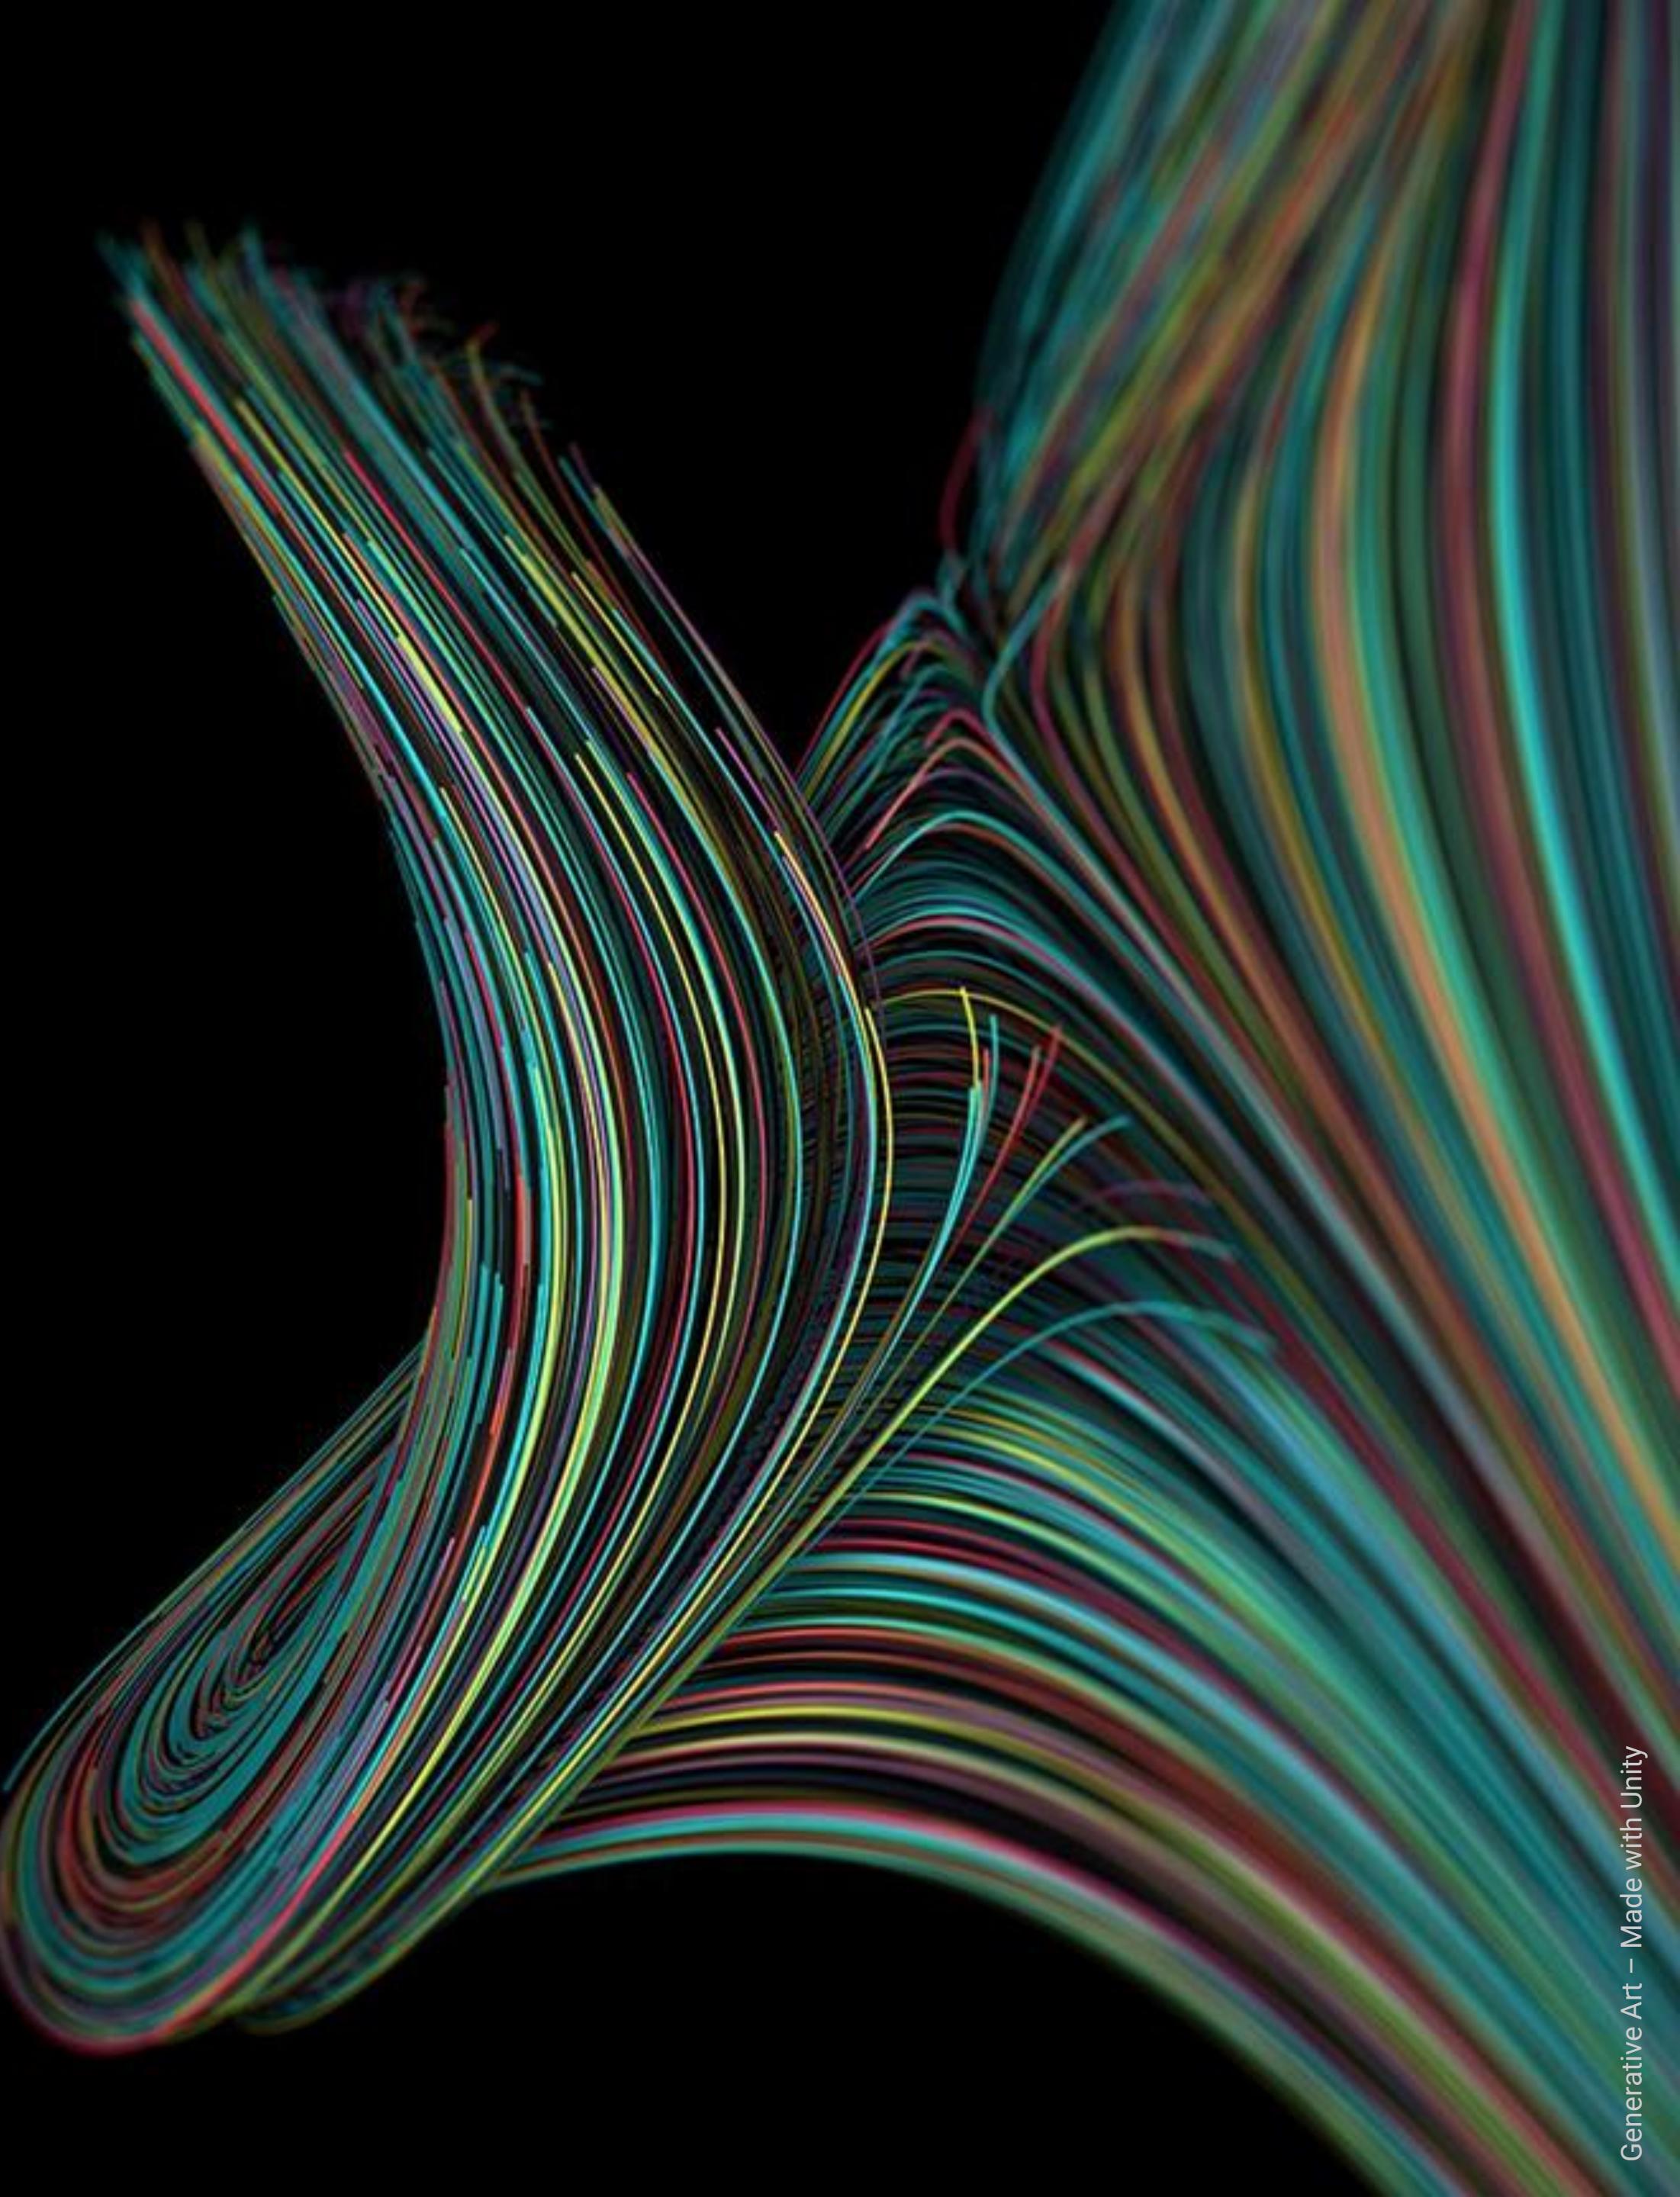
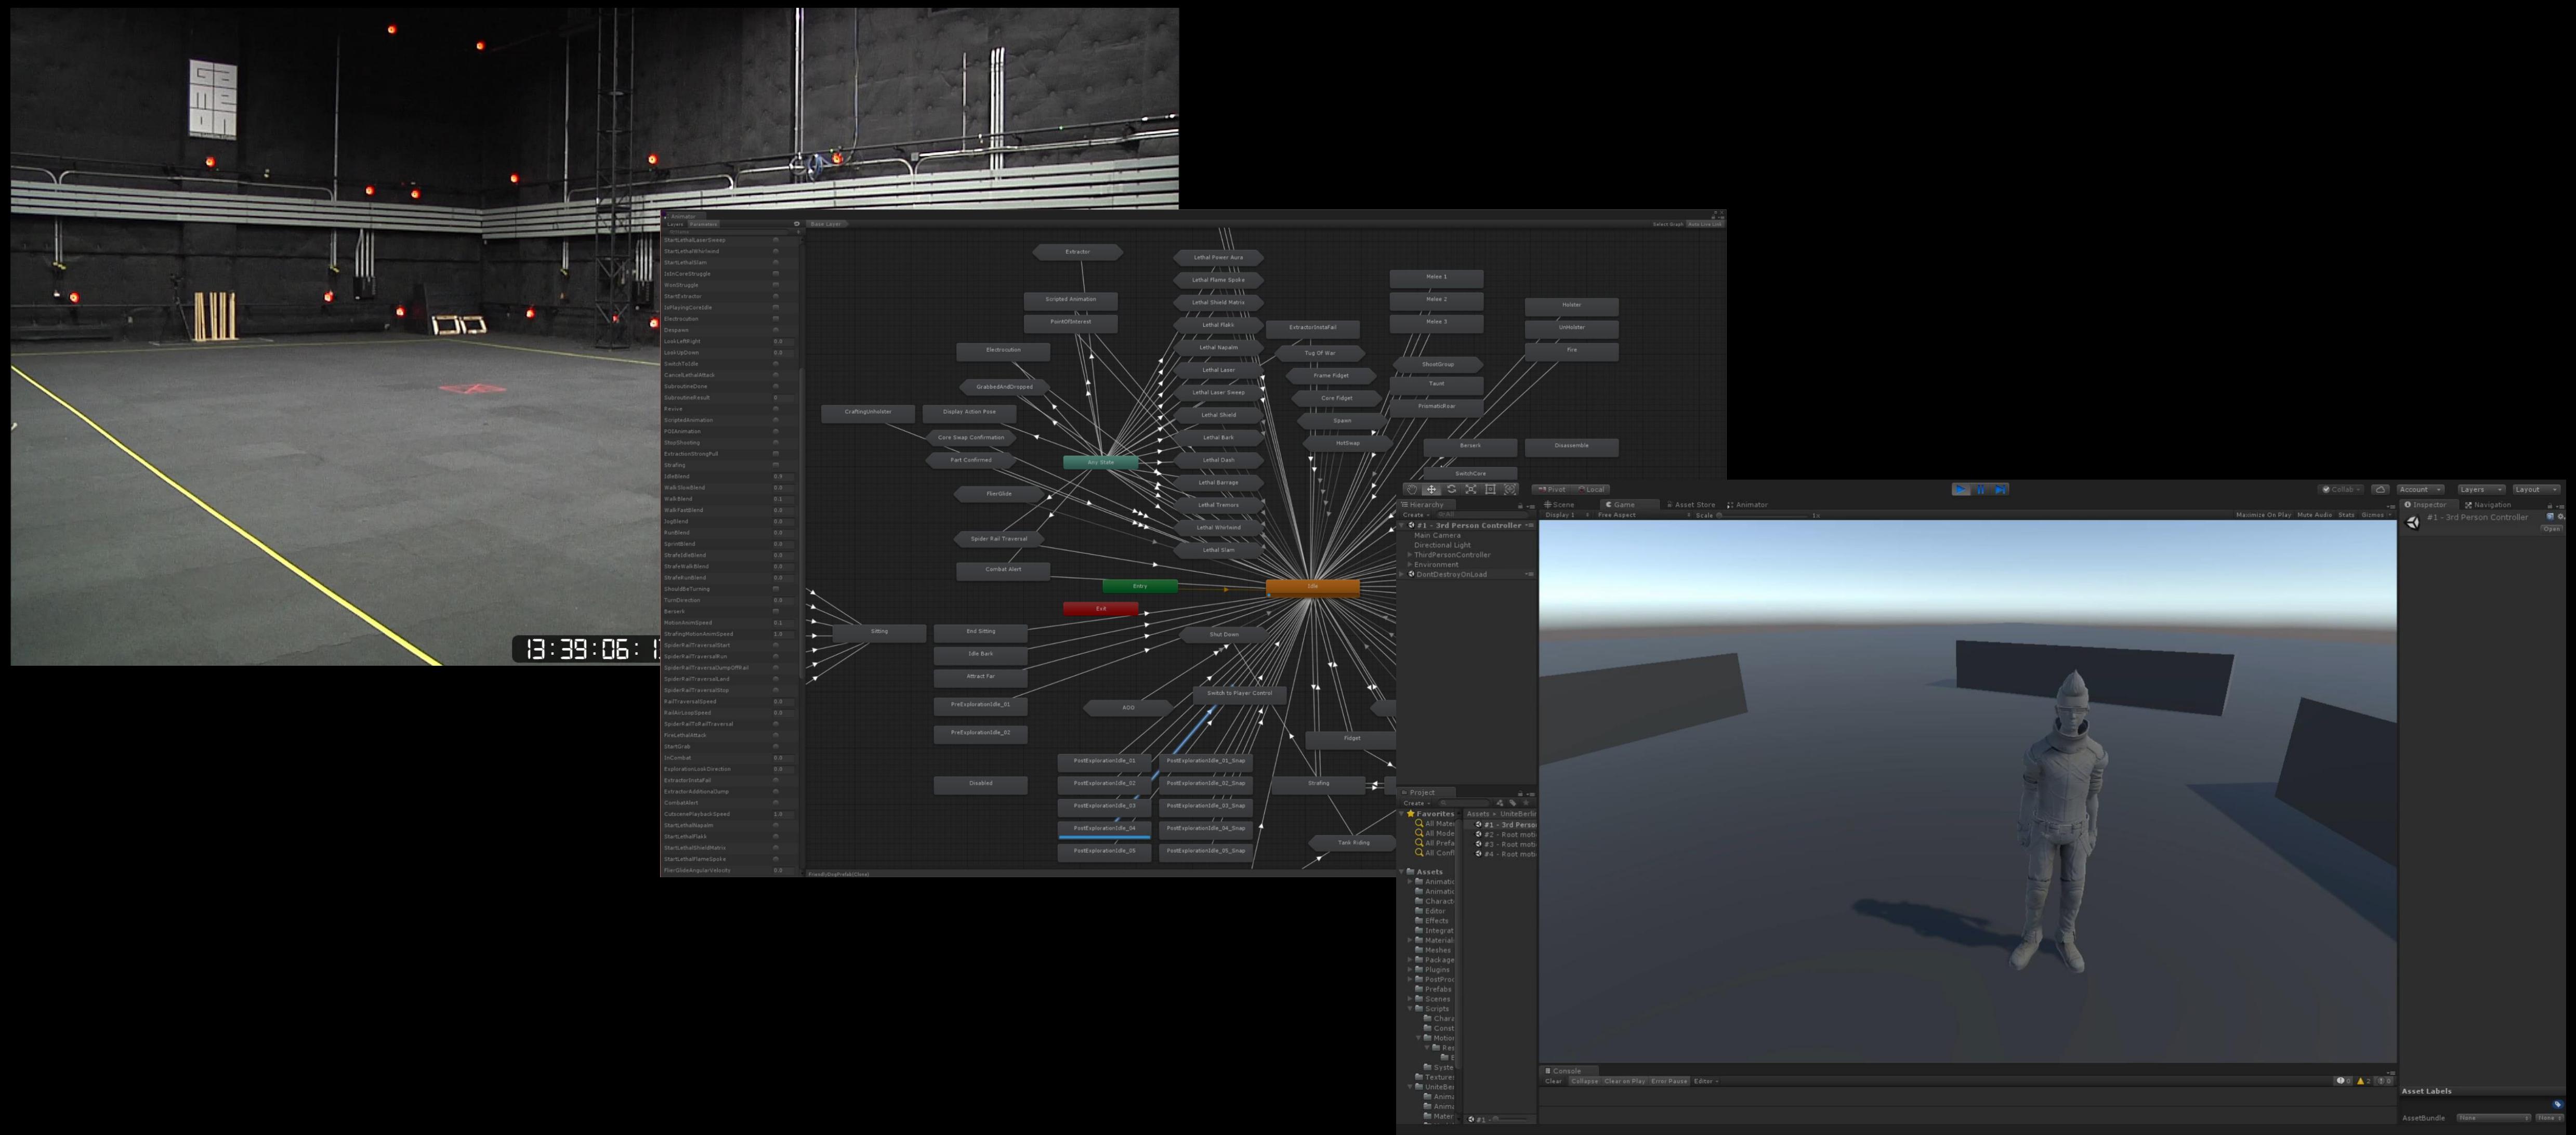
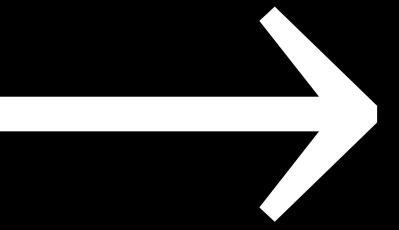
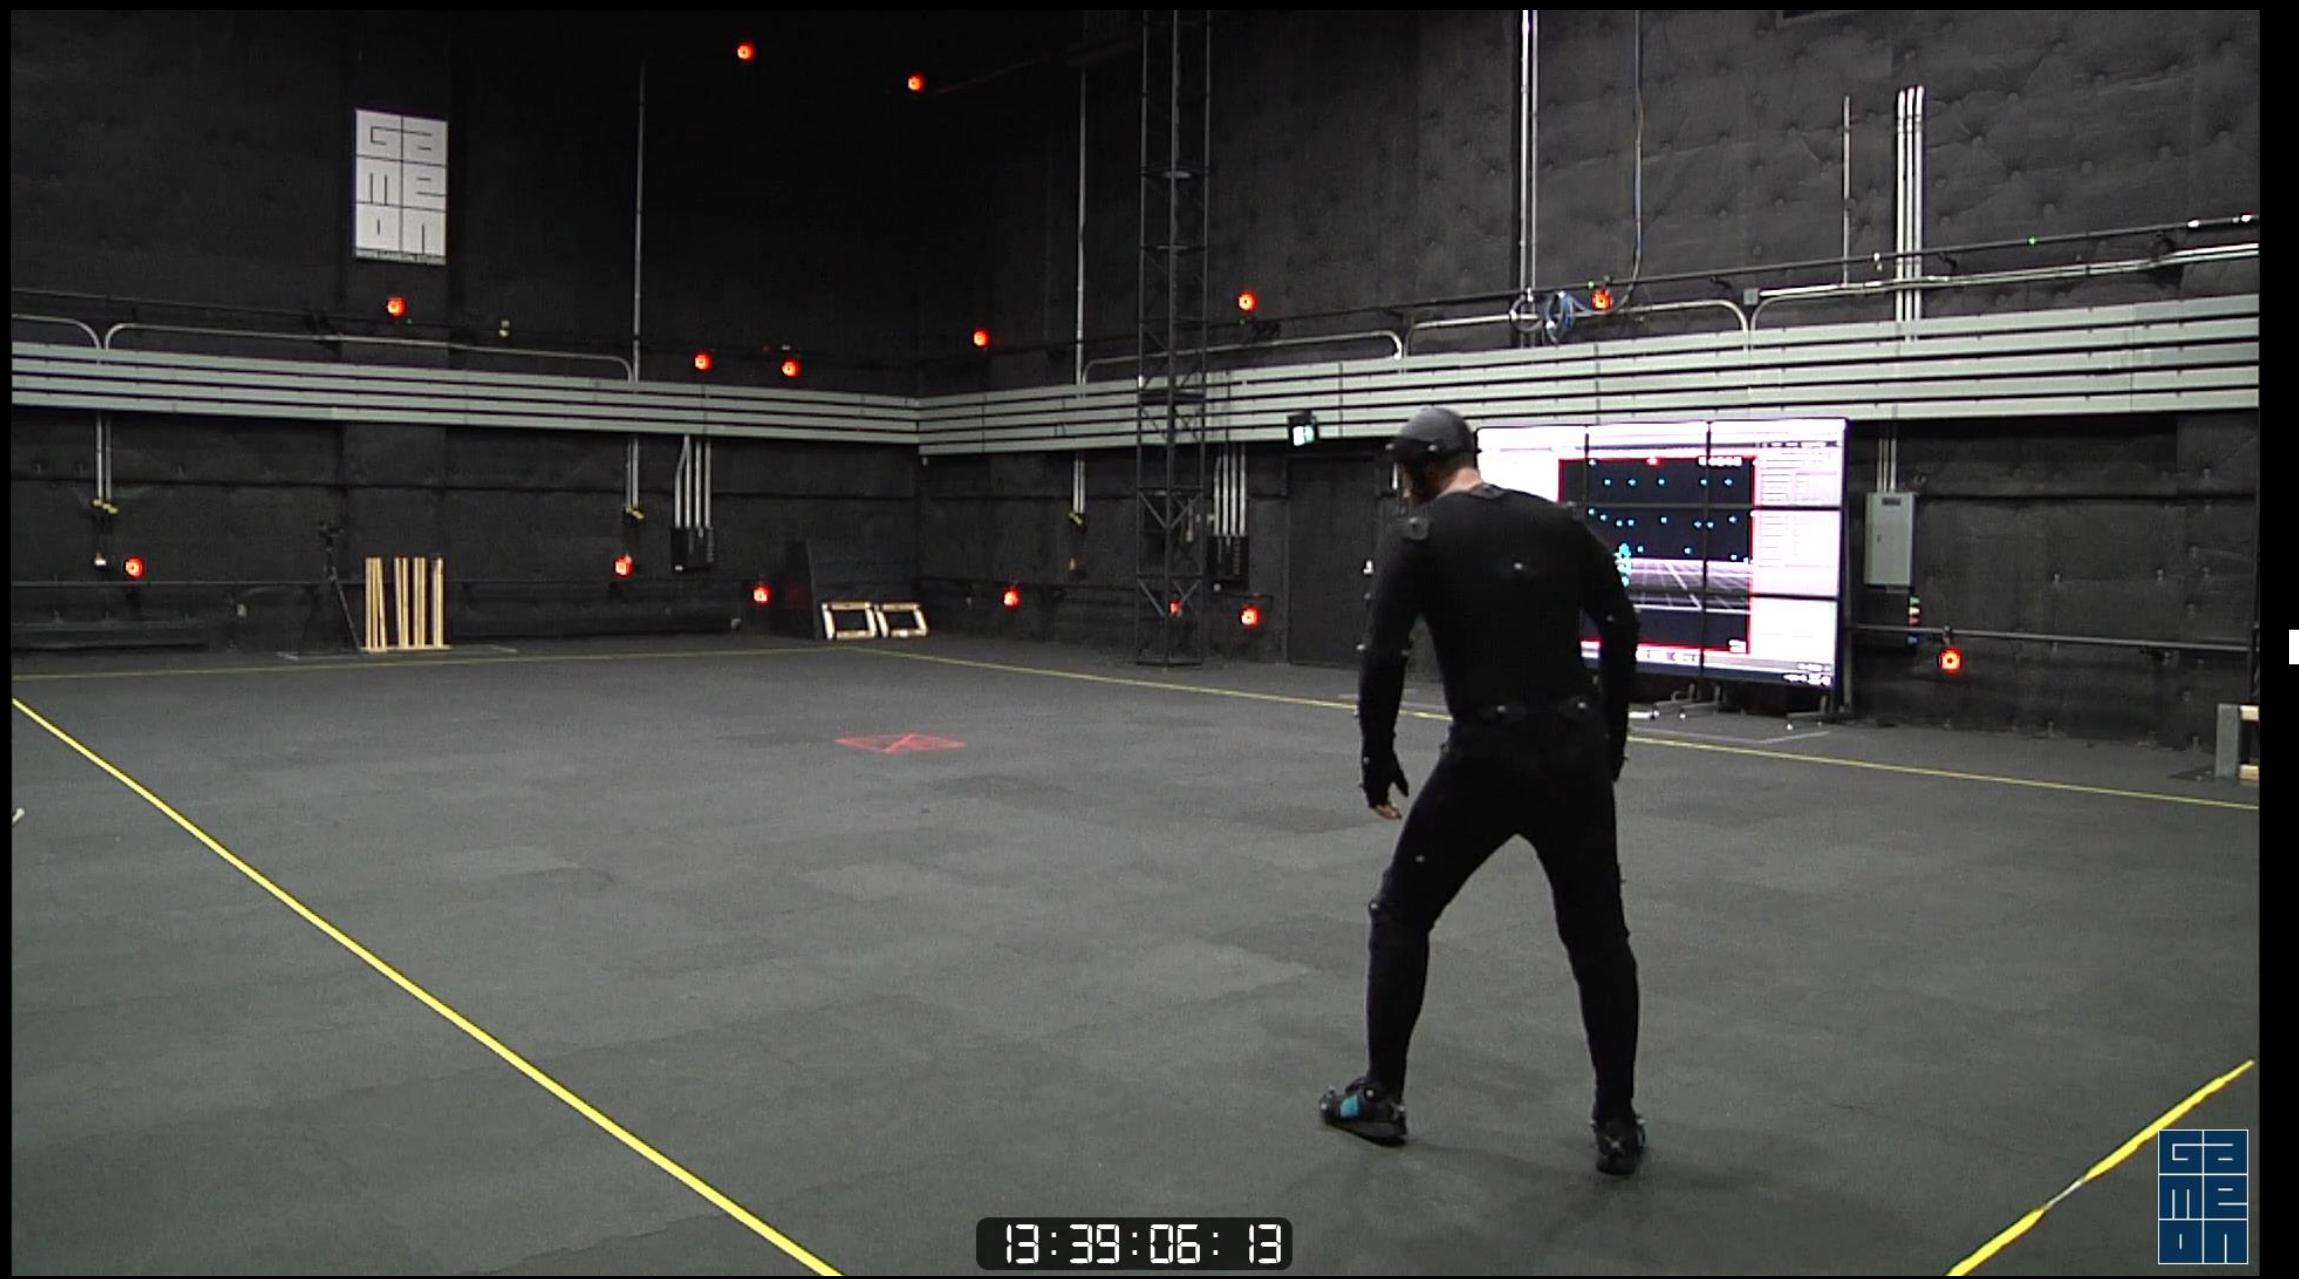


