DEEP LEARNING FOR LOCOMOTION ANIMATION

Gavriel State, Senior Director, Simulation & AI

GTC 2019

March 19, 2019

Deep Learning Animation: PFNN

- Breakthrough 2017 paper on using motion capture + DL to drive locomotion animation

Phase-Functioned Neural Networks for Character Control

DANIEL HOLDEN, University of Edinburgh
TAKU KOMURA, University of Edinburgh
JUN SAITO, Method Studios

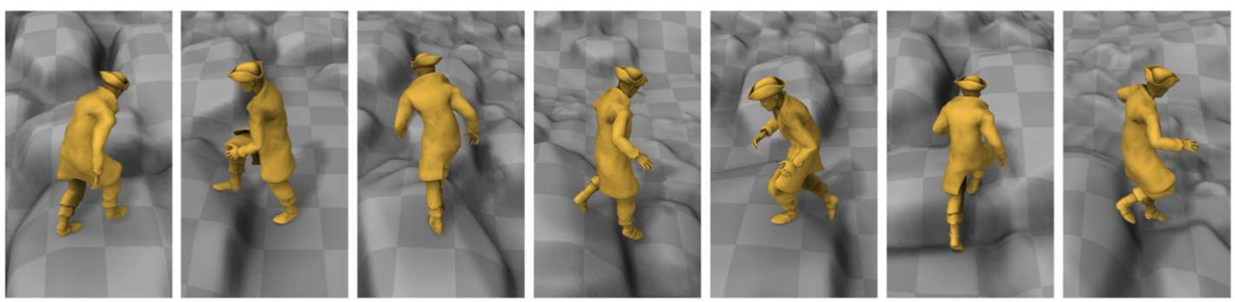

- http://theorangeduck.com/page/phase-functioned-neural-networks-character-control

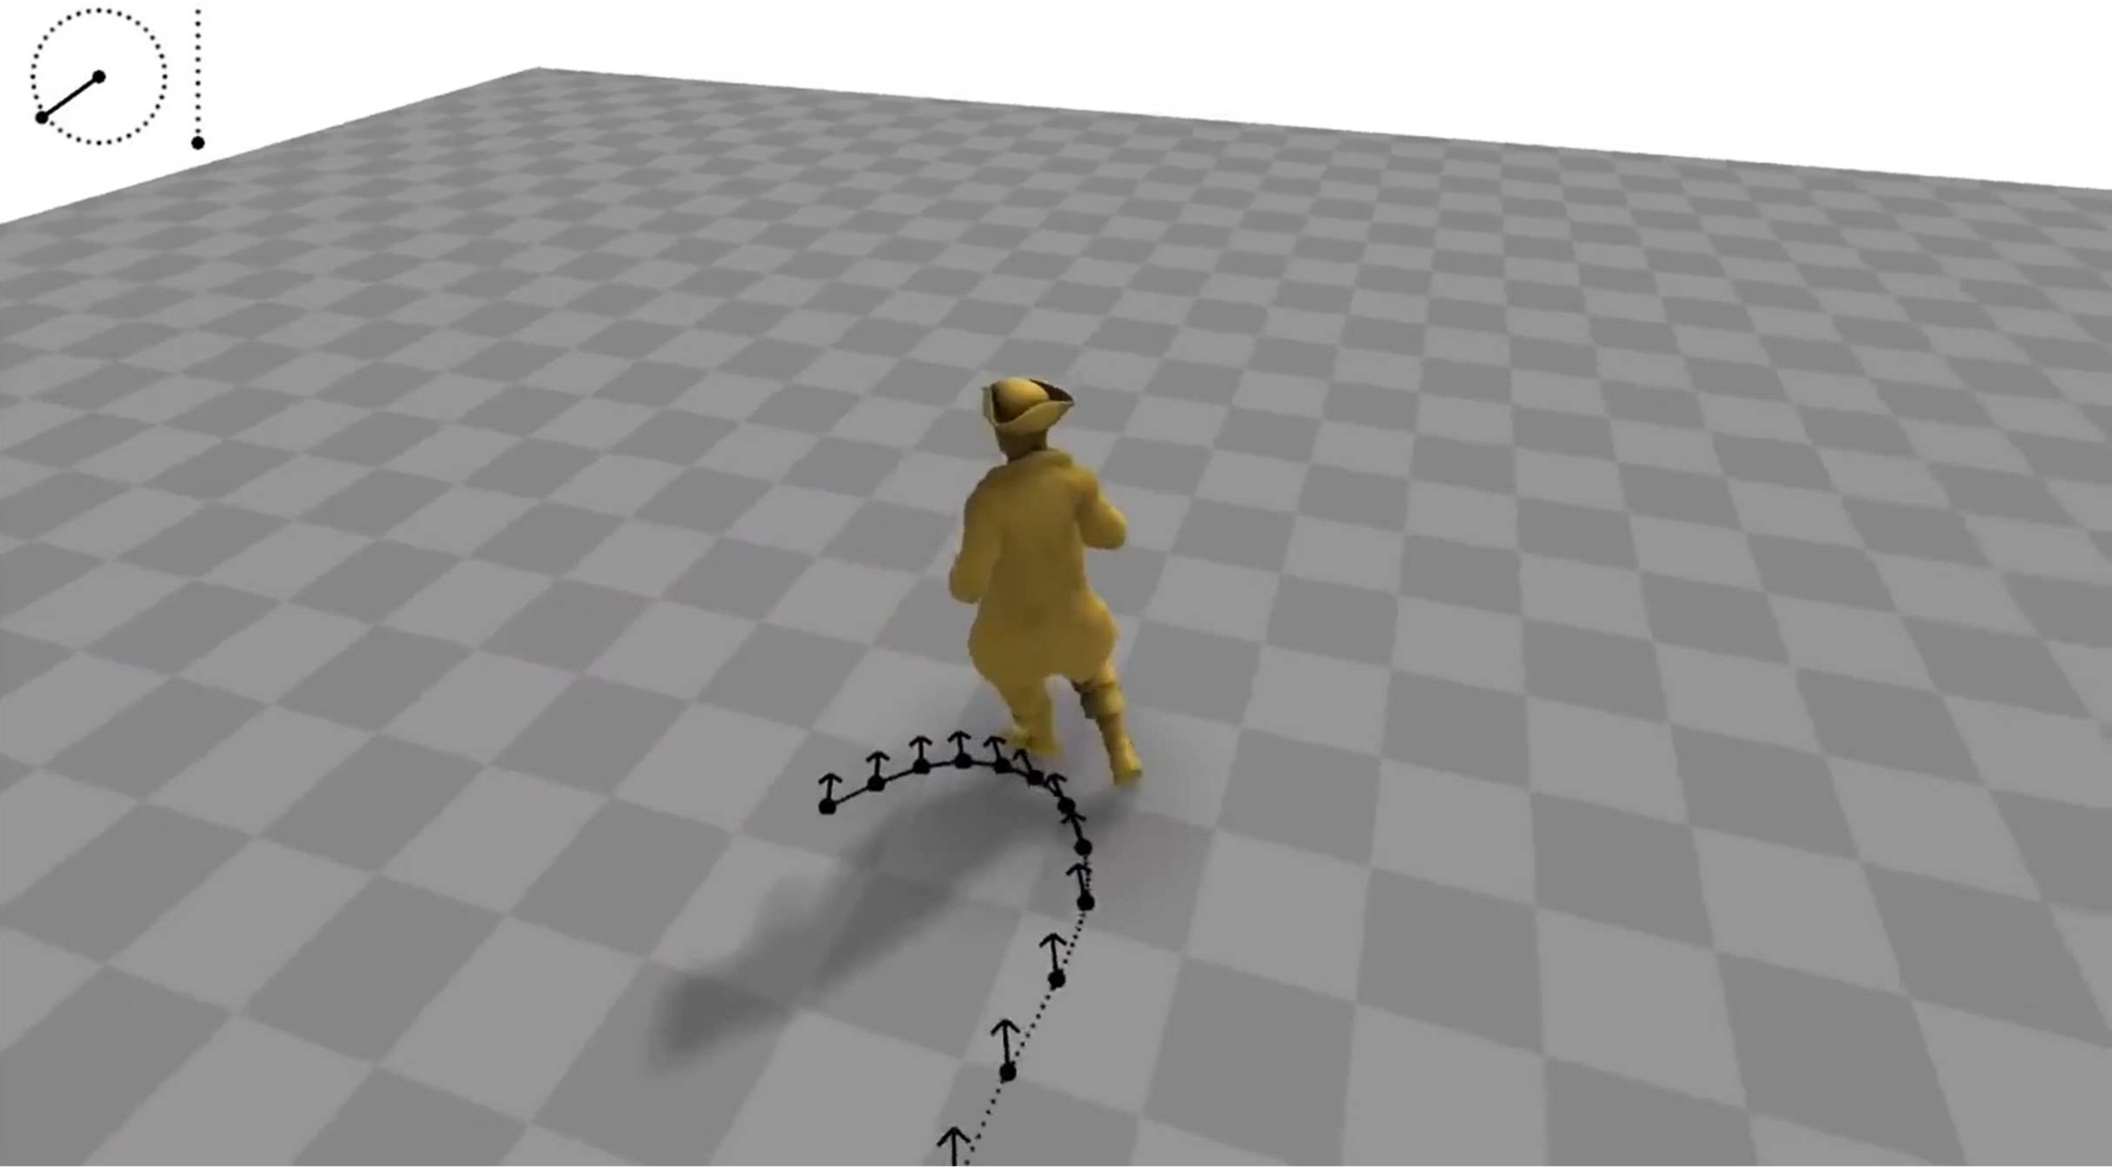

Applications

Games

Applications

VFX Crowd Simulation

Applications

Human/Robot interaction safety

Mimus, Madeline Gannon / ATONATON (2016)

Applications

Holodeck - Before

Applications

Holodeck - after

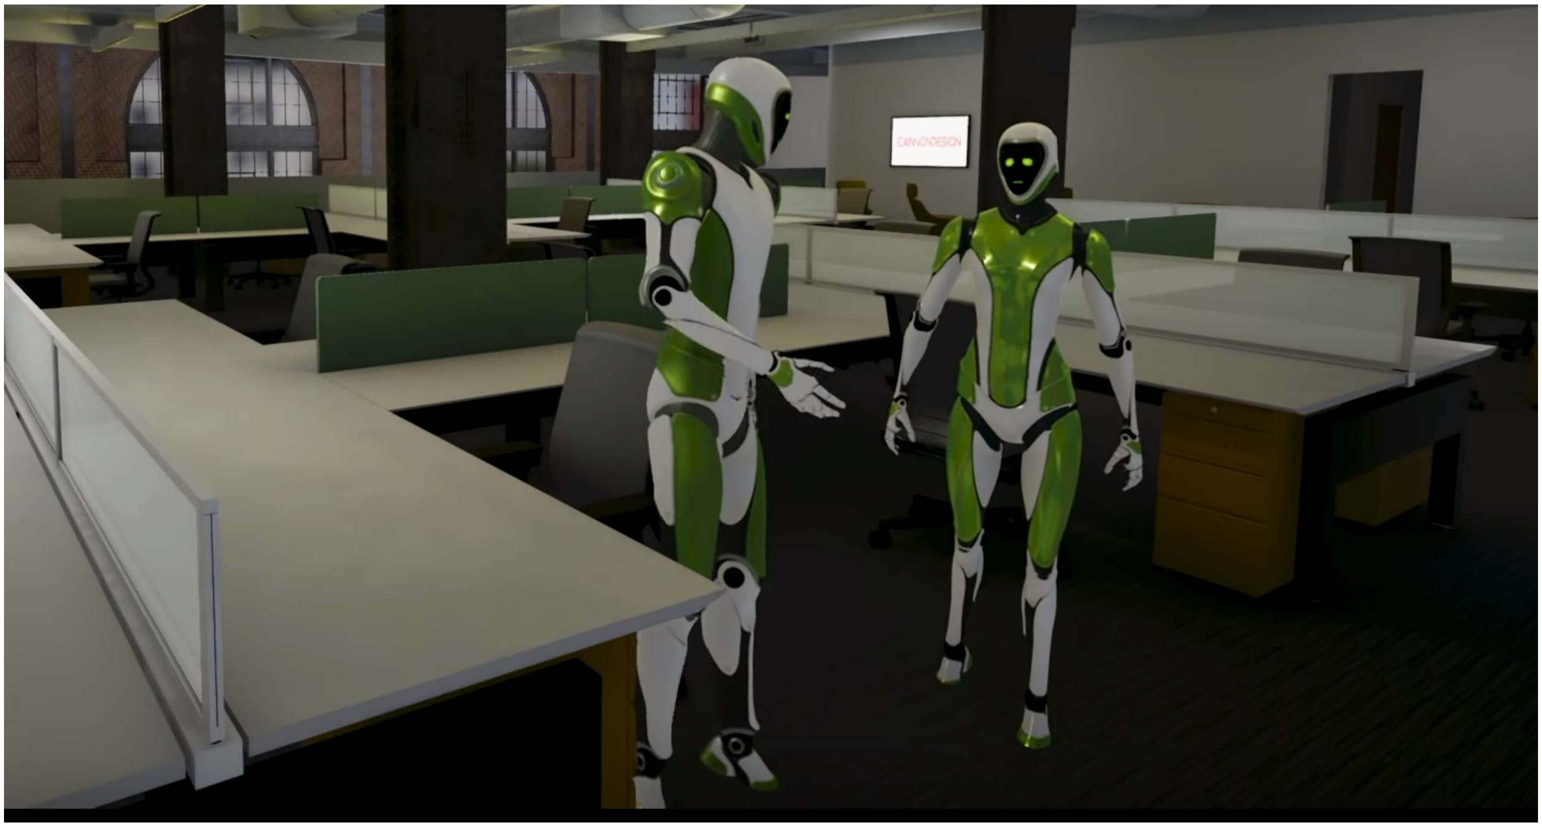

Applications

Auto Simulation

Image from the SYNTHIA dataset

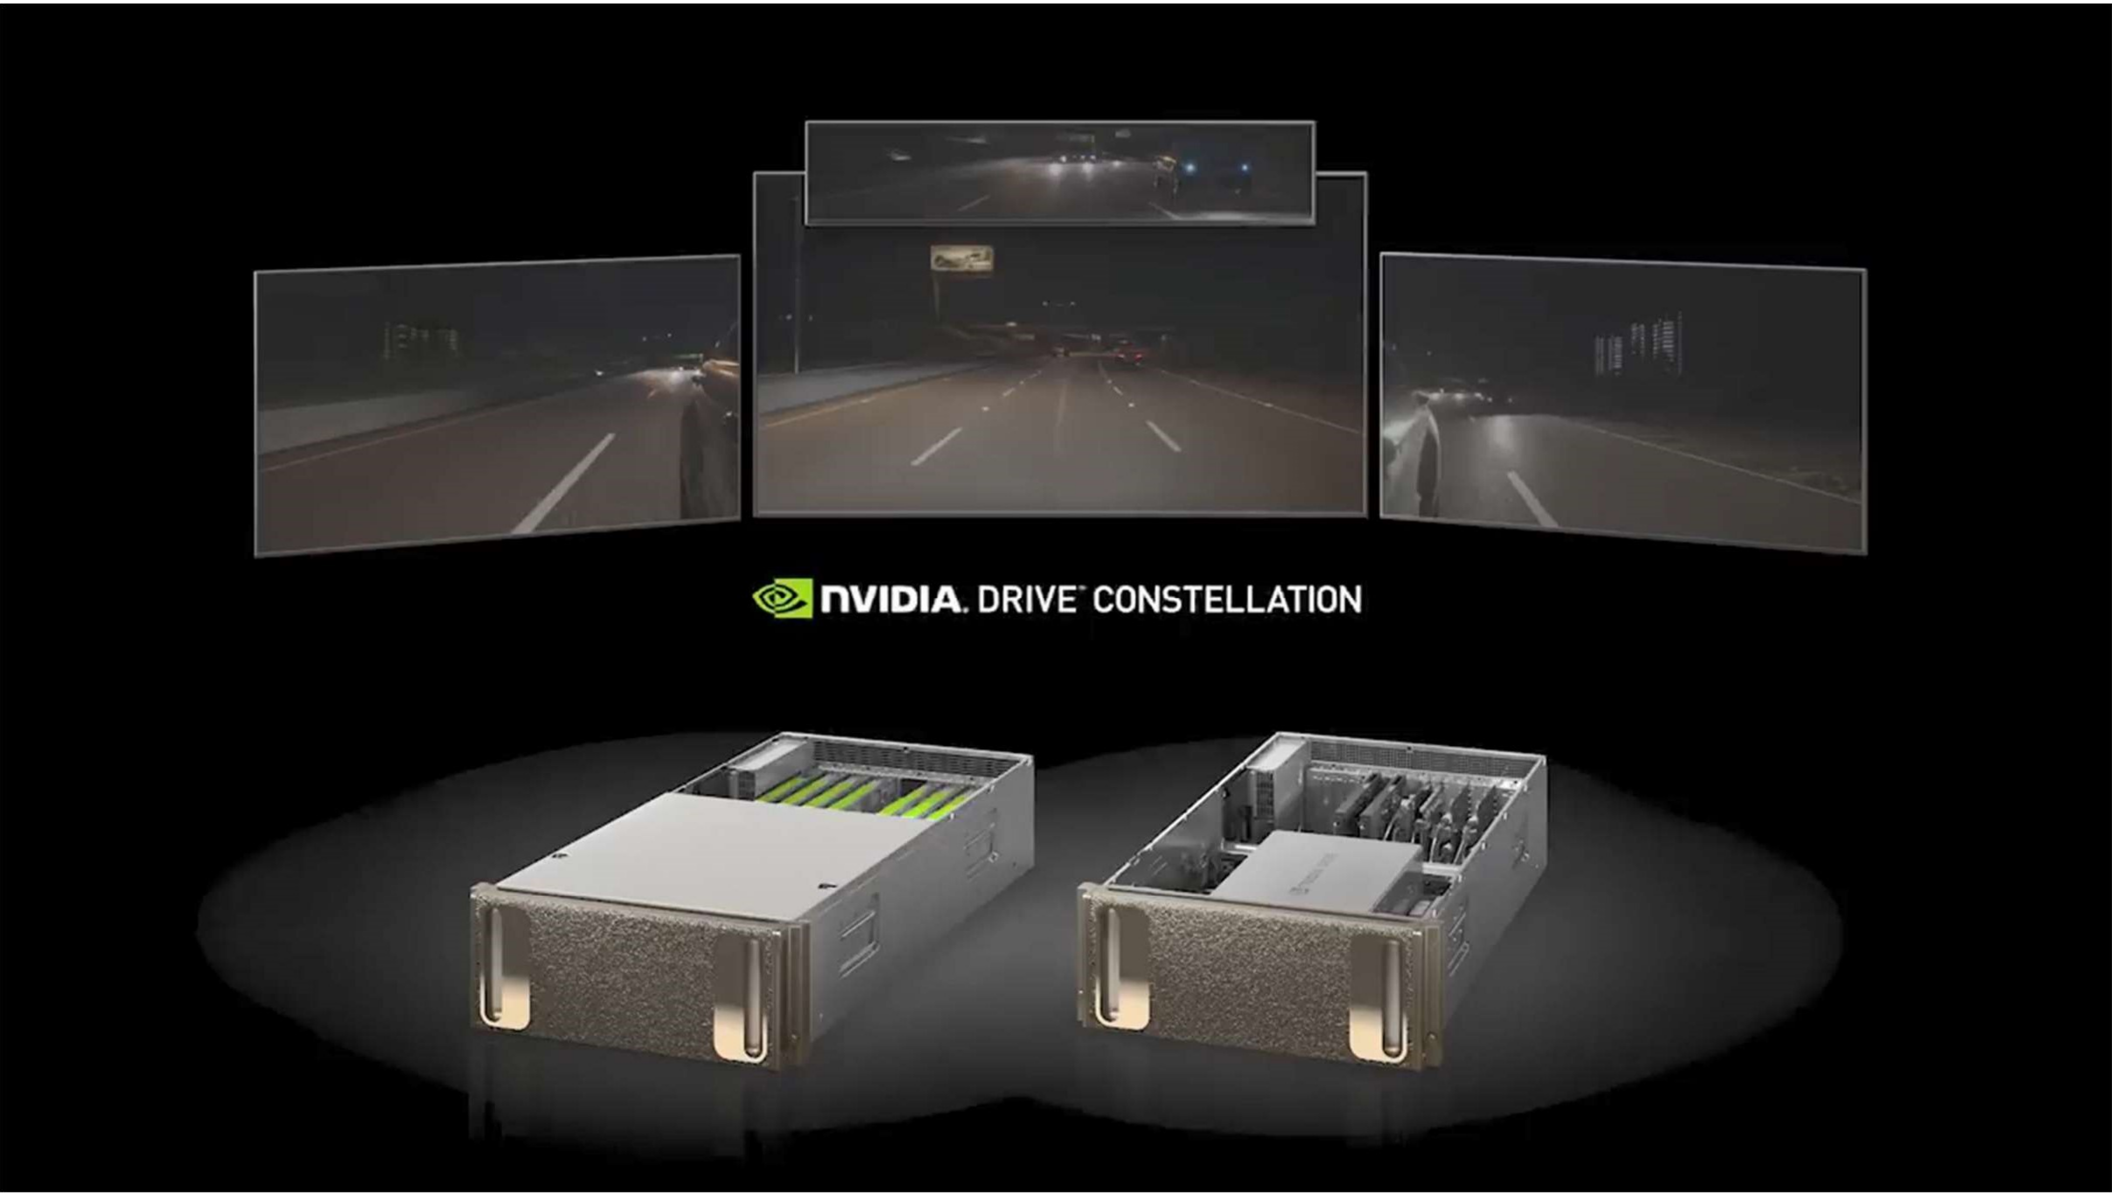
NVIDIA DRIVE CONSTELLATION

PFNN: How does it work?

