

PoseTriplet: Co-evolving 3D Human Pose Estimation, Imitation, and Hallucination under Self-supervision

Kehong Gong^{1,3*}
Jing Huang³

Bingbing Li^{2*}
Michael Bi Mi³

Jianfeng Zhang^{1*}
Jiashi Feng¹

Tao Wang^{1*}
Xinchao Wang^{1†}

¹National University of Singapore ²Nanyang Technological University ³Huawei International Pte Ltd

<https://github.com/Garfield-kh/PoseTriplet>.

Abstract

Existing self-supervised 3D human pose estimation schemes have largely relied on weak supervisions like consistency loss to guide the learning, which, inevitably, leads to inferior results in real-world scenarios with unseen poses. In this paper, we propose a novel self-supervised approach that allows us to explicitly generate 2D-3D pose pairs for augmenting supervision, through a self-enhancing dual-loop learning framework. This is made possible via introducing a reinforcement-learning-based imitator, which is learned jointly with a pose estimator alongside a pose hallucinator; the three components form two loops during the training process, complementing and strengthening one another. Specifically, the pose estimator transforms an input 2D pose sequence to a low-fidelity 3D output, which is then enhanced by the imitator that enforces physical constraints. The refined 3D poses are subsequently fed to the hallucinator for producing even more diverse data, which are, in turn, strengthened by the imitator and further utilized to train the pose estimator. Such a co-evolution scheme, in practice, enables training a pose estimator on self-generated motion data without relying on any given 3D data. Extensive experiments across various benchmarks demonstrate that our approach yields encouraging results significantly outperforming the state of the art and, in some cases, even on par with results of fully-supervised methods. Notably, it achieves 89.1% 3D PCK on MPI-INF-3DHP under self-supervised cross-dataset evaluation setup, improving upon the previous best self-supervised method [16, 26] by 8.6%.

1. Introduction

Video-based 3D human pose estimation aims to infer 3D pose sequences from videos, and therefore plays a crucial role in many applications such as action recognition [47,

*Equal contribution. Email: {gongkehong, zhangjianfeng, taowang}@u.nus.edu, l.libingbing@gmail.com

†Corresponding author. Email: xinchao@nus.edu.sg

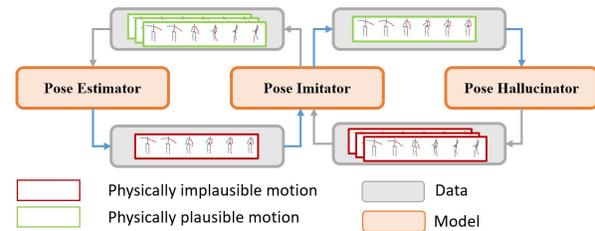


Figure 1. **Overview of our PoseTriplet framework.** The pose estimator, imitator, hallucinator are trained jointly in a dual-loop strategy. In the first loop, the estimator provides physically implausible motion information, which is then enhanced by the imitator via enforcing physical constraints to generate physically plausible motion. In the second loop, the hallucinator generates more diverse motion patterns given motion sequence from previous loop, and sends them to the imitator again for further refinement. This dual-loop paradigm facilitates tight co-evolution of the three components and enables iterative self-improving training of the estimator with the generated diverse and plausible motion data.

58], virtual try-on [31], and mixed reality [5, 20, 34]. Existing methods [22, 33, 37, 38, 48] mainly rely on the fully-supervised paradigms, in which the ground truth 3D data are given as input. However, capturing 3D pose data is cost-intensive and time-consuming, as it typically requires a multi-view setup or a motion capturing system [17, 34], making it infeasible under in-the-wild scenarios.

To this end, two categories of methods have been introduced to alleviate the 3D data availability issue. The first category explores the semi-supervised settings, in which only a small amount of the 3D annotations are given [30, 36, 68]. The second category, on the other hand, assumes no 3D data are available at all and only 2D poses are provided. Under this setup, state-of-the-art methods have mainly focused on imposing weak supervision signals to guide the training, such as aligning the projection of an inferred 3D pose with a 2D pose [4, 16, 61]. Due to the lack of 3D data and hence the missing of 2D-3D pairs, these methods are, by nature, brittle to the challenging scenarios such as unseen poses inherent to the in-the-wild tasks.

In this paper, we propose a novel self-supervised approach termed as PoseTriplet, which allows for explicitly generating physically- and semantically-plausible 2D-3D pose pairs, so that full supervisions can be imposed and further significantly strengthen the self-learning process. This is made possible through introducing a reinforcement-learning-based *imitator*, which is jointly optimized with the pose *estimator* alongside a pose *hallucinator*. Specifically, the imitator takes the form of a physics simulator with non-differentiable dynamics to ensure physical plausibility. The hallucinator helps generate more diverse motion with generative motion completion. These three key components are integrated into a self-contained framework and co-evolve via a dual-loop strategy as the training proceeds. With only 2D pose data as input, PoseTriplet progressively generates, refines and hallucinates 3D data, which in turn reinforces all components in the loop. Once trained, each component of PoseTriplet can be readily taken out and serves as an off-the-shelf tool for its dedicated task, such as pose estimation or imitation.

The key motivation behind co-evolving the pose estimator, imitator and hallucinator, lies in their complementary natures. In particular, pose estimator takes 2D poses as input and generates 3D poses with reasonable semantics (e.g., nature behaviors) but implausible dynamics; such derived 3D poses are then refined through the physics-based imitator that enforces physical constraints. Conversely, the reinforcement-learning-based imitator is possible to generate unnatural behaviors (e.g., overly energetic movements), which can be rectified through the pose estimator to ensure the semantic plausibility. Pose hallucinator, on the other hand, enhances the data diversity by producing realistic 3D pose sequences under both the semantic and physical guidance, which further strengthens data synthesizing and hence improves generalization performance.

We show the overall workflow of PoseTriplet in Fig. 1, which effectively aligns with aforementioned motivation. Unlike prior endeavors that rely on self-consistency-based supervisions or 3D sequences as input, PoseTriplet, through the dual-loop scheme, turns the input 2D poses into dependable 3D poses of realistic semantics and dynamics, thereby lending itself to much stronger supervisions and consequently the co-evolution of the pose estimator, imitator and hallucinator. Experimental results across H36M, 3DHP, and 3DPW datasets demonstrate that, PoseTriplet gives rises to pose estimation results significantly superior to the state-of-the-art self-supervised methods, and sometimes even on par with results from fully-supervised ones. Notably, it achieves 89.1% 3D PCK on MPI-INF-3DHP under self-supervised cross-dataset evaluation setup, improving upon the previous best self-supervised method [16, 26] by 8.6%.