# Machine Learning for Motion Synthesis and Character Control in Games

Michael Buttner, Unity Labs  
Principal Research Engineer  
[michaelbu@unity3d.com](mailto:michaelbu@unity3d.com)







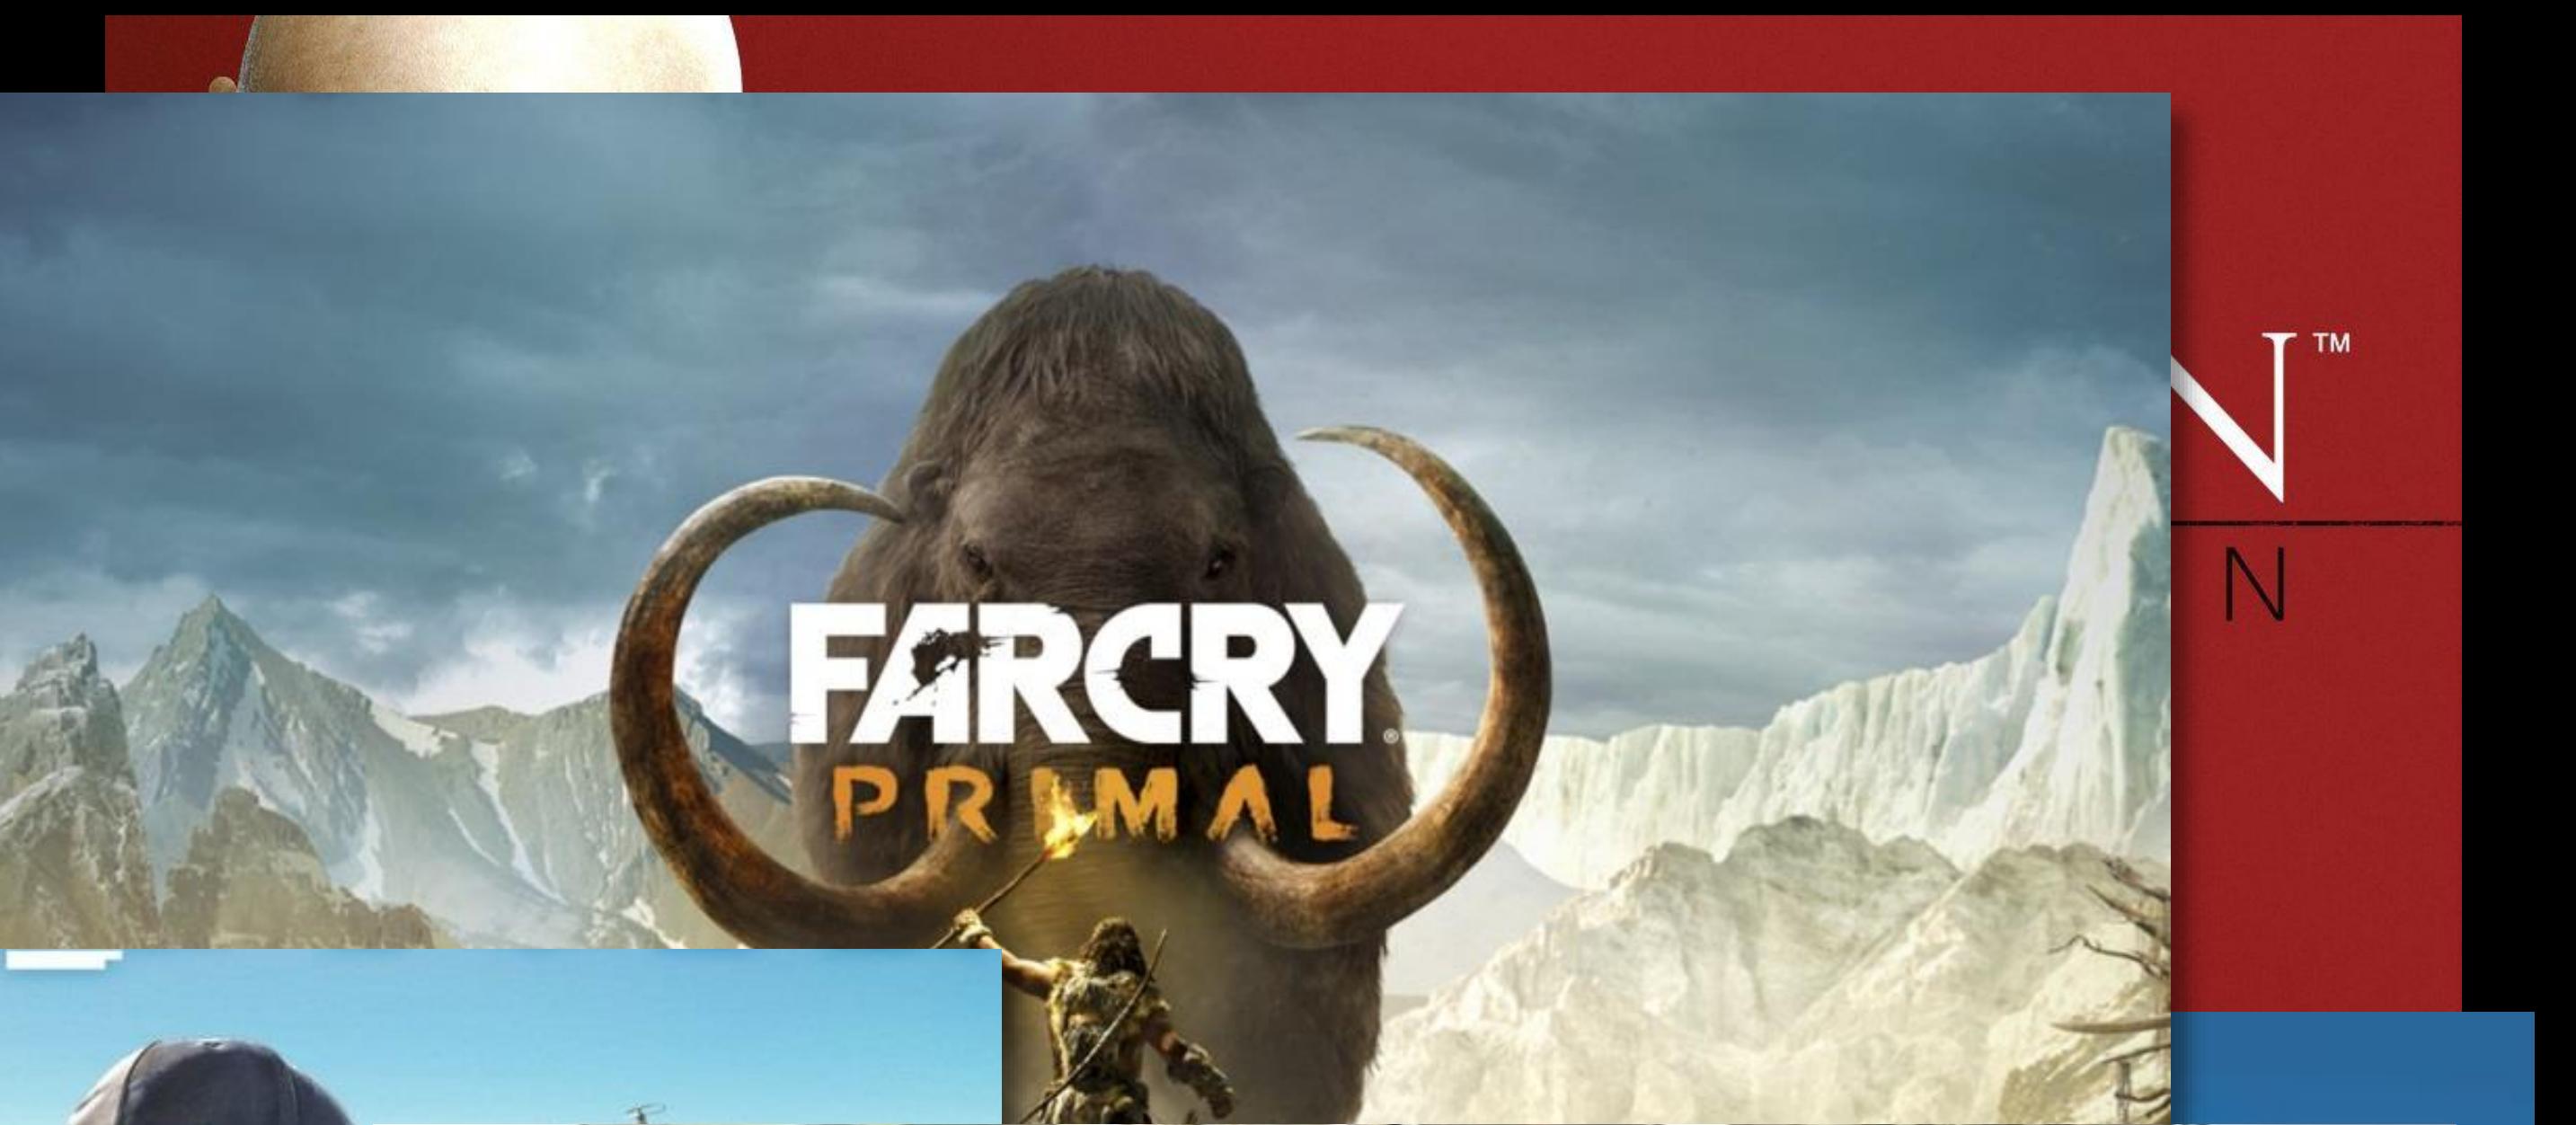
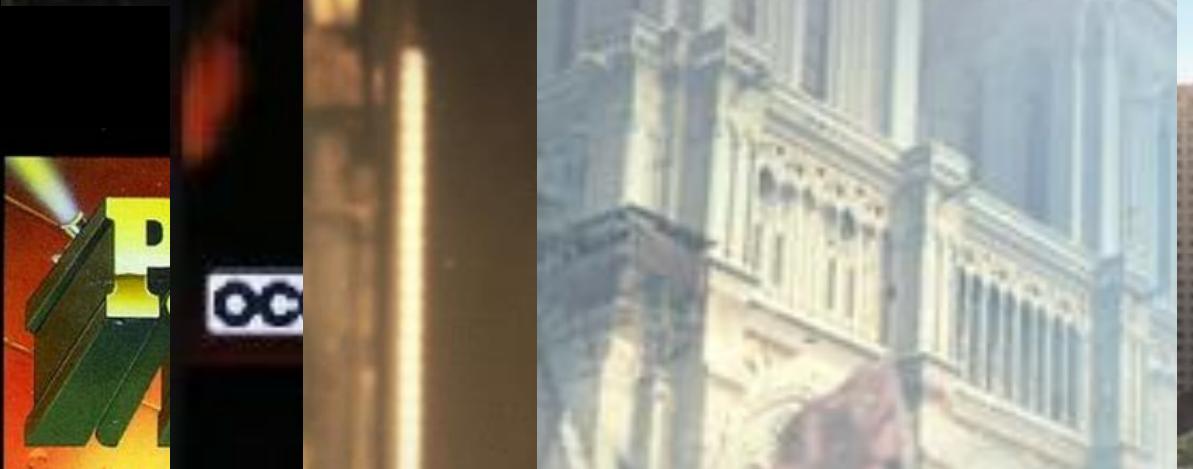
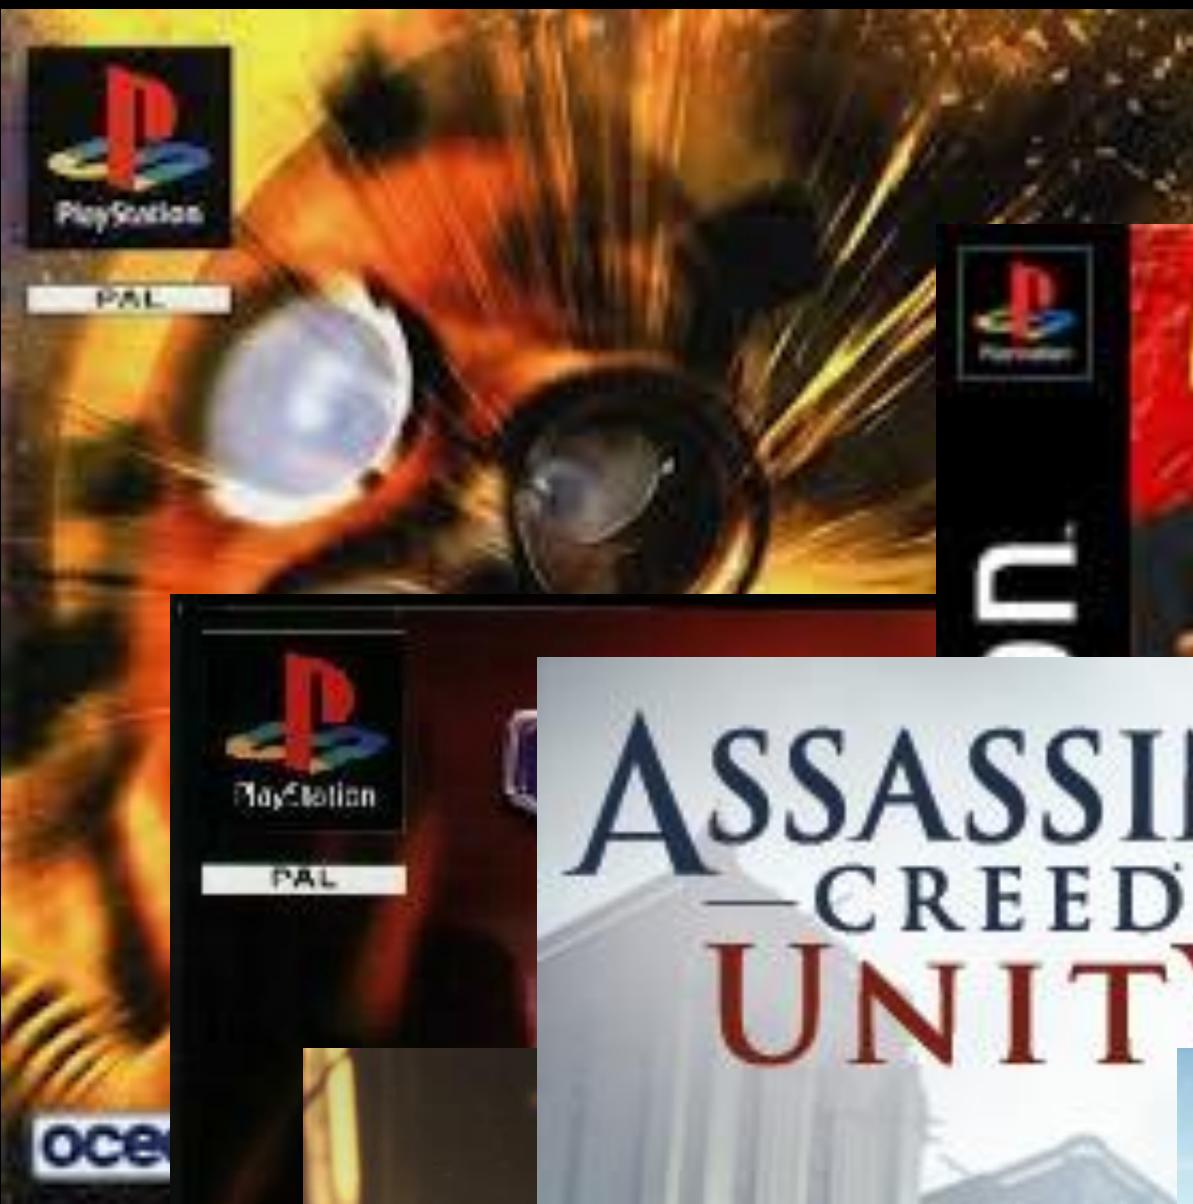




