Motion In-betweening for Physically Simulated Characters

Deepak Gopinath[1], Hanbyul Joo[2], Jungdam Won[1]

[1]Meta AI, [2]Seoul National University

SIGGRAPH ASIA 2022 DAEGU

Abstract

We present a motion in-betweening framework to generate high quality, physically plausible character animation when we are given temporally sparse keyframes as soft animation constraints. More specifically, we learn imitation policies for physically simulated characters by using deep reinforcement learning where the policies can access limited information only. Once learned, the physically simulated characters are capable of adapting to external perturbations while following given sparse input keyframes. We demonstrate the performance of our framework on two different motion datasets and also compare our results with the the results generated by a baseline imitation policy.

Background

Motion in-betweening is a popular method to create skeletal animations, where users (artists) provide keyframe poses with less temporal granularity and the system automatically generates intermediate poses with finer granularity.

We demonstrate a new approach to solve motion in-betweening given sparsely keyframed poses using physically simulated characters for which control policies are trained with deep reinforcement learning (RL)

Using physically simulated characters provides several unique advantages over existing kinematics-based approaches.

- ✓ We can generate physically plausible motions even if bad input poses are given.
- ✓ Motions adapting to external perturbation can emerge during motion generation.

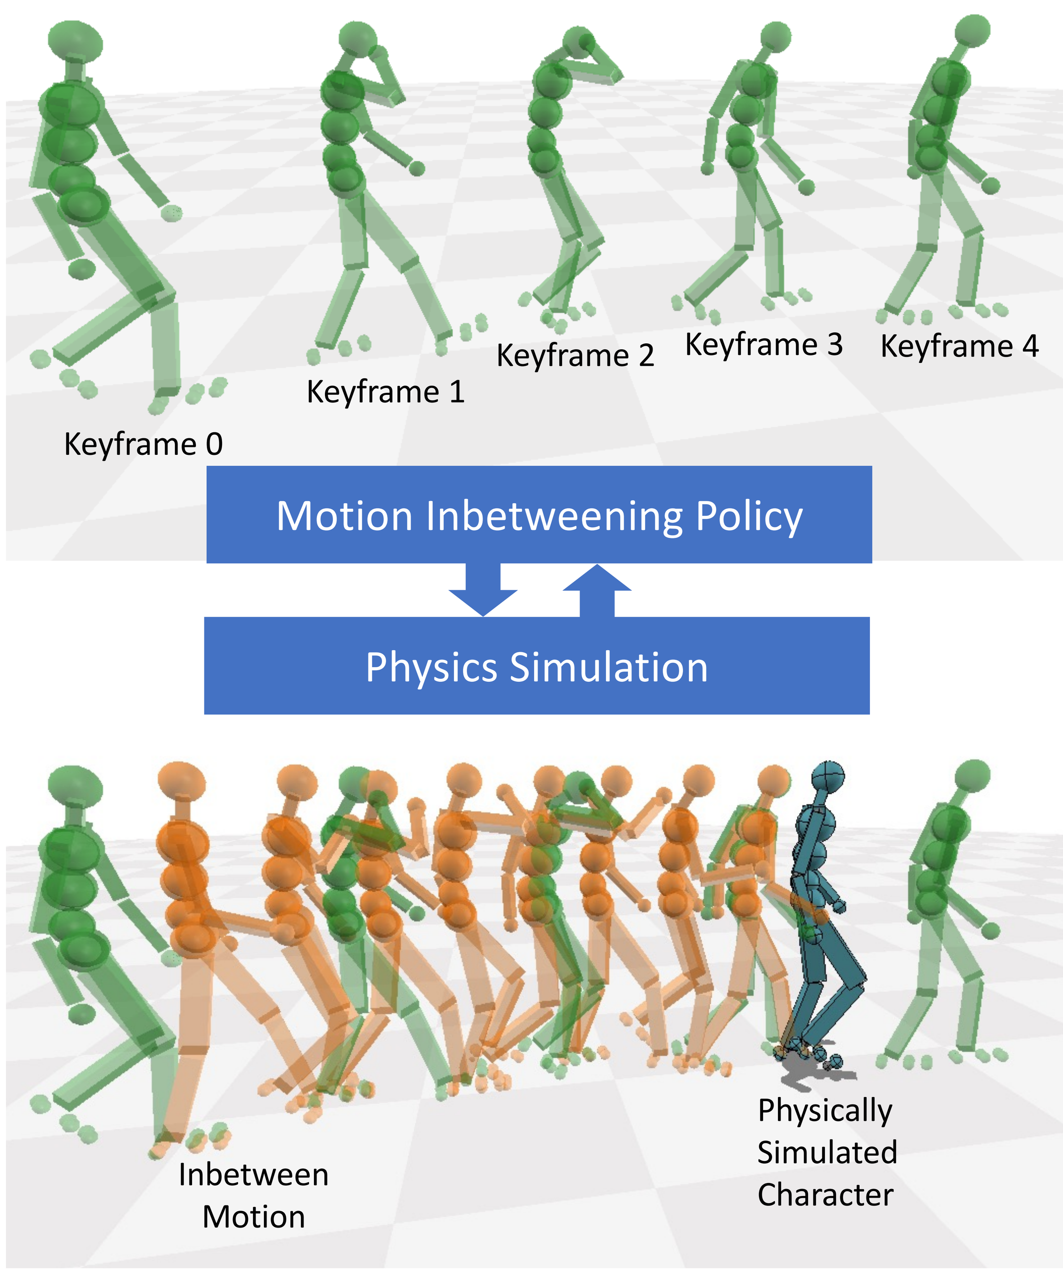

Fig 1. An overview of our system

Method

Our system takes a sequence of keyframe poses $(P_0^{key}, P_k^{key}, P_{2k}^{key}, \cdots)$ with a fixed coarse time interval as input ($k = 1$ in our model), then outputs a motion $(P_0, P_m, \cdots, P_{k-m}, P_k, \cdots, P_{2k-m}, P_{2k}, \cdots)$ at a desired dense time interval ($m = \frac{1}{30}$ in our model). Our framework learns an RL based imitation policy.

The state $s_t = (s_t^{key}, s_t^{sim})$ i.e., the keyframe state $s_t^{key} = (P_t^{key}, P_{t+1}^{key}, F_{offset}^{key}, t_{tta})$ and the simulated state $s_t^{sim} = (S_{body}^{sim}, F_{offset}^{sim})$

- $(P_t^{key}, P_{t+1}^{key})$: current and next input keyframes
- F_{offset}^{key} : relative facing frame of next keyframe wrt current keyframe
- t_{tta} : time-to-arrival feature (time remaining to reach the next keyframe and temporal embedding of the current time step [1]
- S_{body}^{sim} : dynamic state of the simulated character [2]
- F_{offset}^{sim} : relative facing frame of the simulated character wrt the current keyframe.