Our contribution is therefore a novel scheme dedicated

for self-supervised 3D pose estimation, achieved by the co-evolution of a pose estimator, imitator, and hallucinator. The three components complement and benefit one another, together leading to a self-contained system that enables realist 3D pose sequences and further the 2D-3D augmented supervisions. By taking only 2D poses as input, PoseTriplet delivers truly encouraging results across various benchmarks, largely outperforming the state of the art and even approaching full-supervised results.

2. Related works

3D pose estimation 3D pose estimation have been wildly explored under fully supervised, semi-supervised, self-supervised. Various approaches have been explored under fully supervised setting [22, 33, 34, 37, 38, 48, 52, 53, 59, 60, 66]. Through offering impressive results, those approaches highly rely on accurate motion capture data, which are hard to collect. To address high cost of data collection, semi-supervised methods [30, 36, 68] are proposed to utilize the information from unlabeled data. Besides semi-supervised approach, augmentation based methods [9, 29] are proposed to enlarge the data amount through evolution strategy [29] or learnable approach [9].

Different from the above schemes, self-supervised methods, with multi-view data, explore the intrinsic supervision for model training, without requiring ground truth 3D pose [18, 23, 51]. For instance, Kocabas *et al.* [23] utilize the epipolar geometry to generate pseudo label, [18, 51] utilize the 3D pose consistency across different views. Though being effective, those approaches require synchronized multiple cameras, which are not usual in real scenarios. Other methods [4, 7, 16, 61] explore the more challenge single view setting. For example, Drover *et al.* [7] utilize the prior that a random projection of a plausible 3D pose estimation will be plausible in 2D pose distribution through adversary training. Chen *et al.* [4] improves this idea by adding cycle consistency. Yu *et al.* [61] further introduces the scale steps for 2D poses to resolve the ambiguity issue. Zhang *et al.* [65] applies self-supervised learning on test data to adapt model to new scenarios.

Our method belongs to the self-supervised approaches under single view setting. Different from previous self-supervised approaches which implement weak supervision signal through consistency [4] or adversary [7, 61], our method directly uses the strong supervision signal from self-generated data, results in more accurate and stable model performance. The pseudo label strategy [30] under semi-supervised category is close to our approach. However, our approach does not require ground truth data for model pretraining, and our method introduces physical plausibility refinement and diversity enhancement to achieve better performance, which are absent in [30].

Physics-based pose estimation The above methods are all

Kinematics based. Though providing impressive results, they do not consider physical constrains, thus suffering physical implausible artifacts (*e.g.*, foot skating and ground penetration). To ensure physical plausibility, recent works explore physical constraints. Rempe *et al.* [43] introduces physical law to the foot contact and human dynamic, while its iterative optimization is high time costly (*e.g.*, 30 minutes for 2s clip). Later, [45, 46, 57] propose differentiable physical constrain to reduce the time cost. But they only consider foot contact, making them less effective in scenarios with other important contact (*e.g.*, lay down, sitting with chair).

Different from optimization based approaches, physics simulation based methods use physics simulators to provide realistic physical constrains. DeepMimic [40] tries to imitate various motion from reference mocap data in physics engine via reinforcement learning. SFV [41] proposes to refine the low fidelity motion data from video-based pose estimation through imitation learning. However, their adopted imitation learning requires days of training for just one clip. Later, SimPoe [64] addresses this issue by introducing RFC [63] to effectively reduce the time consumption by training one policy for all motion clips. Our method is built upon SimPoe [64] for better generalization and low time cost. However, different from those methods only using physical constraints for post processing, our method propose to involve it in the learning loop. As such, no mocap data for pose estimation training and imitation learning is required.

Motion synthesis Motion synthesis includes non-learning based and learning based approaches. In non-learning based approaches, the motion graph method [25] first builds transition edges between different motion points based on their similarity, and then generates new motion data through traversing the graphs. Motion matching [35] searches proper future frames in motion data based on motion states in real time. In learning based approaches, motion perdition based methods [1, 10, 28, 32, 39, 62, 67] aim to predict future poses conditioned on previous poses. Action generation [2, 3, 55] aims to generate pose sequence conditioned on action labels. Motion completion [8, 11–14, 21] generates realistic transitions between key frames, which most relevant to our work in their aims. The pose hallucination in our framework also aims to generate novel motion sequence, where motion graph and motion match methods are not applicable due to tight restriction in their generated data. We therefore choose motion completion considering it can generate longer sequence with continuously input key frames.

3. Methodology

Given a 2D pose sequence $\mathbf{x}_{1:T} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ of length T , where $\mathbf{x}_t \in \mathbb{R}^{J \times 2}$ is the 2D spatial coordinate of J body

joints at time t , our goal is to estimate the 3D pose sequence $\mathbf{X}_{1:T} = (\mathbf{X}_1, \dots, \mathbf{X}_T)$, where $\mathbf{X}_t \in \mathbb{R}^{J \times 3}$ is the corresponding 3D joint position under the camera coordinate system. Conventionally, a pose estimator $\mathcal{P} : \mathbf{x}_{1:T} \mapsto \mathbf{X}_{1:T}$ with parameter θ is trained with a large set of paired 2D and 3D pose data $\{\mathbf{x}_{1:T}, \mathbf{X}_{1:T}\}$ through fully-supervised learning approaches [22, 33, 34, 38, 48]:

$$\min_{\theta} \mathcal{L}_{\mathcal{P}}(\mathcal{P}_{\theta}(\mathbf{x}_{1:T}), \mathbf{X}_{1:T}). \quad (1)$$

Here $\mathcal{L}_{\mathcal{P}}$ denotes the loss function which is typically defined as mean square errors (MSE) between predicted and ground truth 3D poses sequences. However, ground truth 3D pose data is expensive to capture, which limits the applicability of these approaches. To avoid using 3D data, previous self-supervised approaches typically apply weak 2D re-projection loss [4, 7, 16, 61] to learn the estimator:

$$\min_{\theta} \mathcal{L}_{\mathcal{P}}(\Pi(\mathcal{P}_{\theta}(\mathbf{x}_{1:T})), \mathbf{x}_{1:T}), \quad (2)$$

where Π is the perspective projection function. The re-projection loss only provides weak supervision which tends to induce unstable or unnatural estimations. In this work, we aim to design a self-supervised learning framework of which the core is an iterative self-improving paradigm. Specifically, we propose to enhance the current estimation with some specifically designed transformation \mathcal{T} (*e.g.*, to produce more smooth and diverse motion):

$$\mathbf{X}'_{1:T} = \mathcal{T}(\mathcal{P}_{\theta_n}(\mathbf{x}_{1:T})) \quad (3)$$

The enhanced estimates are then projected to 2D pose to obtain paired training data $\{\mathbf{x}'_{1:T}, \mathbf{X}'_{1:T}\}$, which are used to improve the pose estimator:

$$\theta_{n+1} \leftarrow \min_{\theta} \mathcal{L}_{\mathcal{P}}(\mathcal{P}_{\theta}(\mathbf{x}'_{1:T}), \mathbf{X}'_{1:T}) \quad (4)$$

Here θ_n and θ_{n+1} denote the parameters of the current estimator and improved estimator. The improved estimator can then be utilized to start a new iteration of data enhancement and training. Build on this *self-improving* paradigm, we can train a superior pose estimator starting from only a set of 2D pose sequences $\{\mathbf{x}_{1:T}\}$.