Motion Capture

- Gather Motion Capture data

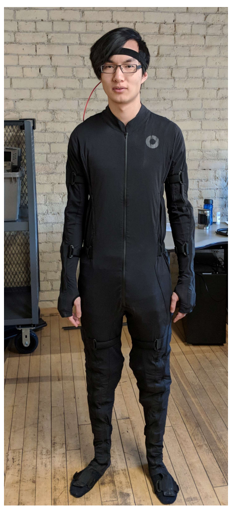

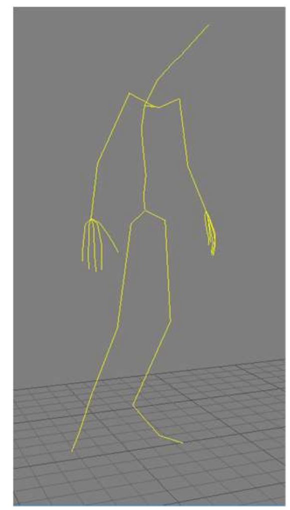

- Lots of free data available from CMU: http://mocap.cs.cmu.edu/
- Many thanks to Fox VFX Lab for our capture above

How does it work?

Metadata labeling

- Additional data needed:
  - Gait (running, walking, crouching, etc)
  - Phase – what point of the walk cycle are we in
  - Footstep positions

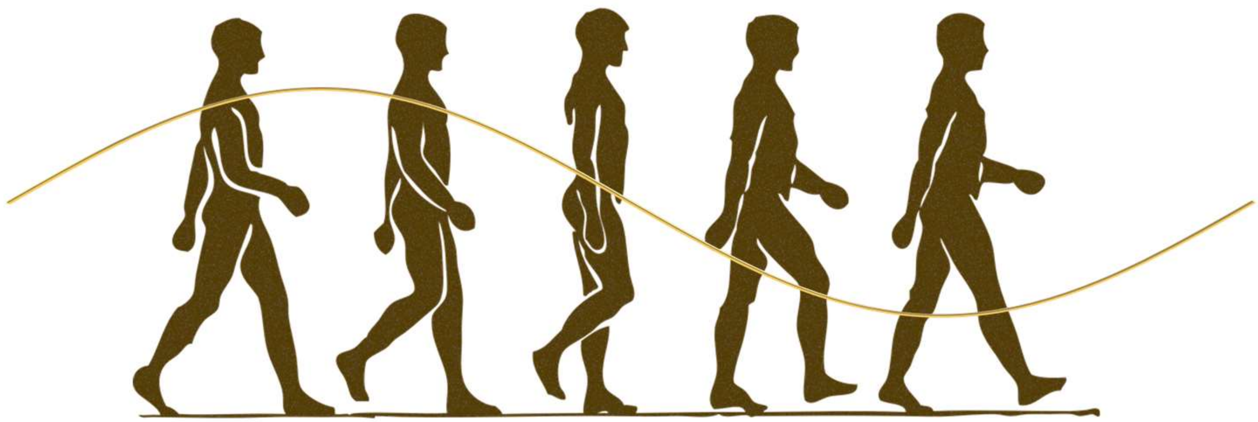

How does it work?

Terrain Fitting

- Generate many different height fields that can fit a given set of character positions
- More robust than just capturing the actual height field, since it gives the network more potential data to fit with

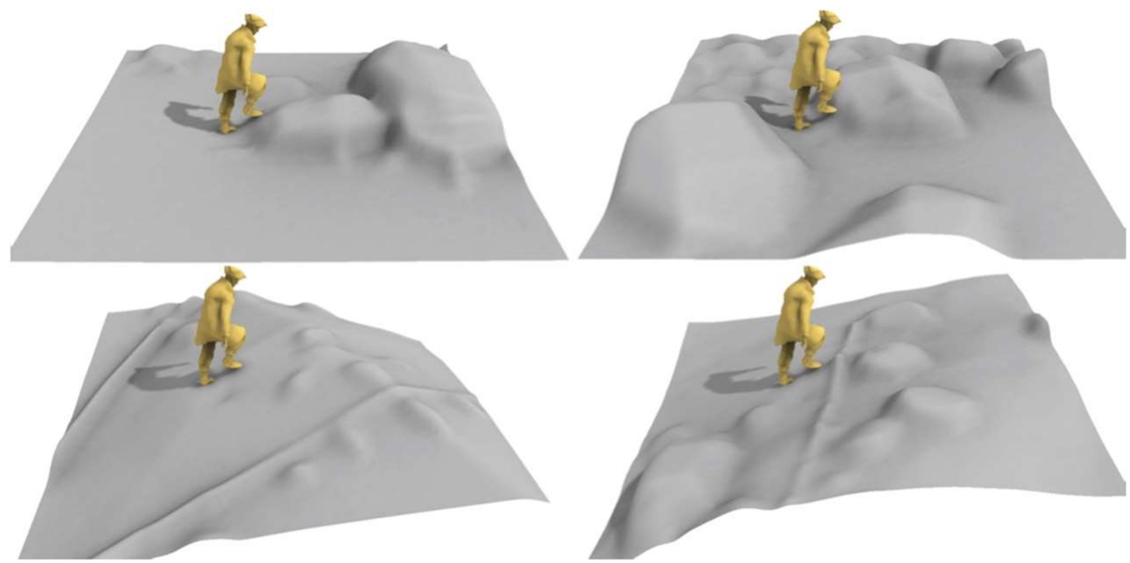

How does it work?

Phase Functioned Neural Network

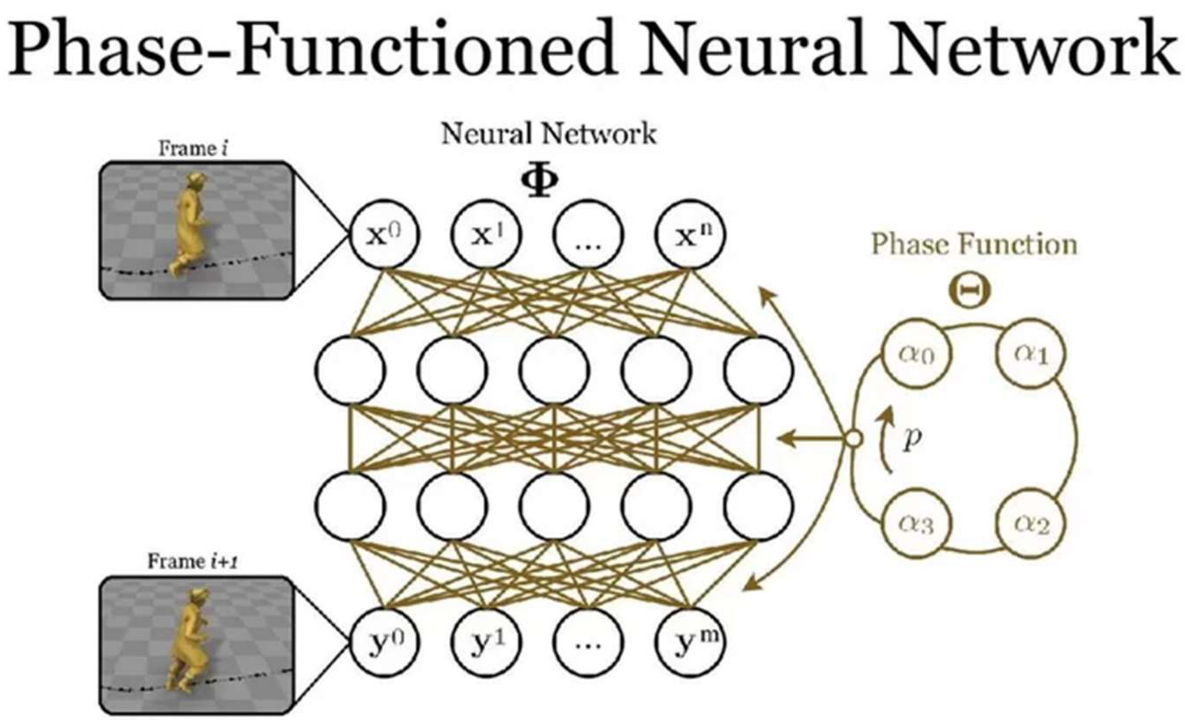

- Weights in the network are different depending on the phase parameter
- Four sets of weights trained
- Mid-cycle weights calculated by spline interpolation or precomputed (requires custom inferencing code or lots of memory)

How does it work?

Runtime Inferencing

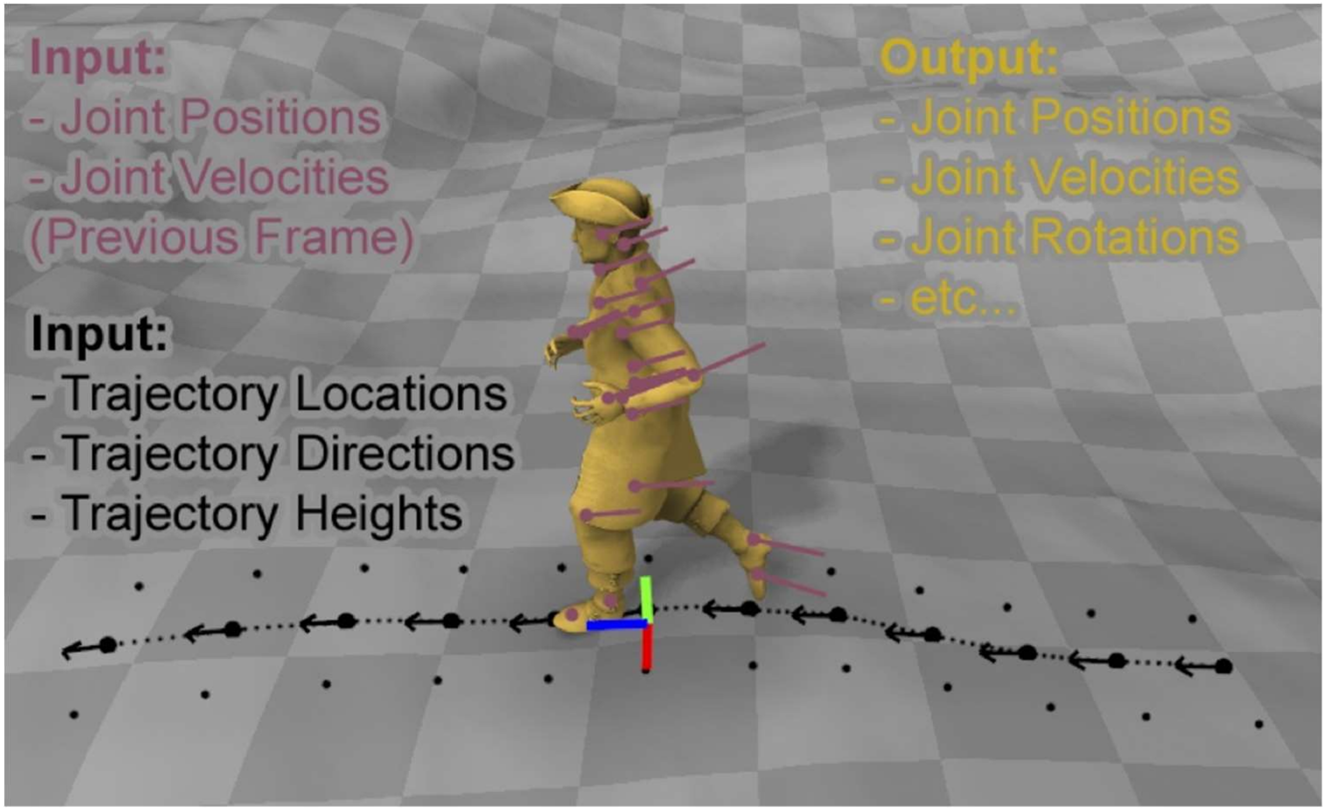


Recording has started

PFNN On GPU

LIVE DEMO: SpaceShip Down

Applications in Robotics

PFNN + Navigation


IsaacSimProject Game Preview Standalone (64-bit/PCD3D_SM5)
Recording has started

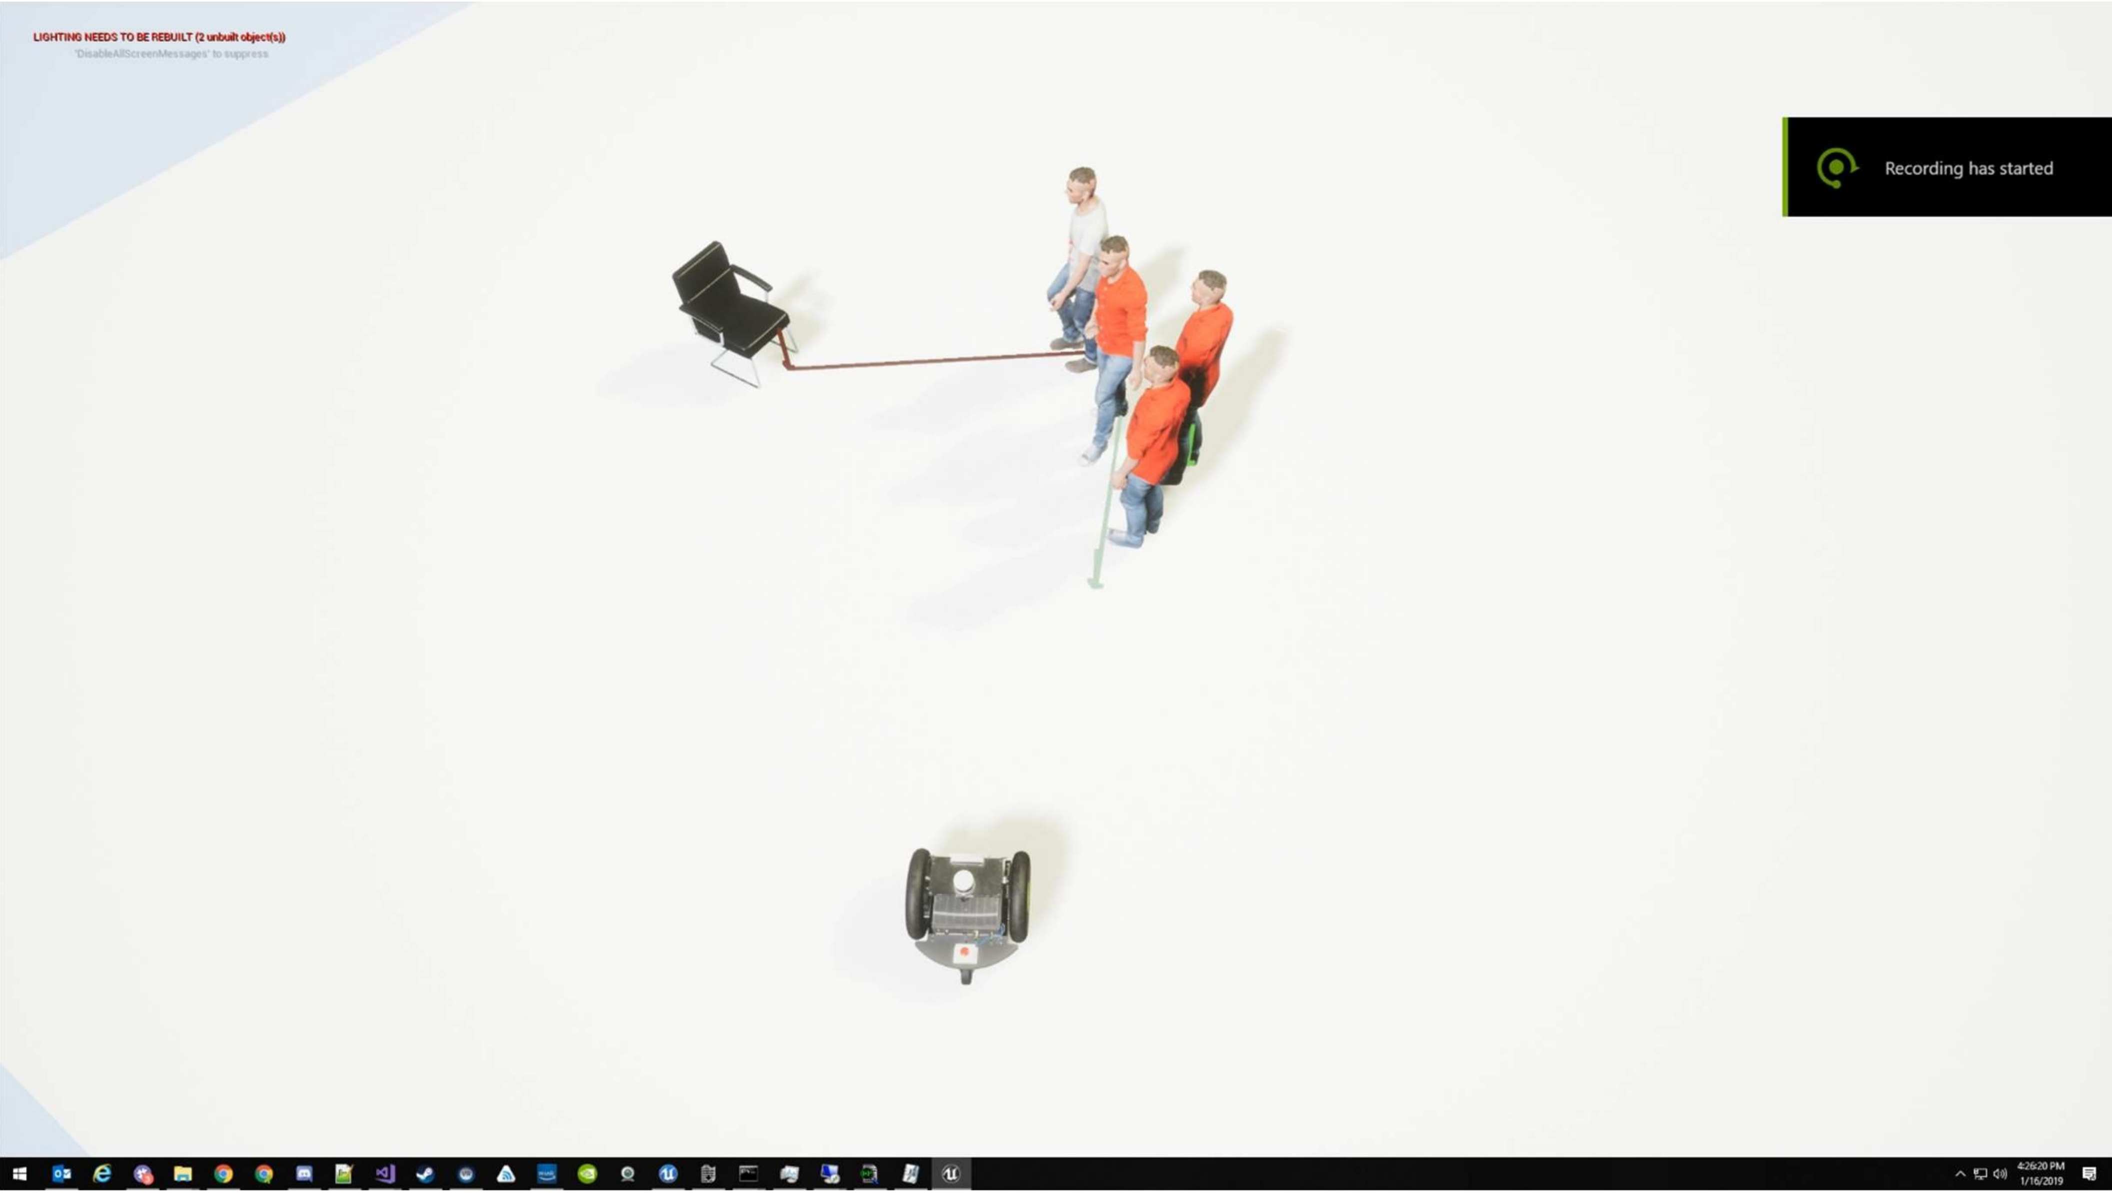
LIGHTING NEEDS TO BE REBUILT (2 unbuilt object(s))
'DisableAllScreenMessages' to suppress
Recording has started
4:26:20 PM
1/16/2019