# Kinematica

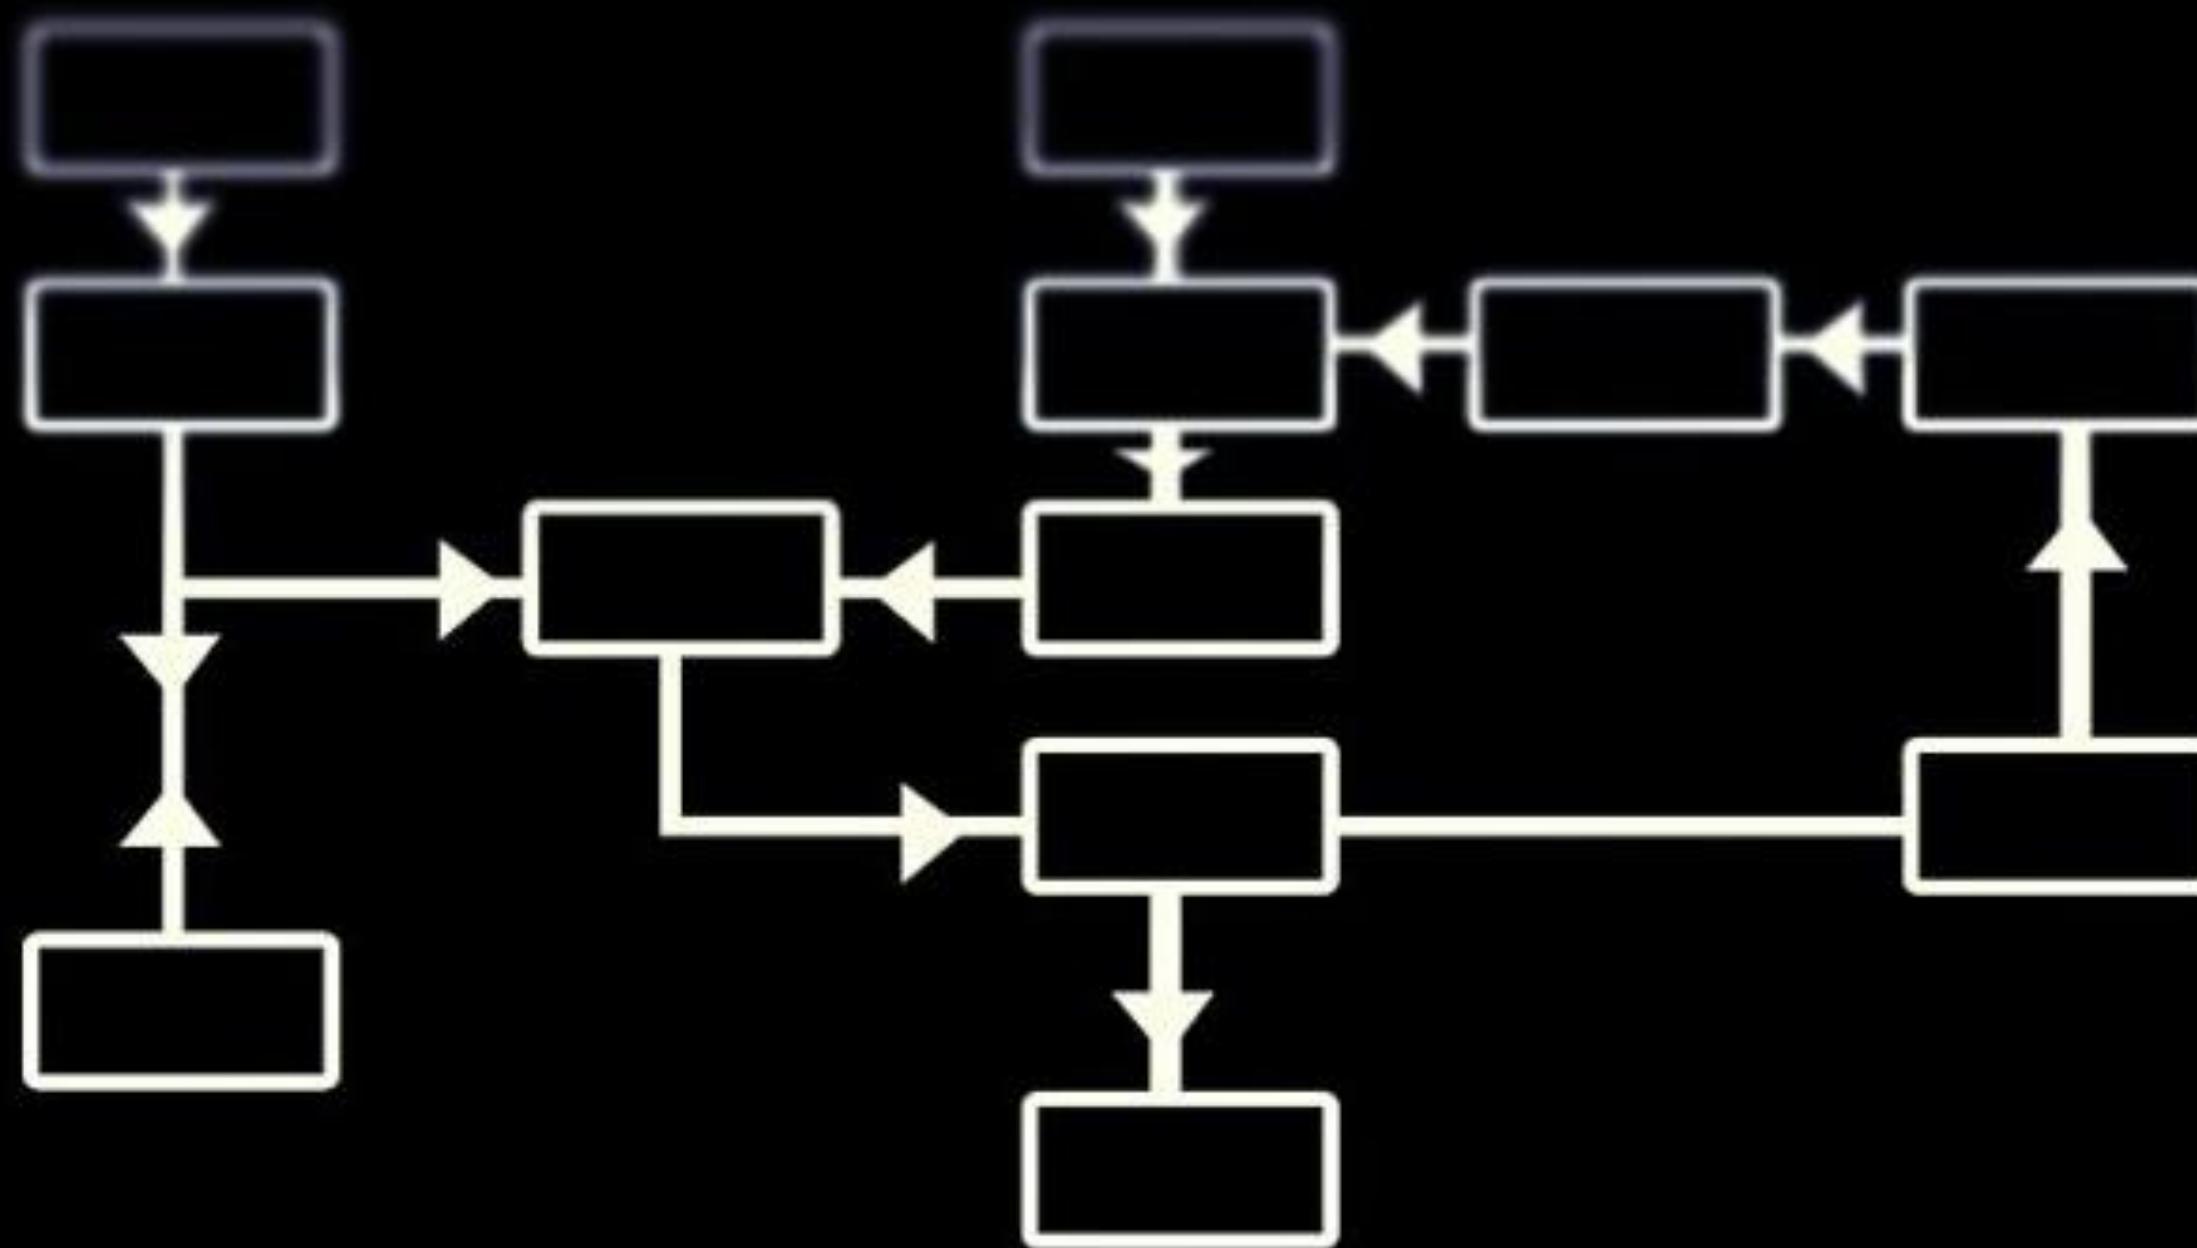
Experimental



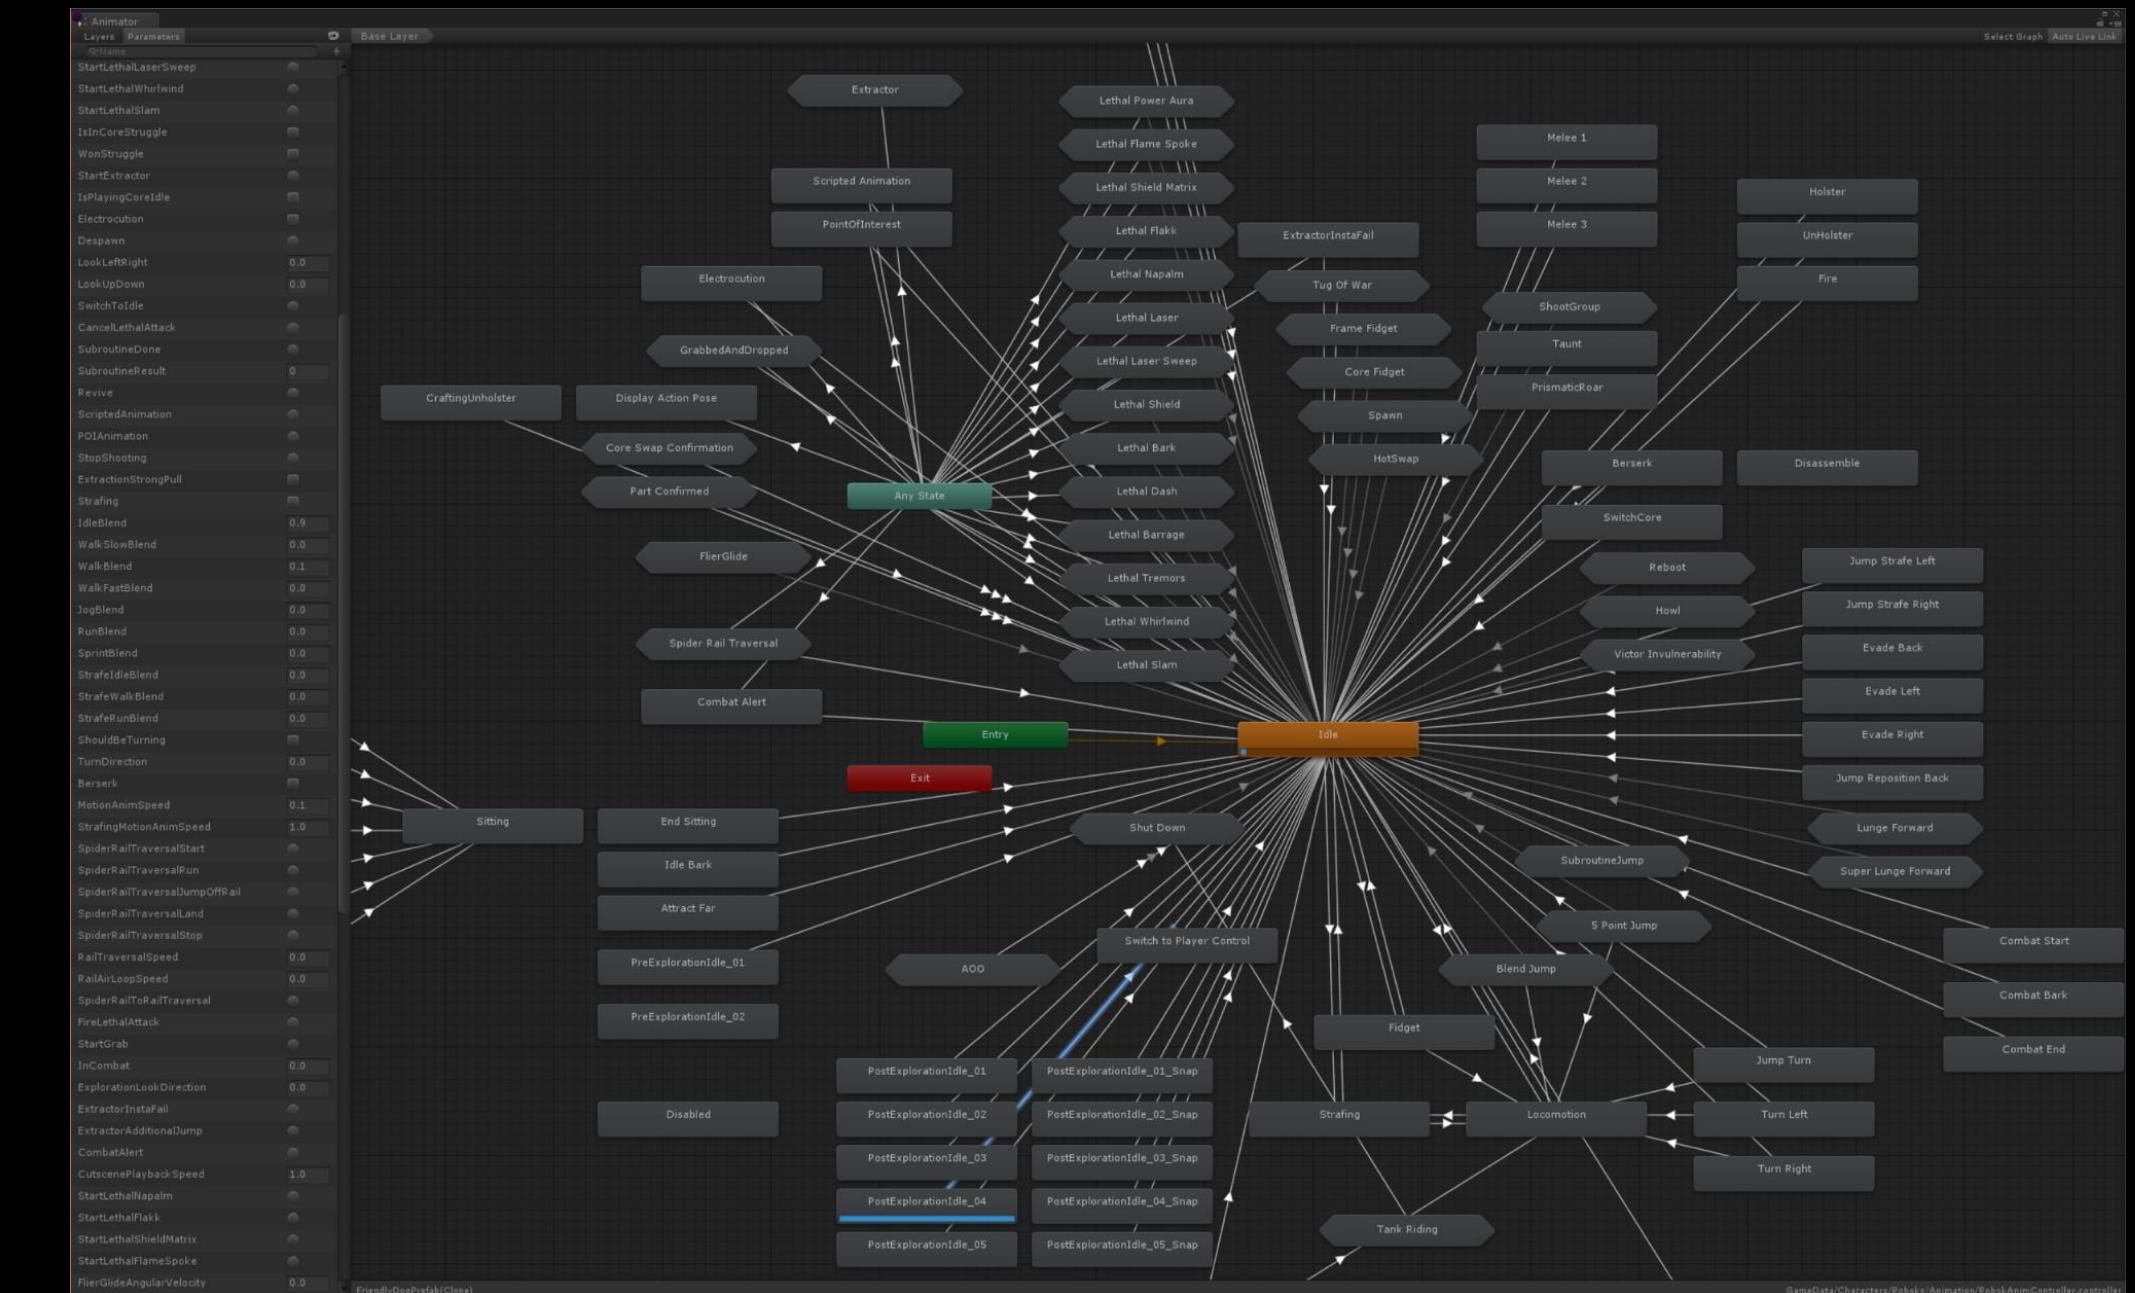


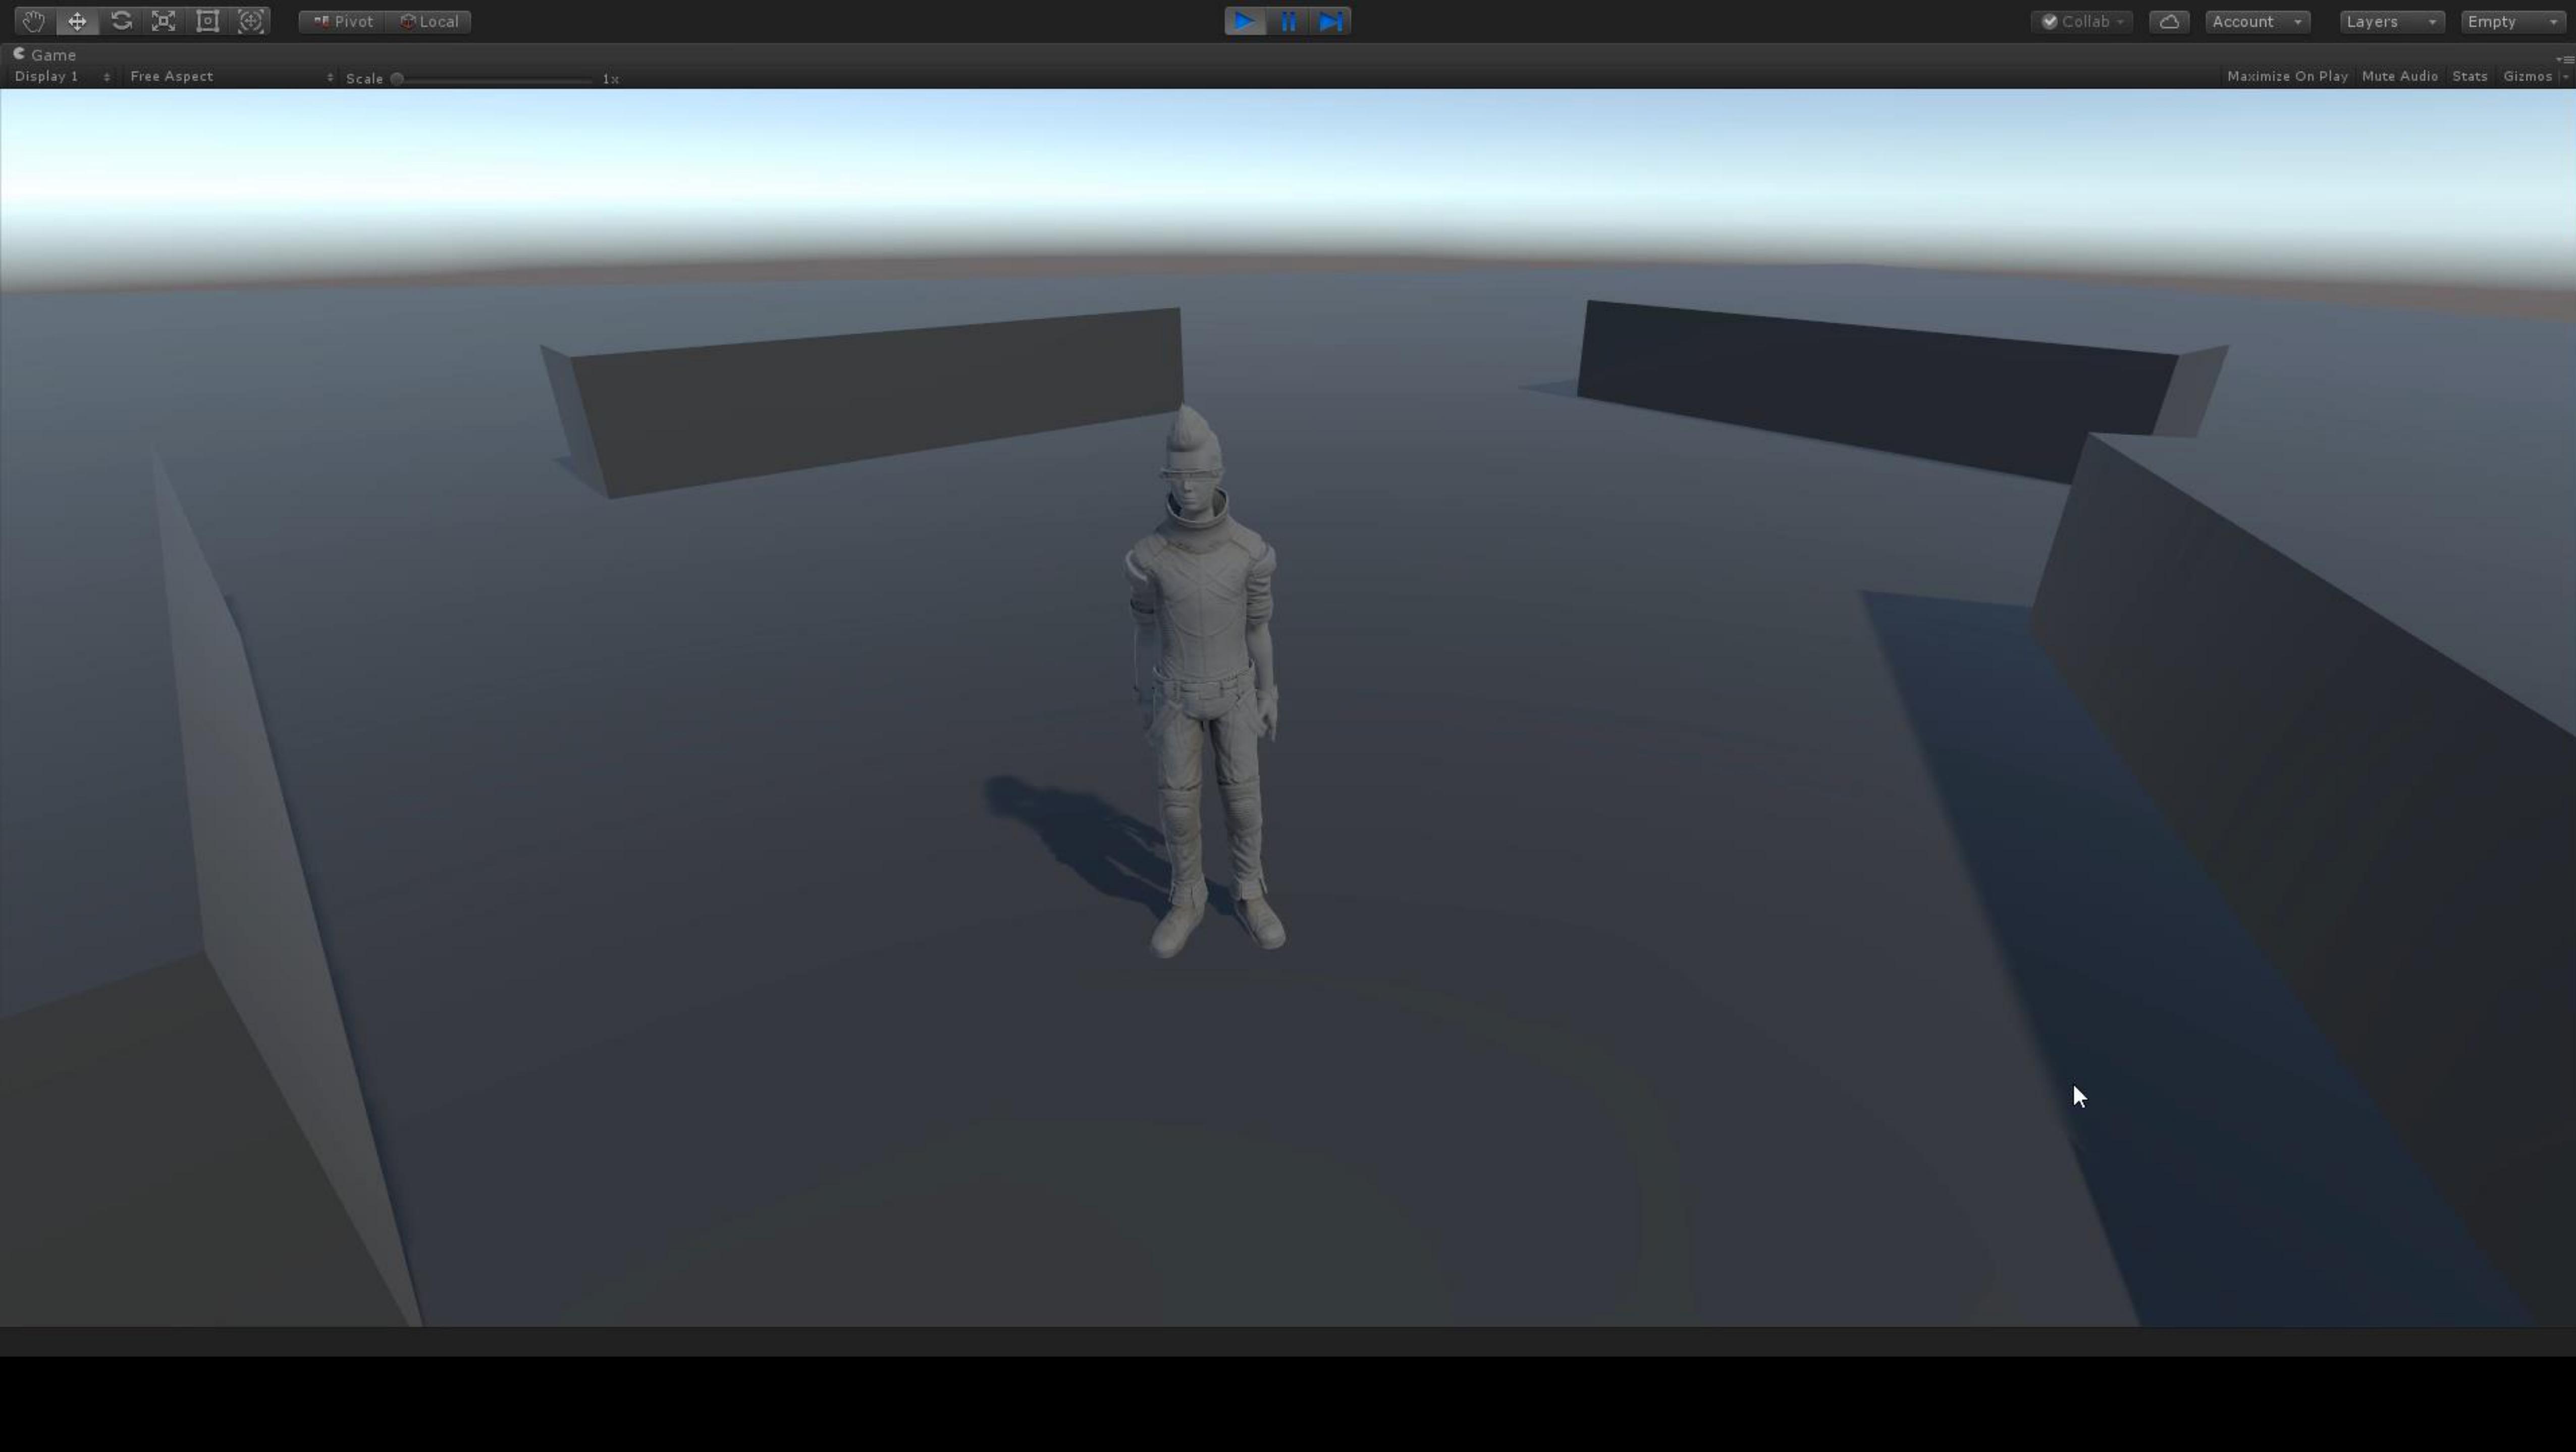
# Game State | Animation Graphs

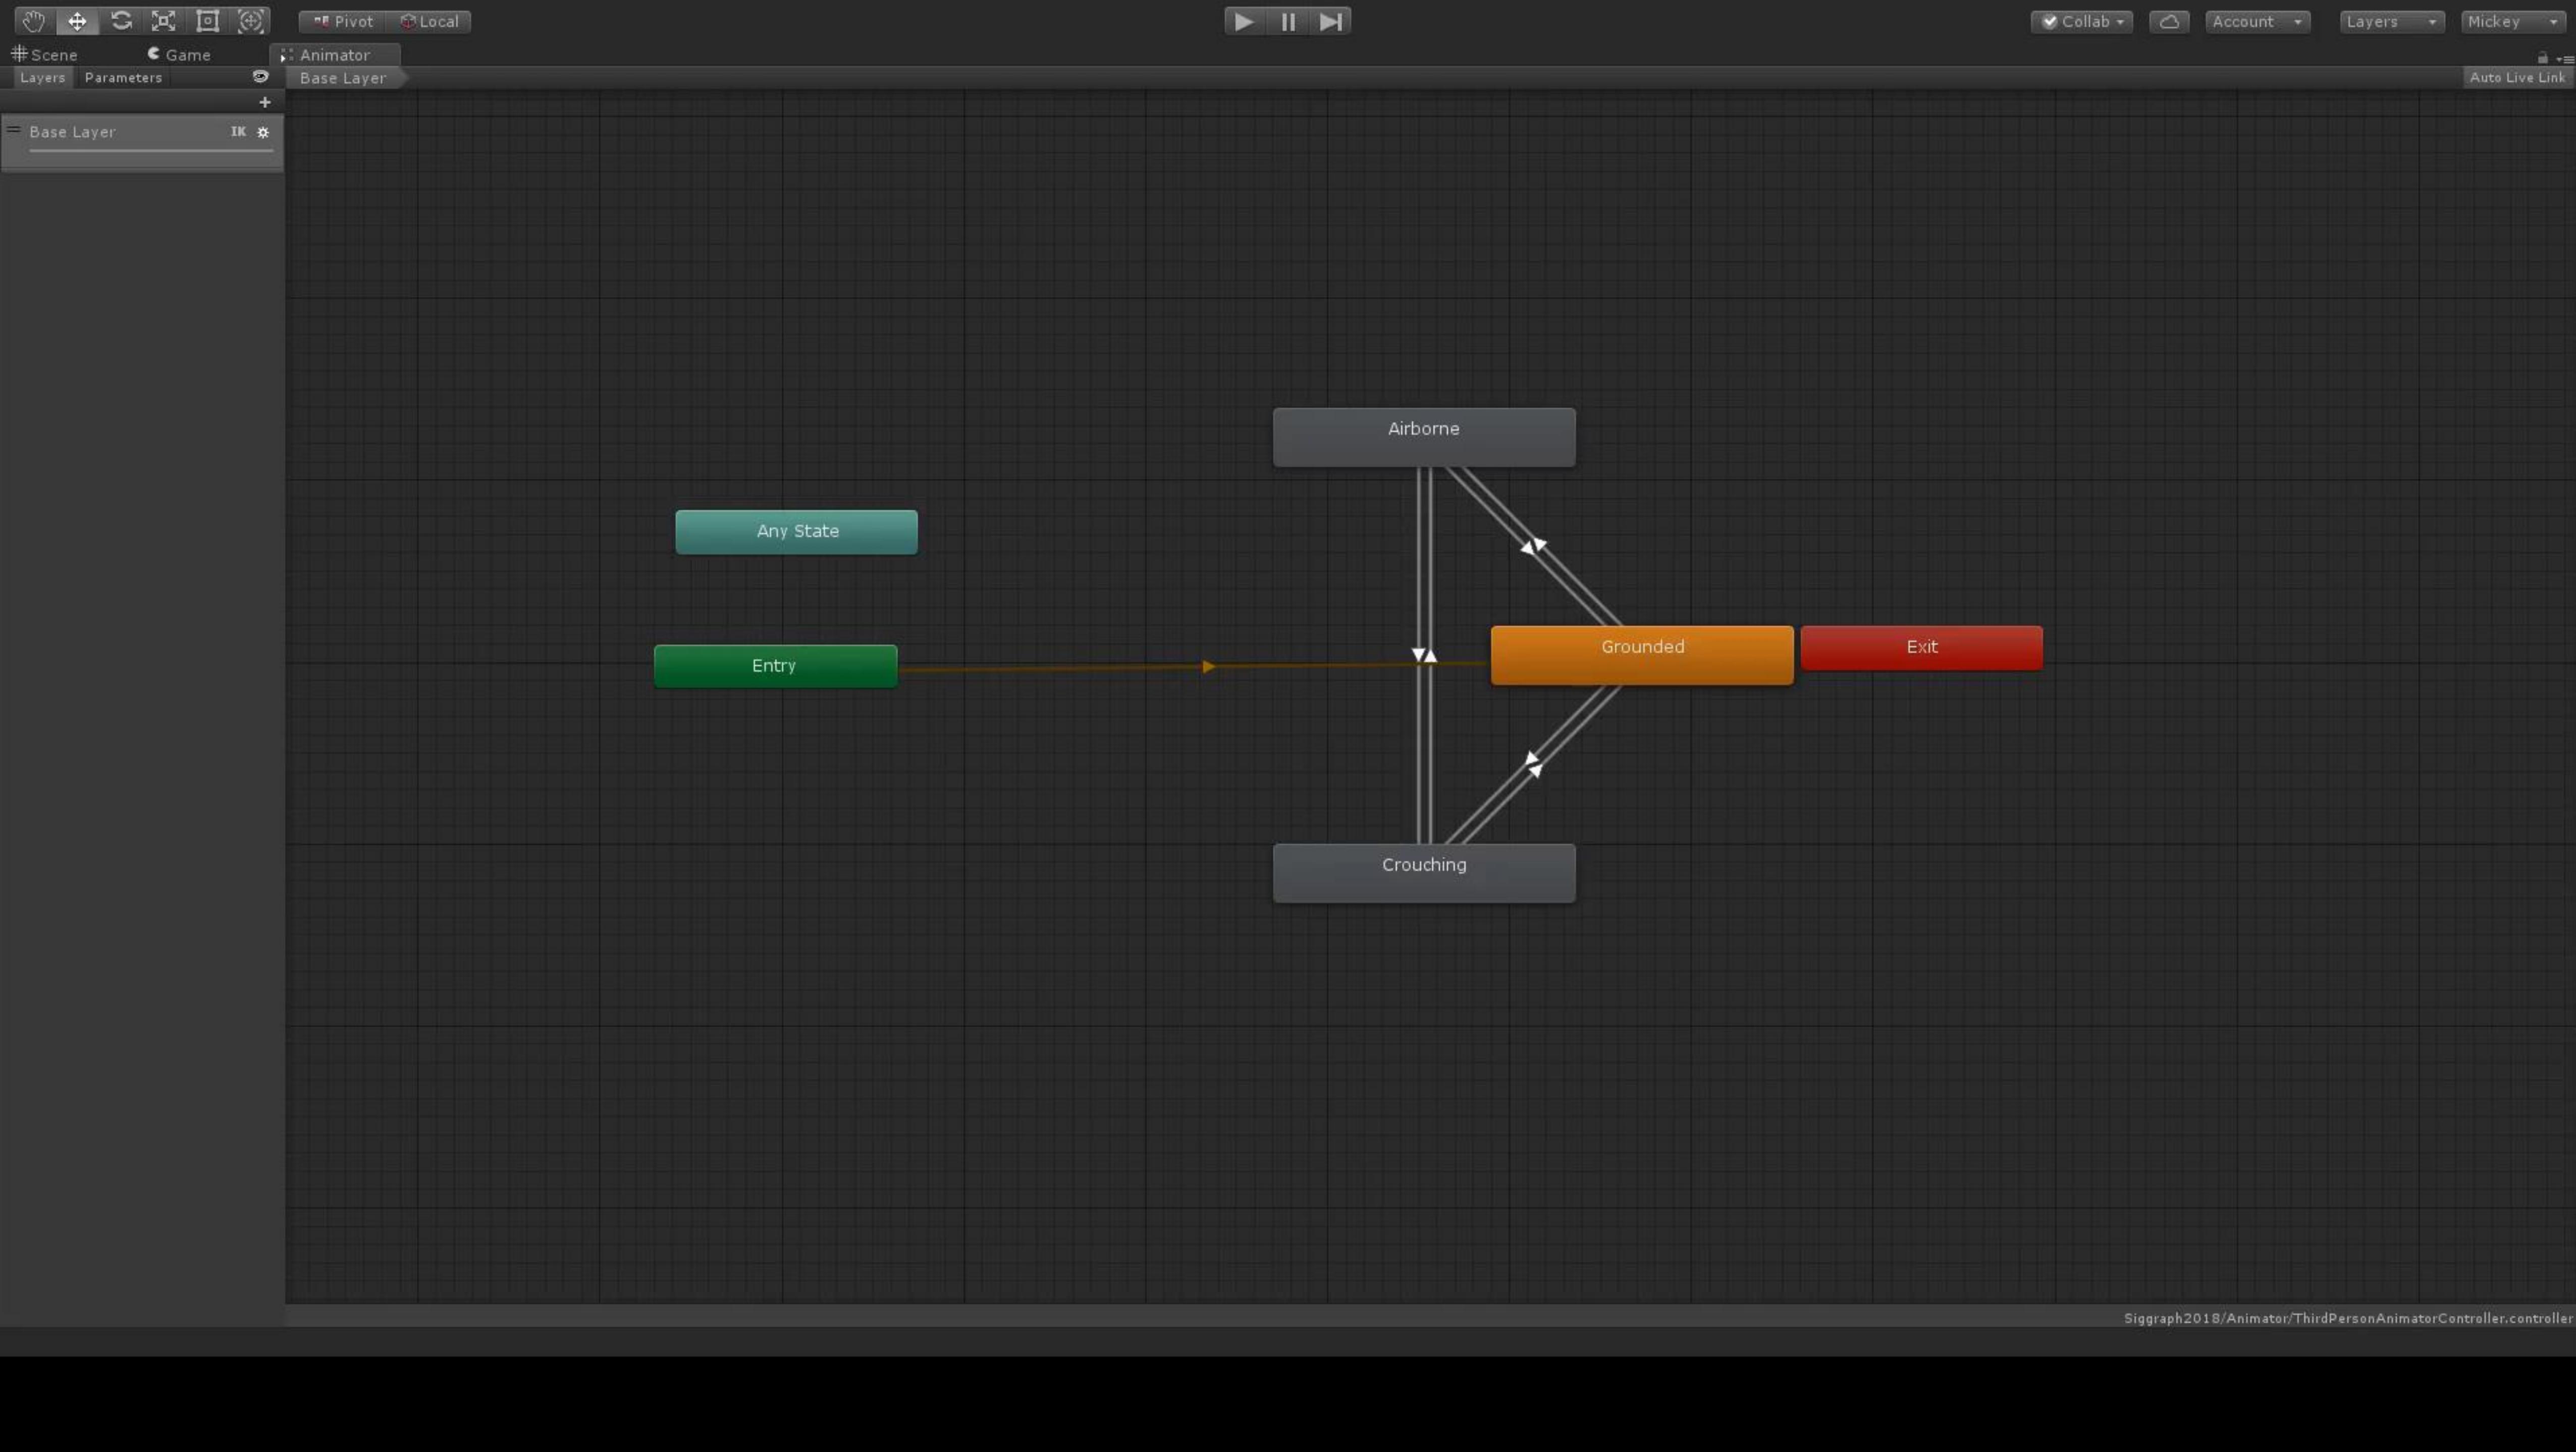
# Game Logic



# Animation Logic







Samples - Microsoft Visual Studio

File Edit View VAssistX Project Build Debug Team Nsight Tools Test Analyze Window Help

Debug Any CPU Attach to Unity

ThirdPersonCharacter.cs

D:\Workspace\ECSAnimationPlayer\Samples\Assets\Siggraph2018\Scripts\ThirdPersonCharacter.cs

Assembly-CSharp

```
1 using UnityEngine;
2
3 namespace ThirdPerson
4 {
5     [RequireComponent(typeof(Rigidbody))]
6     [RequireComponent(typeof(CapsuleCollider))]
7     [RequireComponent(typeof(Animator))]
8     public class ThirdPersonCharacter : MonoBehaviour
9     {
10         void Start()
11         {
12             animator = GetComponent<Animator>();
13             rigidbody = GetComponent<Rigidbody>();
14