3.1. PoseTriplet

To construct an effective self-improving framework, we identify two challenging aspects for enhancing the 3D motion sequence: 1) the pose estimation from the estimator may not be physically plausible due to ignorance of force, mass and contact modeling; 2) existing 2D motion may be limited in diversity and thus the learned model cannot generalize well. To address these challenges, we introduce a pose imitator based on the reinforcement learning aided human motion modeling and a pose hallucinator based on generative motion interpolation accordingly to refine and diversify the 3D motion. The former helps correct the physical artifacts while the latter generates novel pose sequences

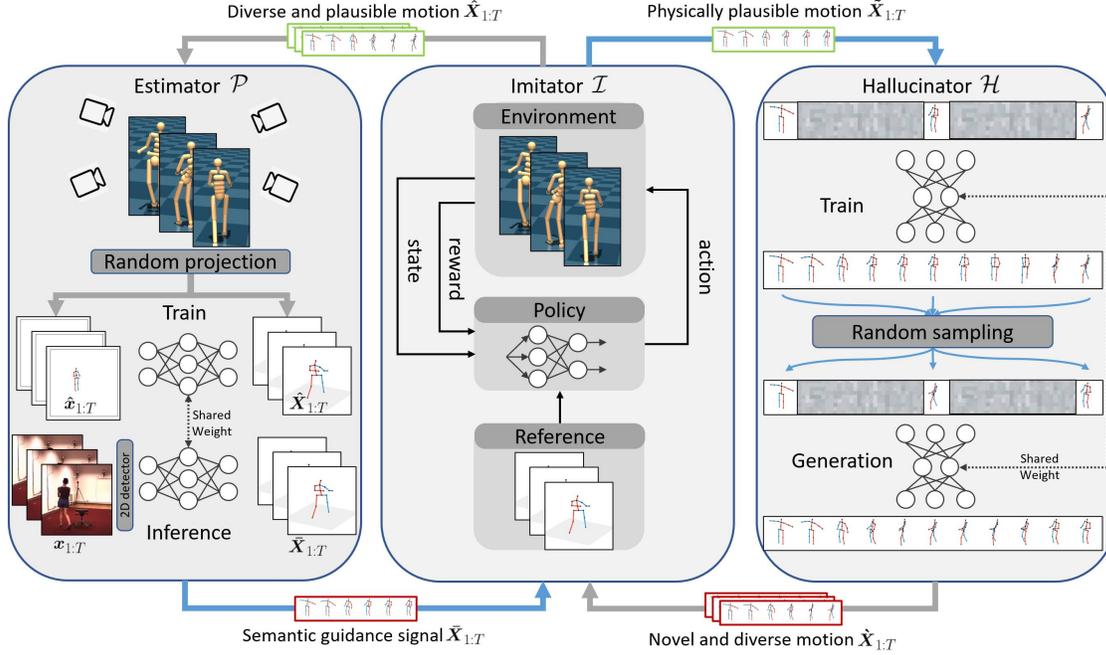


Figure 2. **Detail of our PoseTriplet framework.** Given available 2D pose sequence $x_{1:T}$, the pose estimator \mathcal{P} transforms it to low-fidelity 3D pose sequence $\tilde{X}_{1:T}$. $\tilde{X}_{1:T}$ is then served as semantic guidance signal (*i.e.*, reference motion) for imitator \mathcal{I} to obtain physically plausible motion $\tilde{X}_{1:T}$. The hallucinator \mathcal{H} then generates novel and diverse motion $\hat{X}_{1:T}$ from $\tilde{X}_{1:T}$, which is then refined by the imitator \mathcal{I} to obtain the final enhanced diverse and plausible motion $\hat{X}_{1:T}$. $\hat{X}_{1:T}$ is then projected to 2D-3D pairs to train the estimator. The improved estimator takes the available 2D pose sequence $x_{1:T}$ and start another round of dual-loop optimization.

based on the existing estimates. We find these two aspects in motion are complementary and thus combine them together. The resulting pipeline helps obtain 3D motion data $\{\hat{x}'_{1:T}, \hat{X}'_{1:T}\}$ with significantly improved physical plausibility and motion diversity. Nevertheless, we find naive two-step combination of the two approaches generate inferior-quality 3D pose sequence. The reason is that performing motion diversification first could be ineffective due to implausible estimate while conducting motion diversification later could introduce physical artifacts. Therefore, we further introduce a *dual-loop* scheme and unify the two components with pose estimator into a novel self-supervised framework named PoseTriplet.

Dual-loop architecture Concretely, as shown in Fig. 2, our PoseTriplet introduces a dual-loop architecture to integrate the three modules: a *pose estimator* \mathcal{P} , a *pose imitator* \mathcal{I} , and a *pose hallucinator* \mathcal{H} . Given the set of available 2D pose sequence $x_{1:T}$, the pose estimator first transforms them to low-fidelity 3D pose sequence:

$$\tilde{X}_{1:T} = \mathcal{P}(x_{1:T}) \quad (5)$$

$\{\tilde{X}_{1:T}\}$ is converted to low-fidelity reference motions and served as semantic guidance signal to the pose imitator, which imposes the physical human motion dynamic modeling and obtains physically plausible motion sequence:

$$\tilde{X}_{1:T} = \mathcal{I}(\tilde{X}_{1:T}) \quad (6)$$

By learning a generative motion completion model, the pose hallucinator then generates novel and diverse motion sequences $\{\hat{X}_{1:T}\}$ based on the improved plausible motion from the imitator:

$$\hat{X}_{1:T} = \mathcal{H}(\tilde{X}_{1:T}) \quad (7)$$

Afterwards, instead of closing the loop by treating $\{\hat{X}_{1:T}\}$ as augmented data to the estimator, we introduce another loop. We feed $\{\hat{X}_{1:T}\}$ back into the imitator to correct the induced physical artifacts and obtain the final expected plausible and diverse motion sequences:

$$\hat{X}_{1:T} = \mathcal{I}(\hat{X}_{1:T}) \quad (8)$$

$\{\hat{X}_{1:T}\}$ is then projected to 2D to obtain paired data $\{\hat{x}_{1:T}, \hat{X}_{1:T}\}$ for training the pose estimator.

By jointly optimizing this dual-loop architecture, the three components form a tight co-evolving paradigm: 1) the estimator benefits from the diverse and plausible augmented data to learn more accurate estimation. 2) the imitator learns more robust and physically natural motion based on the improved estimation and diverse data generated from the hallucinator. 3) the hallucinator generates diverse pose sequence of higher quality based on the improved data from the imitator.