Mode Adaptive Neural Network

Quadruped Motion Control

HE ZHANG[†], University of Edinburgh
SEBASTIAN STARKE[†], University of Edinburgh
TAKU KOMURA, University of Edinburgh
JUN SAITO, Adobe Research

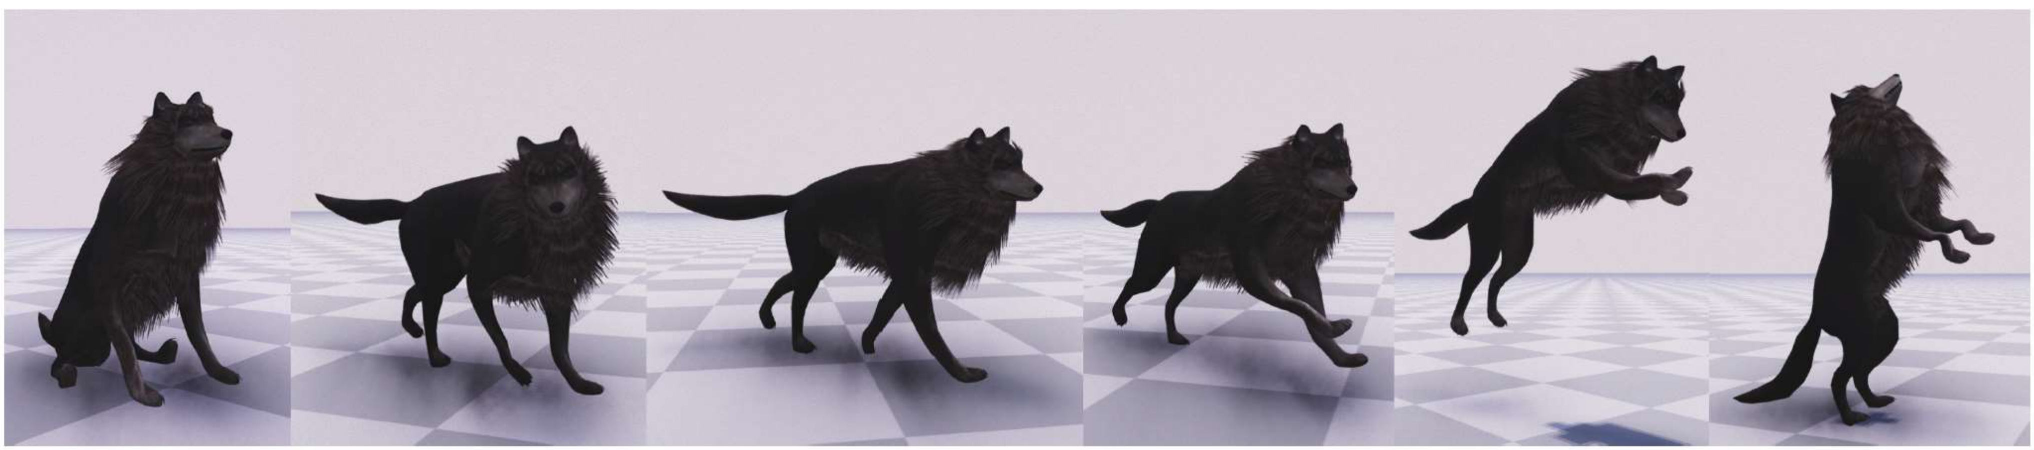


Mode-Adaptive Neural Networks for
Quadruped Motion Control
- SIGGRAPH 2018, Vancouver, Canada -
He Zhang*
Sebastian Starke*
Taku Komura
Jun Saito
*Joint First Authors
THE UNIVERSITY of EDINBURGH
Adobe

MANN for Bipeds - Visualization

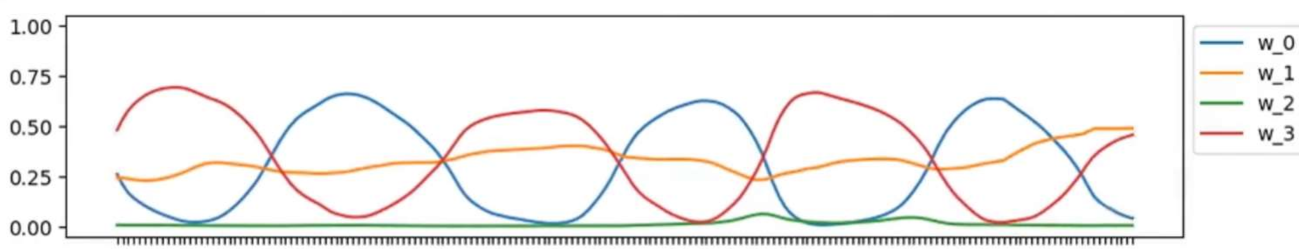

MANN for Bipeds - Hard

No phase information

Standard PFNN

MANN (8 experts)

NVIDIA Improved Biped MANN

What’s Wrong With This Picture?

What's Wrong With This Picture?

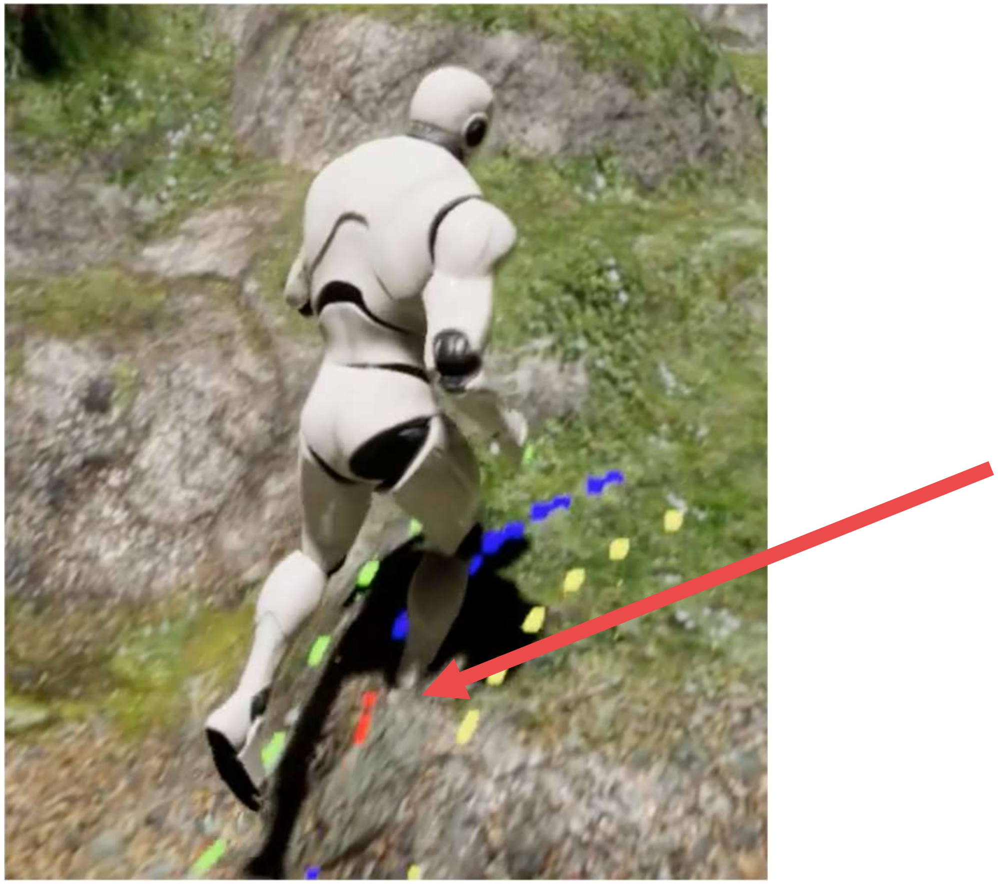

Inverse Kinematics

Traditional Gaming Approach

- Quick and dirty solution
  when encountering obstacles

Inverse Kinematics

Traditional Gaming Approach

- Quick and dirty solution when encountering obstacles
- Just adjust skeleton backwards from intersections
  - First the lower leg

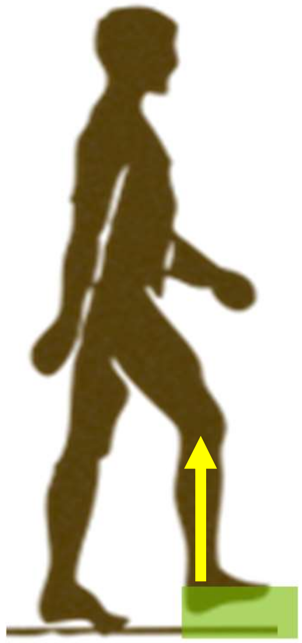

Inverse Kinematics

Traditional Gaming Approach

- Quick and dirty solution when encountering obstacles
- Just adjust skeleton backwards from intersections
    - First the lower leg
    - Then the thigh

Inverse Kinematics

Traditional Gaming Approach

- Quick and dirty solution when encountering obstacles
- Just adjust skeleton backwards from intersections
    - First the lower leg
    - Then the thigh
- Many limitations however

Physics!

The Real Solution

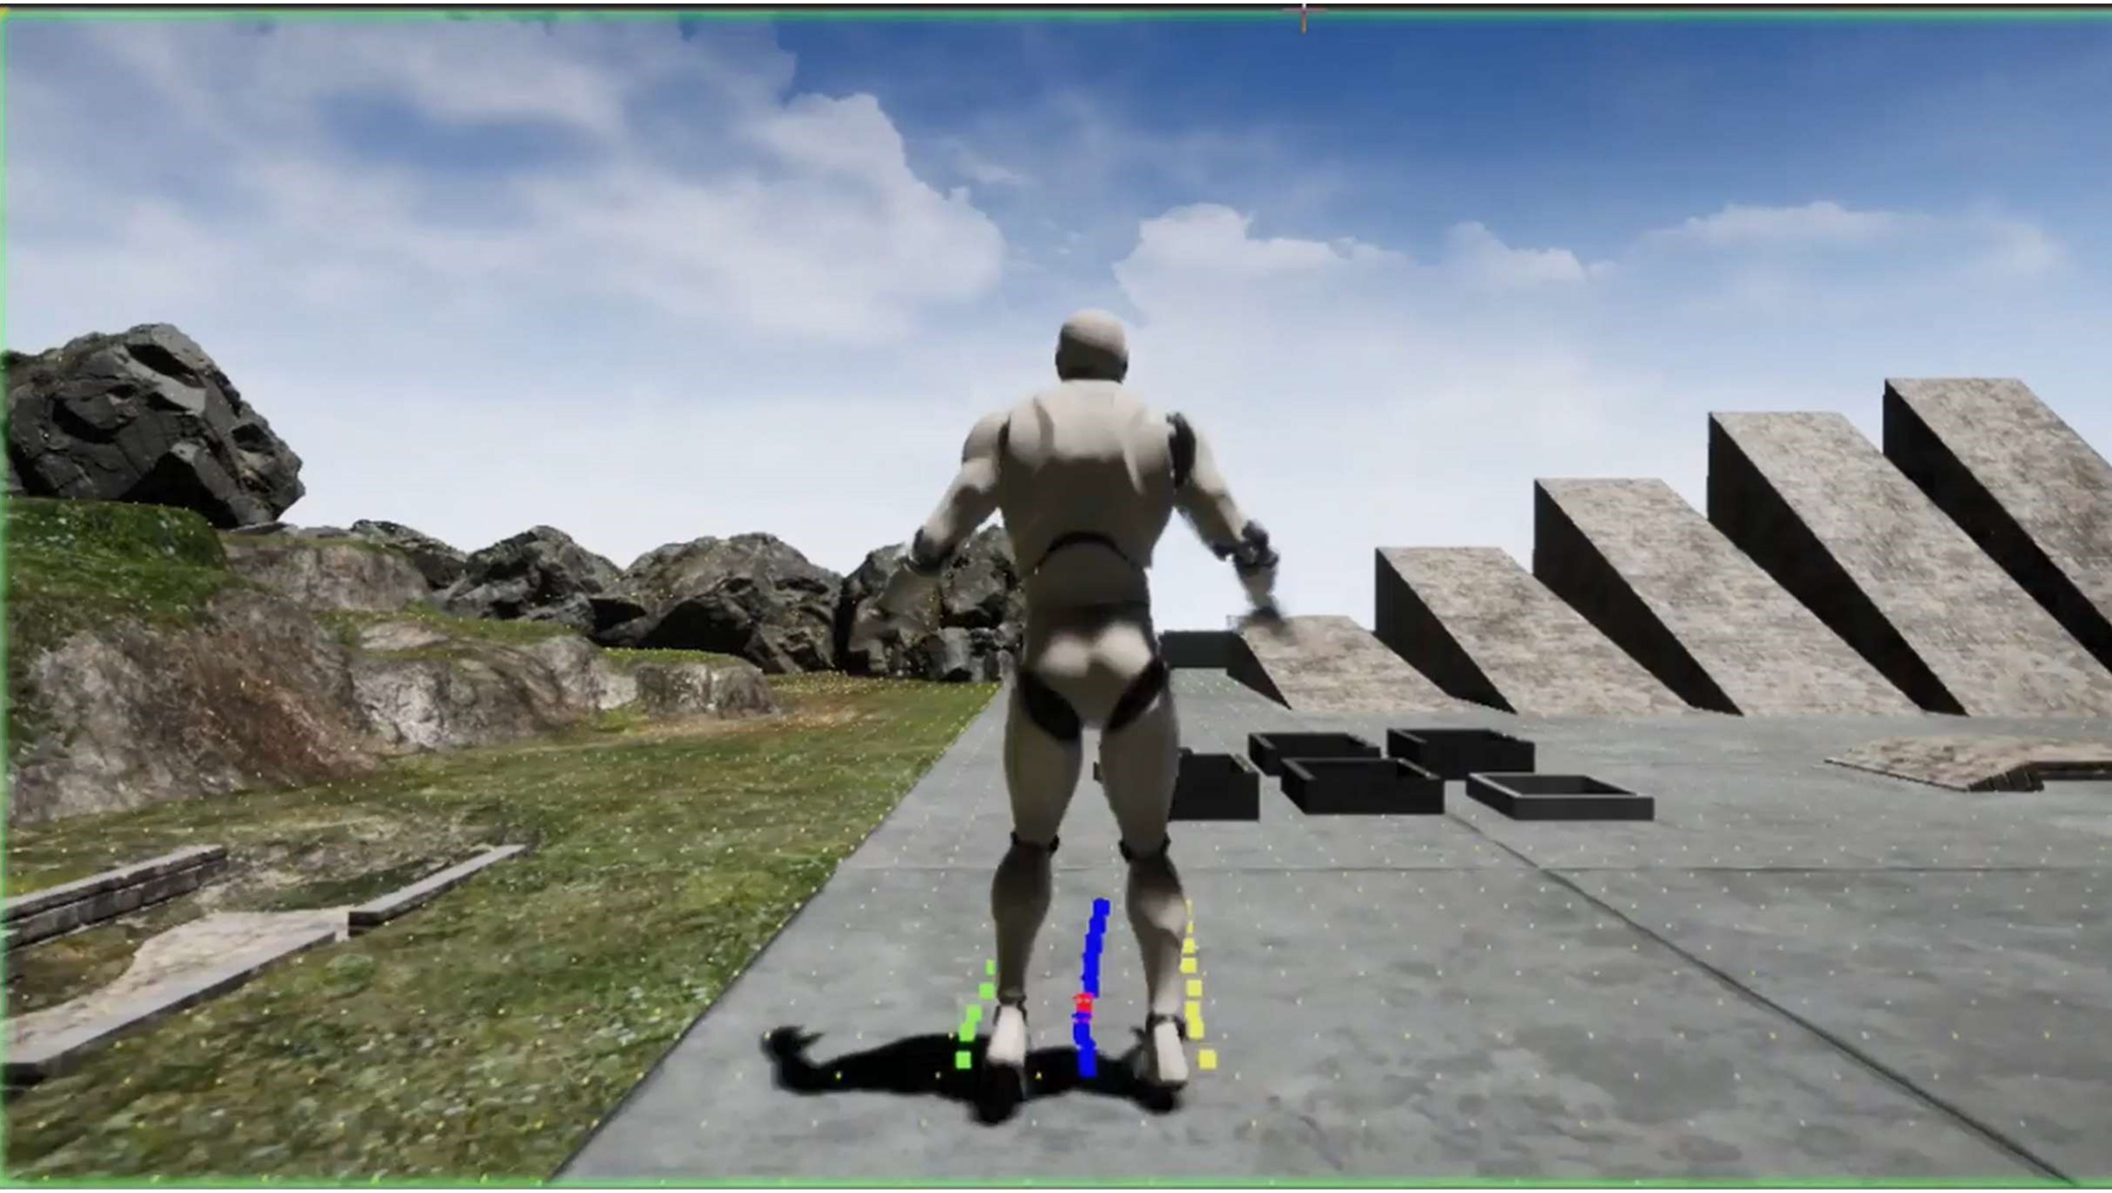

DeepLoco: Physics + RL

- Another major recent work adds physics and high level control:

DeepLoco: Dynamic Locomotion Skills Using Hierarchical Deep Reinforcement Learning

Xue Bin Peng (1) Glen Berseth (1) KangKang Yin (2) Michiel van de Panne (1)
(1)University of British Columbia
(2)National University of Singapore

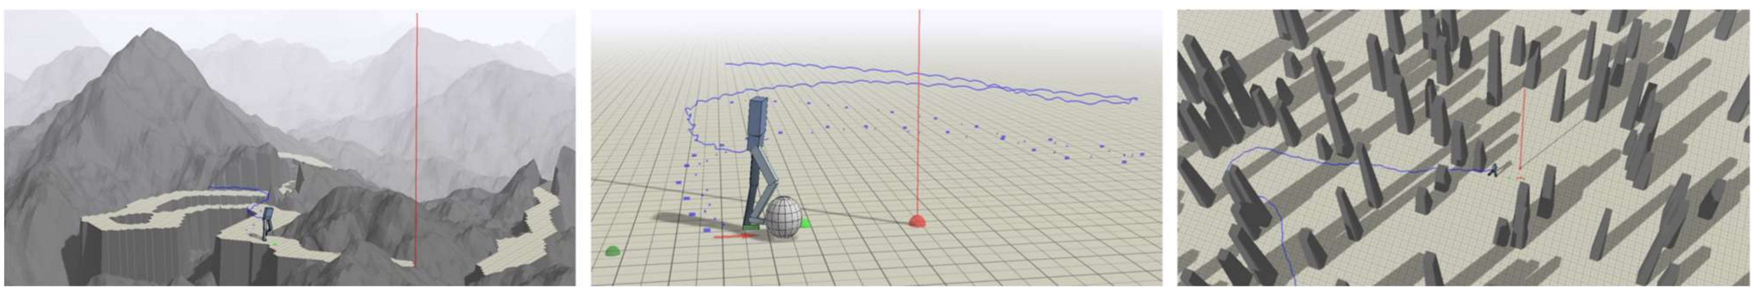

- http://www.cs.ubc.ca/~van/papers/2017-TOG-deepLoco/

Reinforcement Learning

A very very short introduction

- Take a set of states from an environment
- Define a 'Reward' that the agent receives for performing well at a task. For example:
    - Not falling down +
    - Following a motion capture example
- We must learn a policy of how the agent should act to maximize this reward over time
- A difficult problem – especially when acting over long time horizons!