15             rigidbody.constraints = RigidbodyConstraints.FreezeRotationX | RigidbodyConstraints.FreezeRotationY | RigidbodyConstraints.FreezePositionZ;
16             originalGroundCheckDistance = groundCheckDistance;
17         }
18
19         private void FixedUpdate()
20         {
21             // Read current gamepad state
22             // ...
23
24             float h = Input.GetAxis("Left Analog Horizontal");
25             float v = -Input.GetAxis("Left Analog Vertical");
26         }
27     }
28 }
```

210 %

Ready

Ln 2 Col 1 Ch 1 INS ↑ 0 ↗ 1 ECSAnimationPlayer master

# Animation Graphs

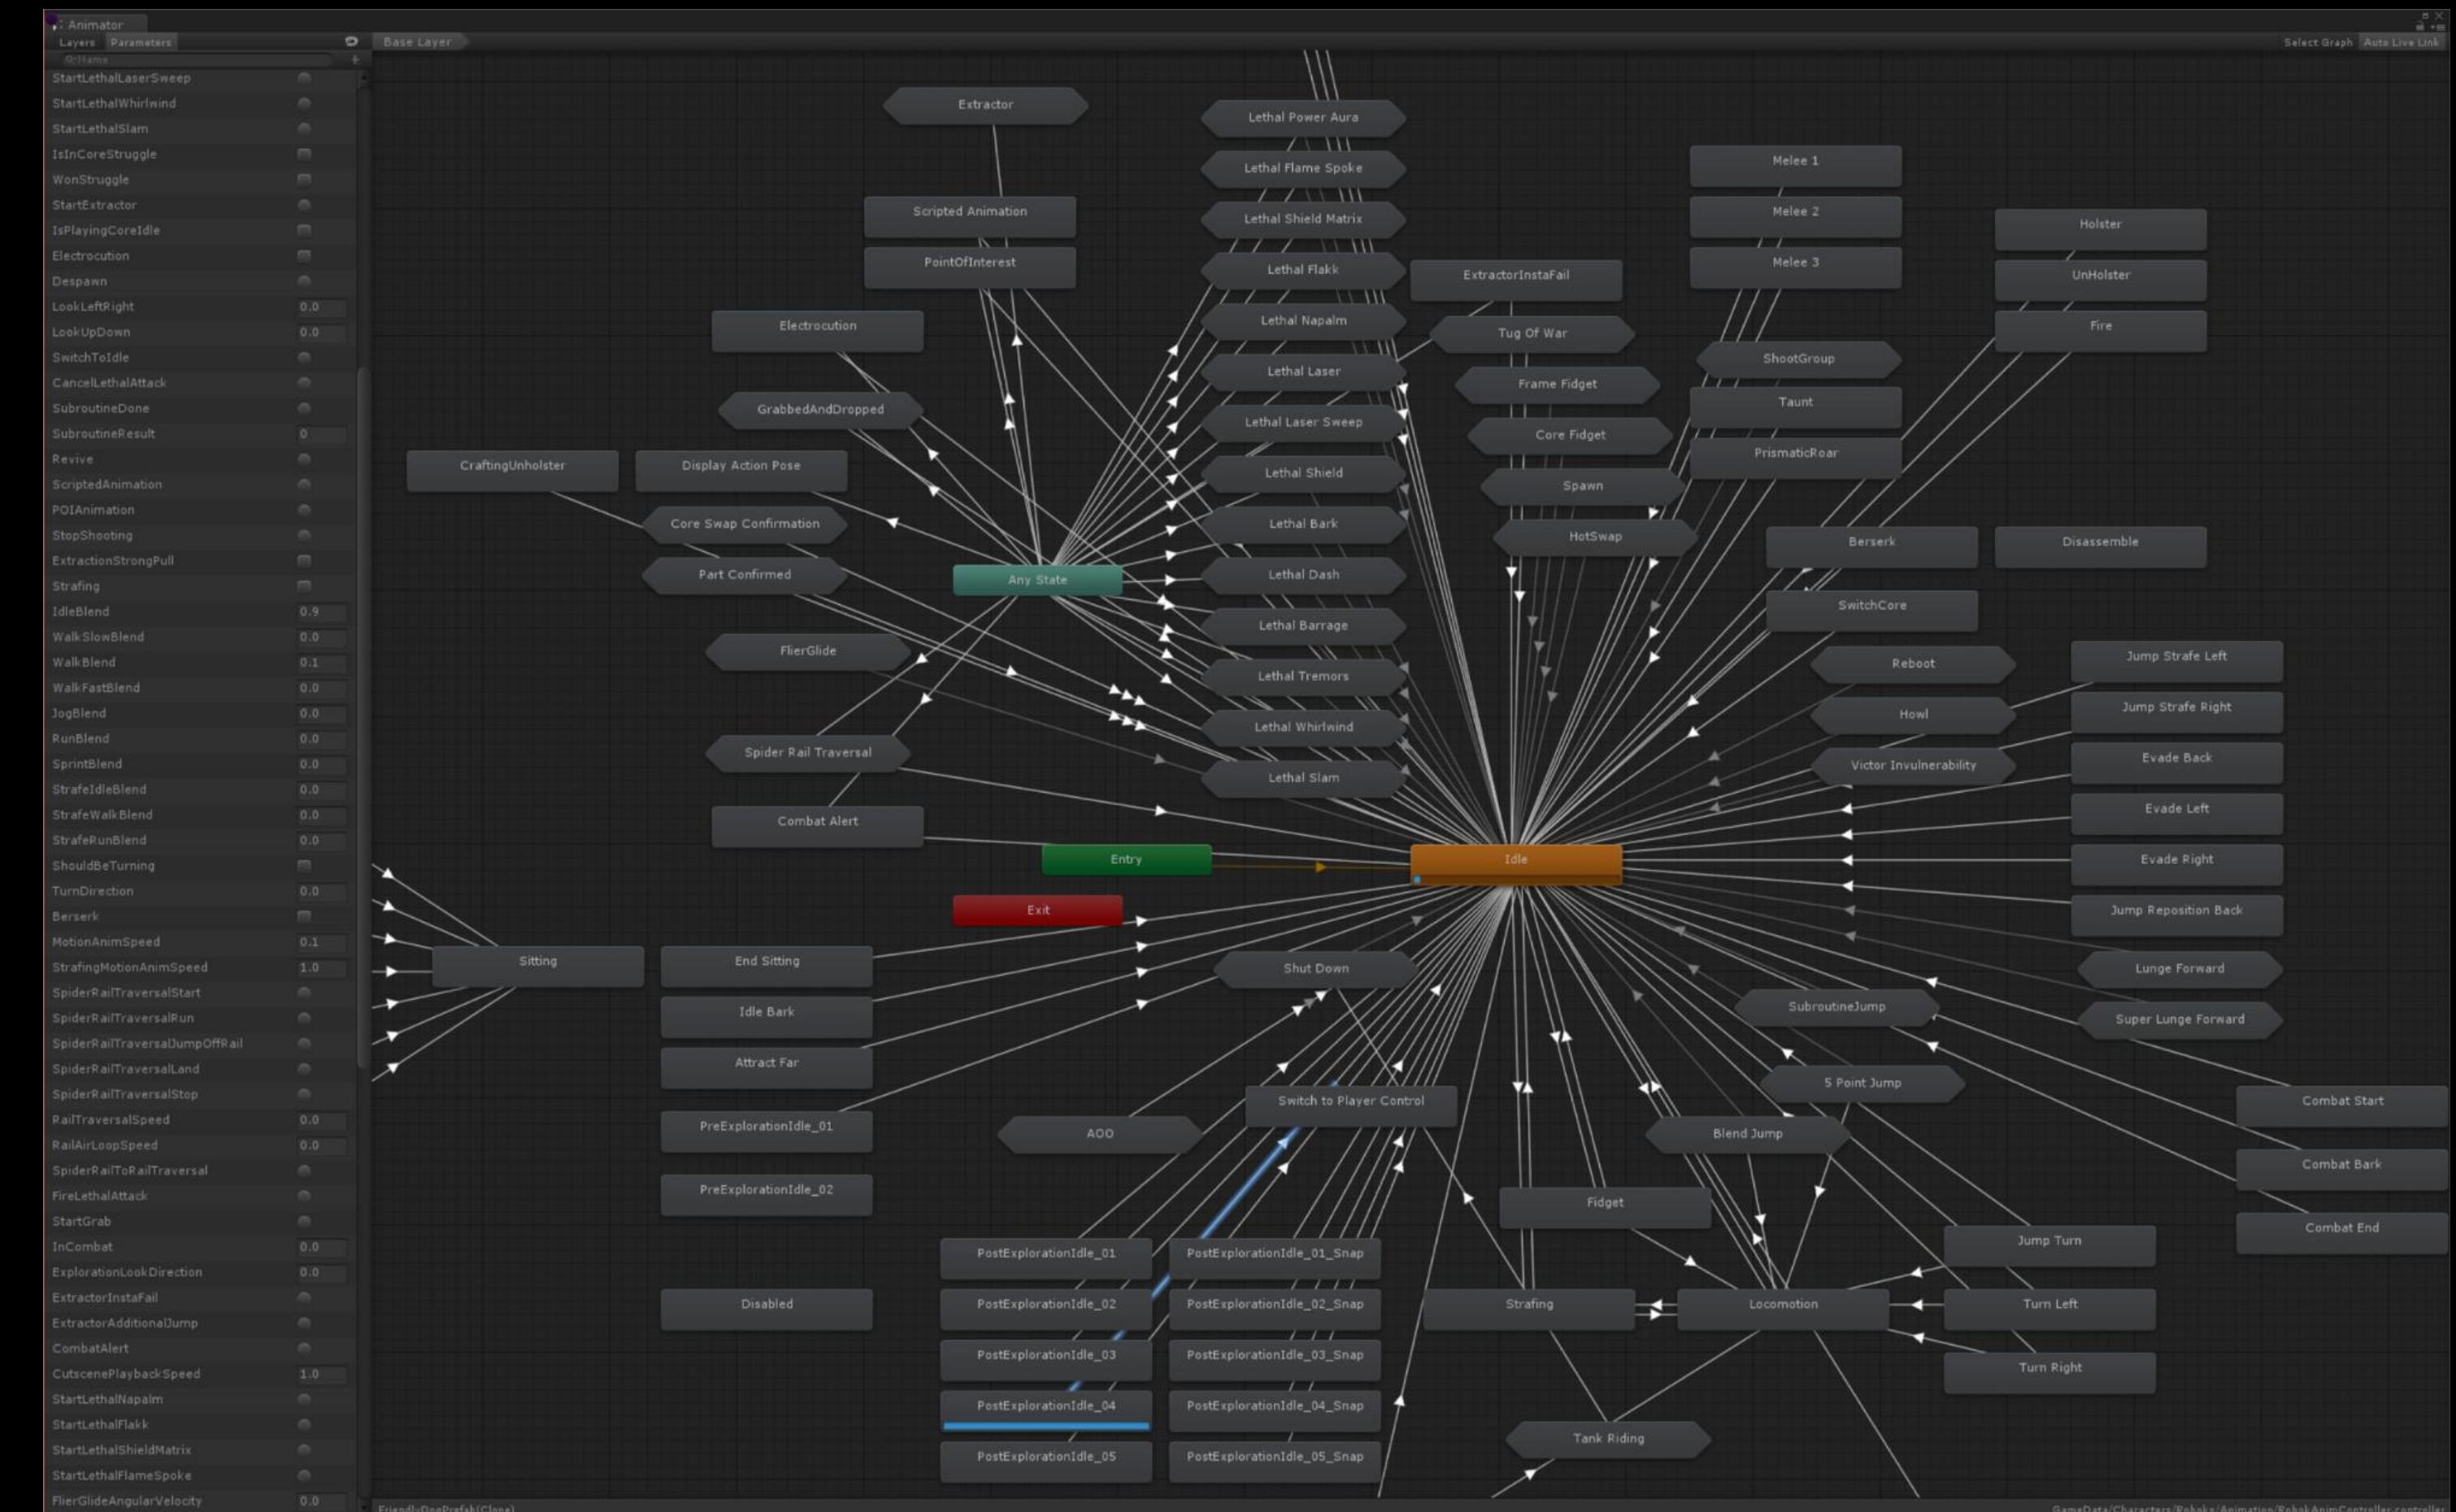
# Provide structure for animation clips

# Allow animations to be “addressable”

# Allow related animation to be grouped

Transitions have to be explicitly spelled out

# Combinatorial explosion





CHINESE NEW YEAR

0

Warp 36,375

Averages FI 89,502 WORLD 82,796



# Unstructured Animations

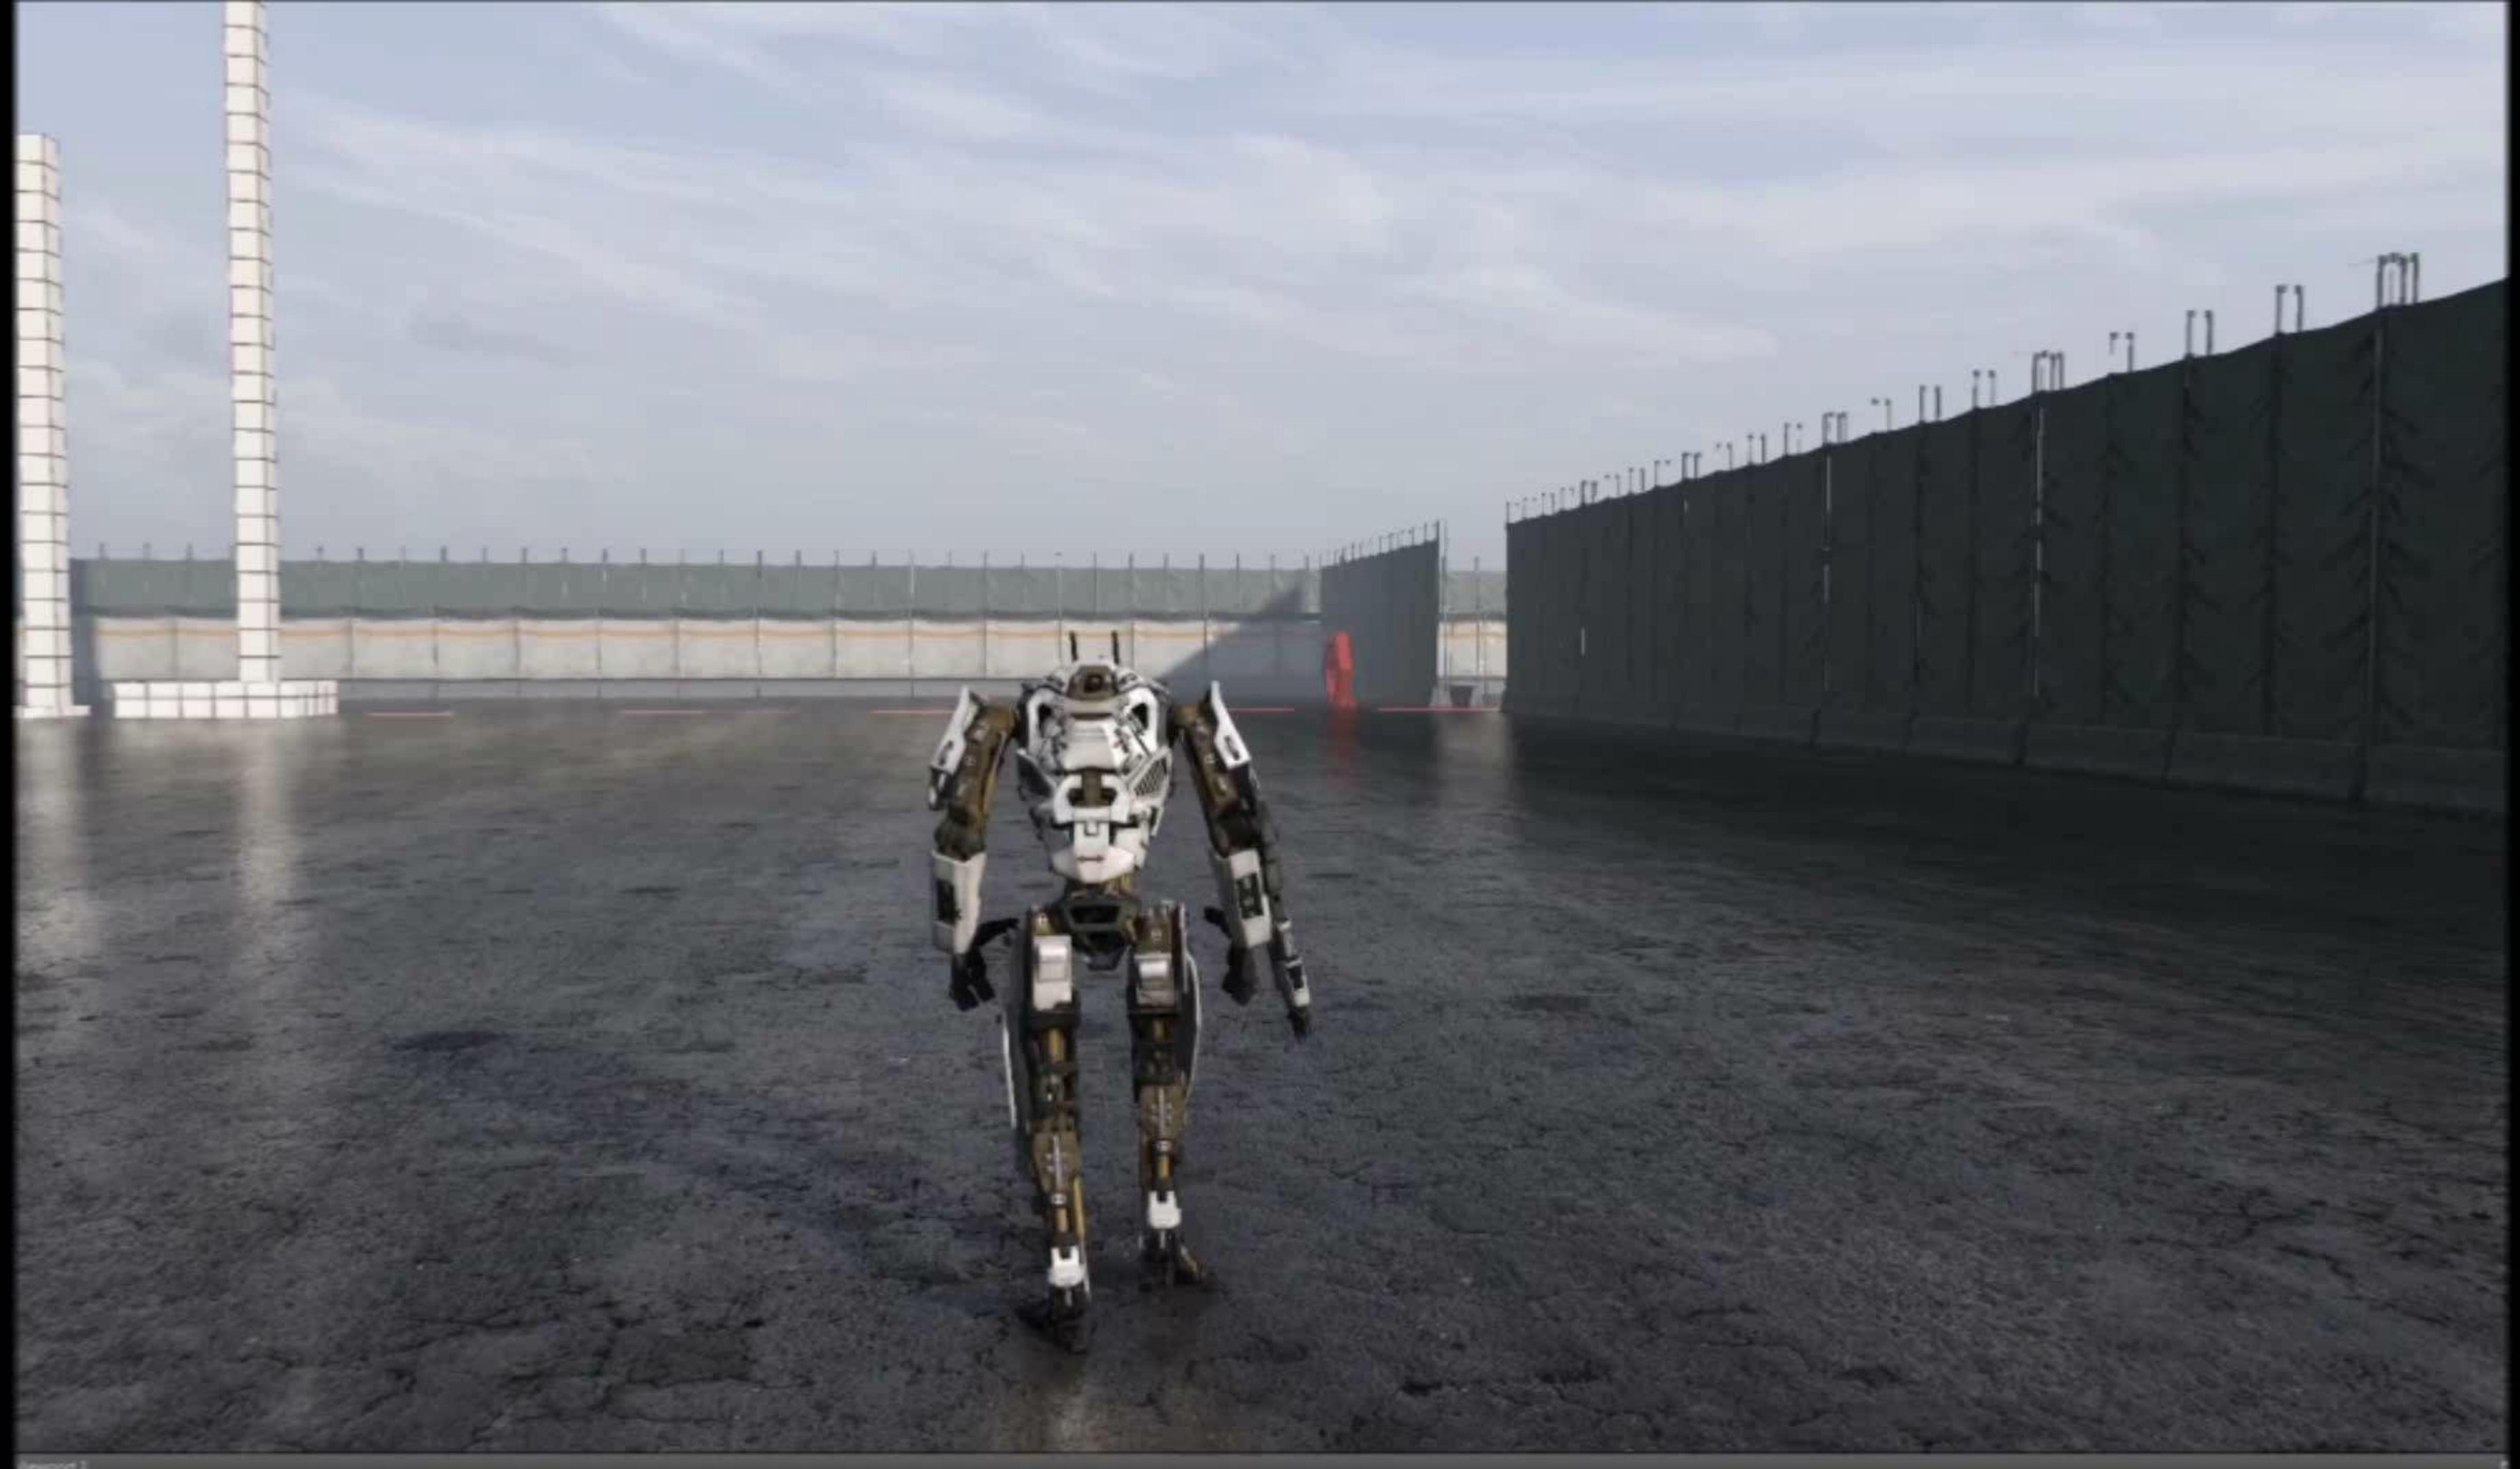
If we would combine all animation clips into a single unstructured library...