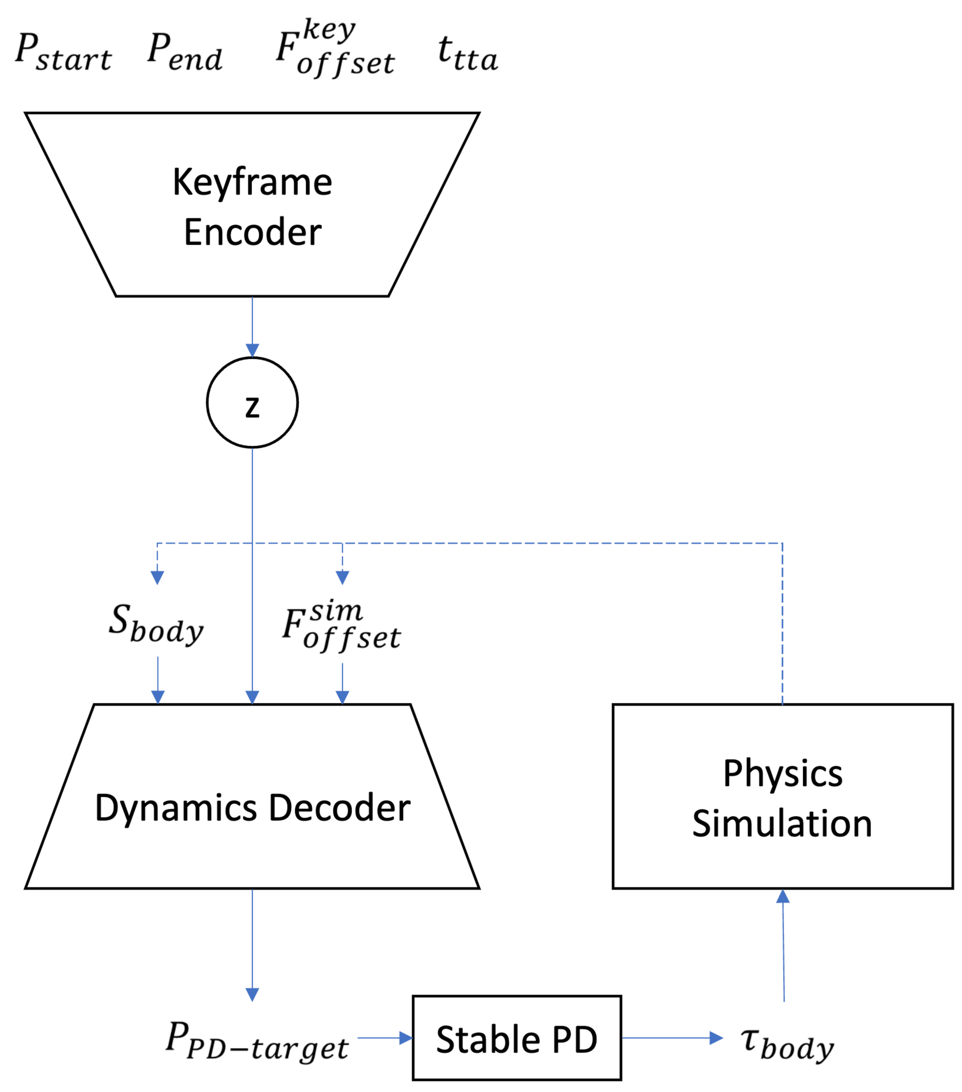

Fig 2. A deep RL control policy for motion in-betweening.

The action a_t in RL is a target pose for the stable PD controller, which computes joint torques τ_{body} to actuate the simulated character. Physics simulation then computes the next state s_{t+1}.
We use the same multiplicative reward function r_t used in ScaDiver [3].

Fig 2 is the encoder-decoder structure that we adopt for our control policy. We employ the keyframe encoder to produce a reduced vector representation z of the in-between pose for the current time step. The output of the encoder is concatenated with the simulation state and fed to the dynamics decoder to produce an action.

Results

Fig 1 shows a snapshot of intermediate postures generated by our model trained with LaFAN1 dataset, where the learned control policy can successfully match input sparse keyframes while generating motions that are physically plausible and resemble ground-truth motions.