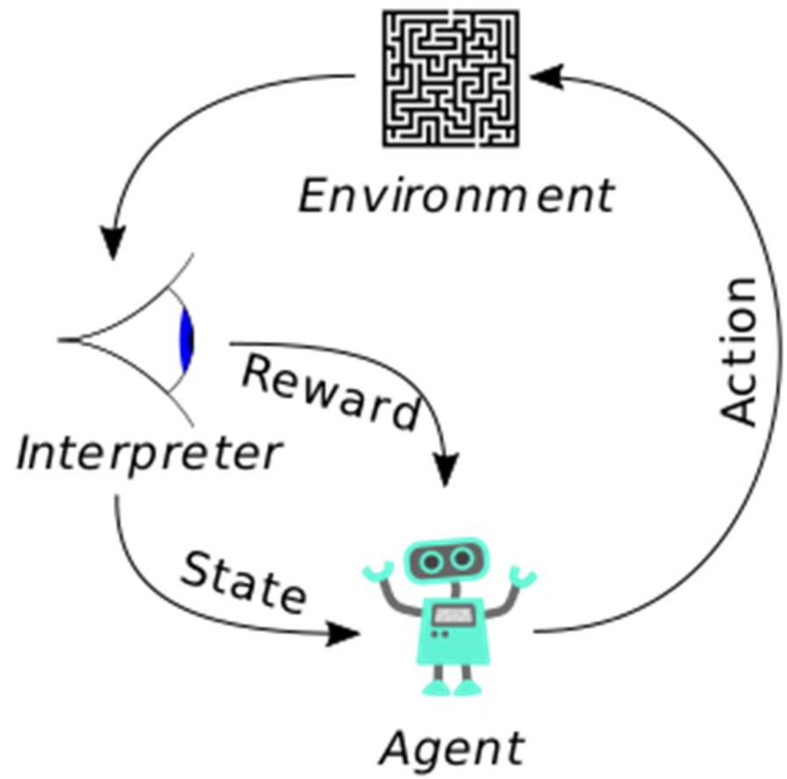

DeepLoco RL System

High level overview

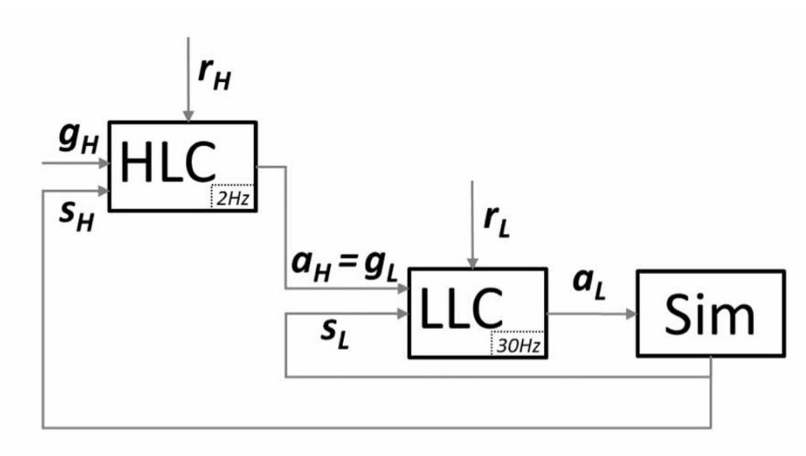

- Simulation engine + RL
  - Bullet Physics Engine, rewards
- Low level controller network
  - Uses phase, like PFNN, but simpler
  - Activates PD controller
- High Level controller network
  - Generates 'footstep plan' based on goals gH
  - Customizable for different tasks

LLC: Walk

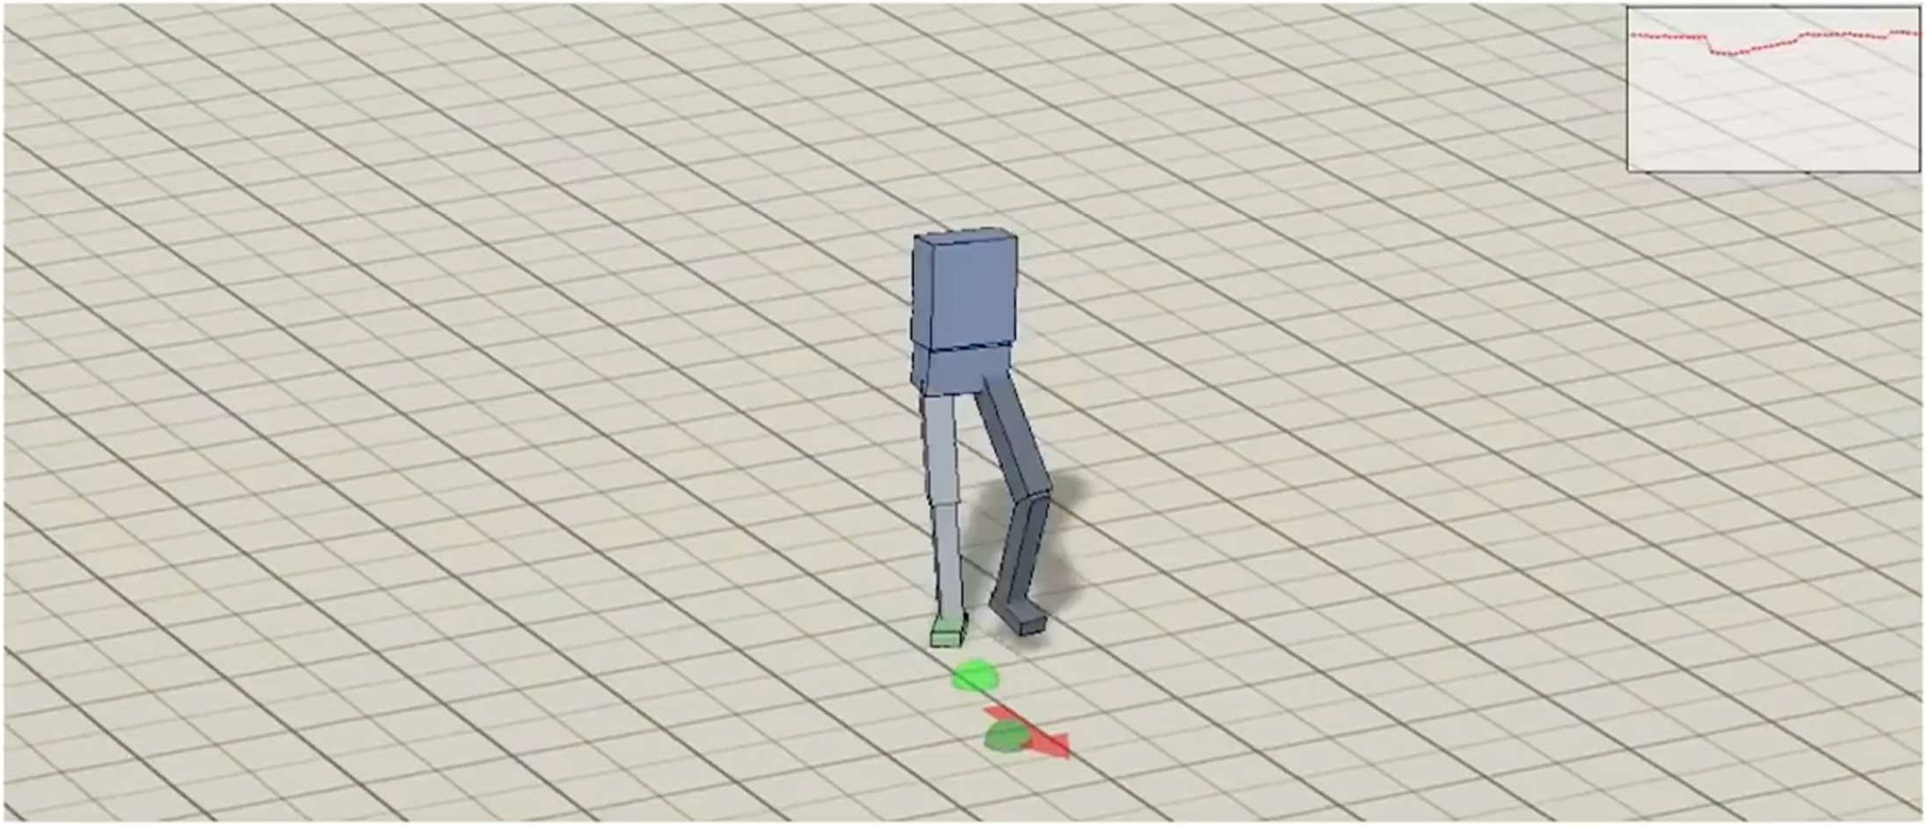

The LLC is first trained to locomote while following random footstep plans.

Early RL Results

DeepLoco-style Reward Function

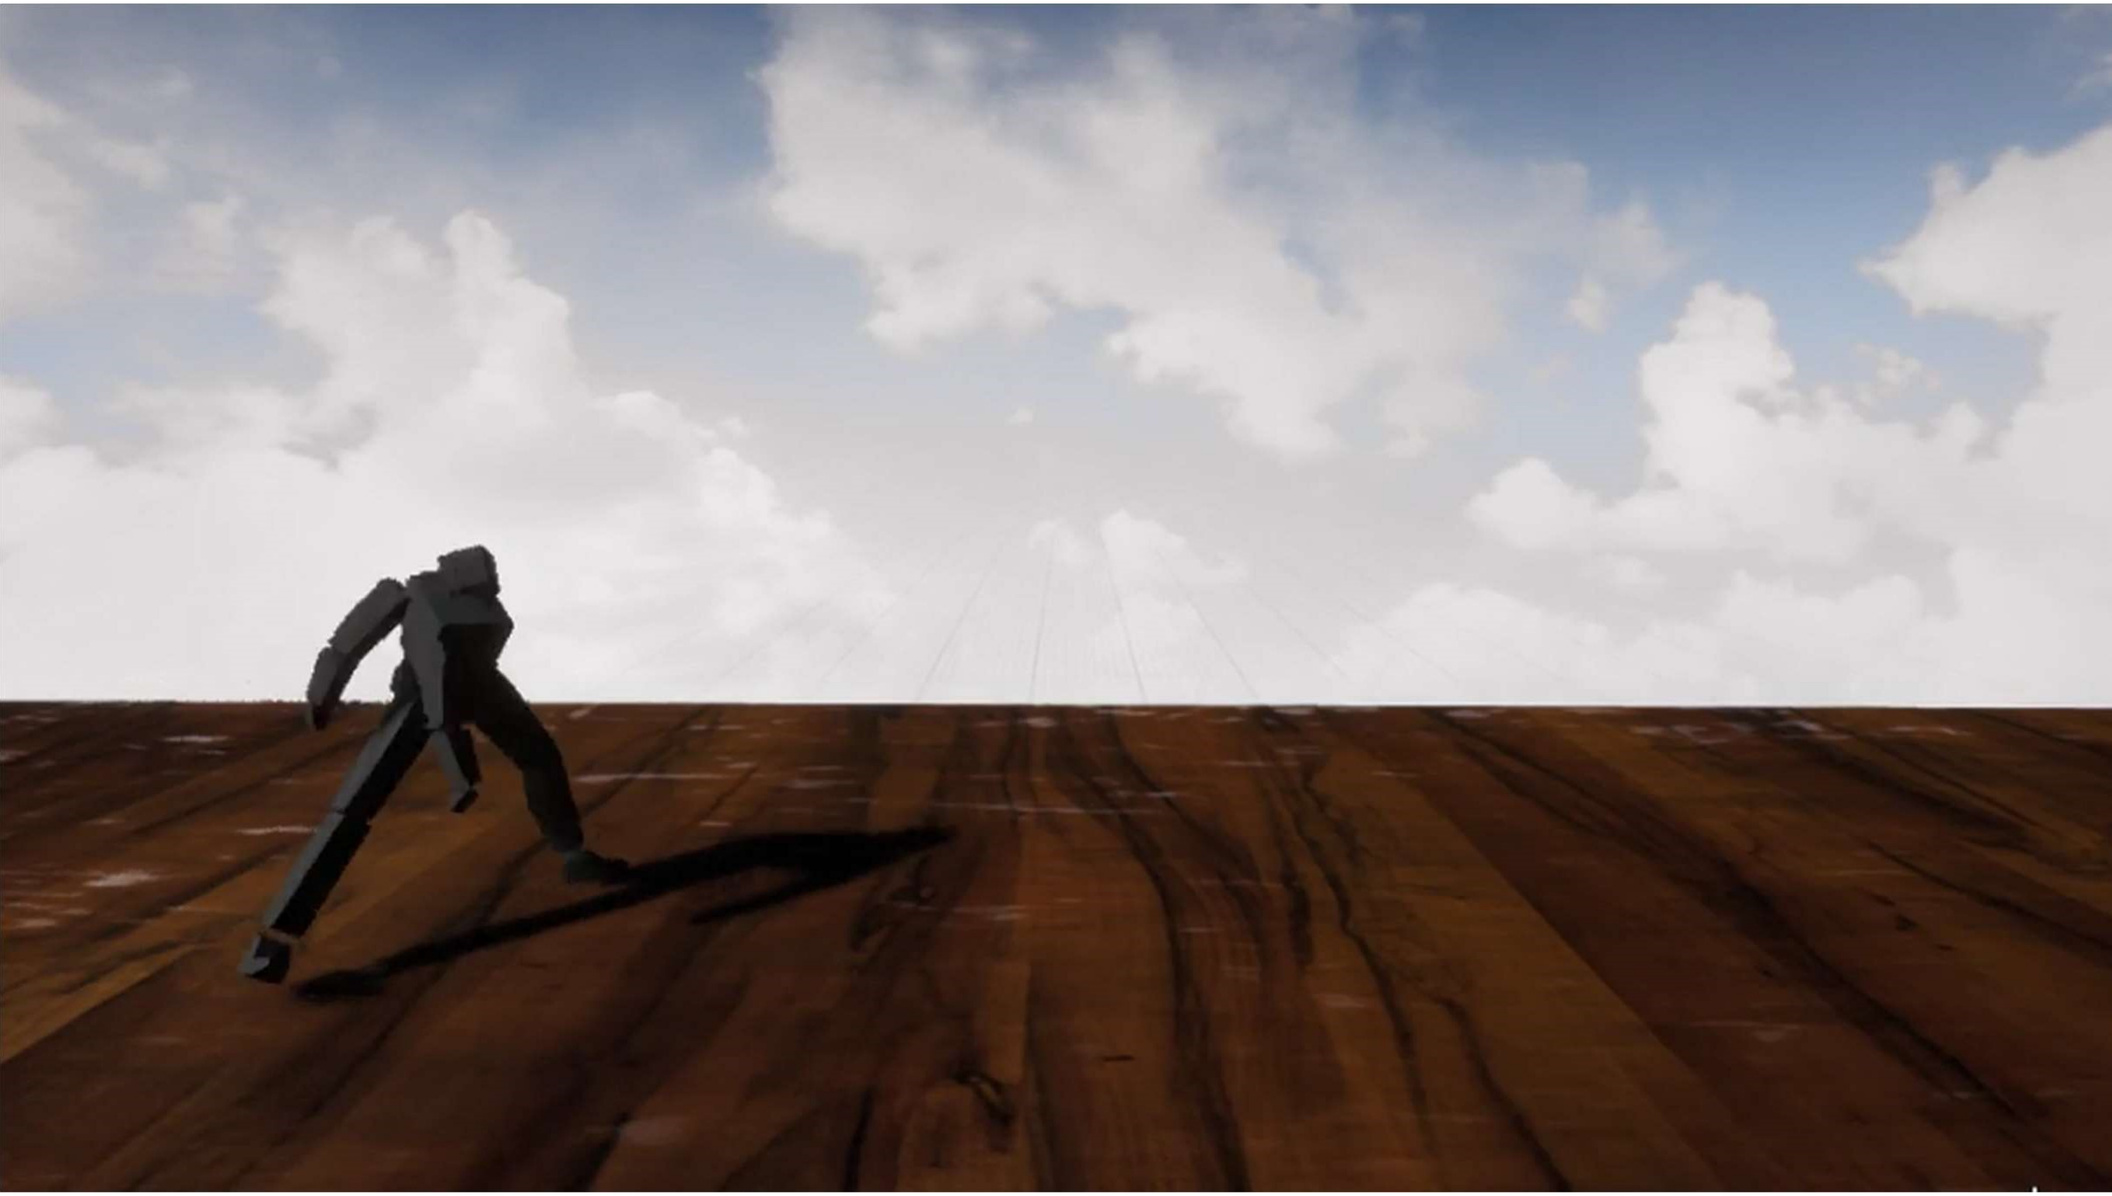

Physics + Mocap + RL

Physics + RL + Uneven Terrain
No Mocap

Ministry of Silly Walks

Physics + RL + Uneven Terrain + Mocap

Ministry of Getting Closer

DeepMimic

Advanced Physics Animation

DeepMimic: Example-Guided Deep Reinforcement Learning of Physics-Based Character Skills

Transactions on Graphics (Proc. ACM SIGGRAPH 2018)

Xue Bin Peng(1) Pieter Abbeel(1) Sergey Levine(1) Michiel van de Panne(2)
(1)University of California, Berkeley (2)University of British Columbia

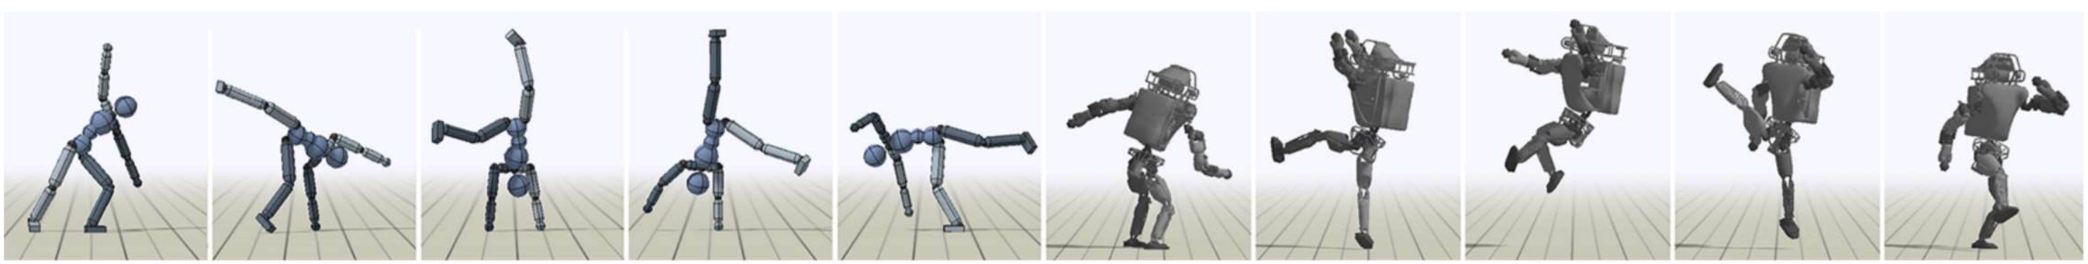

Skill Selector (Flips)