How can we infer information from the gameplay state directly?

Can we extract relevant information from the physical properties of the animations themselves?

How can we achieve the same level of versatility as animation graphs?

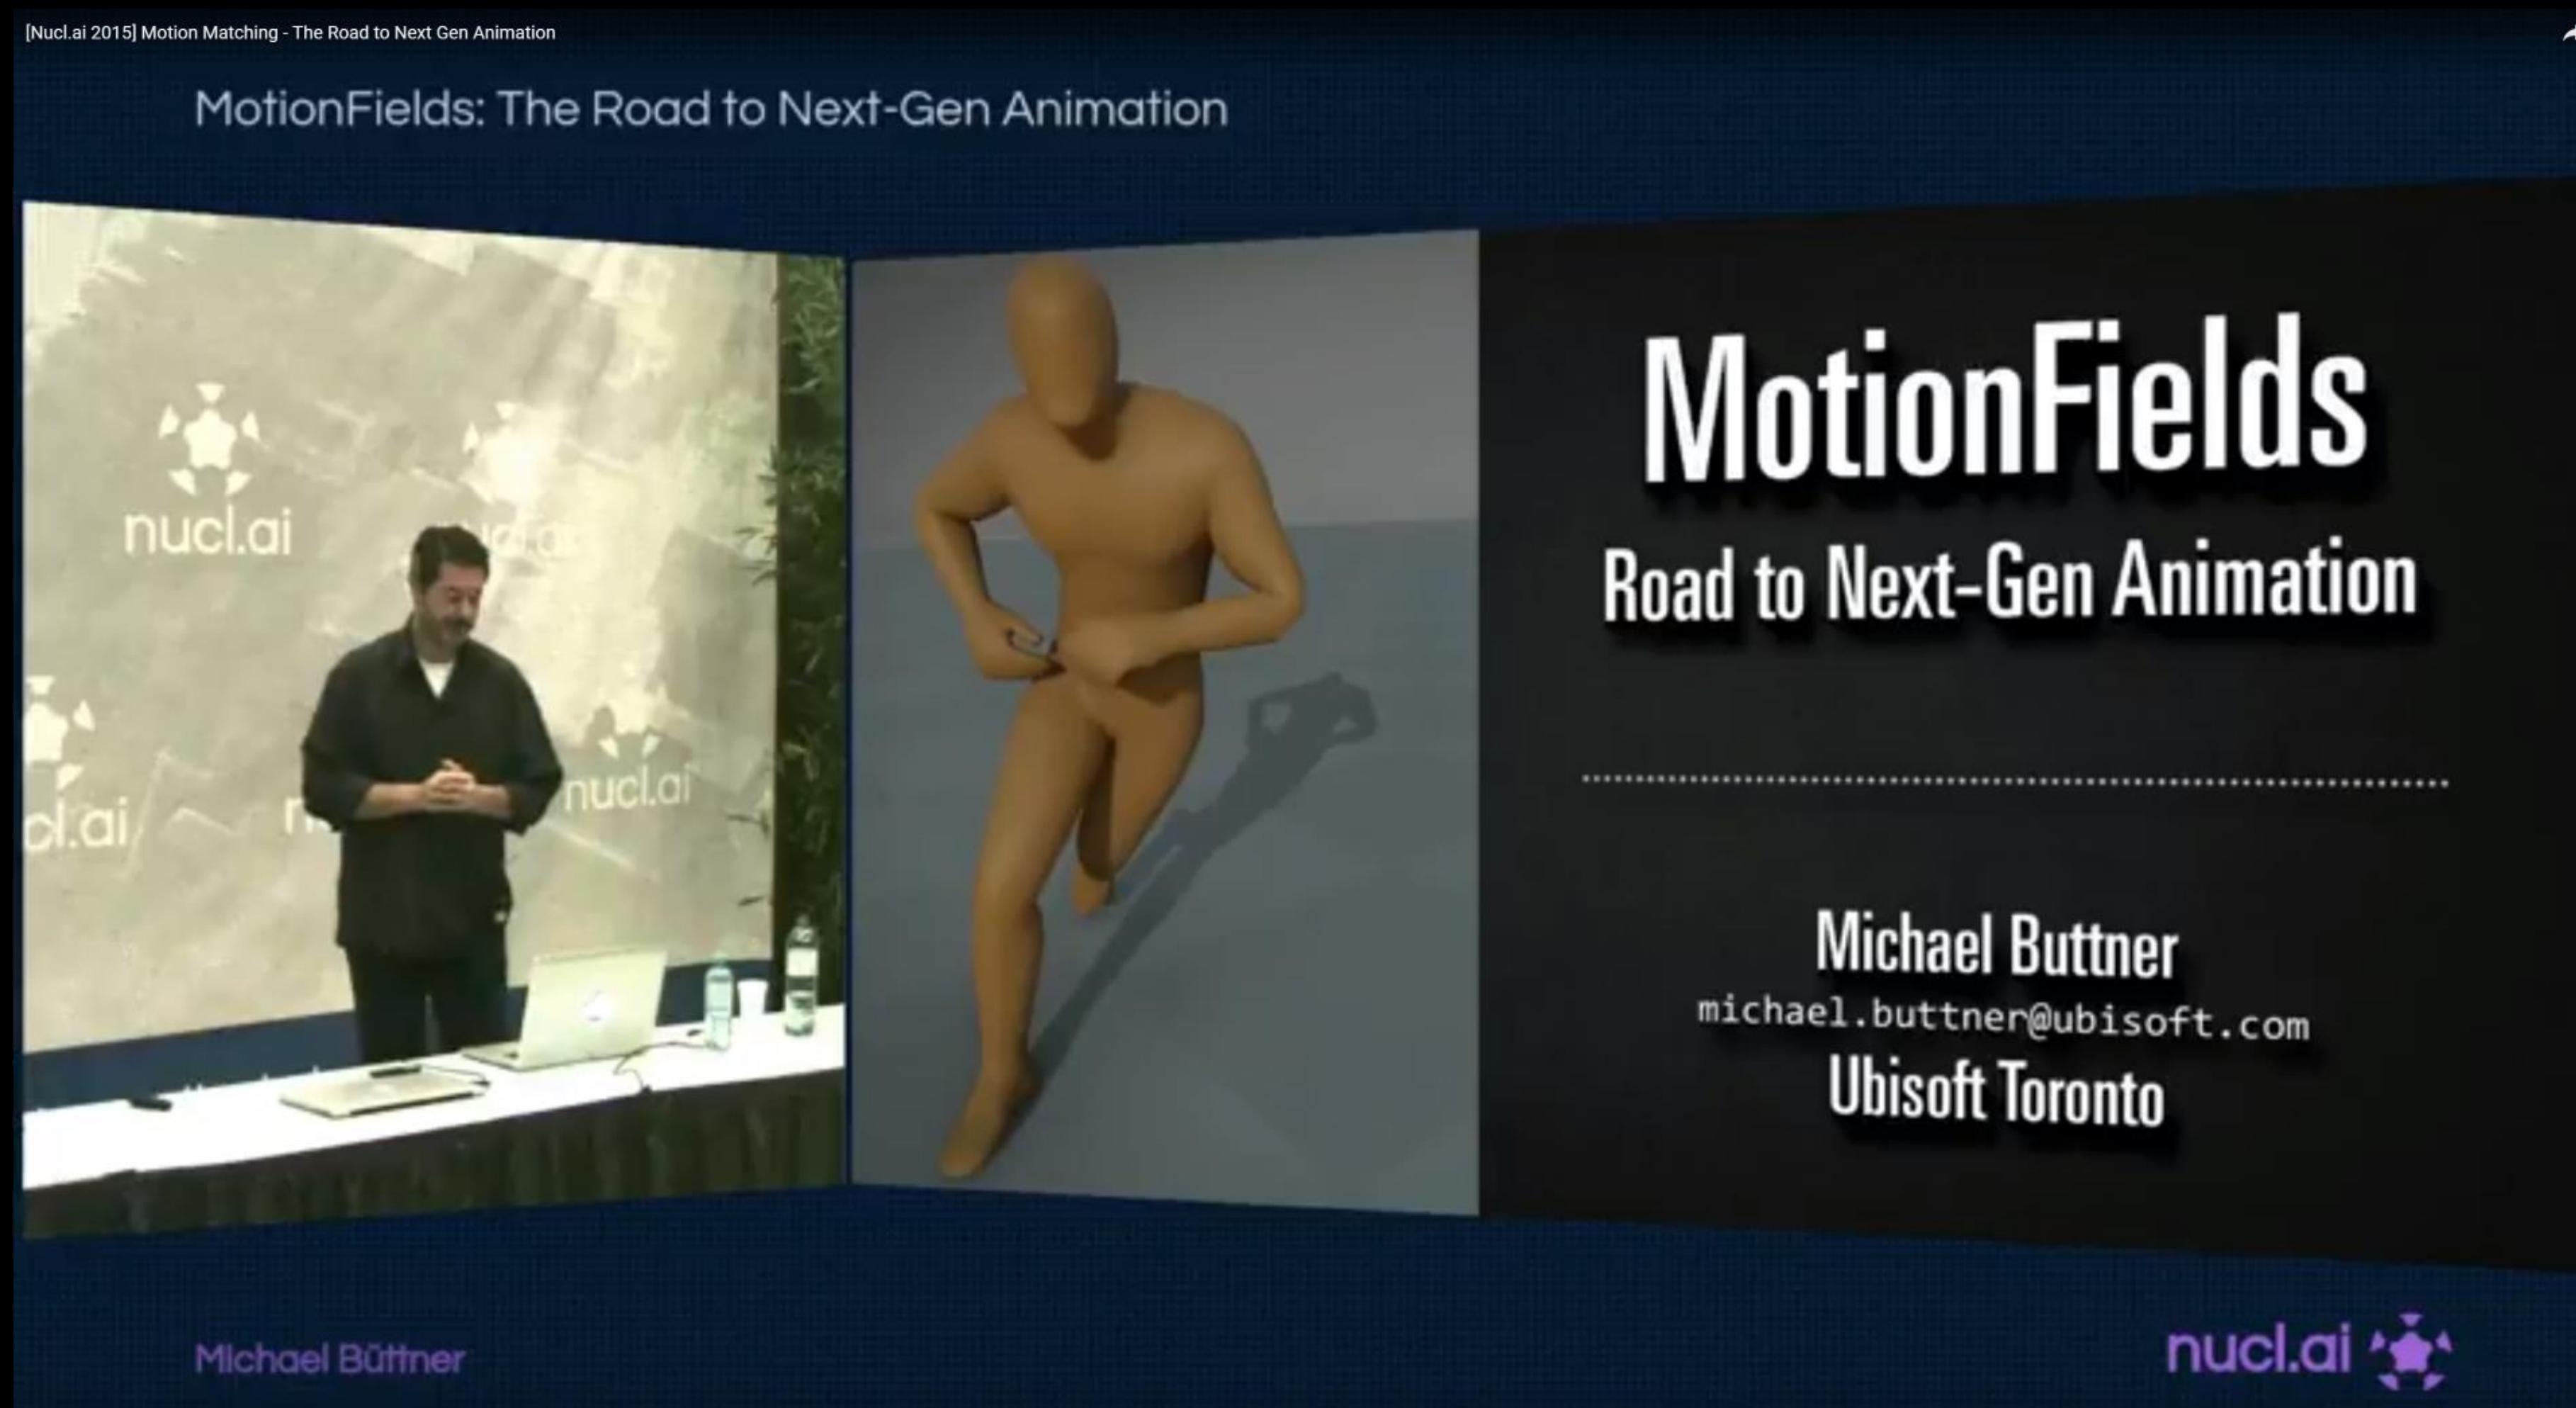




# Motion Matching

[Nucl.ai 2015] Motion Matching - The Road to Next Gen Animation

## MotionFields: The Road to Next-Gen Animation



The image is a composite of two parts. On the left, a man with a beard, Michael Buttner, is standing on a stage, presenting to an audience. He is wearing a dark long-sleeved shirt and dark trousers. He is holding a small device in his hands. In front of him is a white table with a laptop, some papers, and a few small bottles. Behind him is a large projection screen displaying a dense, abstract geometric pattern with the 'nucl.ai' logo repeated across it. On the right side of the image is a 3D rendering of a human-like character in a dynamic, crouching pose, set against a plain grey background. To the right of the character, there is a large title 'MotionFields' in white, bold, sans-serif font, with the subtitle 'Road to Next-Gen Animation' underneath it in a slightly smaller font. Below the title, there is a dotted line separator. Underneath the separator, the name 'Michael Buttner' is written in white, bold font, followed by the email address 'michael.buttner@ubisoft.com'. Below the email, the text 'Ubisoft Toronto' is displayed. At the bottom right of the image is the 'nucl.ai' logo, which consists of the word 'nucl.ai' in a white, lowercase, sans-serif font next to a stylized purple and white geometric icon.

Michael Buttner

michael.buttner@ubisoft.com

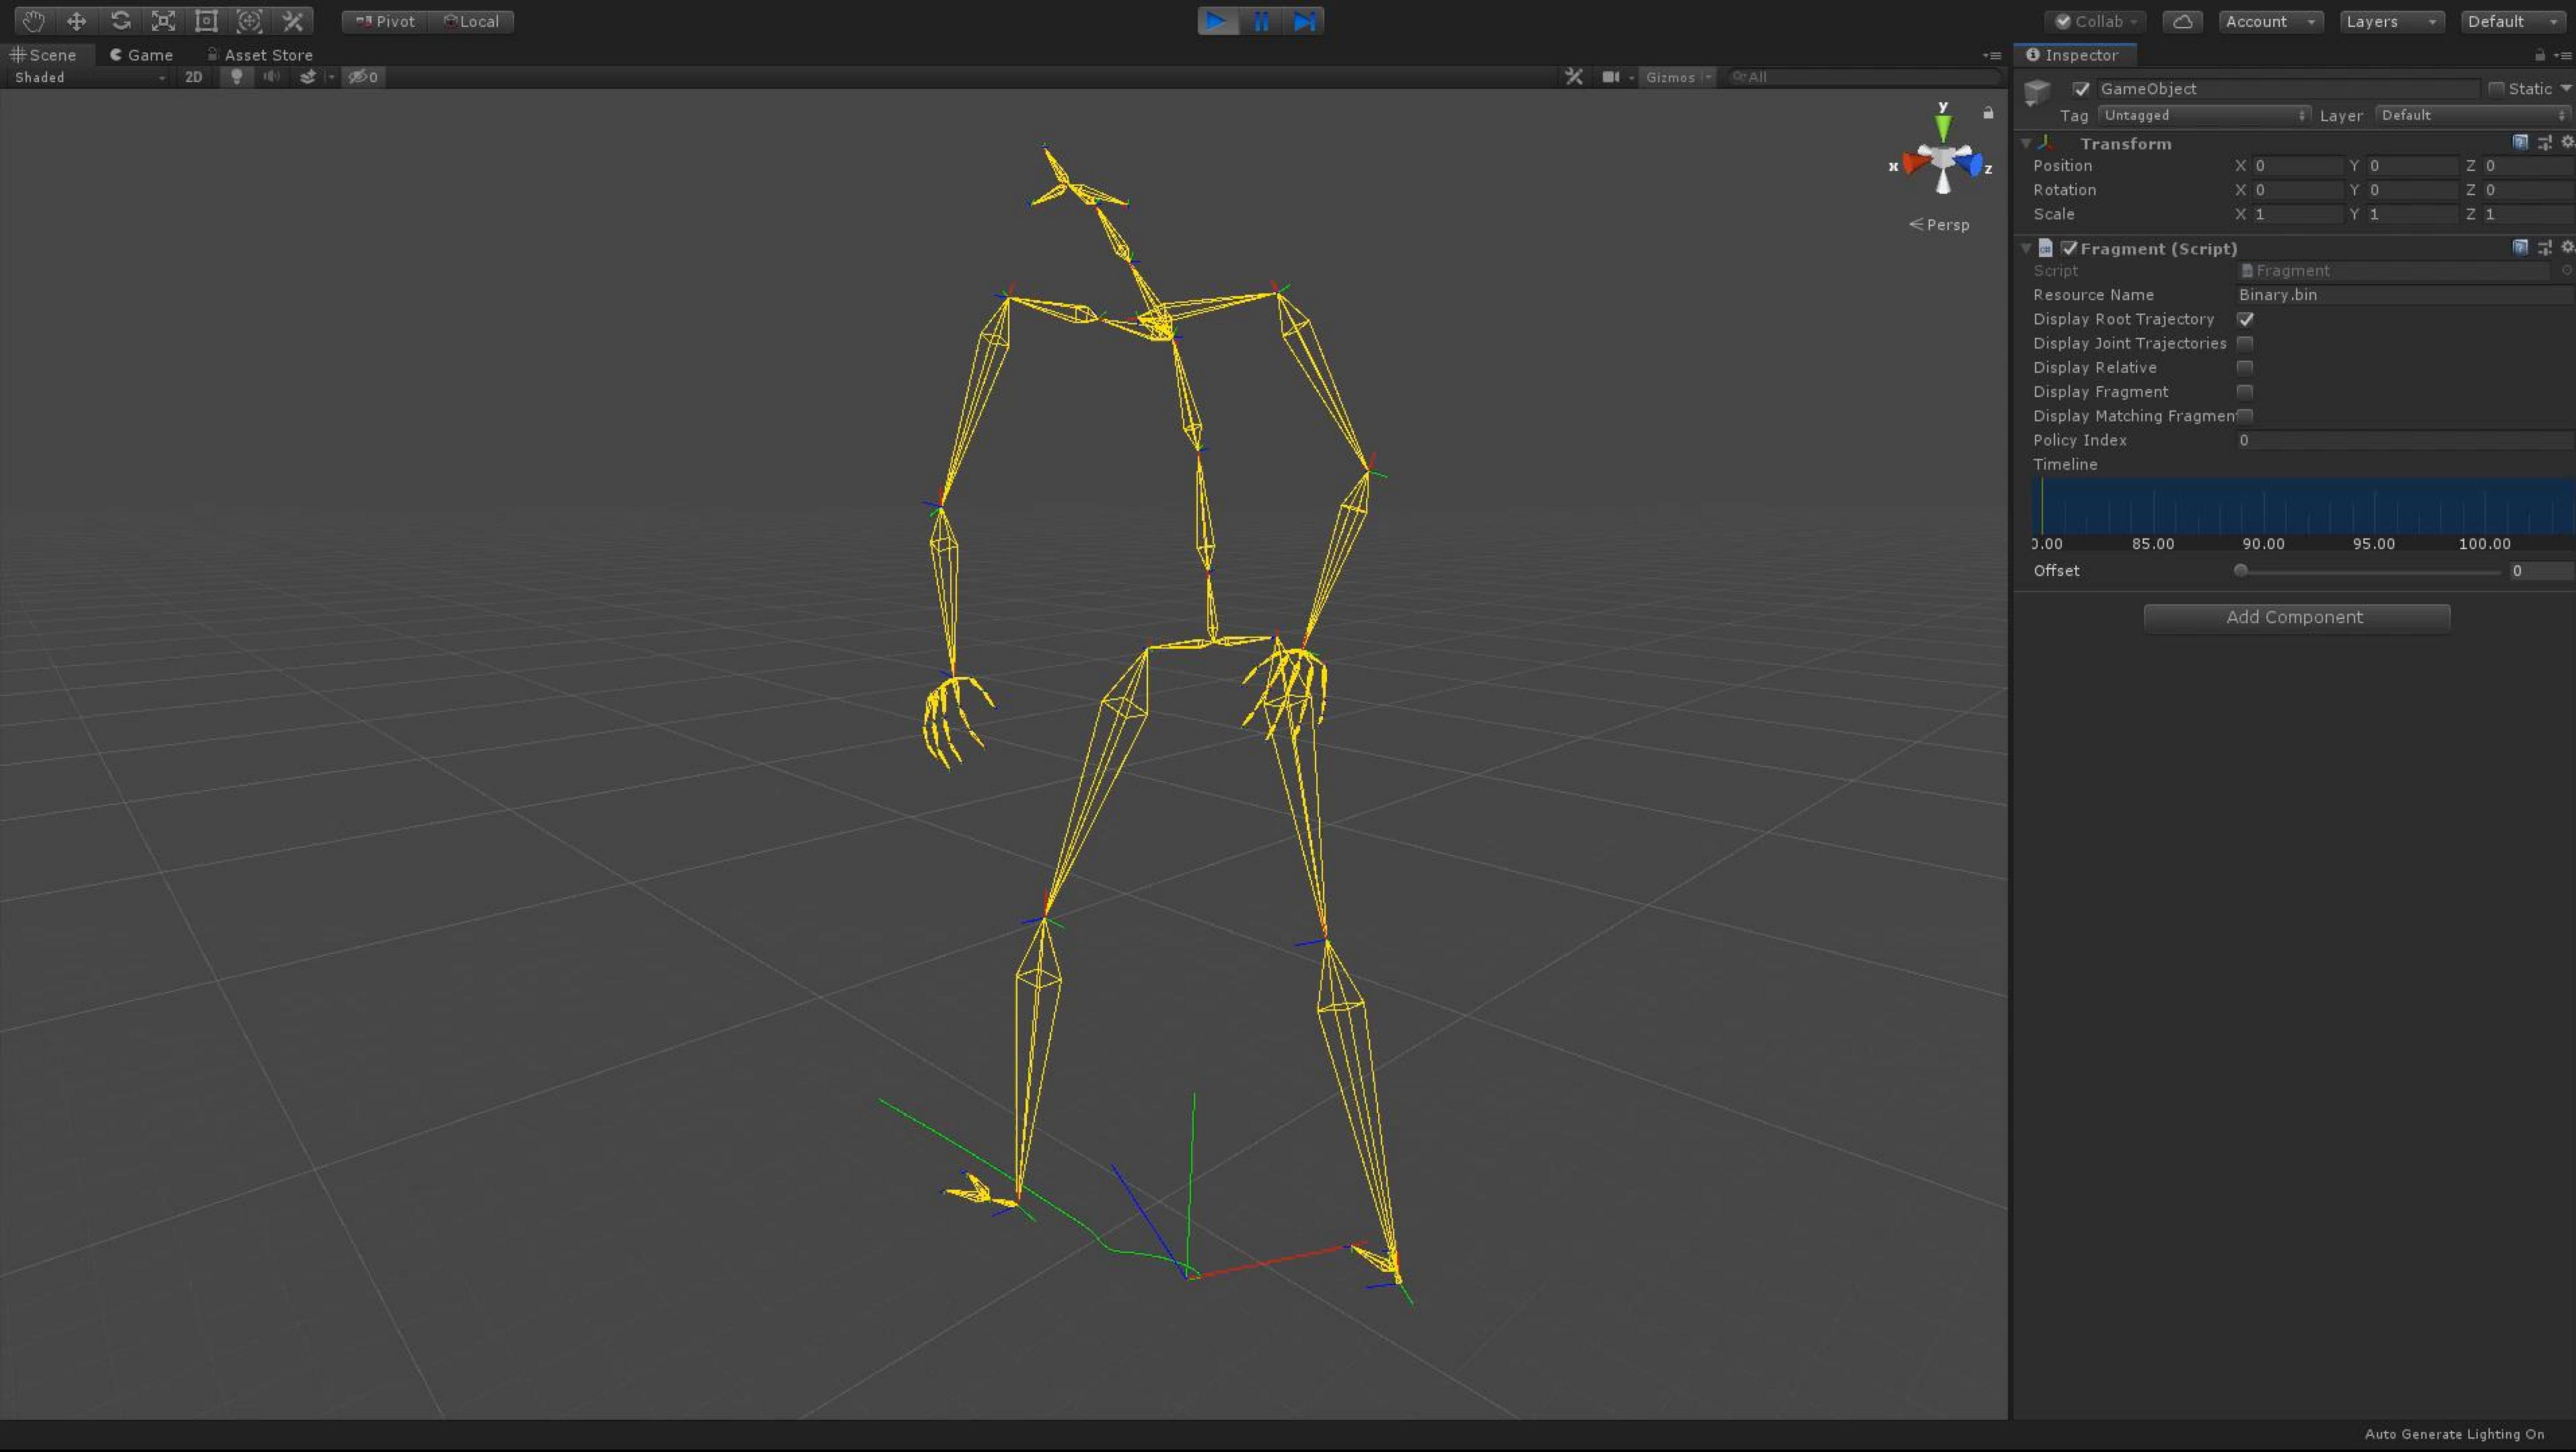
Ubisoft Toronto

nucl.ai

Motion Matching [Michael Buttner, 2015]