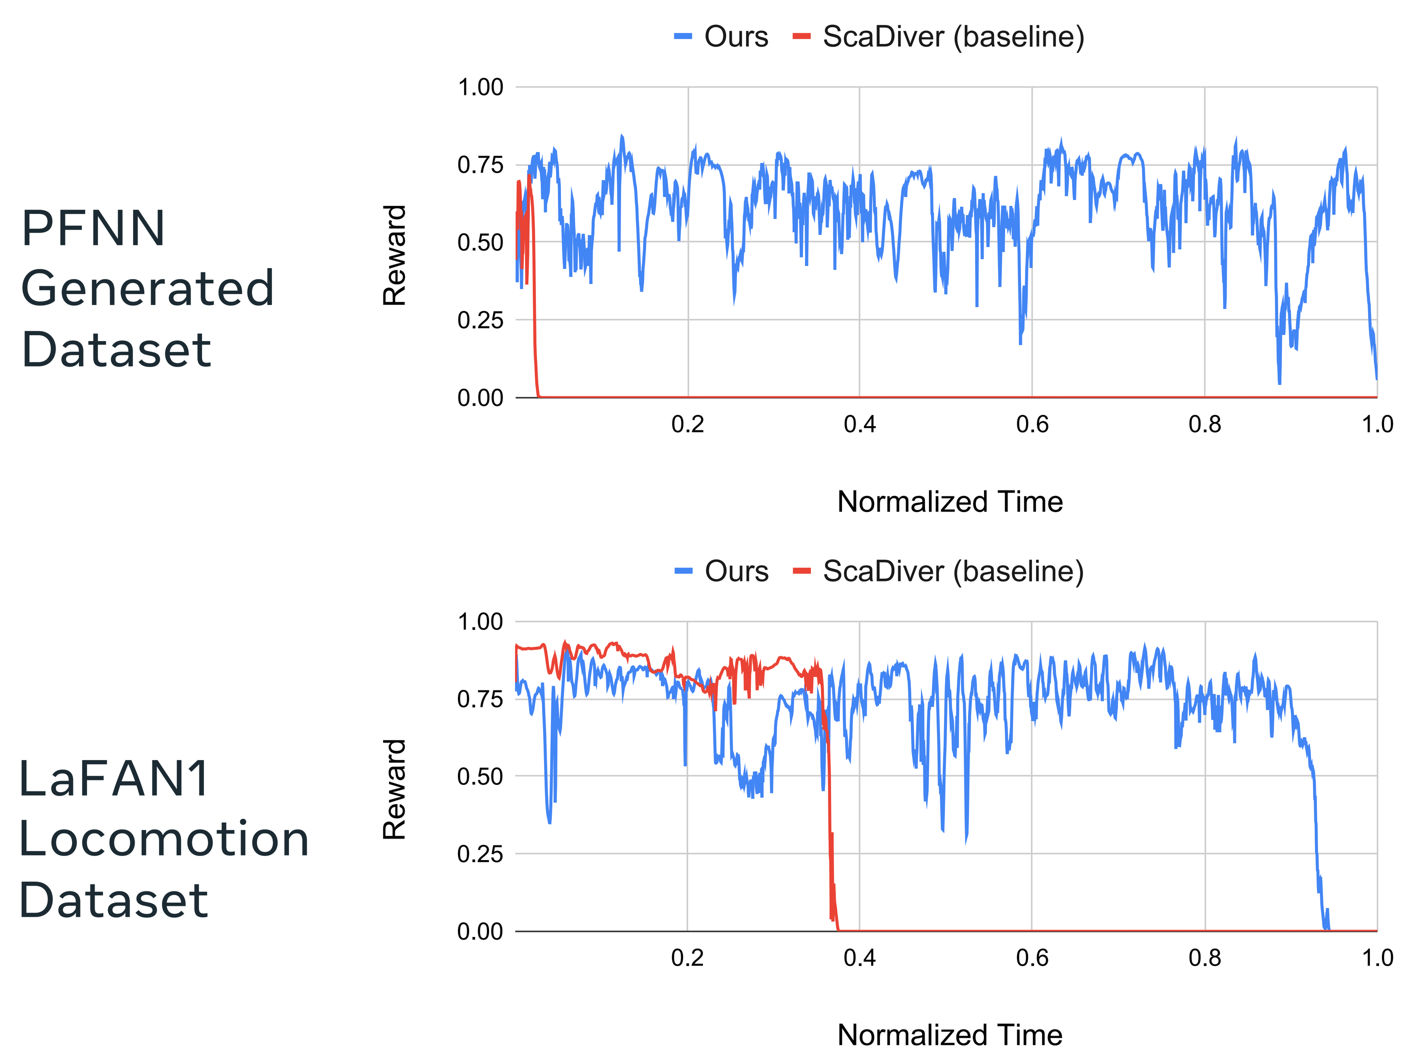

Fig 3. Imitation reward for each frame of in-betweening of keyframes from unseen motions

We trained a baseline imitation policy based on [3], where the policy consumes the densely specified keyframe postures (i.e., future reference motion). Since such information is not available during test time, we provided pseudo-reference motion by linearly interpolating the input sparse keyframe postures. Fig 3 shows the performance comparison over unseen motions (in-distribution). Our policy outperforms the baseline, that suffers from state mismatch between training/test time, by a large margin.

References

[1] Félix G. Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. 2020. Robust Motion In-Betweening. 39, 4 (2020).

[2] Daniel Holden, Taku Komura, and Jun Saito. 2017. Phase-functioned Neural Networks for Character Control. ACM Trans. Graph. 36, 4, Article 42 (July 2017).

[3] Jungdam Won, Deepak Gopinath, and Jessica Hodgins. 2020. A Scalable Approach to Control Diverse Behaviors for Physically Simulated Characters. ACM Trans. Graph. 39, 4, Article 33 (2020).

Take a photo to view supplemental video!