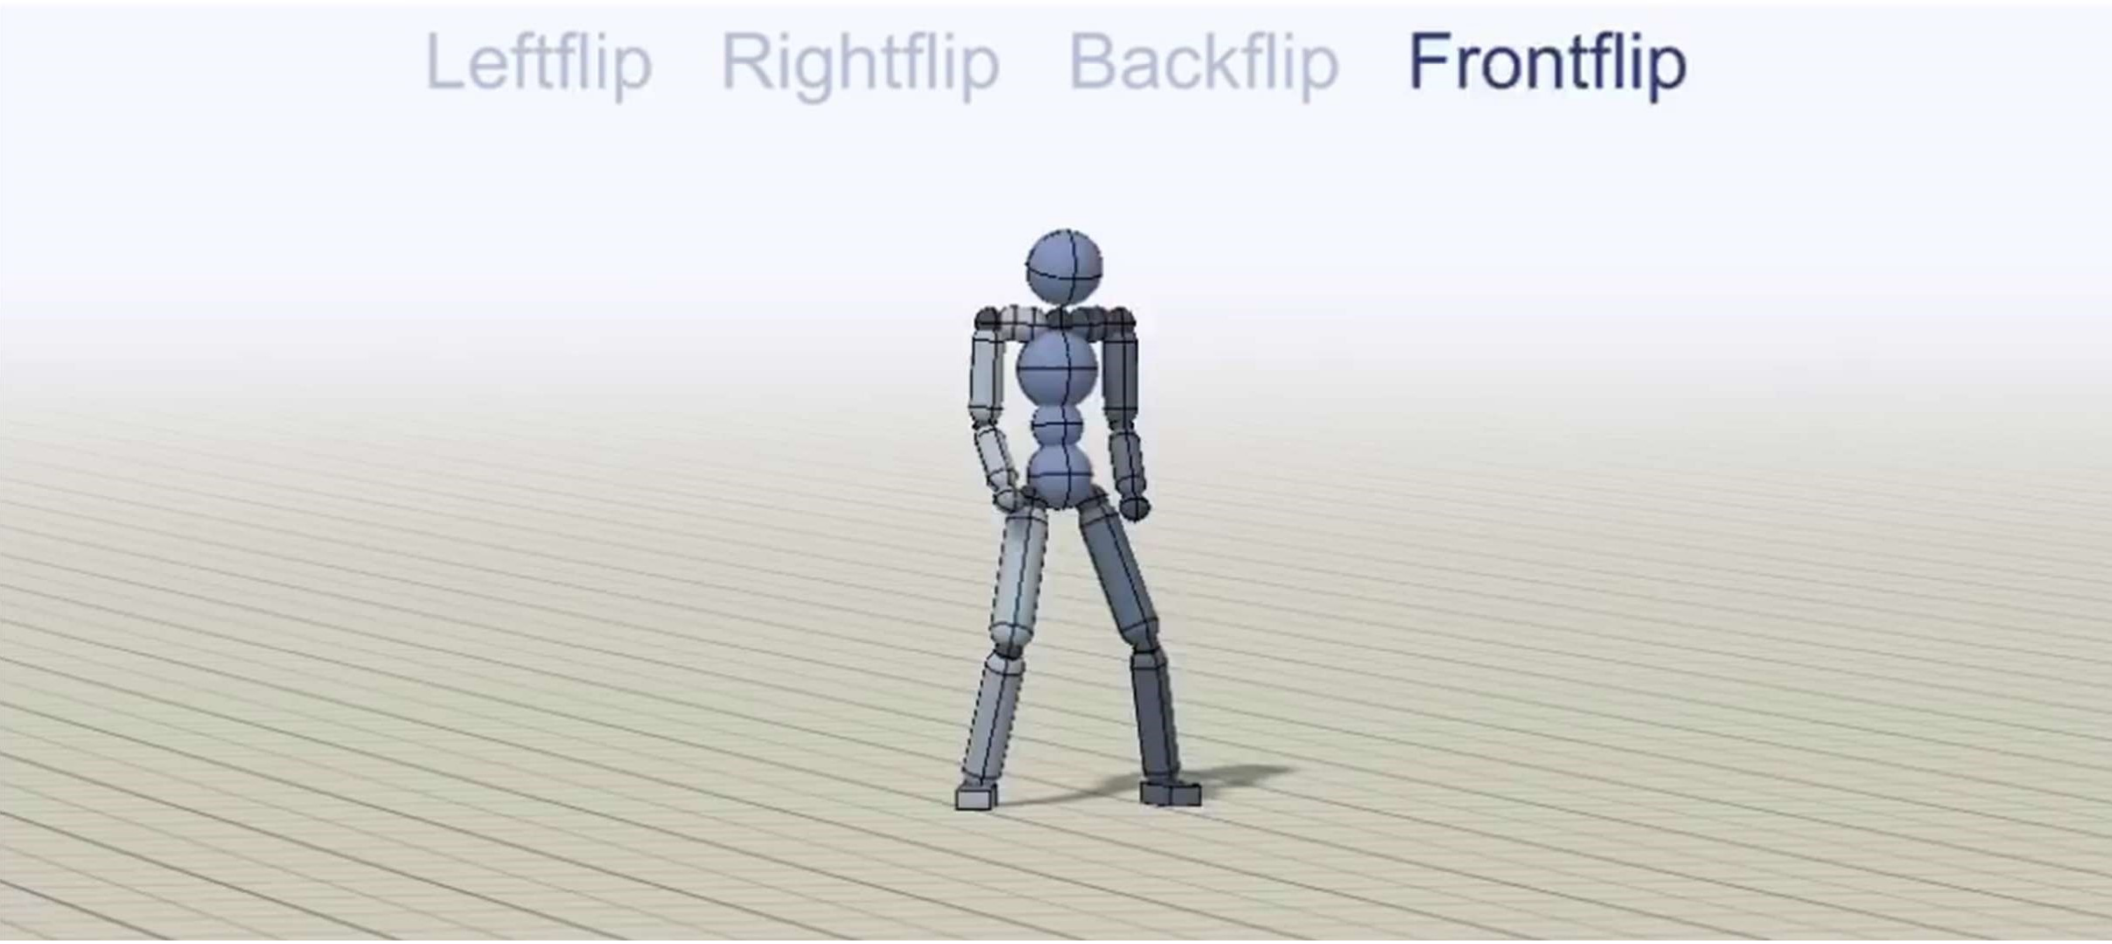

DeepMimic Enhancements

Just a few key tricks!

- Don't always start at the beginning!
  - Reference state initialization from random points in the motion capture clip
  - Simplifies learning hard motions
- Early termination
  - If an agent falls down, start over immediately
  - Don't bother learning how to get up without reference motion

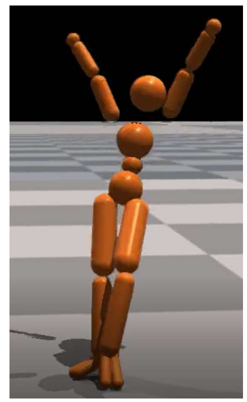

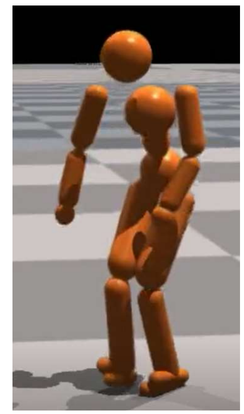

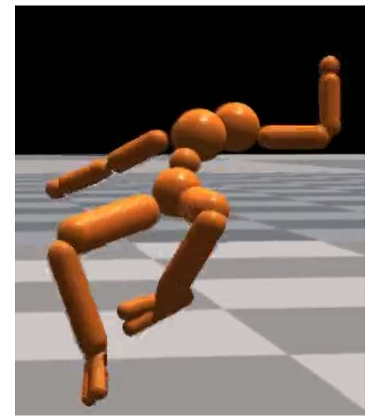

GPU Accelerated Simulation

Apply GPUs to BOTH Sim and RL

GPU-Accelerated Robotic Simulation for Distributed Reinforcement Learning
Conference on Robot Learning (CoRL) 2018

Jacky Liang, Viktor Makoviychuk, Ankur Handa, Nuttapong Chentanez, Miles Macklin, Dieter Fox
NVIDIA

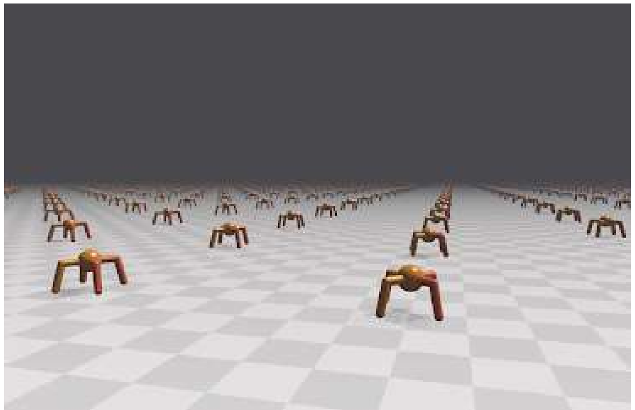

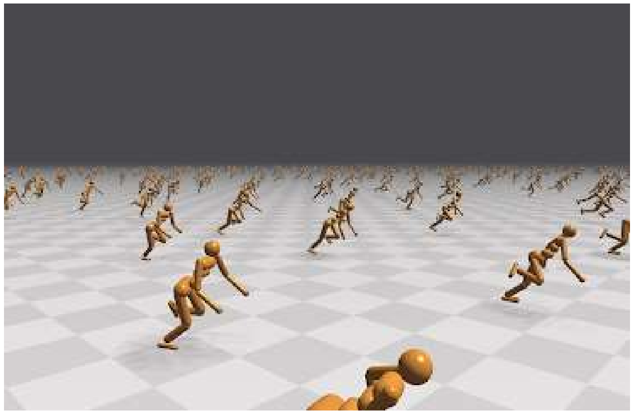

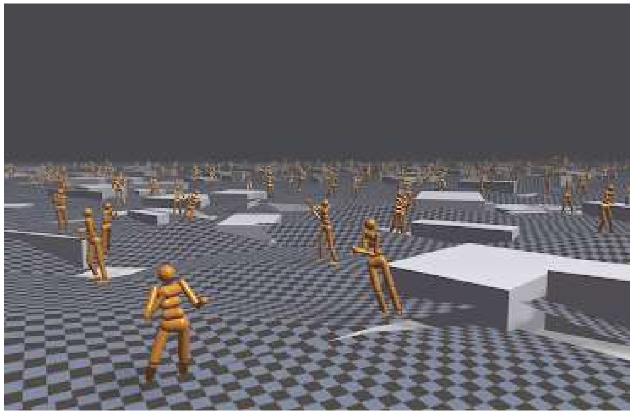

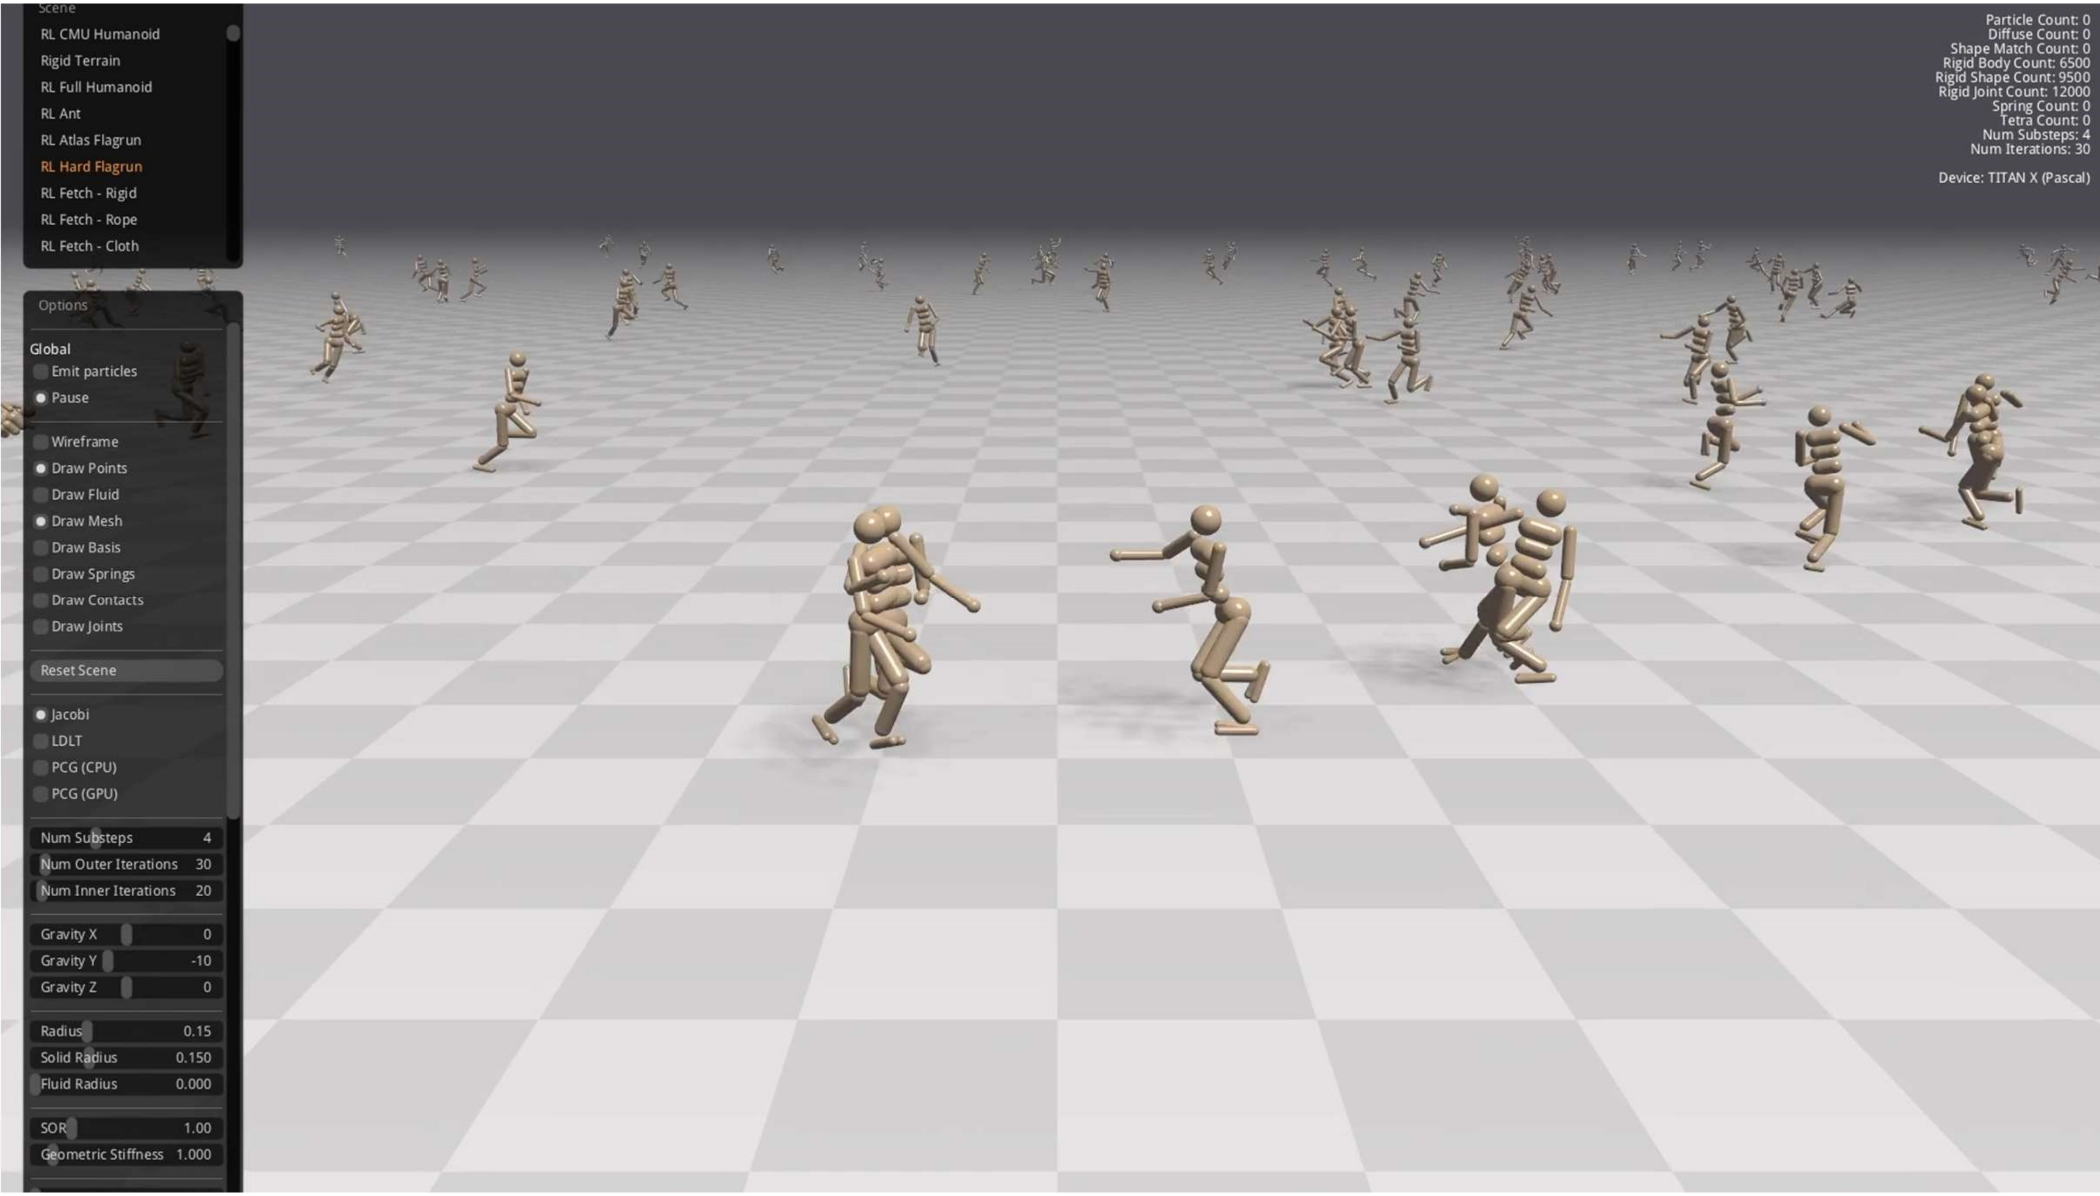
Scene
RL CMU Humanoid
Rigid Terrain
RL Full Humanoid
RL Ant
RL Atlas Flagrun
RL Hard Flagrun
RL Fetch - Rigid
RL Fetch - Rope
RL Fetch - Cloth
Options
Global
Emit particles
Pause
Wireframe
Draw Points
Draw Fluid
Draw Mesh
Draw Basis
Draw Springs
Draw Contacts
Draw Joints
Reset Scene
Jacobi
LDLT
PCG (CPU)
PCG (GPU)
Num Substeps 4
Num Outer Iterations 30
Num Inner Iterations 20
Gravity X 0
Gravity Y -10
Gravity Z 0
Radius 0.15
Solid Radius 0.150
Fluid Radius 0.000
SOR 1.00
Geometric Stiffness 1.000
Particle Count: 0
Diffuse Count: 0
Shape Match Count: 0
Rigid Body Count: 6500
Rigid Shape Count: 9500
Rigid Joint Count: 12000
Spring Count: 0
Tetra Count: 0
Num Substeps: 4
Num Iterations: 30
Device: TITAN X (Pascal)