Loop starting Another challenging aspect of this self-improving learning paradigm is the loop starting. Without

access to 3D motion data, the whole framework cannot start learning. Recall our pose imitator employs physics-based human motion model, we thus develop a zero-data generating strategy that produces initial 3D pose sequence for starting the dual-loop learning. Specifically, we generate root trajectory signal in horizontal plane with random direction and proper velocity. This trajectory is then used for guidance signal for RL agent. By control the agent to follow the generated trajectory, we can generate motion sequences that are physically plausible. These motion sequences is then projected to obtain 2D-3D pose pairs and used to train a initial pose estimator. In this way, the whole dual-loop learning can be started.

3.2. Module detail

3.2.1 Pose estimator

The pose estimator estimates the 3D pose sequence $\mathbf{X}_{1:T}$ from the input sequence $\mathbf{x}_{1:T}$. Specifically, we adopt a similar estimator architecture as VideoPose [38], which predicts both root trajectory and root-relative joint locations. The trajectory can be used as additional movement signal to pose imitator. Meanwhile, the noise in root movement can be corrected by the pose imitator and in turn help the pose estimator. We use Mean Square Error (MSE) loss for the root-related pose estimation and Weighted L1 loss for the trajectory estimation following [38].

Projection for training estimator Given the generated motion sequence data $\{\hat{\mathbf{X}}_{1:T}\}$, we project them to 2D to obtain paired training data. We consider two strategies for the projection: 1) Heuristic random projection. We set the virtual camera with certain elevation, azimuth range, height and distance range to match the indoor capture environment. This is similar to the projection strategy for 3D pose data synthesis as Chen *et al.* [6]; 2) Generative adversarial learning based projection [9]. A generator is used to regress the camera orientation and position for each motion sequence. The regression is learned through a discriminator by distinguishing the real and the projected 2D pose sequences with the generated camera parameters. In this way, reasonable camera viewpoint distribution can be extracted from real 2D pose data, improving the plausibility of generated 2D-3D paired data. The two strategies are combined in our framework to ensure the diversity of camera viewpoints.

3.2.2 Pose imitator

The 3D pose sequences $\{\bar{\mathbf{X}}_{1:T}\}$ predicted from pose estimator \mathcal{P} , due to lack of physical constrain, would suffer unnatural artifacts such as foot skating, floating, floor penetration. Those artifacts prevent it from being used as training data directly for estimator \mathcal{P} or hallucinator \mathcal{H} . To address the issue, motivated by [40, 41, 64], we introduce a reinforcement learning based pose imitator \mathcal{I} to imitate the low fidelity 3D pose sequence $\{\bar{\mathbf{X}}_{1:T}\}$ from pose estima-

tor to generate more physically plausible motion sequence $\{\tilde{\mathbf{X}}_{1:T}\}$.

Background The imitation process can be seen as a Markov decision process. Given a reference motion and current state $s_t \in \mathcal{S}$, the agent interacts with the simulation environment with action $\mathbf{a}_t \in \mathcal{A}$ and receive reward r_t . the action is determined by a policy $\pi(\mathbf{a}_t|s_t)$ conditioned on state $s_t \in \mathcal{S}$; the reward is determined based on how similar the agent behaves like the reference motion. When an action is taken, the current state s_t changes to next state s_{t+1} through transition function $\mathcal{T}(s_{t+1}|s_t, \mathbf{a}_t)$. The goal is to learn a policy that maximizes the average cumulative rewards $\sum_{i=1}^{\infty} \gamma^i r_t$ (*i.e.*, performing similar behavior in physics simulator as reference motion), where γ is the discounting factor. The state, action and rewards are detailed below.

State includes current pose \mathbf{q}_t , current velocity $\dot{\mathbf{q}}_t$, and target pose $\tilde{\mathbf{q}}_{t+1}$ from reference motion. To deal with the noisy reference motion from the pose estimator, we introduce an extra encoded feature ϕ by concatenating and fusing the past and future motion information. In this way, the control policy is aware of past and future reference motion, and is thus more robust to the noise.

Action involves two kinds of forces: internal force and external force. The internal force is applied by actuator on the non-root joints (*e.g.*, elbow, knee). Following previous work [42], we use PD (proportional–derivative) control for internal force control. The external force $\boldsymbol{\eta}_t$ is a virtual force applied on root joint(*i.e.*, hip) [63] for extra interaction (*e.g.*, sitting on the chair) and is regressed by the policy network.

Rewards measure the motion differences between the agent and reference motion. These differences capture pose related (pose, velocity), root related (root height, root velocity) and body end factors (position, velocity). Besides, a regulation loss on virtual force is applied to avoid unnecessary external force following [63]. As we find that the agent is hard to move with the above setting due to the noisy reference motion, we further introduce a feet relative position into the motion characteristics to enhance the feet motion.

3.2.3 Pose hallucinator

The pose hallucinator aims to generate novel and diverse motion sequence based on the refined data from pose imitator. In this work, we choose motion interpolation technique to generate novel pose motions. Specifically, we sample key-frames from the refined pose sequence, and interpolate the missing frames via neural networks to generate new motion data. In details, the pose hallucinator is constructed by a recurrent neural network (RNN) structure. The inputs are the sampled temporal key-frames (we sample key-frames with a certain frame interval). Conditioned on these sampled key-frames, the model predicts the intermediate frames

in sequential manner. A reconstruction loss and an adversary loss is used to train this model. The reconstruction loss measures the L_2 distance between the ground truth and predicted poses. The adversary loss provides temporal supervision to avoid RNN collapse (*i.e.*, predicting average motion). In the inference stage, we randomly select frames from different motion clips and generate novel motion sequences based on these sampled key-frames.

4. Experiments

We study three questions in experiments. 1) Is PoseTriplet able to improve performance of video pose estimator for both intra- and cross-dataset scenarios? 2) How does the performance improves with the round of co-evolving process? 3) How does the amount of training data affects model performance? We conduct experiments with H36M (source dataset) and 3DHP/3DPW (for cross-dataset evaluation). Throughout the experiments, we adopt VideoPose [38] (T=27) as our pose estimator. We report results from estimator for comparison. Please refer to supplementary for more implementation details.

4.1. Dataset

H36M [17] is the most popular 3D pose benchmark captured by marker-based motion capture system. It contains 3.6 million video frames for 11 subjects and 15 scenarios. Following previous works [4, 61], we use the 2D poses of subject S1, S5, S6, S7, S8 as our training set and evaluate the performance on S9 and S11. The two standard metrics Mean Per Joint Position Error (MPJPE) in millimeters and Procrustes Aligned Mean Per Joint Position Error (PA-MPJPE) are used for evaluation.