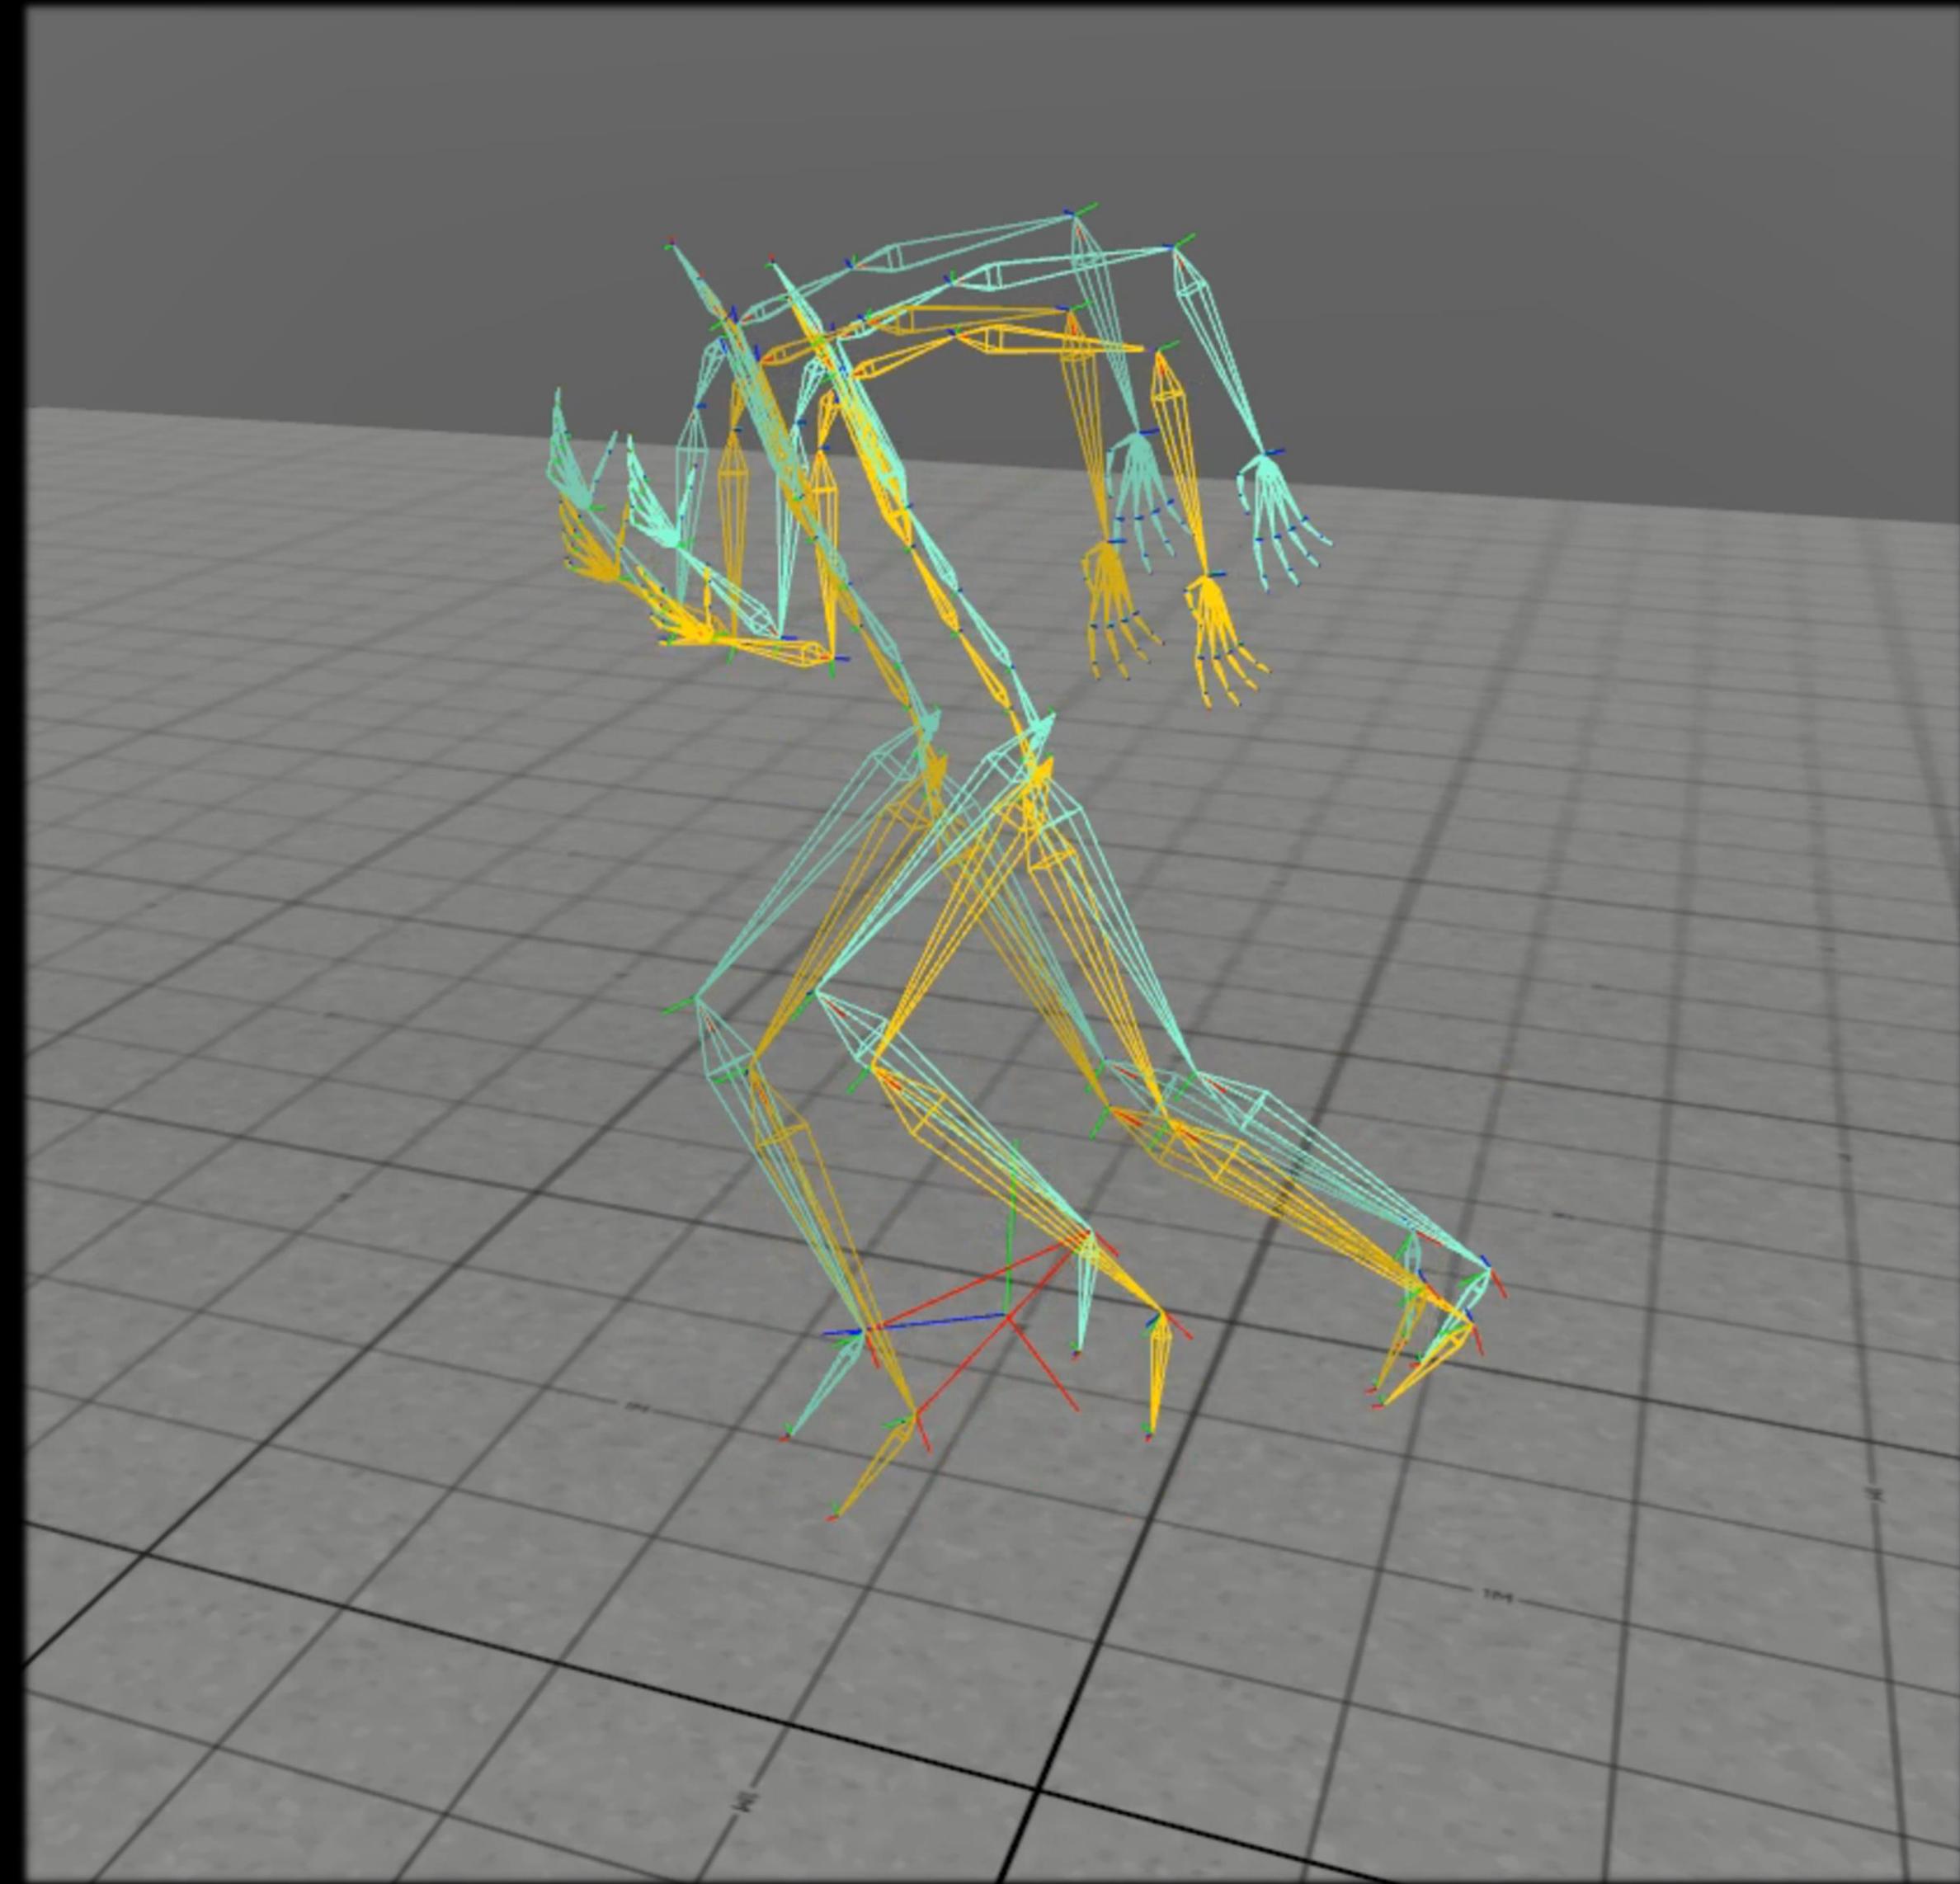
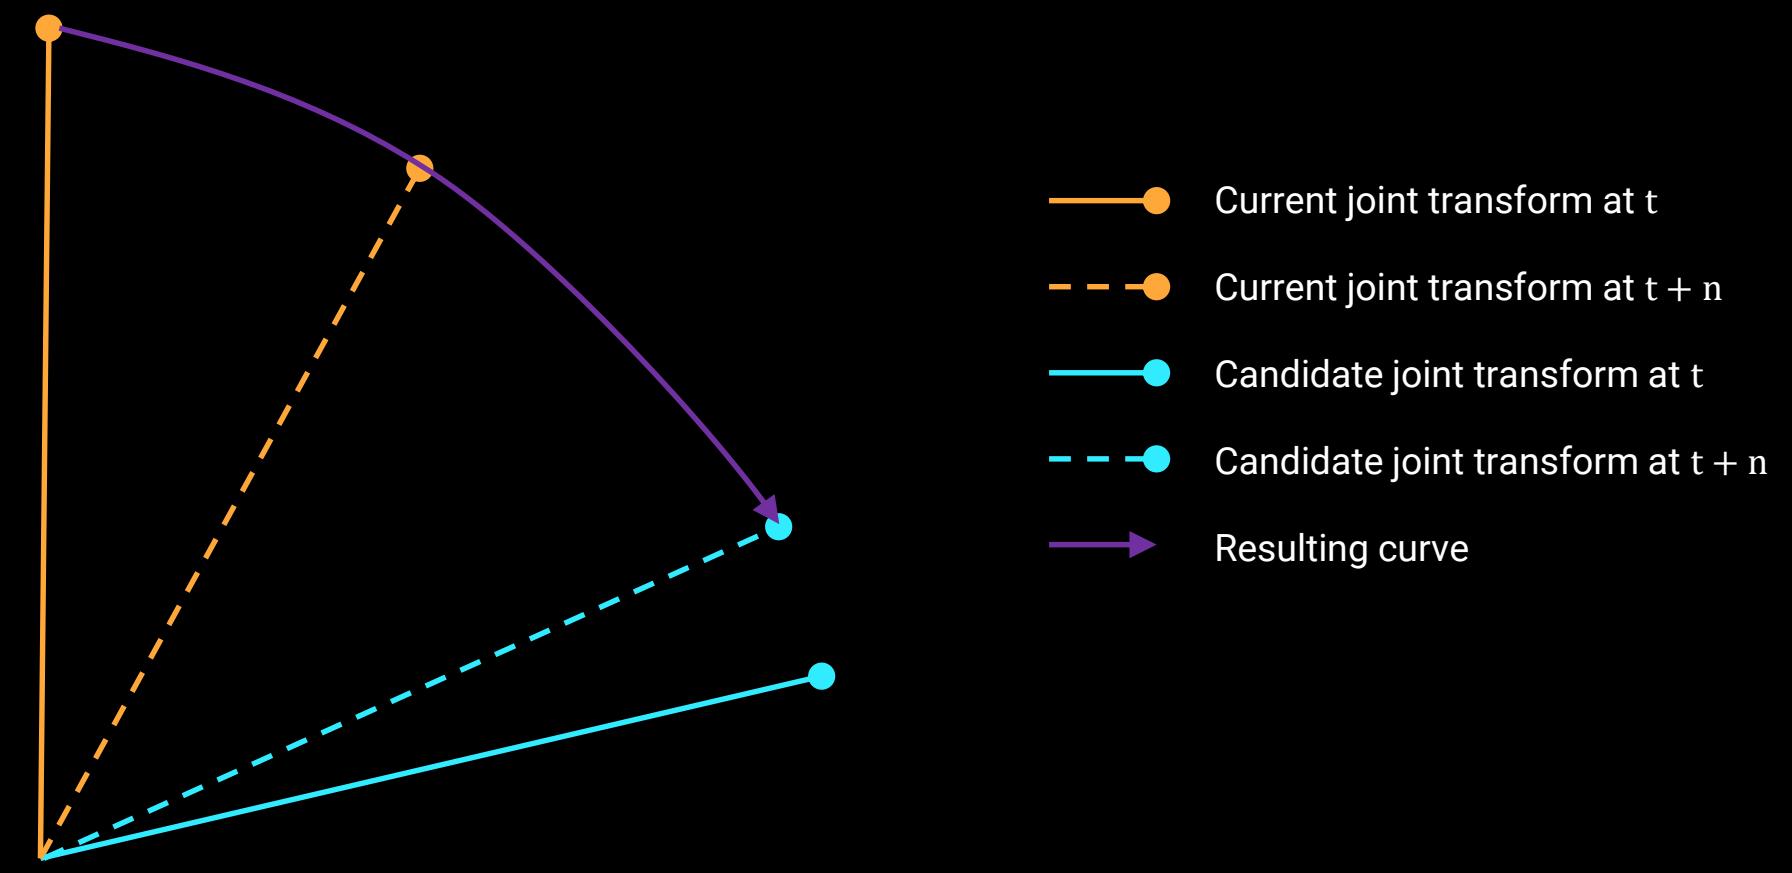


13:39:06:13

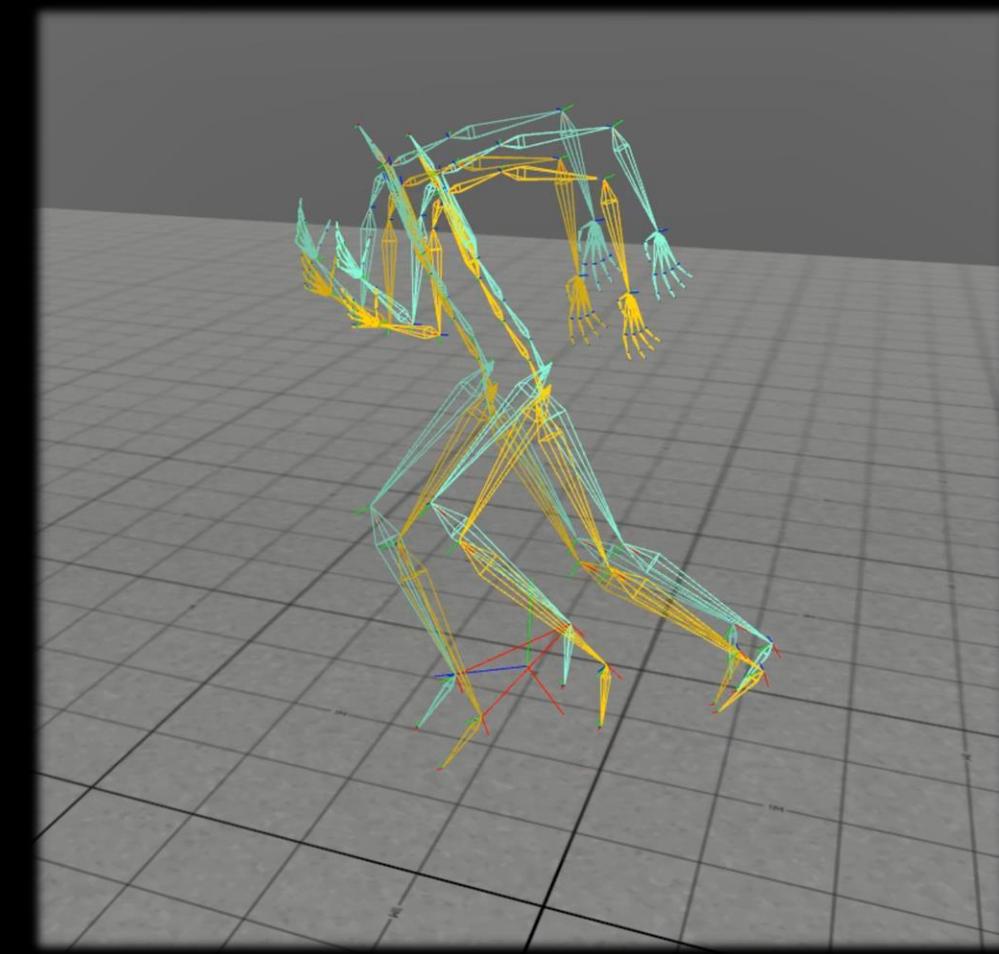


# Geometric Pose Comparison

- Minimum squared distance between joint positions
- Using only “relevant” joints
- Minimum local joint rotation delta is flawed