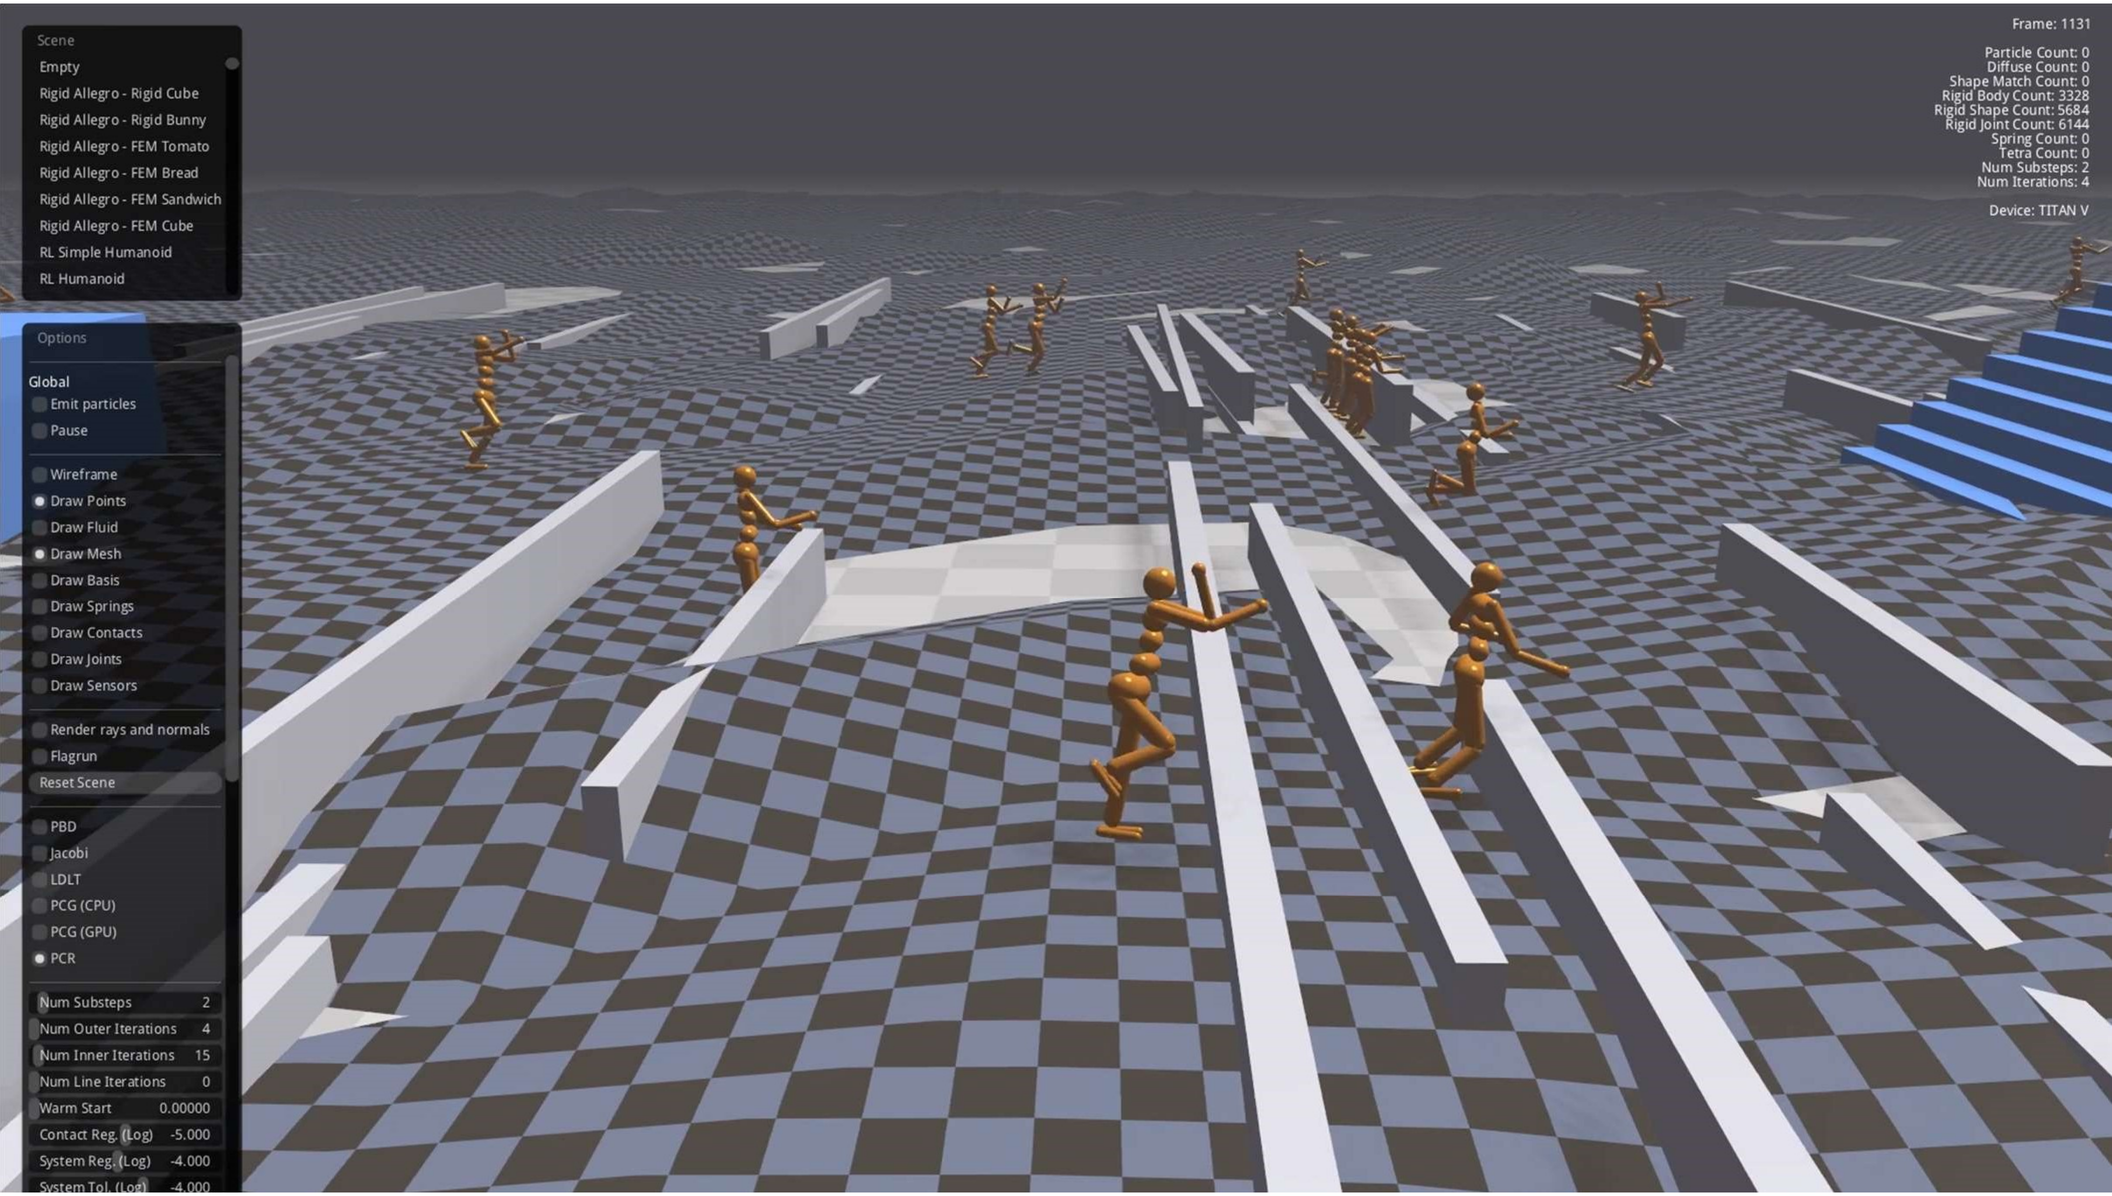
Scene
Empty
Rigid Allegro - Rigid Cube
Rigid Allegro - Rigid Bunny
Rigid Allegro - FEM Tomato
Rigid Allegro - FEM Bread
Rigid Allegro - FEM Sandwich
Rigid Allegro - FEM Cube
RL Simple Humanoid
RL Humanoid
Options
Global
Emit particles
Pause
Wireframe
Draw Points
Draw Fluid
Draw Mesh
Draw Basis
Draw Springs
Draw Contacts
Draw Joints
Draw Sensors
Render rays and normals
Flagrun
Reset Scene
PBD
Jacobi
LDLT
PCG (CPU)
PCG (GPU)
PCR
Num Substeps 2
Num Outer Iterations 4
Num Inner Iterations 15
Num Line Iterations 0
Warm Start 0.00000
Contact Reg. (Log) -5.000
System Reg. (Log) -4.000
System Tol. (Log) -4.000
Frame: 1131
Particle Count: 0
Diffuse Count: 0
Shape Match Count: 0
Rigid Body Count: 3328
Rigid Shape Count: 5684
Rigid Joint Count: 6144
Spring Count: 0
Tetra Count: 0
Num Substeps: 2
Num Iterations: 4
Device: TITAN V

Arbitrary Motion Imitation with Physics

Single Network, Thousands of Clips

Physics-based Motion Capture Imitation with Deep Reinforcement Learning

Motion, Interaction, and Games (MIG) 2018

Nuttapong Chentanez, Matthias Müller, Miles Macklin, Viktor Makoviychuk, Stefan Jeschke

NVIDIA

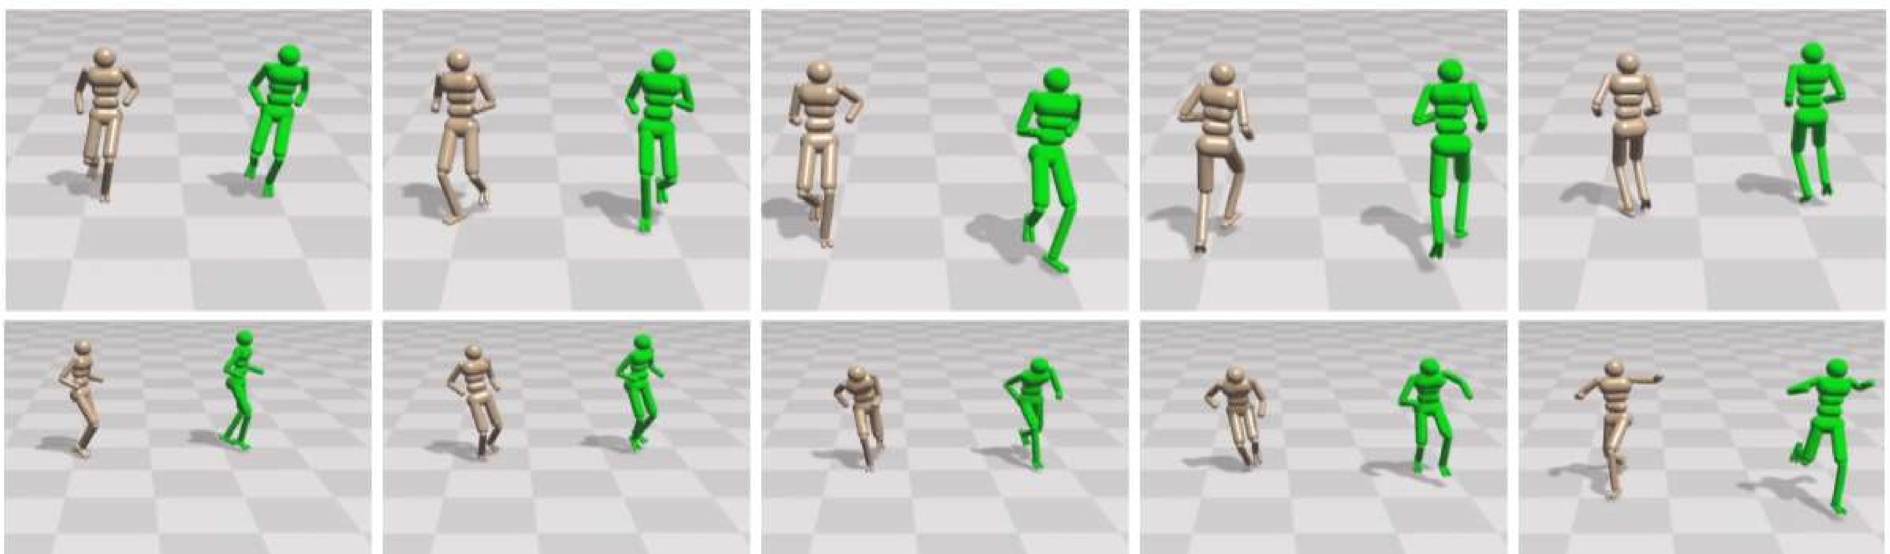

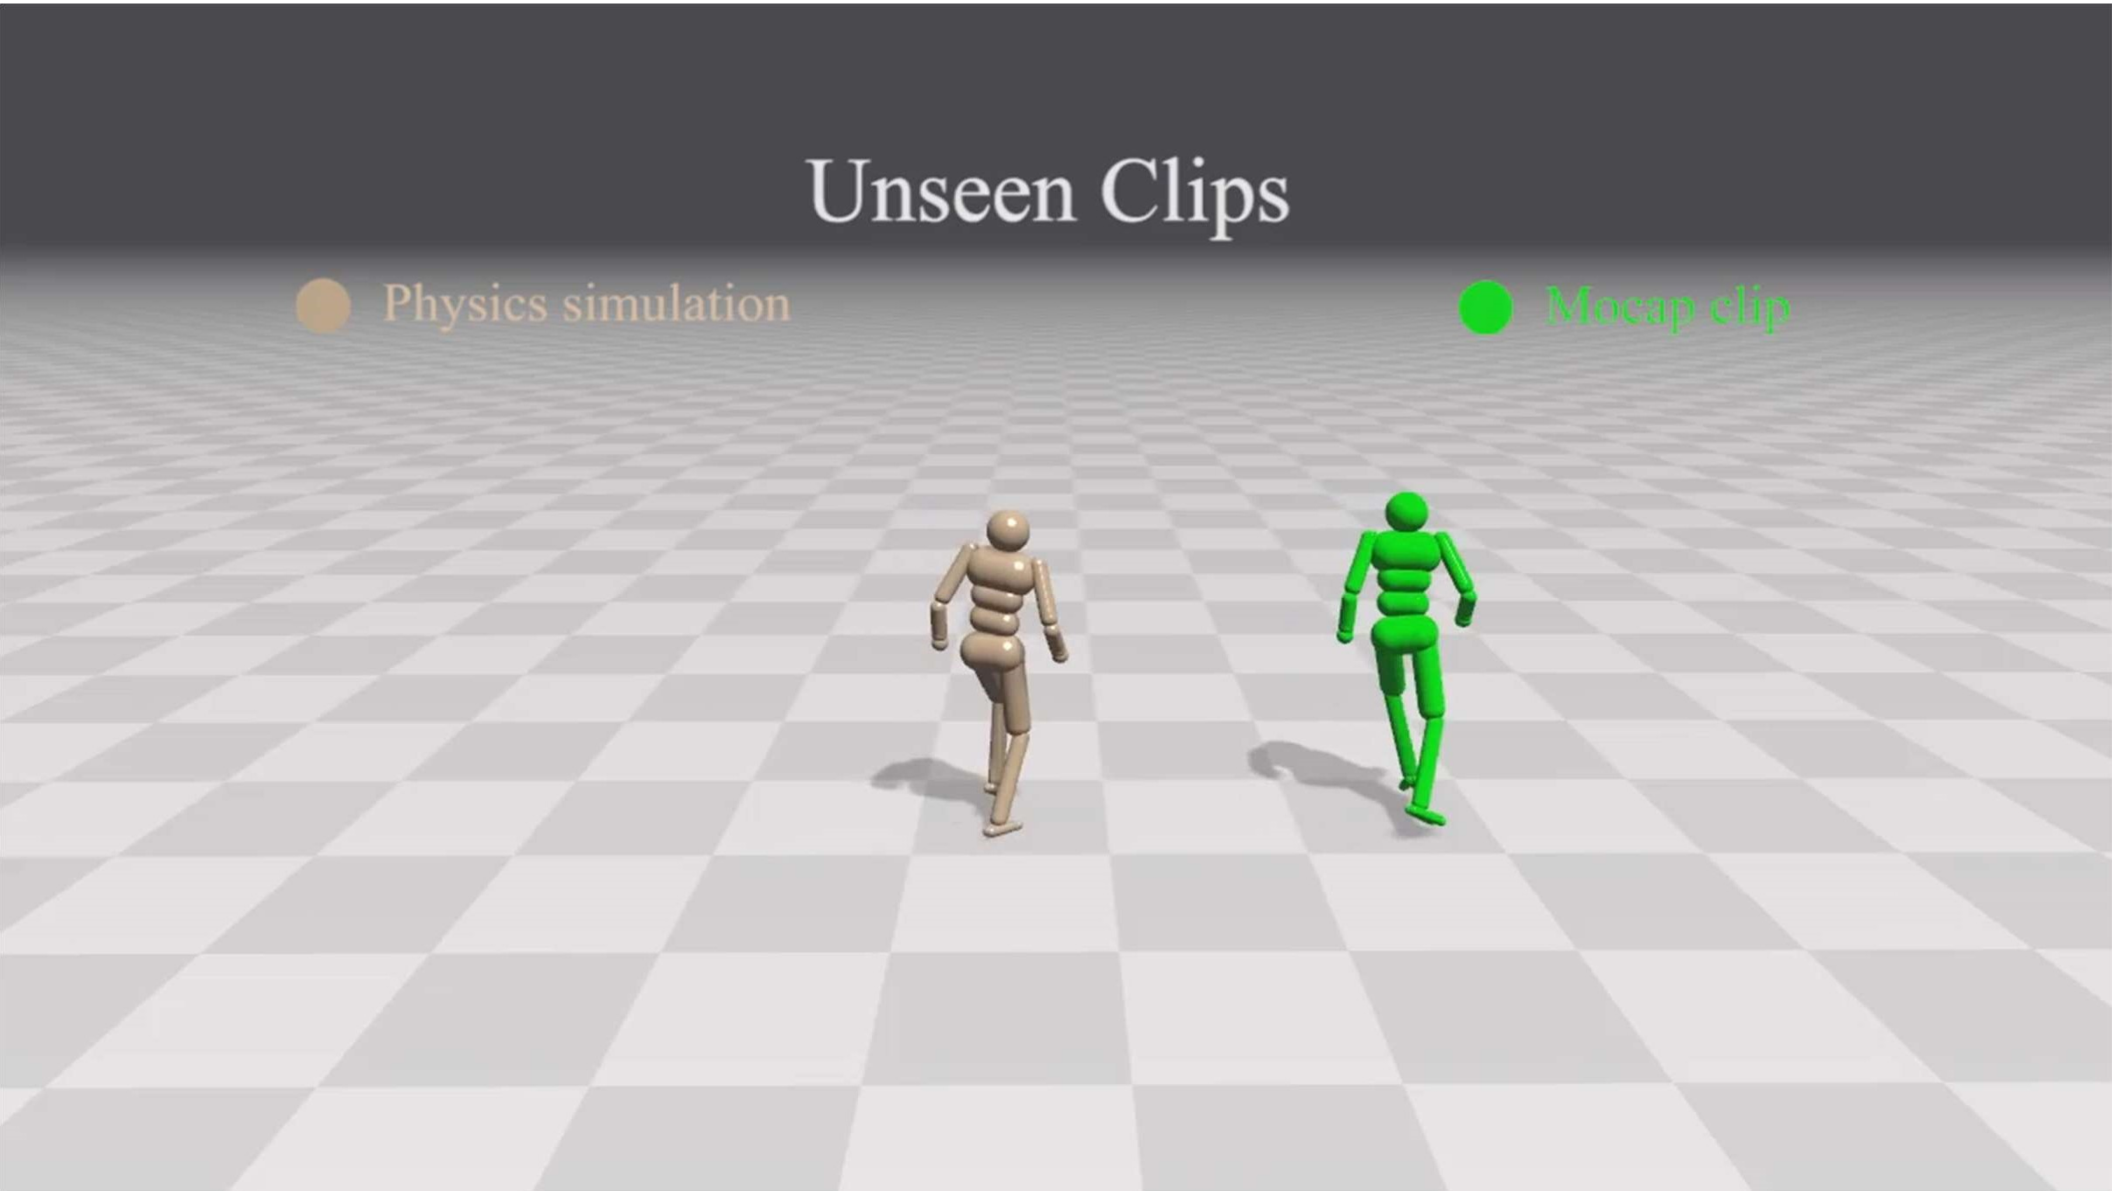
Unseen Clips
Physics simulation
Mocap clip

IsaacGym

Advanced Physics RL Training Environment

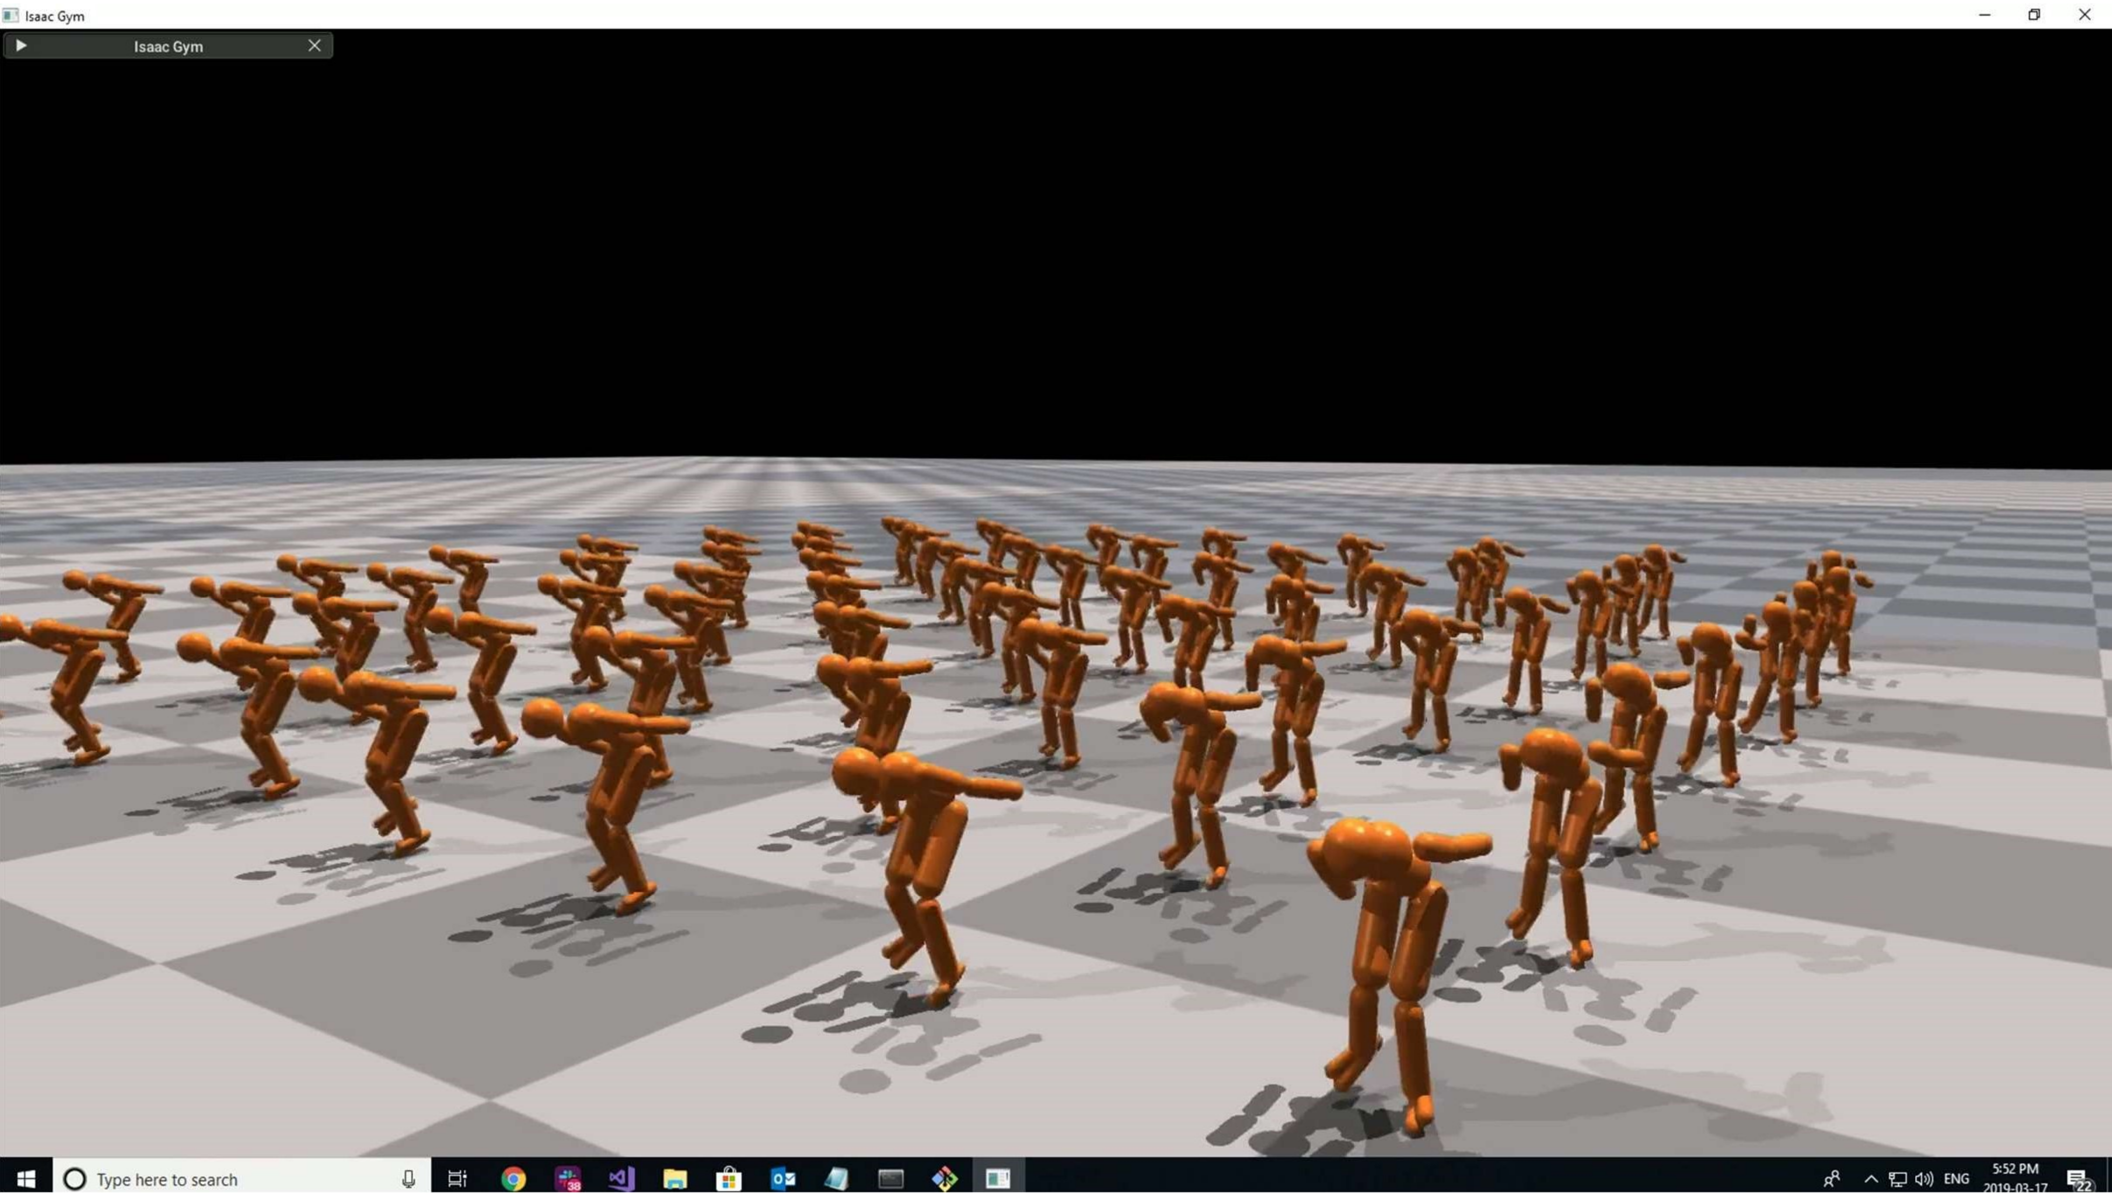
Isaac Gym
Isaac Gym
Type here to search
ENG
5:52 PM