3DHP [34] is a large 3D pose dataset. It contains both indoor and outdoor scenarios. Following previous works [4, 24], we report the metrics of MPJPE, Percentage of Correct Keypoints (PCK) and Area Under the Curve (AUC) after scale and rigid alignment for evaluation. We only use its test set to evaluate the model’s generalization performance.

3DPW [50] is a more challenging in-the-wild dataset. It contains more complicated activities and scenarios. Same as 3DHP, we only use its test set to evaluate model’s generalization performance. Follow previous work [24], we report MPJPE and PA-MPJPE for 3DPW.

4.2. Quantitative results

Results on H36M We compare our PoseTriplet with other state-of-the-art self-supervised methods [4, 16, 26, 44, 61] under GT (ground truth 2D poses) and Det (detected 2D poses) settings as shown in Table 1. Among which, [4, 26, 44] implement weak supervision (*i.e.*, consistency supervision), [16, 61] utilize temporal information through adversary learning [61] and smoothness constrains [16].

Our method outperforms the best of them by a large margin in MPJPE for both GT (85.3 vs. 68.2) and Det (82.1 vs. 78.0) settings. The result verifies that our method with co-evolving strategy and augmented supervision performs better compared with previous approaches. Moreover, our method also outperforms some weakly-supervised approaches [15, 18, 30, 56] which involve ground truth data during training. Especially, comparing with Li *et al.* [30] which implements low rank representation and temporal smoothing for pseudo 3D label generation, our approach, utilizing the advantage of physics simulator, provides better refinement and outperforms [30] by a large margin in MPJPE (88.8 Vs 78.8) even it use ground truth data (*i.e.*, subject 1). This verifies the effectiveness of our co-evolving strategy on reducing reliance on 3D data.

Mode	Method	GT		Det	
		P1 (↓)	P2 (↓)	P1 (↓)	P2 (↓)
Full	Martinez <i>et al.</i> [33]	45.5	37.1	62.9	47.7
Full	Pavlo <i>et al.</i> [38]	37.2	27.2	46.8	36.5
Weak	3DInterpreter [56]	-	88.6	-	98.4
Weak	AIGN [15]	-	79.0	-	97.4
Weak	Drover <i>et al.</i> [7]	-	38.2	-	64.6
Weak	Li <i>et al.</i> [30]	-	-	88.8	66.5
Weak	Umar <i>et al.</i> [18]	-	-	-	55.9
Self	Rhodin <i>et al.</i> [44]	-	-	131.7	98.2
Self	Chen <i>et al.</i> [4]	-	51.0	-	68.0
Self	Kundu <i>et al.</i> [26]	-	-	-	62.4
Self	Kundu <i>et al.</i> [27]	-	-	-	63.8
Self	Yu <i>et al.</i> [61]	85.3	42.0	92.4	52.3
Self	Hu <i>et al.</i> [16]	-	-	82.1	-
Self	Wandt <i>et al.</i> [51]*	-	-	81.9	53.0
Self	Ours	68.2	45.1	78.0	51.8

Table 1. **Results on H36M** in terms of MPJPE (P1) and PA-MPJPE (P2). * uses multi-view setting. Best results are shown in **bold** under self supervised setting.

Results on 3DHP We then evaluate the generalization performance of our method on cross dataset 3DHP. We compare our PoseTriplet with state-of-the-art methods, including fully supervised, weakly supervised, and self supervised approaches [4, 24, 26, 27, 34, 49, 61]. As shown in Table 2, under cross-data evaluation, our method overruns previous self-supervised methods [4, 26, 61] significantly in PCK (82.2 vs. 89.1) and MPJPE (103.8 vs. 79.5). The result indicates that the diverse and plausible motion generated by our PoseTriplet improves generalization. Exception is that Kundu *et al.* [26], which uses extra data and unpaired 3D poses for model training and thus achieves slightly better performance in AUC (56.3 vs. 53.1). Our method also outperforms self-supervised methods [4, 26, 27, 61] trained on 3DHP dataset directly. In addition, our method achieves better performance than weakly supervised approaches [24, 49] in all metrics, even though they use un-

paired images and 3D poses for supervision during the training process. Notably, our method even achieves comparable performance with fully supervised approaches [24, 34, 49]). In summary, the cross-dataset performance of our self-supervised framework PoseTriplet is comparable with the intra-dataset result from fully/semi-supervision. This indicates a good generalization performance of our PoseTriplet.

Mode	Method	CE	PCK (\uparrow)	AUC (\uparrow)	MPJPE (\downarrow)
Full	VNect [34]		83.9	47.3	98.0
Full	HMR [49]		86.3	47.8	89.8
Full	SPIN [24]		92.5	55.6	67.5
Weak	HMR [49]		77.1	40.7	113.2
Weak	SPIN [24]		87.0	48.5	80.4
Self	Chen <i>et al.</i> [4]		71.1	36.3	-
Self	Kundu <i>et al.</i> [27]		80.2	44.8	97.1
Self	Kundu <i>et al.</i> [26]		84.6	60.8	93.9
Self	Yu <i>et al.</i> [61]		86.2	51.7	-
Self	Chen <i>et al.</i> [4]	✓	64.3	31.6	-
Self	Kundu <i>et al.</i> [26]*	✓	82.1	56.3	103.8
Self	Yu <i>et al.</i> [61]	✓	82.2	46.6	-
Self	Ours	✓	89.1	53.1	79.5

Table 2. **Results on 3DHP** in terms of PCK, AUC, and MPJPE. CE denotes cross-data evaluation. * uses extra unpaired 2D/3D dataset for training. Best results are shown in **bold**.

Results on 3DPW We further evaluate the generalization performance of our method on in-the-wild 3DPW dataset. Note that there is few works evaluated on 3DPW under the self-supervised cross dataset setting. Therefore, we compare to the supervised approaches [24, 38, 49, 54] directly. From Table 3, we can observe our method achieves comparable results with the fully supervised baseline without relying on any 3D data. This demonstrates that our method performs well on complicated and challenging in-the-wild scenarios.

Mode	Method	CE	MPJPE (\downarrow)	P-MPJPE (\downarrow)
Full	Wang <i>et al.</i> [54]	✓	124.2	-
Full	DSD-SATN [49]	✓	-	69.5
Full	CRMH [19]	✓	105.3	62.3
Full	BMP [66]	✓	104.1	63.8
Full	VideoPose [38]	✓	101.8	63.0
Self	Ours	✓	115.0	69.5

Table 3. **Results on 3DPW** in terms of MPJPE and PA-MPJPE. CE denotes cross-data evaluation.

