphics, 34(6):1– 16, 2015.
- [20] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable person image synthesis with attribute-decomposed gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [21] Matthias Minderer, Chen Sun, Ruben Villegas, Forrester Cole, Kevin P Murphy, and Honglak Lee. Unsupervised learning of object structure and dynamics from videos. In Advances in Neural Information Processing Systems, pages 92–102, 2019.
- [22] Natalia Neverova, Riza Alp Guler, and Iasonas Kokkinos. Dense pose transfer. In *Proceedings of the European Conference on Computer Vision*, pages 123–138, 2018.
- [23] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. ACM Transactions on Graphics, 37(4):1–14, 2018.
- [24] Xue Bin Peng, Angjoo Kanazawa, Jitendra Malik, Pieter Abbeel, and Sergey Levine. Sfv: Reinforcement learning of physical skills from videos. ACM Transactions On Graphics, 37(6):1–14, 2018.
- [25] Yurui Ren, Xiaoming Yu, Junming Chen, Thomas H Li, and Ge Li. Deep image spatial transformation for person image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7690–7699, 2020.
- [26] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017.
- [27] Soshi Shimada, Vladislav Golyanik, Weipeng Xu, and Christian Theobalt. Physically plausible monocular 3d

motion capture in real time. ACM Transactions on Graphics, 39(6):1–16, 2020.

- [28] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2377–2386, 2019.
- [29] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In Advances in Neural Information Processing Systems, volume 32, pages 7137–7147, 2019.
- [30] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. arXiv preprint arXiv:1812.01717, 2018.
- [31] Kevin Xie, Tingwu Wang, Umar Iqbal, Yunrong Guo, Sanja Fidler, and Florian Shkurti. Physics-based human motion estimation and synthesis from videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11532–11541, 2021.
- [32] Ye Yuan, Shih-En Wei, Tomas Simon, Kris Kitani, and Jason Saragih. Simpoe: Simulated character control for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 7159– 7169, 2021.
- [33] Jason Y Zhang, Panna Felsen, Angjoo Kanazawa, and Jitendra Malik. Predicting 3d human dynamics from video. In Proceedings of the IEEE International Conference on Computer Vision, pages 7114–7123, 2019.
- [34] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [35] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [36] Zhen Zhu, Tengteng Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. Progressive pose attention transfer for person image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2347–2356, 2019.