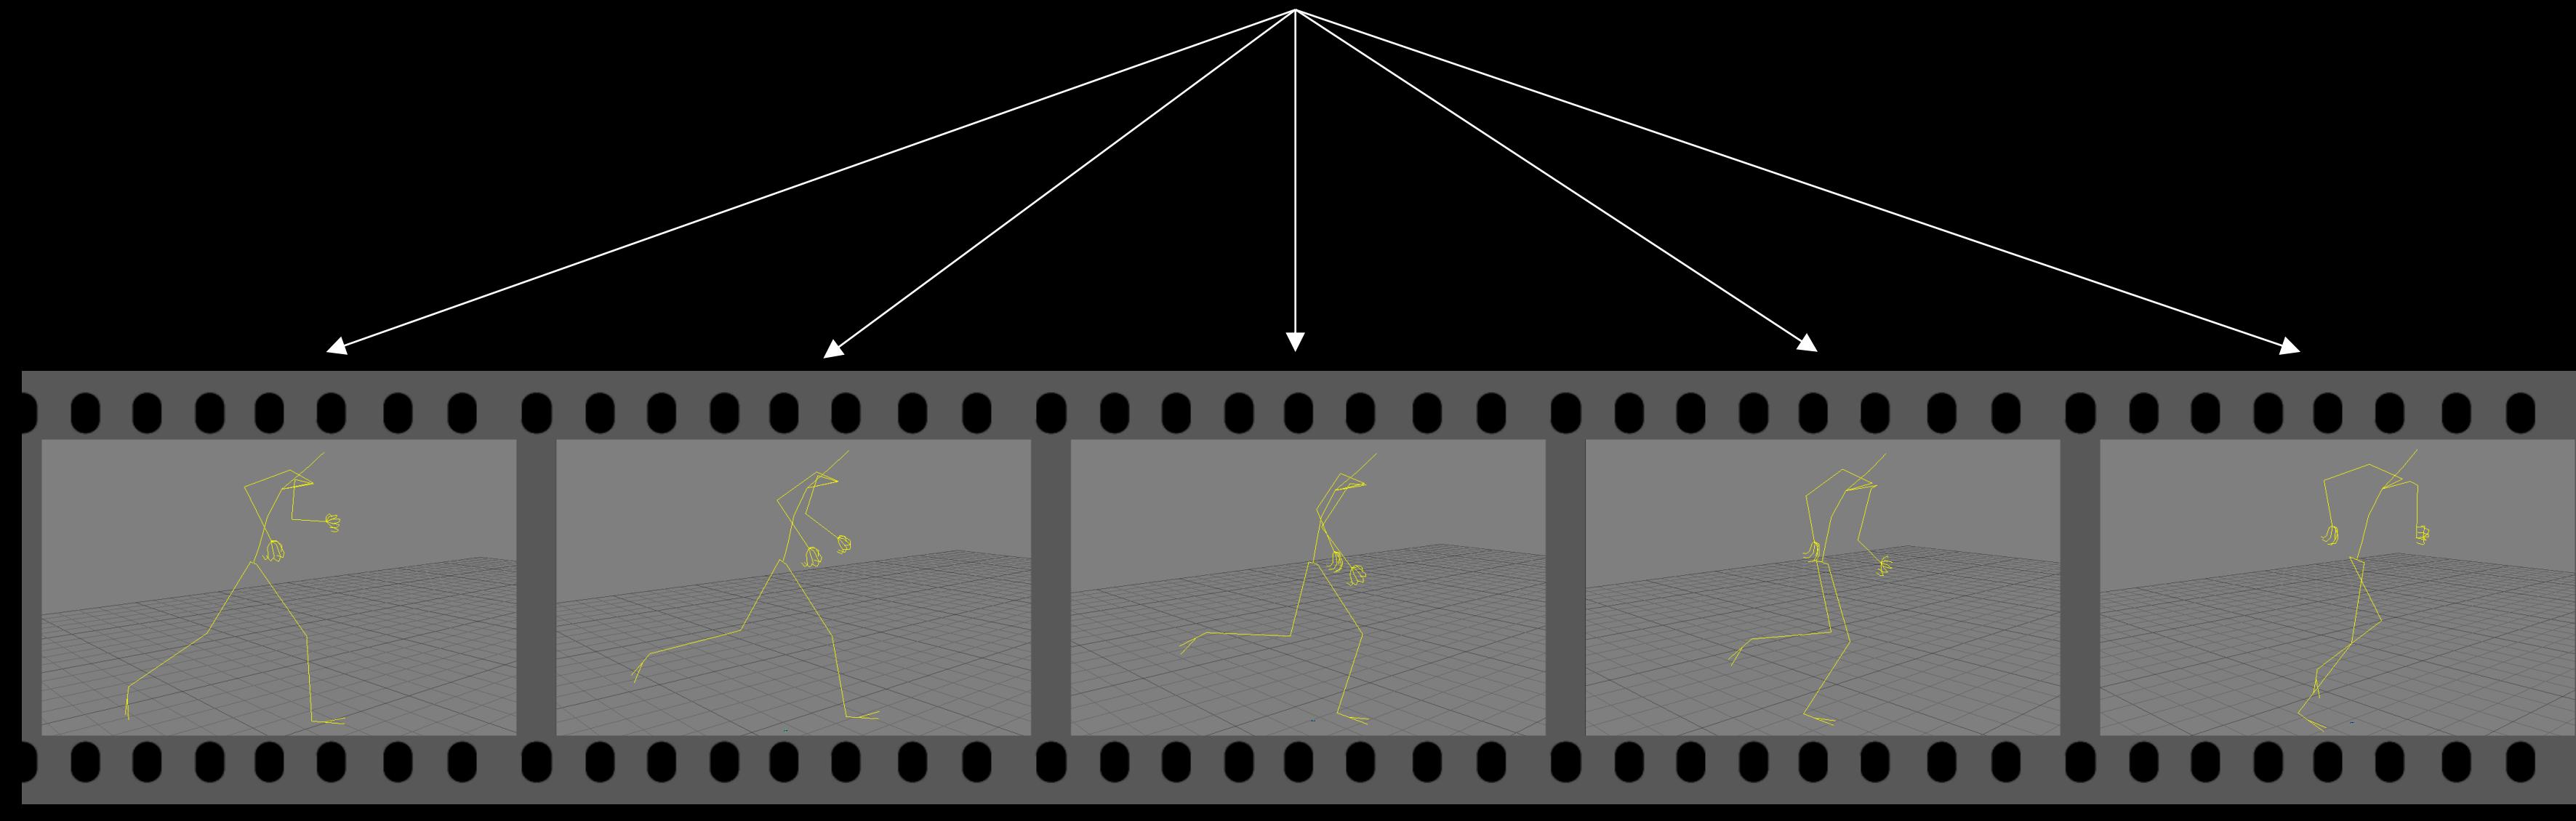


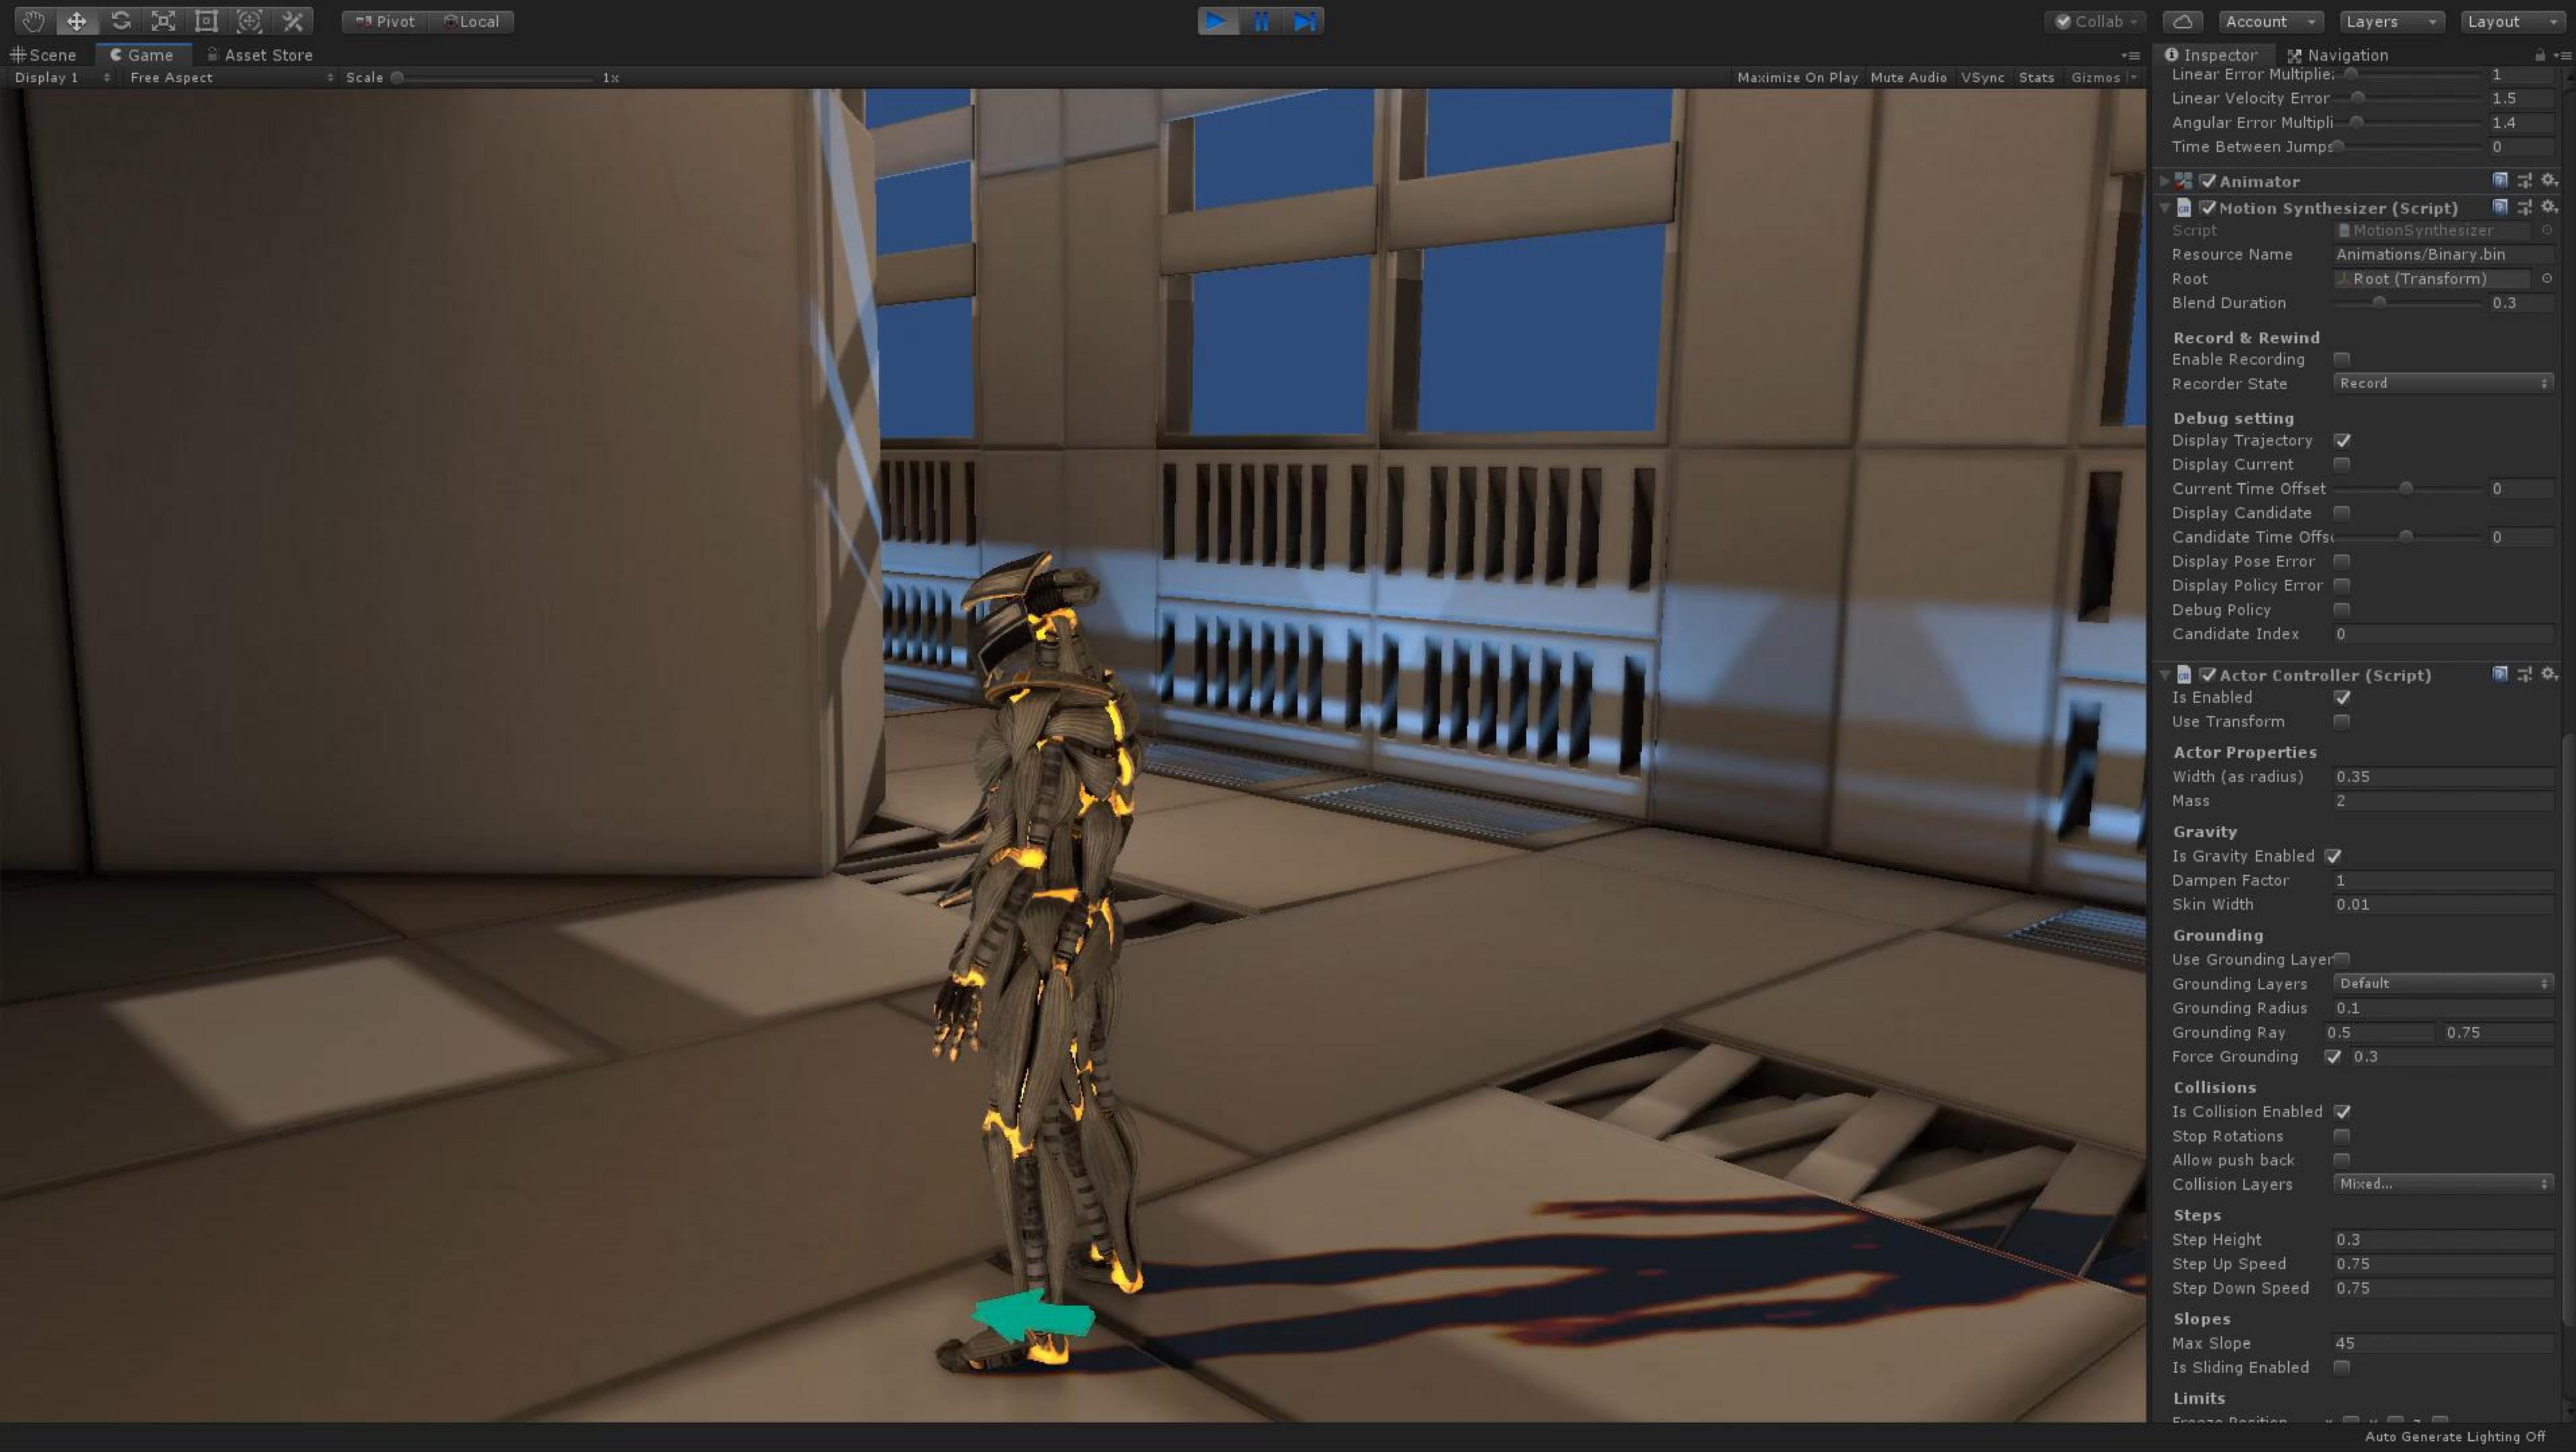
# Motion Matching Algorithm



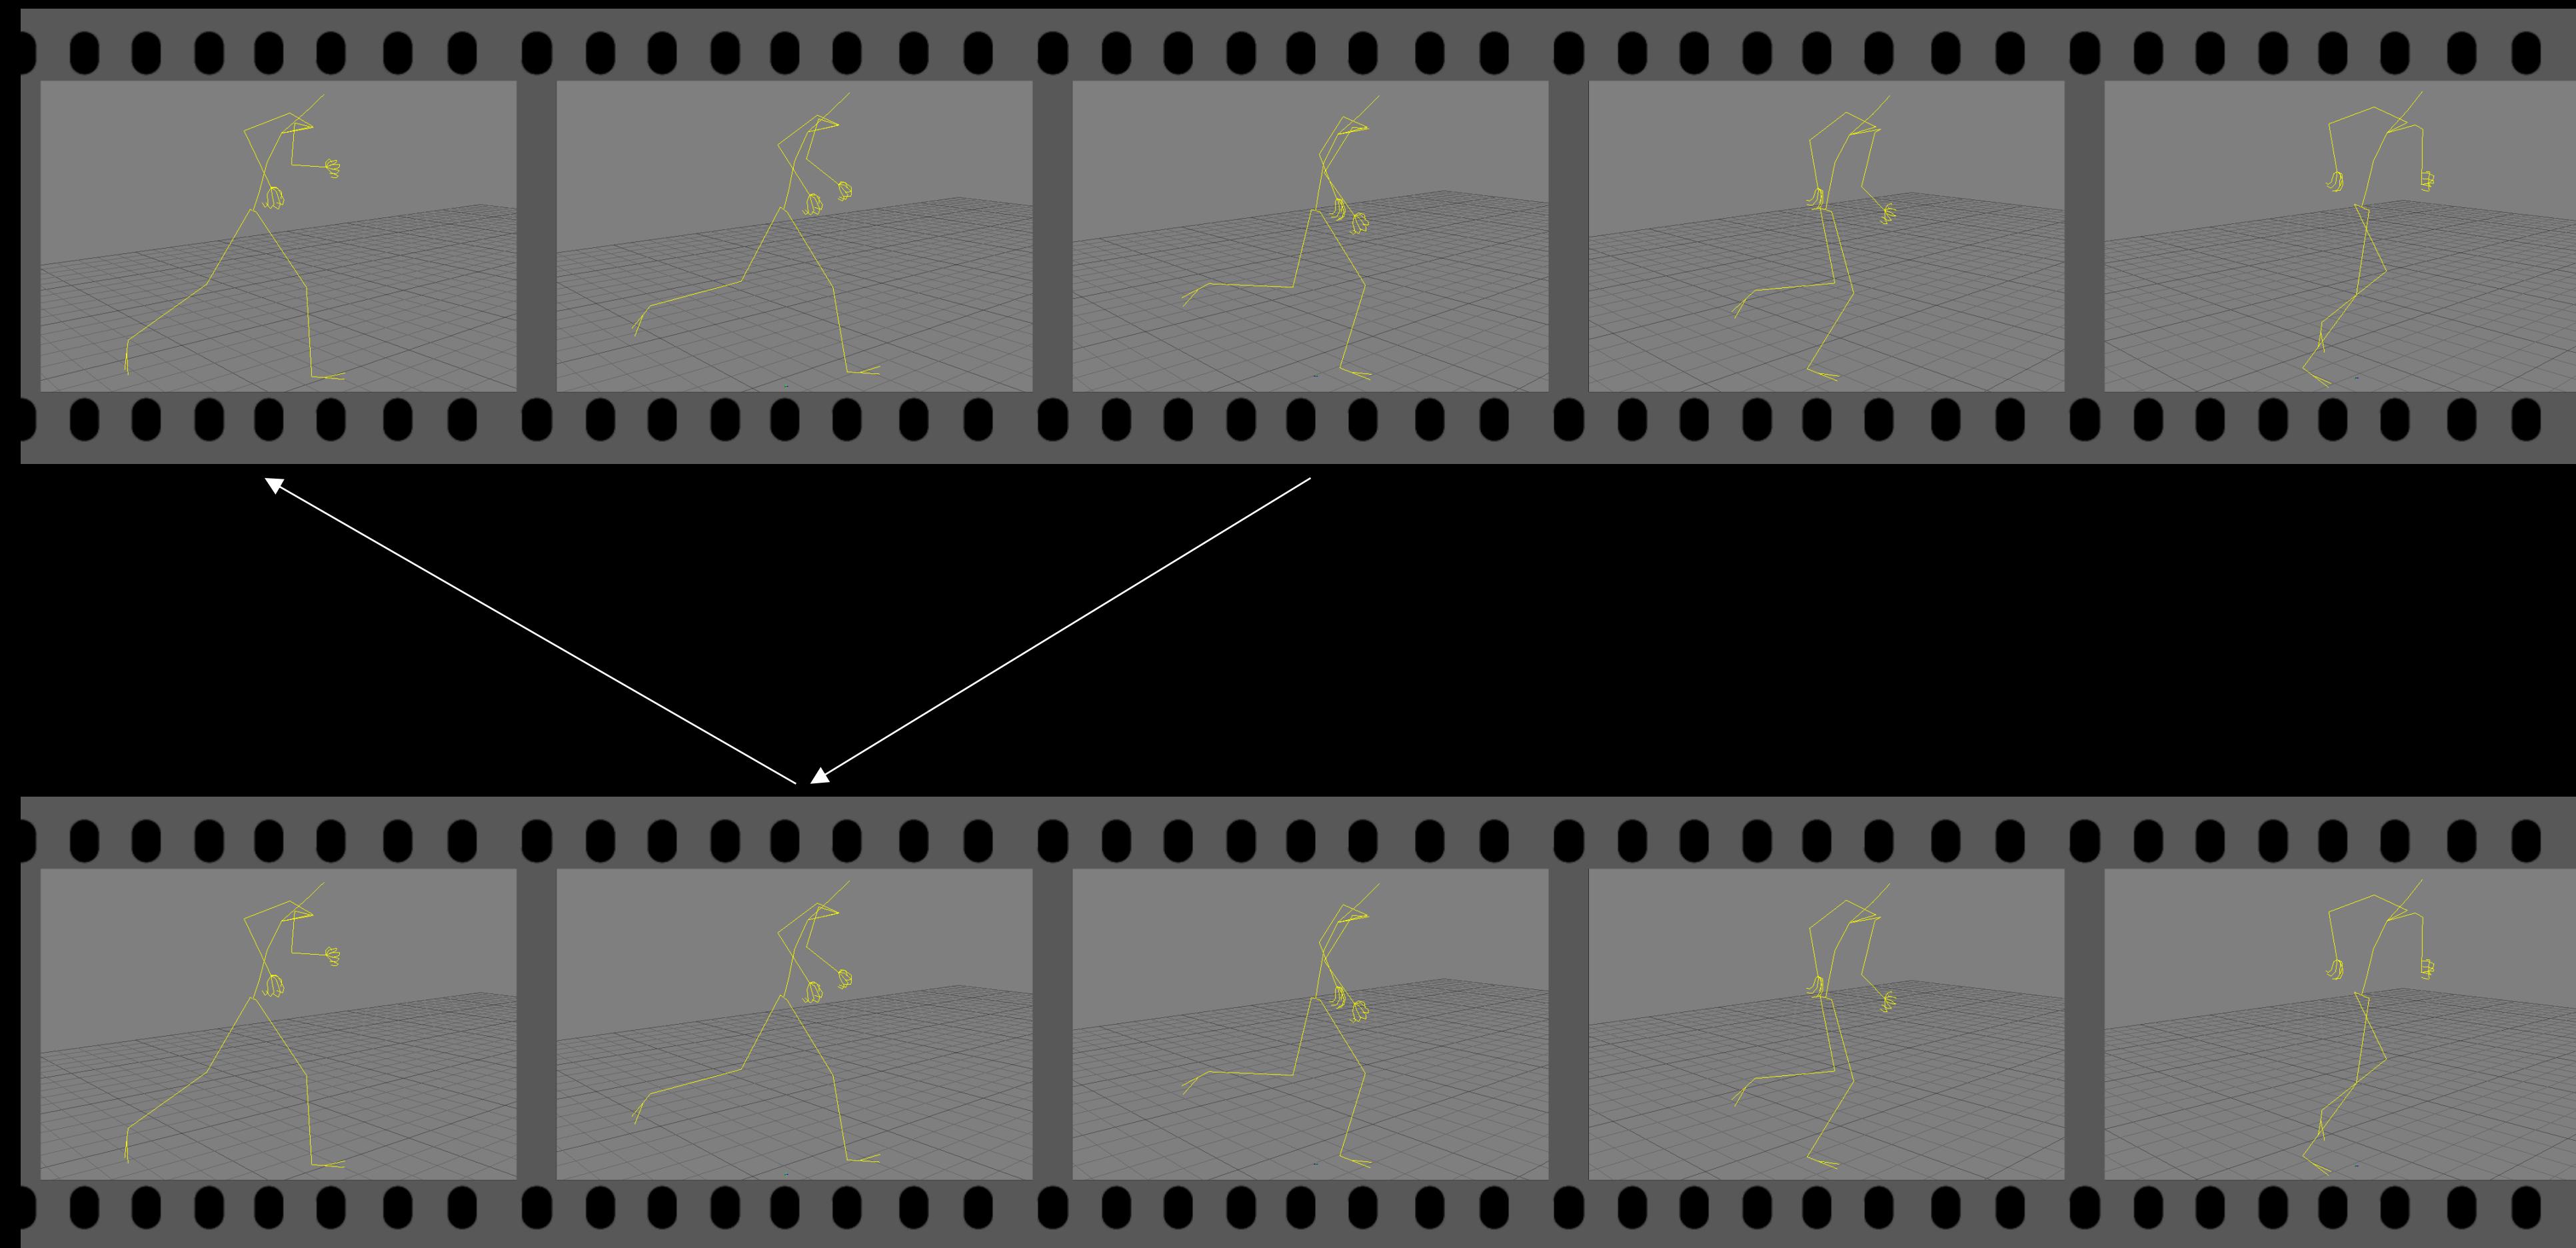
Query  $\rightarrow F(y)$

$$\arg \min x s(F(x), F(y))$$





# Back-in-time Problem



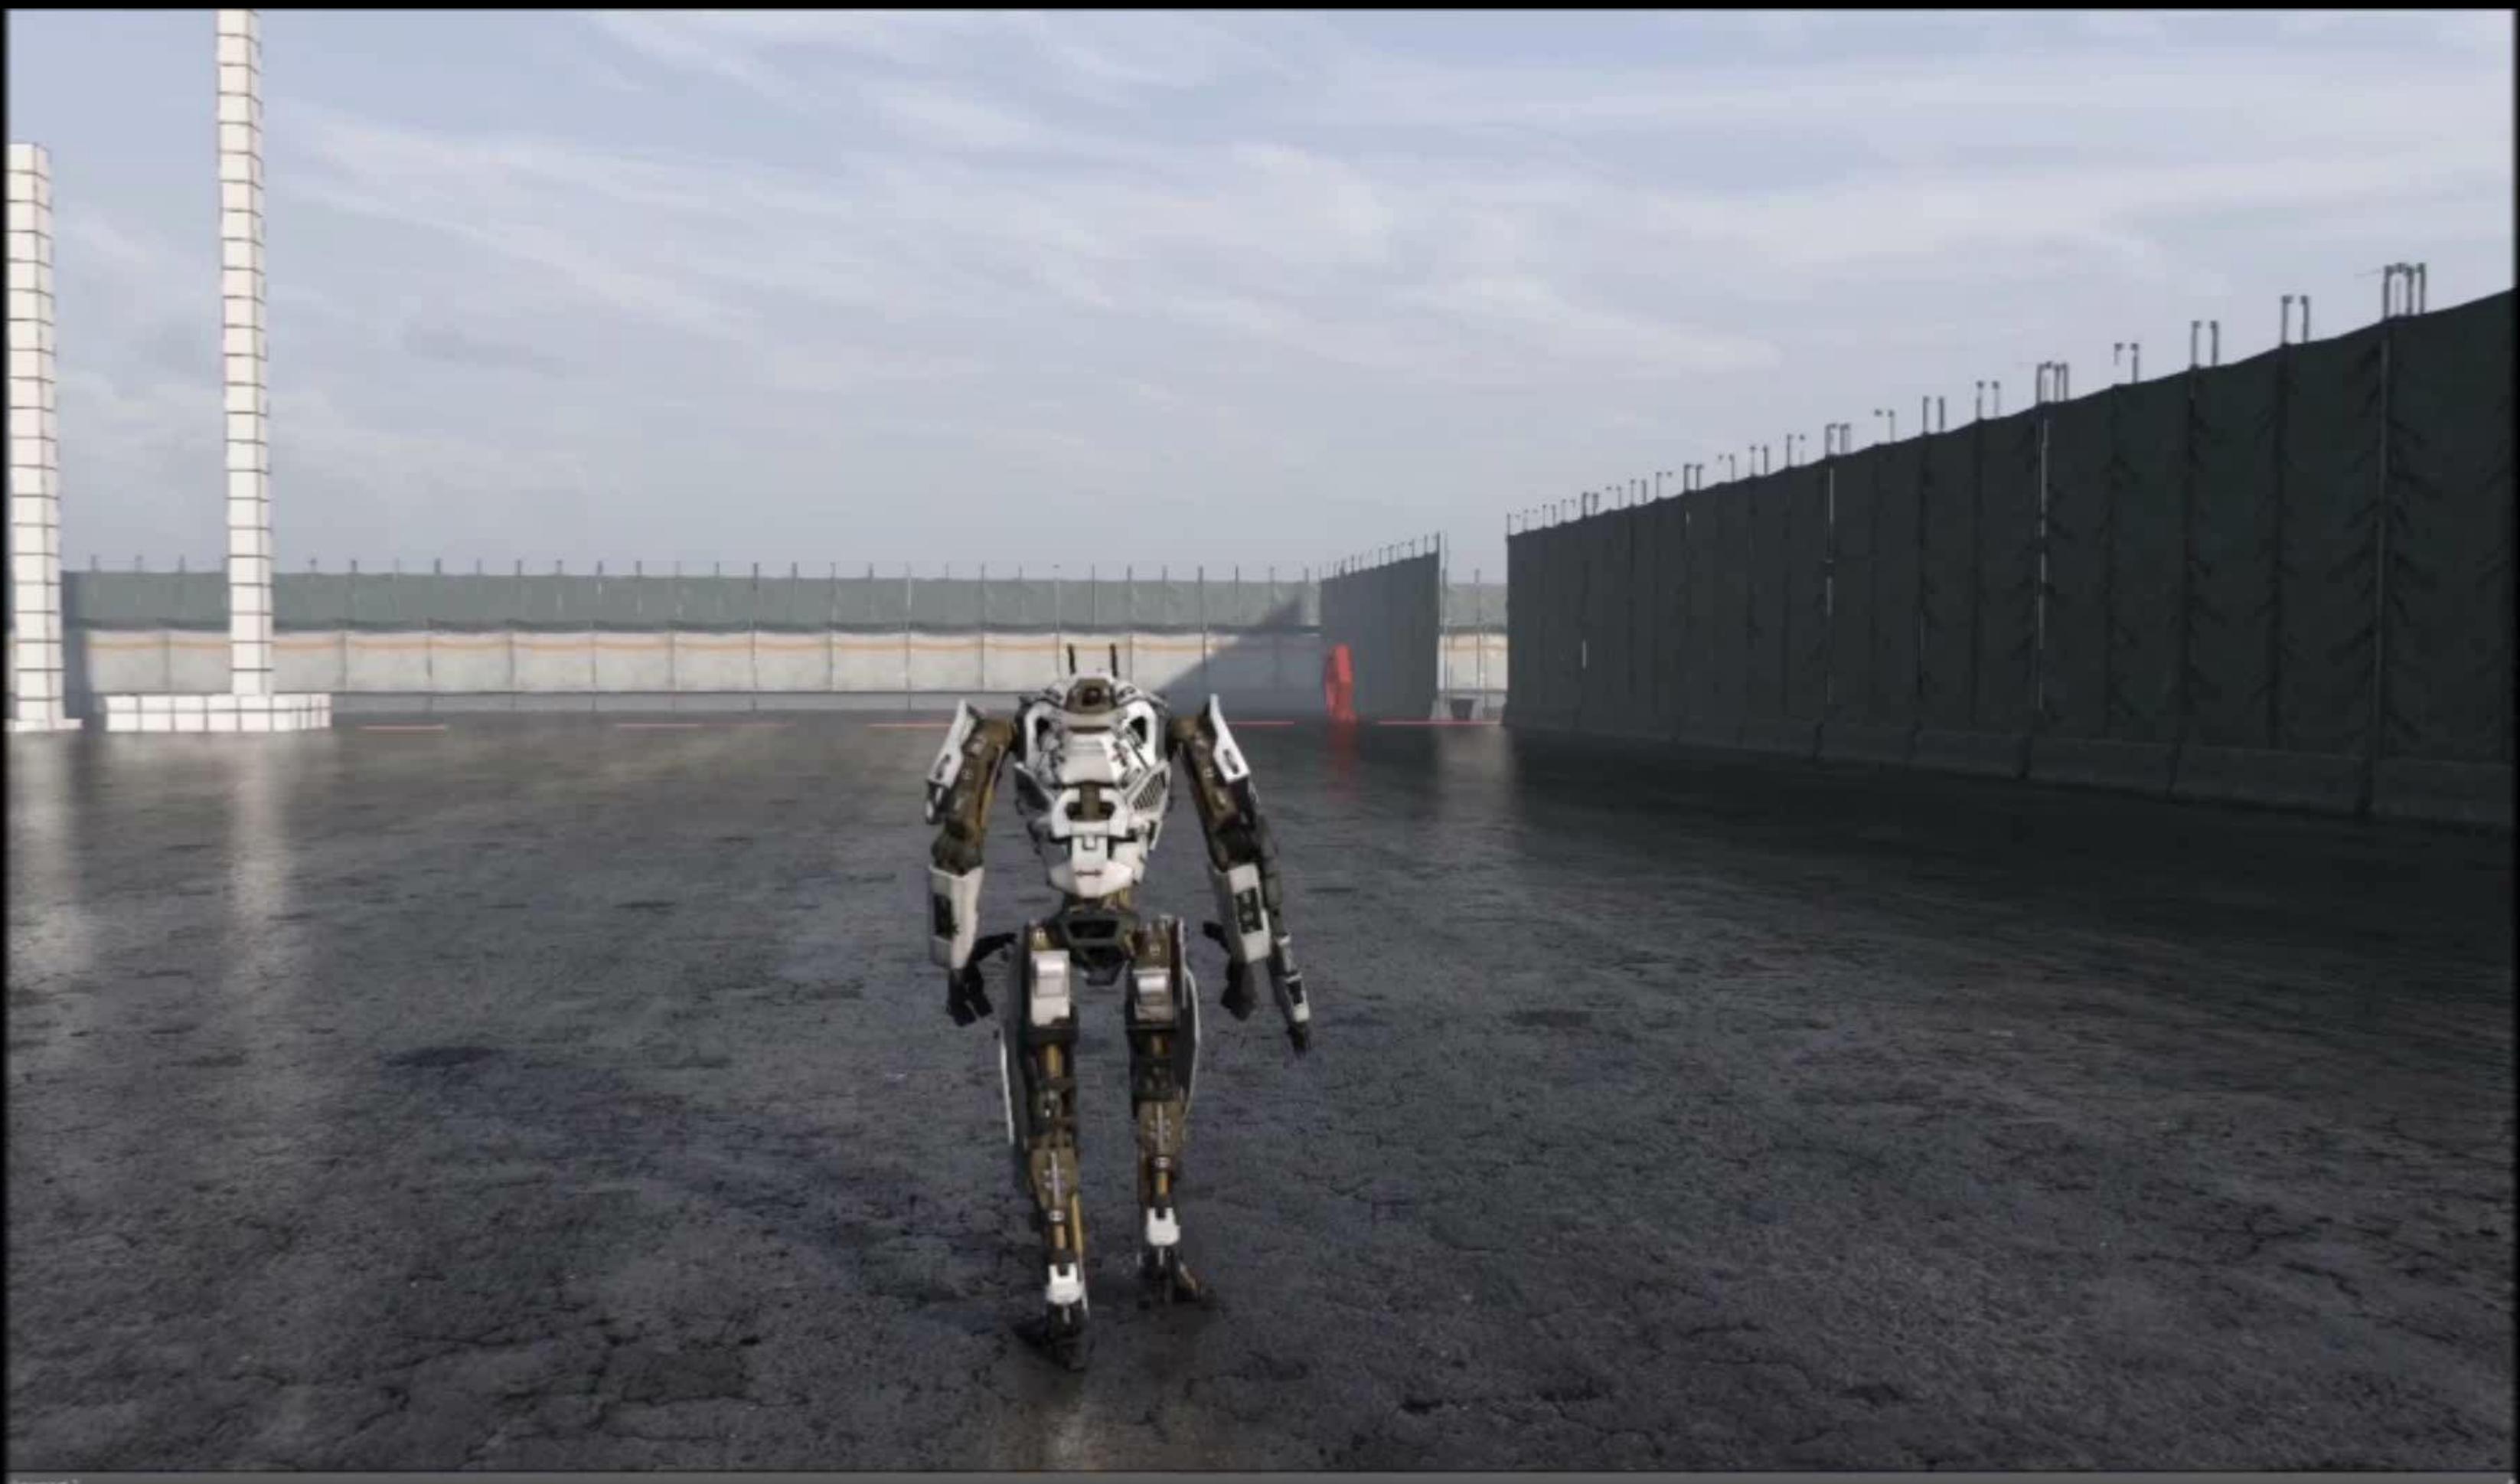
# Motion Matching

## Pros

- Preserves high quality result
- Does not rely on phases
- Relatively easy to implement

## Cons

- Prediction must match data