4.3. Qualitative results

While previous self-supervised methods rely on weak supervision signal (*e.g.*, consistency loss), our method trains the pose estimator with augmenting supervision from the self-generated data, resulting in more stable, plausible, and

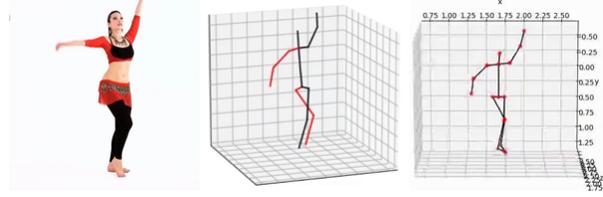


Figure 3. **Result on UID [16] comparison with Hu *et al.* [16]**. The figure includes: input (left), ours (middle), Hu *et al.* [16] (right).

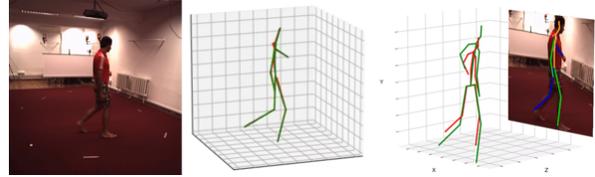


Figure 4. **Result on H36M comparison with Yu *et al.* [61]**. The figure includes: input (left), ours (middle), Yu *et al.* [61] (right). Red skeleton is prediction, green skeleton is ground truth.

accurate estimation¹. As shown in Fig. 3, although Hu *et al.* [16] implements temporal smoothness prior during the training process, jittering effect is still obvious. While our result, learned from co-evolving approach, is much smoother. Yu *et al.* [61] introduce a scale estimation strategy for 2D pose to reduce the scale ambiguity. Through the weak supervision from the bone length consistency and scale distribution, his result still contains scale ambiguity (*i.e.*, the body size varies) as shown in Fig. 4. Ours result maintains stable and accurate in term of body size compared with it. We further demonstrate the result from 3DHP (Fig. 5) and 3DPW (Fig. 6). These results demonstrate that our method perform well on unseen poses for in-the-wild scenarios. More in-the-wild examples can be viewed in the supplementary material in video format.

4.4. Ablation study

4.4.1 Ablation on round of co-evolution

We then analysis how the co-evolving round improves the performance of each component (estimator \mathcal{P} , imitator \mathcal{I} , hallucinator \mathcal{H}). To demonstrate the improvement, we select three evaluation metrics for each component. For estimator, we evaluate the trained model \mathcal{P} on H36M test set and report the MPJPE as evaluation metric. For imitator, we evaluate the trained policy \mathcal{I} on GT 3D reference motion (H36M) to measure the number of termination (*e.g.*, fall down) as evaluation metric. For hallucinator, we evaluate the trained model \mathcal{H} on GT 3D data (Walking scenario [12] in H36M) for intermediate pose completion. We measure

¹Fig.3-8 are video figure in arxiv version that are best viewed in Adobe Reader (click and play), and videos are in supplementary materials.

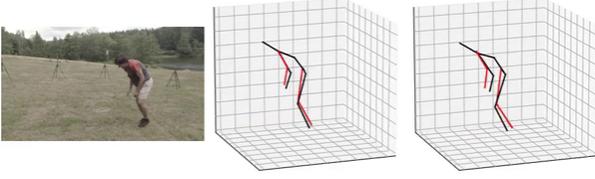


Figure 5. **Result on 3DHP compared with Ground Truth.** The figure includes: input (left), ours (middle), ground truth (right).

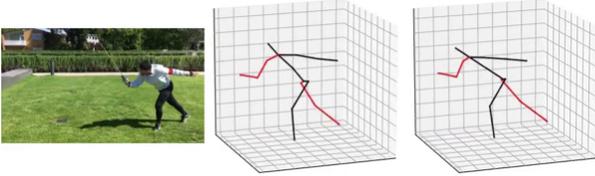


Figure 6. **Result on 3DPW compared with Ground Truth.** The figure includes: input (left), ours (middle), ground truth (right).

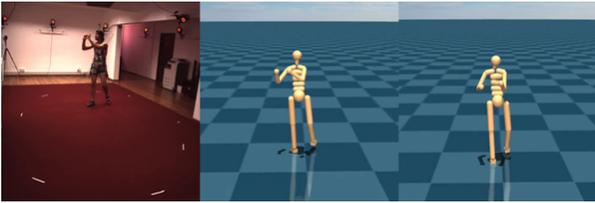


Figure 7. **Results on co-evolving for imitator \mathcal{I} .** The figure includes: video source (left), our co-evolving result (middle), oracle trained with ground truth data(right).

the MPJPE of pose and root position as evaluation metric. We involve an oracle by training each model using the GT data directly as showed in the last row of Table 4. Through iterative co-evolving, the performance of estimator \mathcal{P} , imitator \mathcal{I} , hallucinator \mathcal{H} are improved and getting closer to the result which is trained with GT data. We further provide visualization result for imitator \mathcal{I} (Fig. 7), hallucinator \mathcal{H} (Fig. 8). This result shows that the imitator \mathcal{I} and hallucinator \mathcal{H} co-evolved by our PoseTriplet without using 3D data achieve a comparable performance compared with the oracle trained with GT 3D data.

Round	\mathcal{P}	\mathcal{I}	\mathcal{H}	
Num.	P1 (pose)	Termination Num.	P1 (pose)	P1 (root)
0	193.6	-	-	-
1	112.2	928	-	-
2	77.8	280	71.4	62.6
3	68.2	132	67.3	54.0
Oracle	37.2	81	53.0	33.7

Table 4. **Results on co-evolving for estimator \mathcal{P} , imitator \mathcal{I} , hallucinator \mathcal{H} .** Note that round 0 is the **Loop starting**, and we involve hallucinator \mathcal{H} after round one to ensure the quality of initial pose estimation.

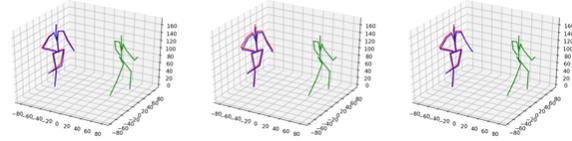


Figure 8. **Results on co-evolving for hallucinator \mathcal{H} .** The figure includes: ground truth (left), our co-evolving result (middle), oracle trained with ground truth data(right).

4.4.2 Ablation on amount of data usage

To study how the amount of data affects the performance, we construct an ablation experiment with limited 2D pose data. As shown in Table 5, we gradually involve more data in our method, (*i.e.*, S1, S1+S5, S1+S5+S6+S7+S8). Result shows that the performance of PoseTriplet can be improved gradually by adding more 2D pose data in both intra and cross-dataset scenarios.

Mode	Sub	H36M	3DHP	3DPW
Self	S1	89.2	94.0	135.8
Self	S1,S5	81.9	83.5	128.6
Self	S1,S5,S6,S7,S8	68.2	79.5	115.0

Table 5. **Results on ablation amount of data** in terms of MPJPE.