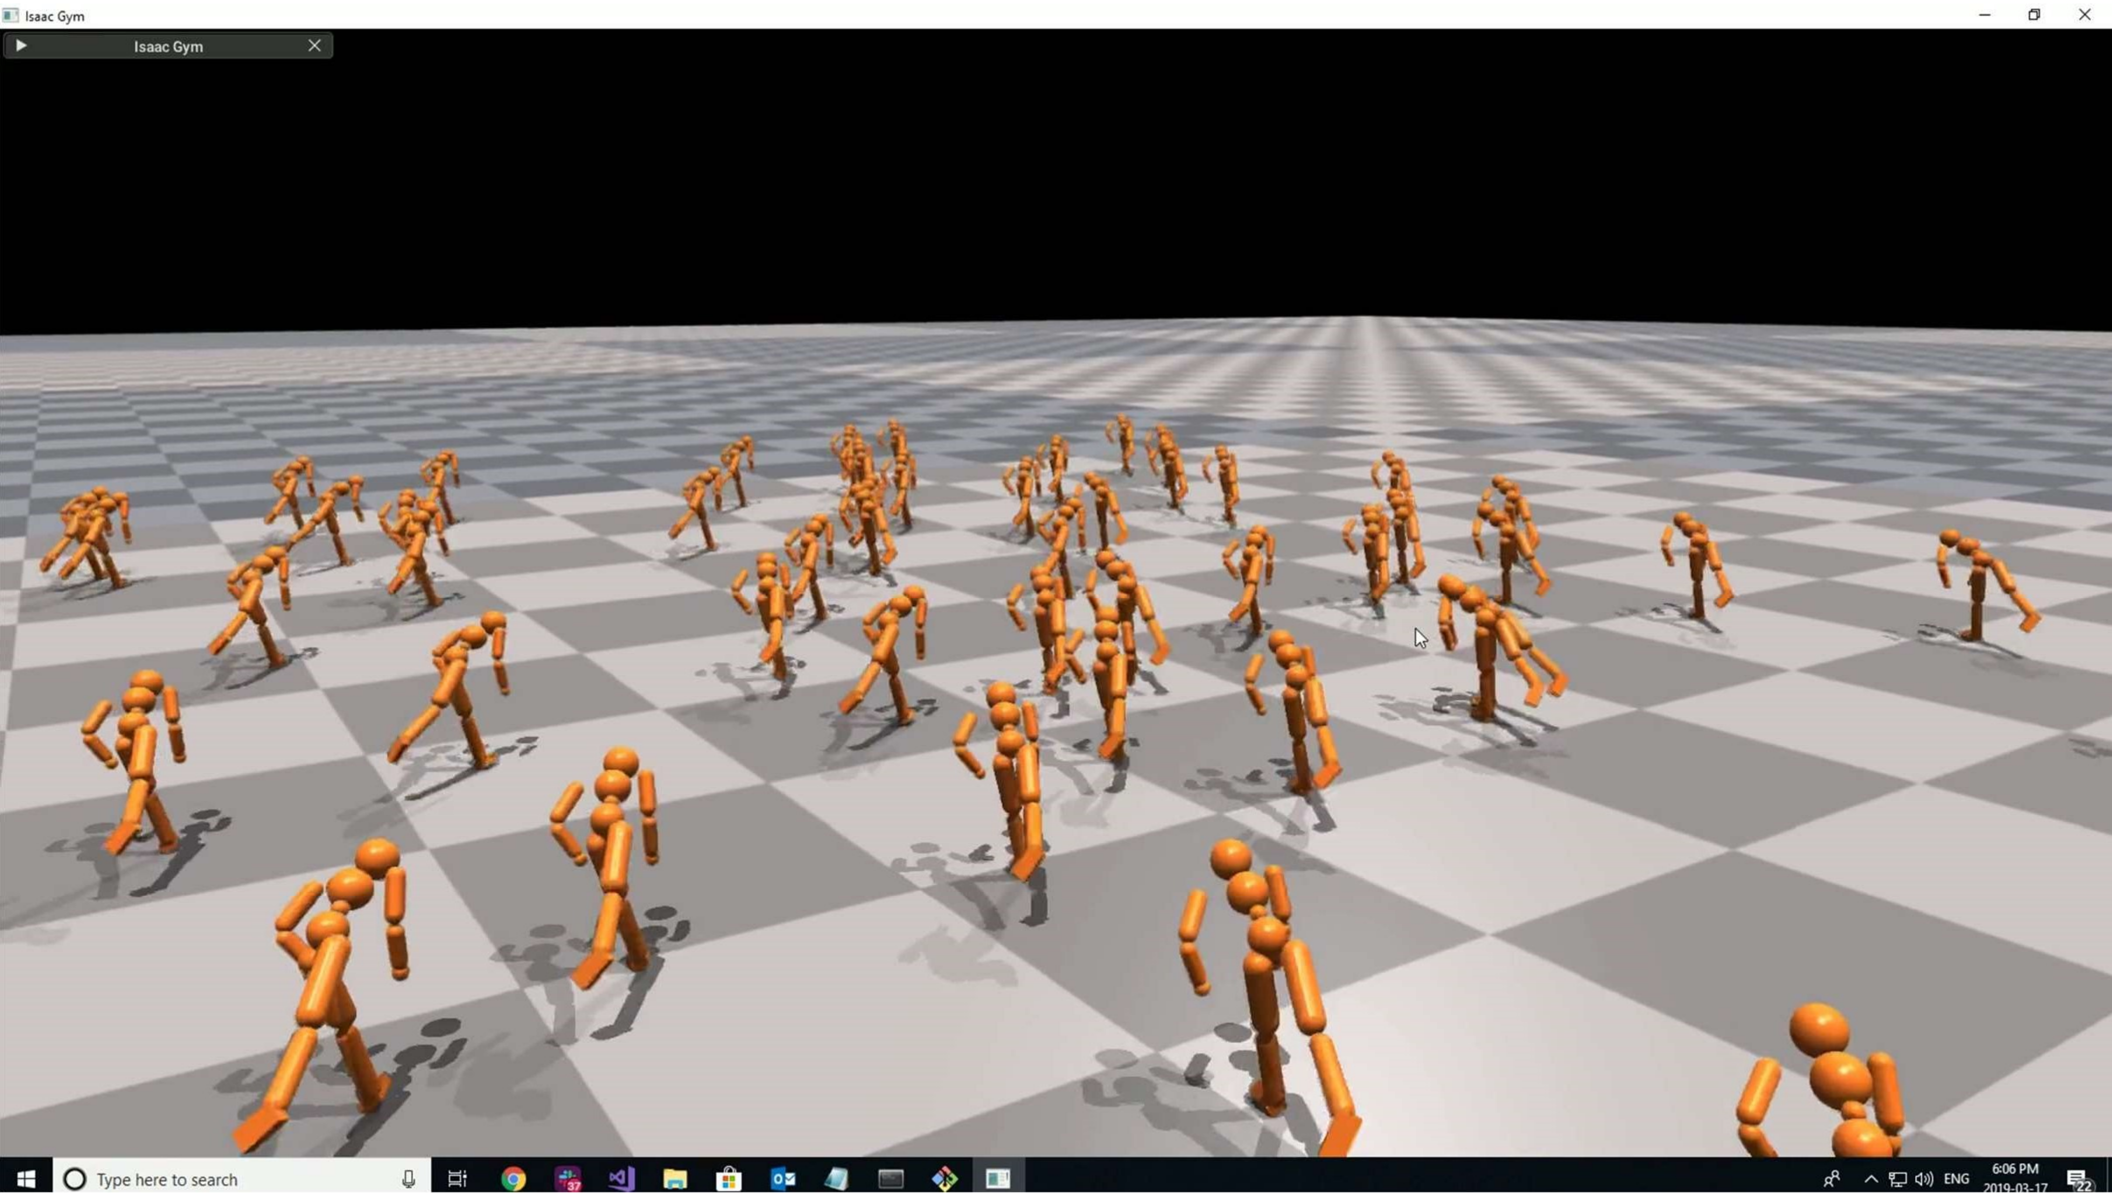
Isaac Gym
Isaac Gym
Type here to search
ENG
6:06 PM
2019-03-17

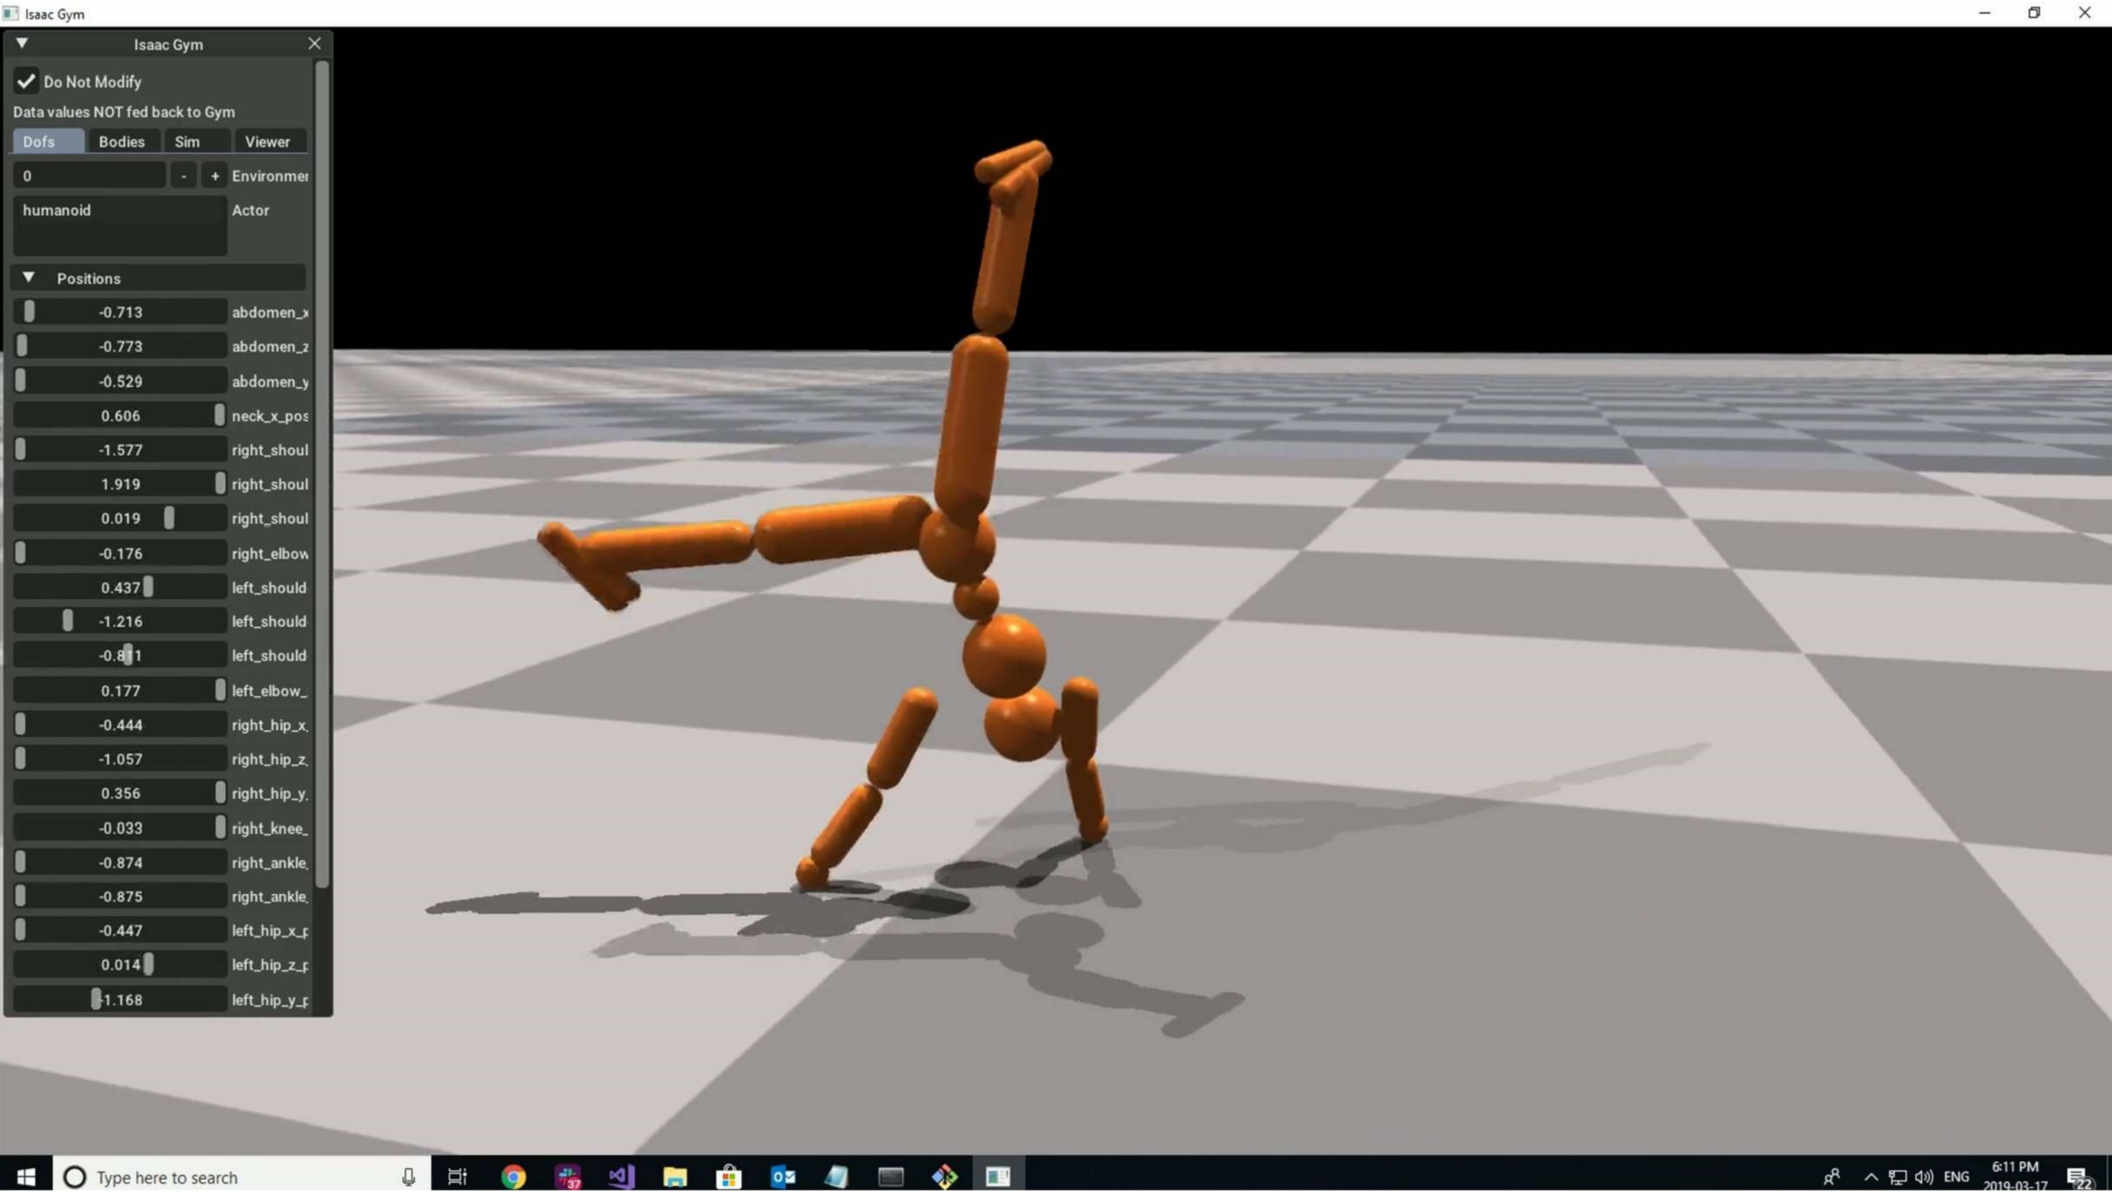
Isaac Gym
Isaac Gym
Do Not Modify
Data values NOT fed back to Gym
Dofs
Bodies
Sim
Viewer
0
Environmer
humanoid
Actor
Positions
-0.713 abdomen_x
-0.773 abdomen_z
-0.529 abdomen_y
0.606 neck_x_pos
-1.577 right_shoul
1.919 right_shoul
0.019 right_shoul
-0.176 right_elbow
0.437 left_should
-1.216 left_should
0.177 left_elbow_
-0.444 right_hip_x_
-1.057 right_hip_z_
0.356 right_hip_y_
-0.033 right_knee_
-0.874 right_ankle_
-0.875 right_ankle_
-0.447 left_hip_x_p
0.014 left_hip_z_p
left_hip_y_p
Type here to search
ENG
6:11 PM
2019-03-17

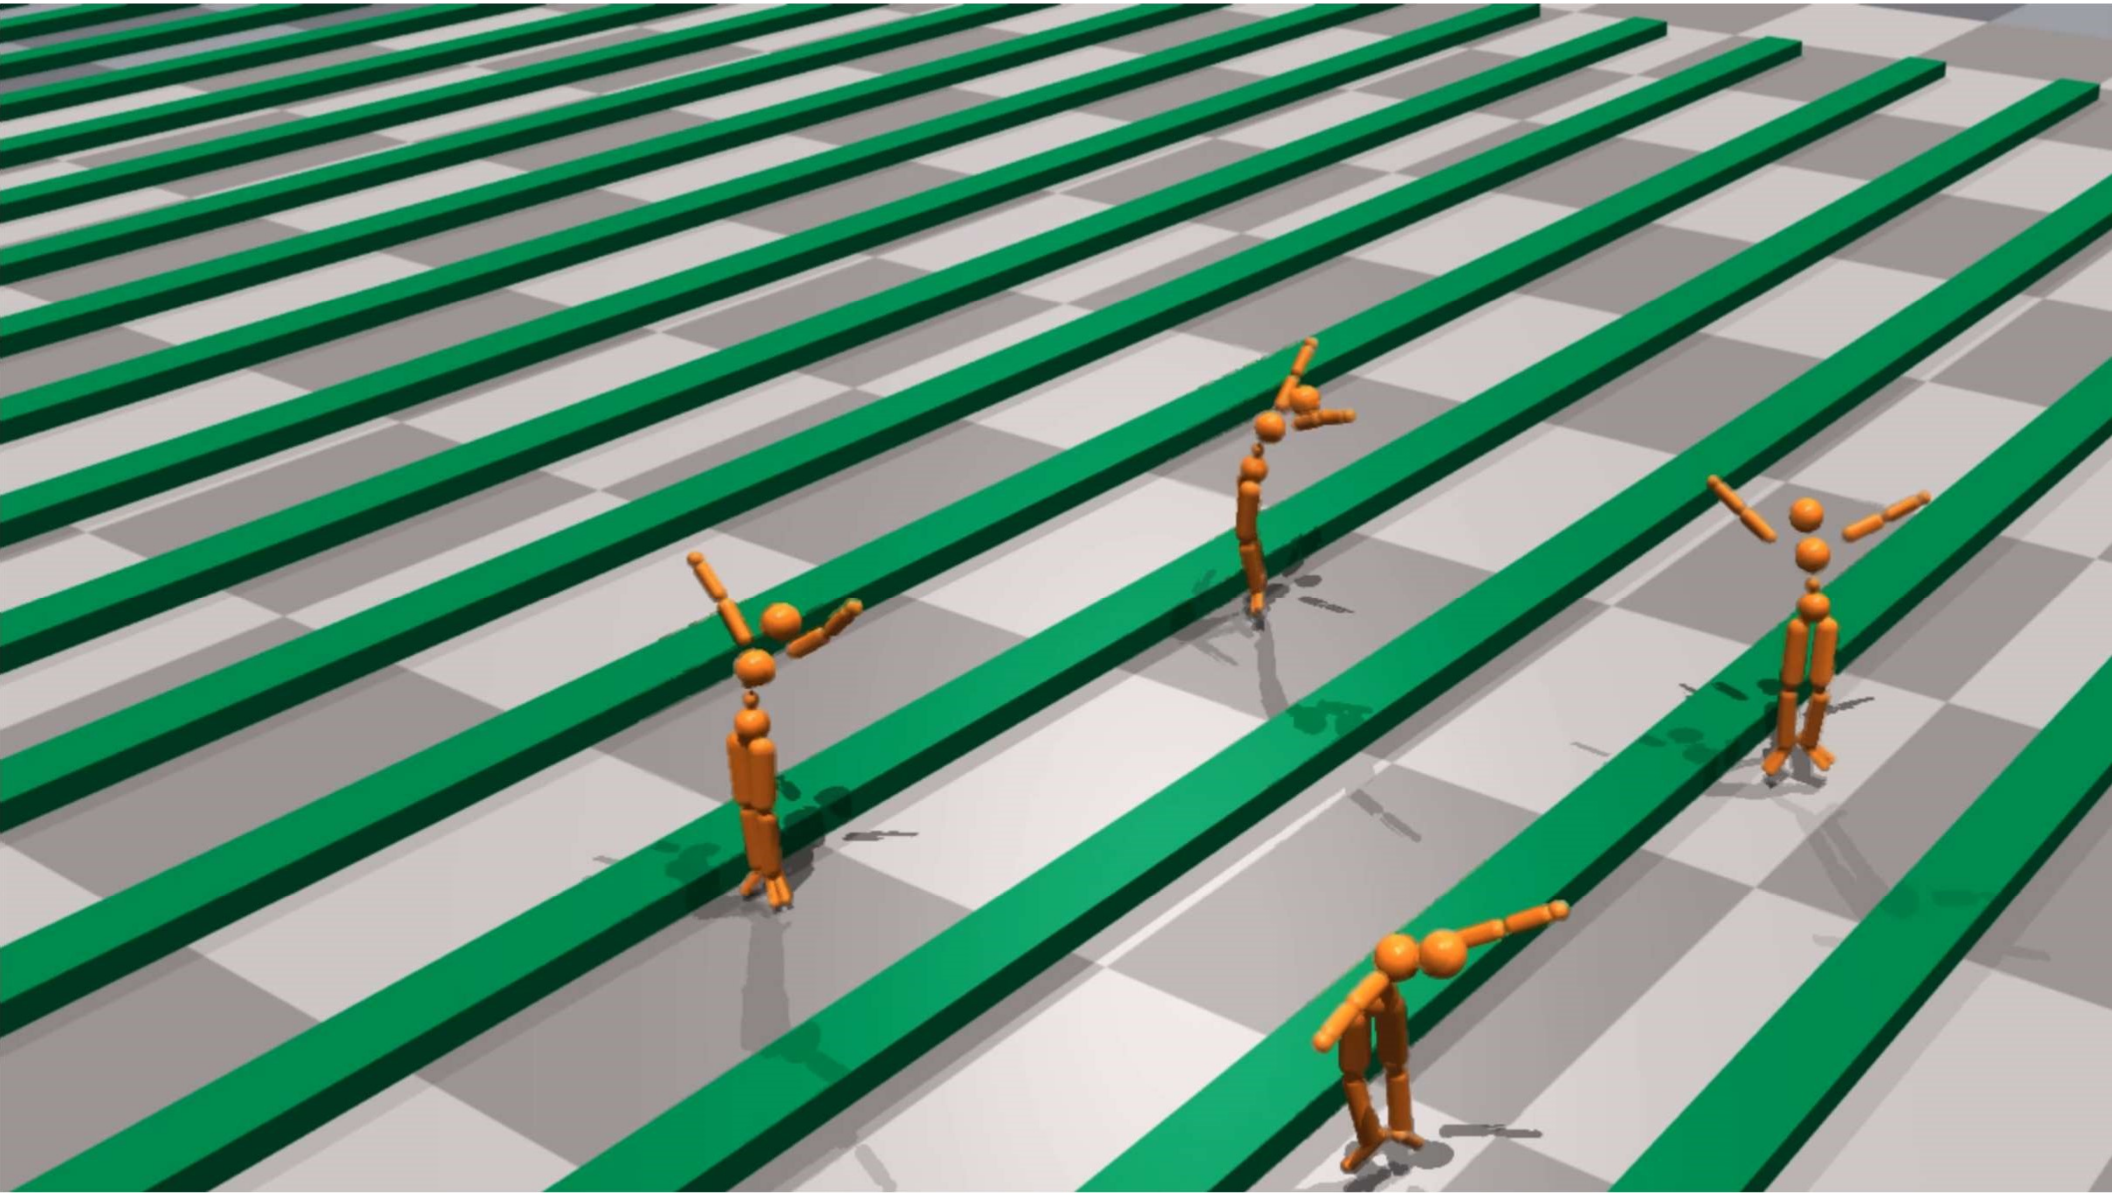

High Level Behavior

Exciting times ahead!

LLC: Walk

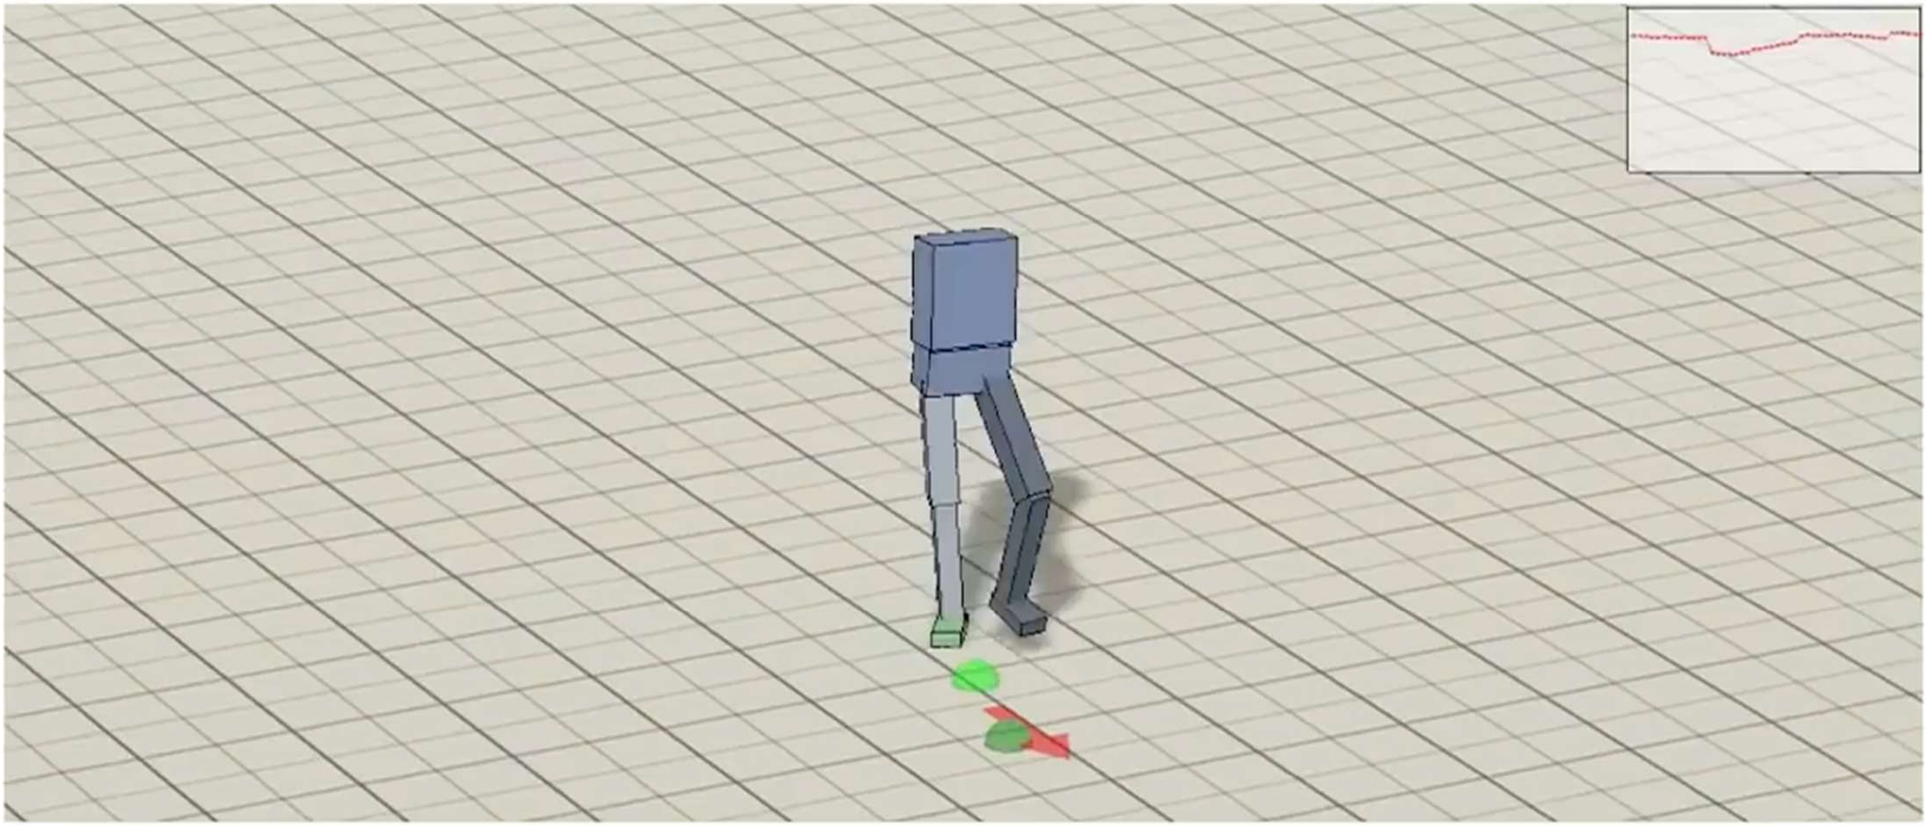

The LLC is first trained to locomote while
following random footstep plans.

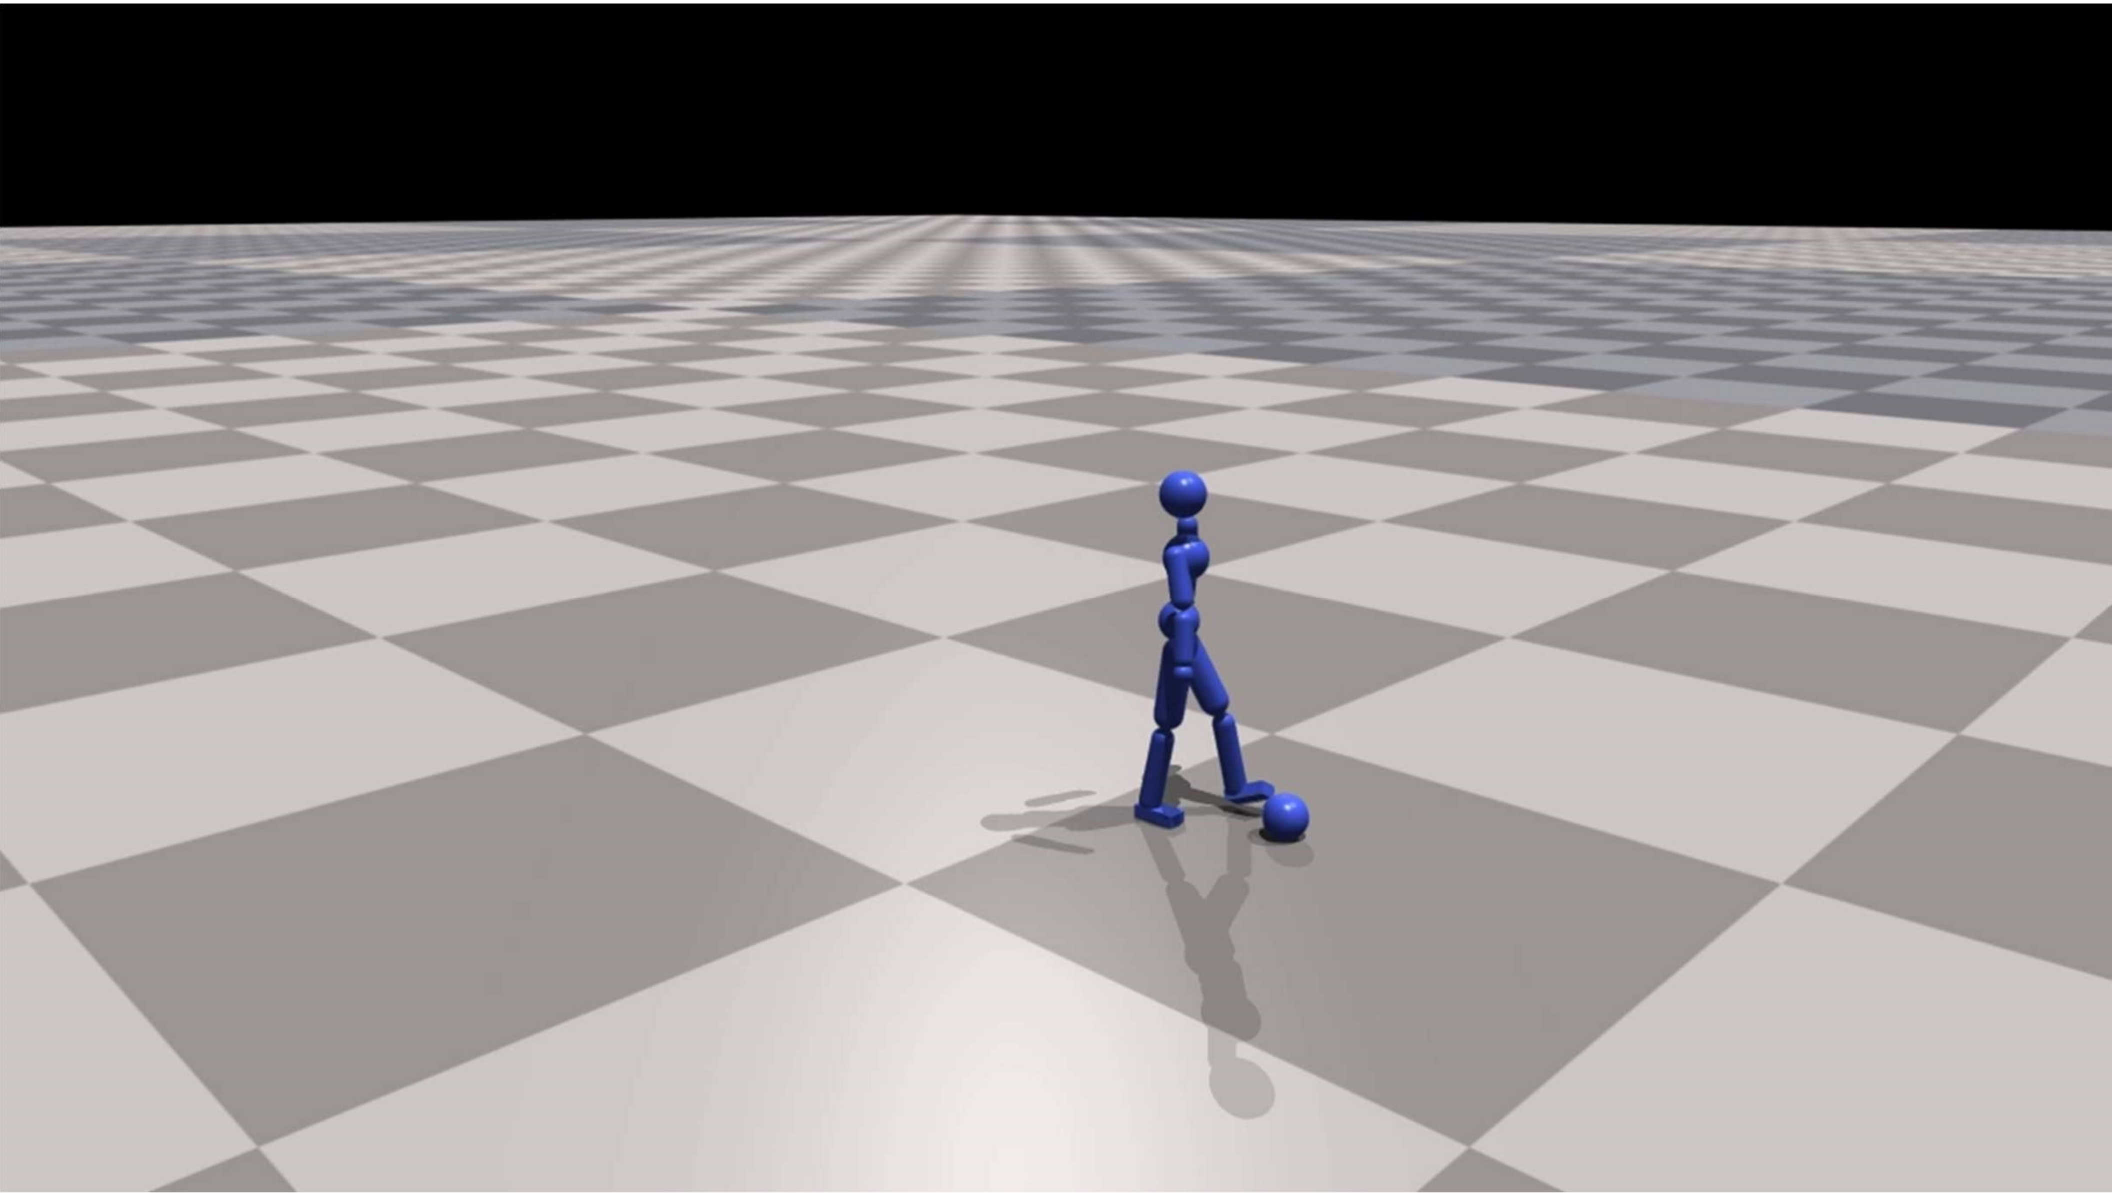


Isaac Gym
Type here to search
ENG
10:43 PM

QUESTIONS?