5. Conclusion

In this work we present a novel framework PoseTriplet for self-supervised 3D pose estimation, which is achieved by a co-evolution strategy of a pose estimator, imitator, and hallucinator. These three components, complement and strength one another through a dual-loop strategy as the training procedure. The framework enables generating diverse and plausible motion data, which help train superior pose estimator. Experiments on varies benchmarks demonstrate that PoseTriplet yields encouraging results. It outperforms the state of the art self-supervised approaches and even competes with fully-supervised approaches.

Limitations The major limitation is that our pipeline suffers low training efficiency, *e.g.*, it takes 7 days to train for 3 rounds on a machine with a Intel Xeon Gold 6278C CPU and a Tesla T4 GPU. The reason is that the imitator (\mathcal{I}) is implemented with CPU-based reinforcement learning (RL) and the hallucinator (\mathcal{H}) is instantiated with RNN architecture. In the future, we will explore GPU-based RL implementation and more efficient hallucinator architecture (*e.g.*, transformer) to speed up the training process.

Acknowledgement This project is supported by NUS Faculty Research Committee Grant (WBS: A-0009440-00-00) and NUS Advanced Research and Technology Innovation Centre (Project Reference ECT-RP2). Kehong would like to thank Ye Yuan for the discussion.

References

- [1] Emad Barsoum, John Kender, and Zicheng Liu. Hp-gan: Probabilistic 3d human motion prediction via gan. In *CVPRw*, 2018. 3
- [2] Neeraj Battan, Yudhik Agrawal, Sai Soorya Rao, Aman Goel, and Avinash Sharma. Glocalnet: Class-aware long-term human motion synthesis. In *WACV*, 2021. 3
- [3] Judith Butepage, Michael J Black, Danica Kragic, and Hedvig Kjellstrom. Deep representation learning for human motion prediction and classification. In *CVPR*, 2017. 3
- [4] Ching-Hang Chen, Amrith Tyagi, Amit Agrawal, Dylan Drover, Rohith MV, Stefan Stojanov, and James M Rehg. Unsupervised 3d pose estimation with geometric self-supervision. In *CVPR*, 2019. 1, 2, 3, 6, 7
- [5] Mingfei Chen, Jianfeng Zhang, Xiangyu Xu, Lijuan Liu, Jiashi Feng, and Shuicheng Yan. Geometry-guided progressive nerf for generalizable and efficient neural human rendering. *arXiv*, 2021. 1
- [6] Wenzheng Chen, Huan Wang, Yangyan Li, Hao Su, Zhenhua Wang, Changhe Tu, Dani Lischinski, Daniel Cohen-Or, and Baoquan Chen. Synthesizing training images for boosting human 3d pose estimation. In *3DV*, 2016. 5
- [7] Dylan Drover, Rohith MV, Ching-Hang Chen, Amit Agrawal, Amrith Tyagi, and Cong Phuoc Huynh. Can 3d pose be learned from 2d projections alone? In *ECCVw*, 2018. 2, 3, 6
- [8] Yinglin Duan, Tianyang Shi, Zhengxia Zou, Yenan Lin, Zhehui Qian, Bohan Zhang, and Yi Yuan. Single-shot motion completion with transformer. *arXiv*, 2021. 3
- [9] Kehong Gong, Jianfeng Zhang, and Jiashi Feng. Poseaug: A differentiable pose augmentation framework for 3d human pose estimation. In *CVPR*, 2021. 2, 5
- [10] Liang-Yan Gui, Yu-Xiong Wang, Xiaodan Liang, and José MF Moura. Adversarial geometry-aware human motion prediction. In *ECCV*, 2018. 3
- [11] Félix G Harvey and Christopher Pal. Recurrent transition networks for character locomotion. In *ACM Transactions on Graphics*, 2018. 3
- [12] Félix G Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. In *ACM Trans. on Graphics*, 2020. 3, 7
- [13] Alejandro Hernandez, Jurgen Gall, and Francesc Moreno-Noguer. Human motion prediction via spatio-temporal inpainting. In *ICCV*, 2019. 3
- [14] Daniel Holden, Jun Saito, and Taku Komura. A deep learning framework for character motion synthesis and editing. In *ACM Trans. on Graphics*, 2016. 3
- [15] Tung Hsiao-Yu Fish, Harley Adam W, Seto William, and Fragkiadaki Katerina. Adversarial inverse graphics networks: Learning 2d-to-3d lifting and image-to-image translation from unpaired supervision. In *ICCV*, 2017. 6
- [16] Xiaodan Hu and Narendra Ahuja. Unsupervised 3d pose estimation for hierarchical dance video recognition. In *ICCV*, 2021. 1, 2, 3, 6, 7
- [17] 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 Trans. on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014. 1, 6
- [18] Umar Iqbal, Pavlo Molchanov, and Jan Kautz. Weakly-supervised 3d human pose learning via multi-view images in the wild. In *CVPR*, 2020. 2, 6
- [19] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *CVPR*, 2020. 7
- [20] Hanbyul Joo, Natalia Neverova, and Andrea Vedaldi. Exemplar fine-tuning for 3d human pose fitting towards in-the-wild 3d human pose estimation. *3DV*, 2020. 1
- [21] Manuel Kaufmann, Emre Aksan, Jie Song, Fabrizio Pece, Remo Ziegler, and Otmar Hilliges. Convolutional autoencoders for human motion infilling. *3DV*, 2020. 3
- [22] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *CVPR*, 2020. 1, 2, 3
- [23] Muhammed Kocabas, Salih Karagoz, and Emre Akbas. Self-supervised learning of 3d human pose using multi-view geometry. In *CVPR*, 2019. 2
- [24] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019. 6, 7
- [25] Lucas Kovar, Michael Gleicher, and Frédéric Pighin. Motion graphs. In *ACM Transactions on Graphics*, 2008. 3
- [26] Jogendra Nath Kundu, Siddharth Seth, Varun Jampani, Mugalodi Rakesh, R Venkatesh Babu, and Anirban Chakraborty. Self-supervised 3d human pose estimation via part guided novel image synthesis. In *CVPR*, 2020. 1, 2, 6, 7
- [27] Jogendra Nath Kundu, Siddharth Seth, MV Rahul, Mugalodi Rakesh, Venkatesh Babu Radhakrishnan, and Anirban Chakraborty. Kinematic-structure-preserved representation for unsupervised 3d human pose estimation. In *AAAI*, 2020. 6, 7
- [28] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *CVPR*, 2018. 3
- [29] Shichao Li, Lei Ke, Kevin Pratama, Yu-Wing Tai, Chi-Keung Tang, and Kwang-Ting Cheng. Cascaded deep monocular 3d human pose estimation with evolutionary training data. In *CVPR*, 2020. 2
- [30] Zhi Li, Xuan Wang, Fei Wang, and Peilin Jiang. On boosting single-frame 3d human pose estimation via monocular videos. In *ICCV*, 2019. 1, 2, 6
- [31] Ting Liu, Jianfeng Zhang, Xuecheng Nie, Yunchao Wei, Shikui Wei, Yao Zhao, and Jiashi Feng. Spatial-aware texture transformer for high-fidelity garment transfer. In *IEEE Trans. on Image Processing*, 2021. 1
- [32] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *CVPR*, 2017. 3
- [33] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *ICCV*, 2017. 1, 2, 3, 6
- [34] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel,

- Weipeng Xu, Dan Casas, and Christian Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Trans. on Graphics*, 36(4):44, 2017. 1, 2, 3, 6, 7
- [35] Buttner Michael. Motion matching-the road to next-gen animation. In *Nucl. ai*, 2015. 3
- [36] Rahul Mitra, Nitesh B Gundavarapu, Abhishek Sharma, and Arjun Jain. Multiview-consistent semi-supervised learning for 3d human pose estimation. In *CVPR*, 2020. 1, 2
- [37] Xuecheng Nie, Jianfeng Zhang, Shuicheng Yan, and Jiashi Feng. Single-stage multi-person pose machines. In *ICCV*, 2019. 1, 2
- [38] Dario Pavllo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *CVPR*, 2019. 1, 2, 3, 5, 6, 7
- [39] Dario Pavllo, David Grangier, and Michael Auli. Quaternet: A quaternion-based recurrent model for human motion. *arXiv*, 2018. 3
- [40] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. In *ACM Trans. on Graphics*, 2018. 3, 5
- [41] Xue Bin Peng, Angjoo Kanazawa, Jitendra Malik, Pieter Abbeel, and Sergey Levine. Sfv: Reinforcement learning of physical skills from videos. In *ACM Trans. on Graphics*, 2018. 3, 5
- [42] Xue Bin Peng and Michiel van de Panne. Learning locomotion skills using deeprl: Does the choice of action space matter? In *ACM*, 2017. 5
- [43] Davis Remppe, Leonidas J Guibas, Aaron Hertzmann, Bryan Russell, Ruben Villegas, and Jimei Yang. Contact and human dynamics from monocular video. In *ECCV*, 2020. 3
- [44] Helge Rhodin, Mathieu Salzmann, and Pascal Fua. Unsupervised geometry-aware representation for 3d human pose estimation. In *ECCV*, 2018. 6
- [45] Soshi Shimada, Vladislav Golyanik, Weipeng Xu, Patrick Pérez, and Christian Theobalt. Neural monocular 3d human motion capture with physical awareness. 2021. 3
- [46] Soshi Shimada, Vladislav Golyanik, Weipeng Xu, and Christian Theobalt. Physcap: Physically plausible monocular 3d motion capture in real time. In *ACM Trans. on Graphics*, 2020. 3
- [47] Chenyang Si, Wentao Chen, Wei Wang, Liang Wang, and Tieniu Tan. An attention enhanced graph convolutional lstm network for skeleton-based action recognition. In *CVPR*, 2019. 1
- [48] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *ECCV*, 2018. 1, 2, 3
- [49] Yu Sun, Yun Ye, Wu Liu, Wenpeng Gao, Yili Fu, and Tao Mei. Human mesh recovery from monocular images via a skeleton-disentangled representation. In *ICCV*, 2019. 6, 7
- [50] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *ECCV*, 2018. 6
- [51] Bastian Wandt, Marco Rudolph, Petrisa Zell, Helge Rhodin, and Bodo Rosenhahn. Canonpose: Self-supervised monocular 3d human pose estimation in the wild. In *CVPR*, 2021. 2, 6
- [52] Jue Wang, Shaoli Huang, Xinchao Wang, and Dacheng Tao. Not all parts are created equal: 3d pose estimation by modelling bi-directional dependencies of body parts. In *ICCV*, 2019. 2
- [53] Tao Wang, Jianfeng Zhang, Yujun Cai, Shuicheng Yan, and Jiashi Feng. Direct multi-view multi-person 3d pose estimation. In *NeurIPS*, 2021. 2
- [54] Zhe Wang, Daeyun Shin, and Charless C Fowlkes. Predicting camera viewpoint improves cross-dataset generalization for 3d human pose estimation. In *ECCVw*, 2020. 7
- [55] Zhenyi Wang, Ping Yu, Yang Zhao, Ruiyi Zhang, Yufan Zhou, Junsong Yuan, and Changyou Chen. Learning diverse stochastic human-action generators by learning smooth latent transitions. In *AAAI*, 2020. 3
- [56] Jiajun Wu, Tianfan Xue, Joseph J Lim, Yuandong Tian, Joshua B Tenenbaum, Antonio Torralba, and William T Freeman. Single image 3d interpreter network. In *ECCV*, 2016. 6
- [57] Kevin Xie, Tingwu Wang, Umar Iqbal, Yunrong Guo, Sanja Fidler, and Florian Shkurti. Physics-based human motion estimation and synthesis from videos. In *CVPR*, 2021. 3
- [58] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018. 1
- [59] Yiding Yang, Jiayan Qiu, Mingli Song, Dacheng Tao, and Xinchao Wang. Distilling knowledge from graph convolutional networks. In *CVPR*, 2020. 2
- [60] Yiding Yang, Zhou Ren, Haoxiang Li, Chunlun Zhou, Xinchao Wang, and Gang Hua. Learning dynamics via graph neural networks for human pose estimation and tracking. In *CVPR*, 2021. 2
- [61] Zhenbo Yu, Bingbing Ni, Jingwei Xu, Junjie Wang, Chenglong Zhao, and Wenjun Zhang. Towards alleviating the modeling ambiguity of unsupervised monocular 3d human pose estimation. In *ICCV*, 2021. 1, 2, 3, 6, 7
- [62] Ye Yuan and Kris Kitani. Dlow: Diversifying latent flows for diverse human motion prediction. In *ECCV*, 2020. 3
- [63] Ye Yuan and Kris Kitani. Residual force control for agile human behavior imitation and extended motion synthesis. In *NeurIPS*, 2020. 3, 5
- [64] Ye Yuan, Shih-En Wei, Tomas Simon, Kris Kitani, and Jason Saragih. Simpoe: Simulated character control for 3d human pose estimation. In *CVPR*, 2021. 3, 5
- [65] Jianfeng Zhang, Xuecheng Nie, and Jiashi Feng. Inference stage optimization for cross-scenario 3d human pose estimation. In *NeurIPS*, 2020. 2
- [66] Jianfeng Zhang, Dongdong Yu, Jun Hao Liew, Xuecheng Nie, and Jiashi Feng. Body meshes as points. In *CVPR*, 2021. 2, 7
- [67] Yan Zhang, Michael J Black, and Siyu Tang. We are more than our joints: Predicting how 3d bodies move. *CVPR*, 2021. 3

- [68] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *ICCV*, 2017. [1](#), [2](#)