- Construction of cost functions
- Requires a lot of tweaking
- Doesn't scale well
- Duplicate data problem
- "Back-in-time" problem



# Motion Synthesis Research

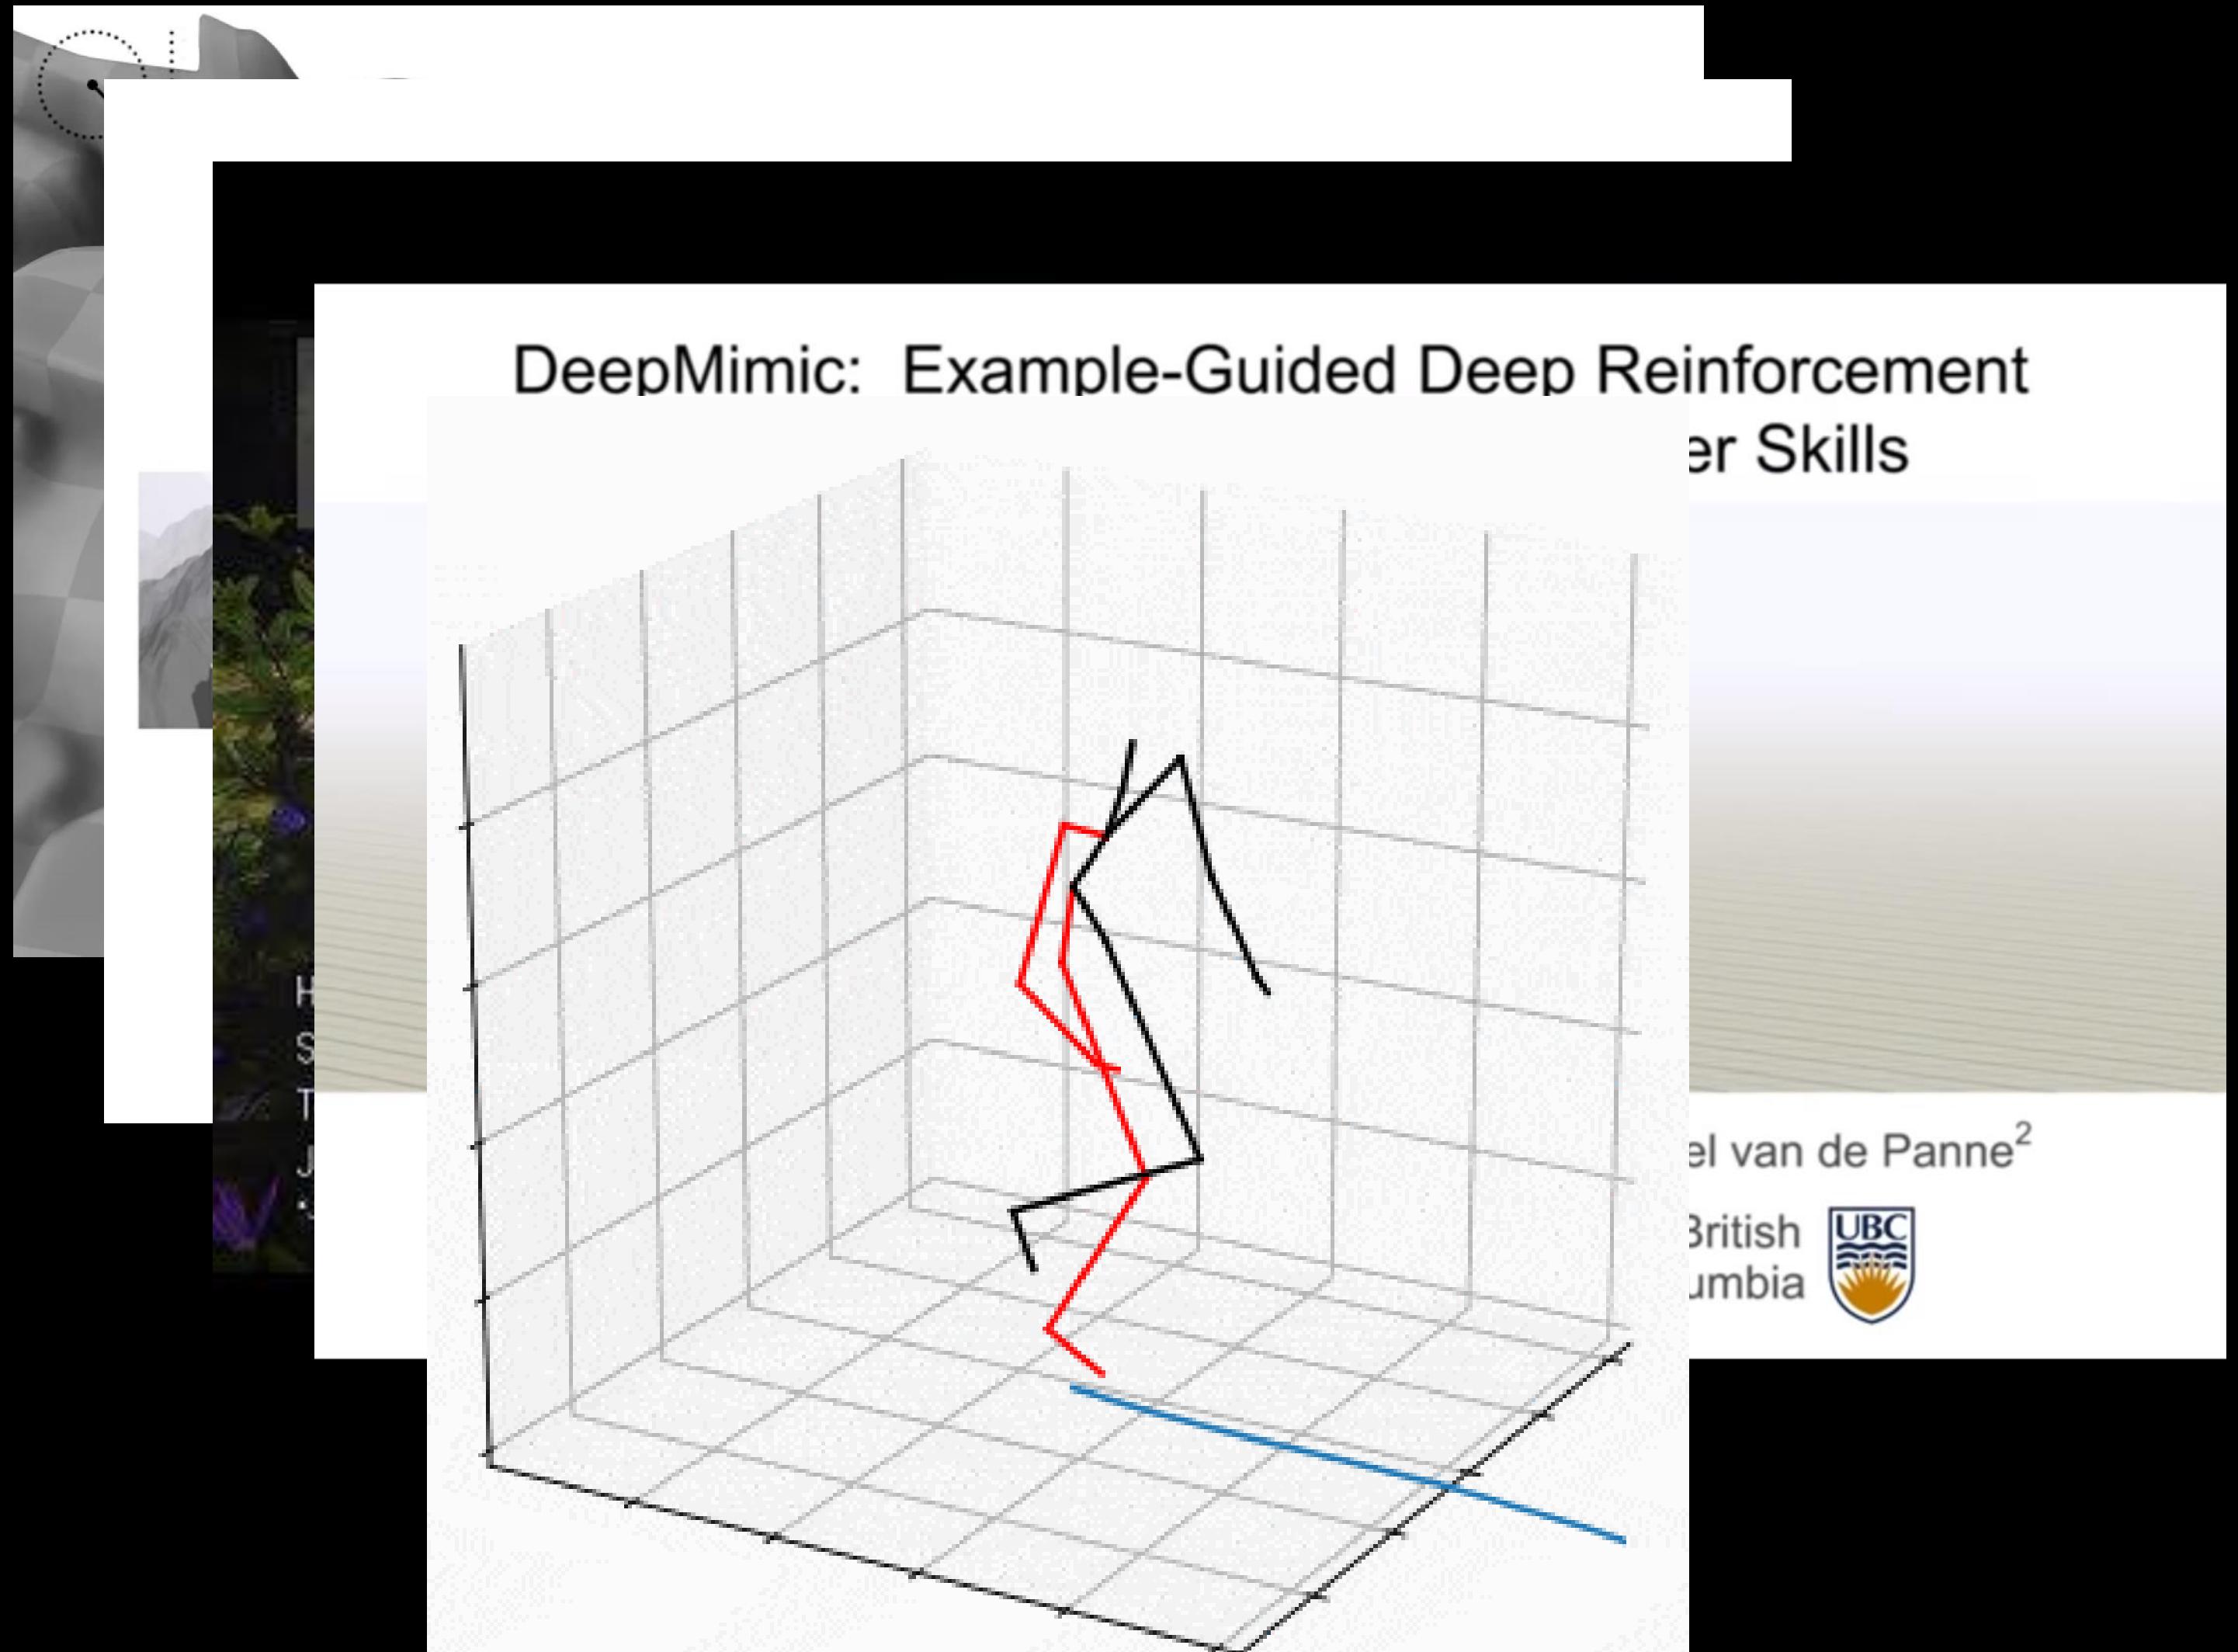
Phase Function NN [Daniel Holden, 2017]

Deep Loco [Peng, 2017]

Mode Adaptive NN [Sebastian Starke, 2018]

Deep Mimic [Peng, 2018]

QuaterNet [Pavvlo, 2018]



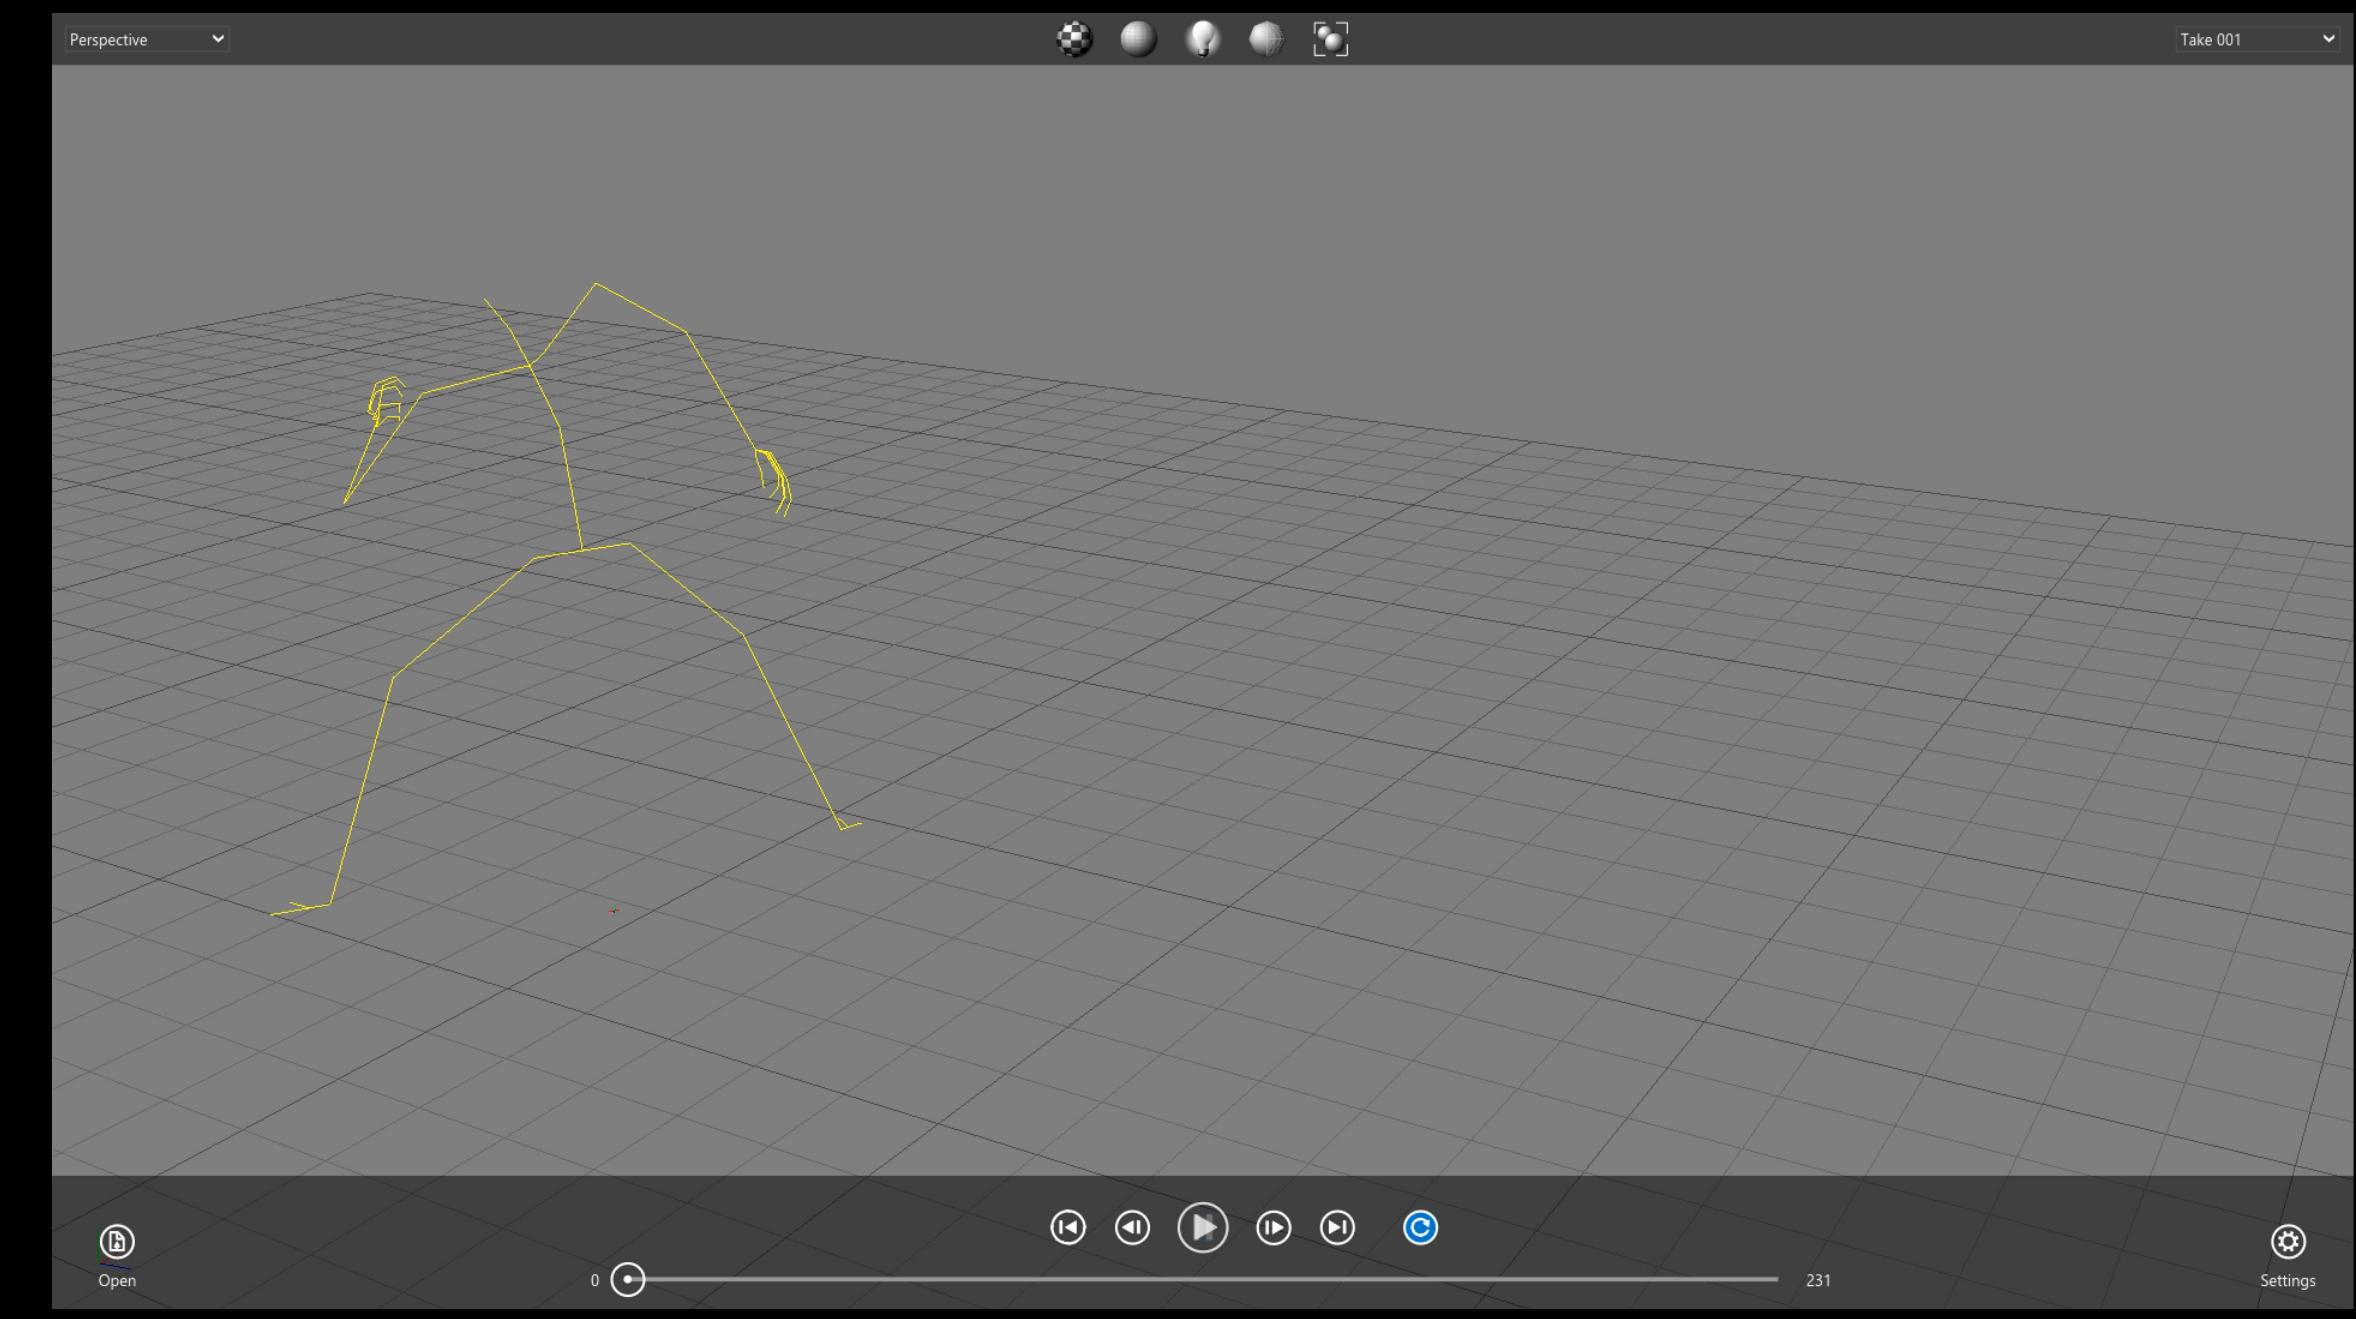
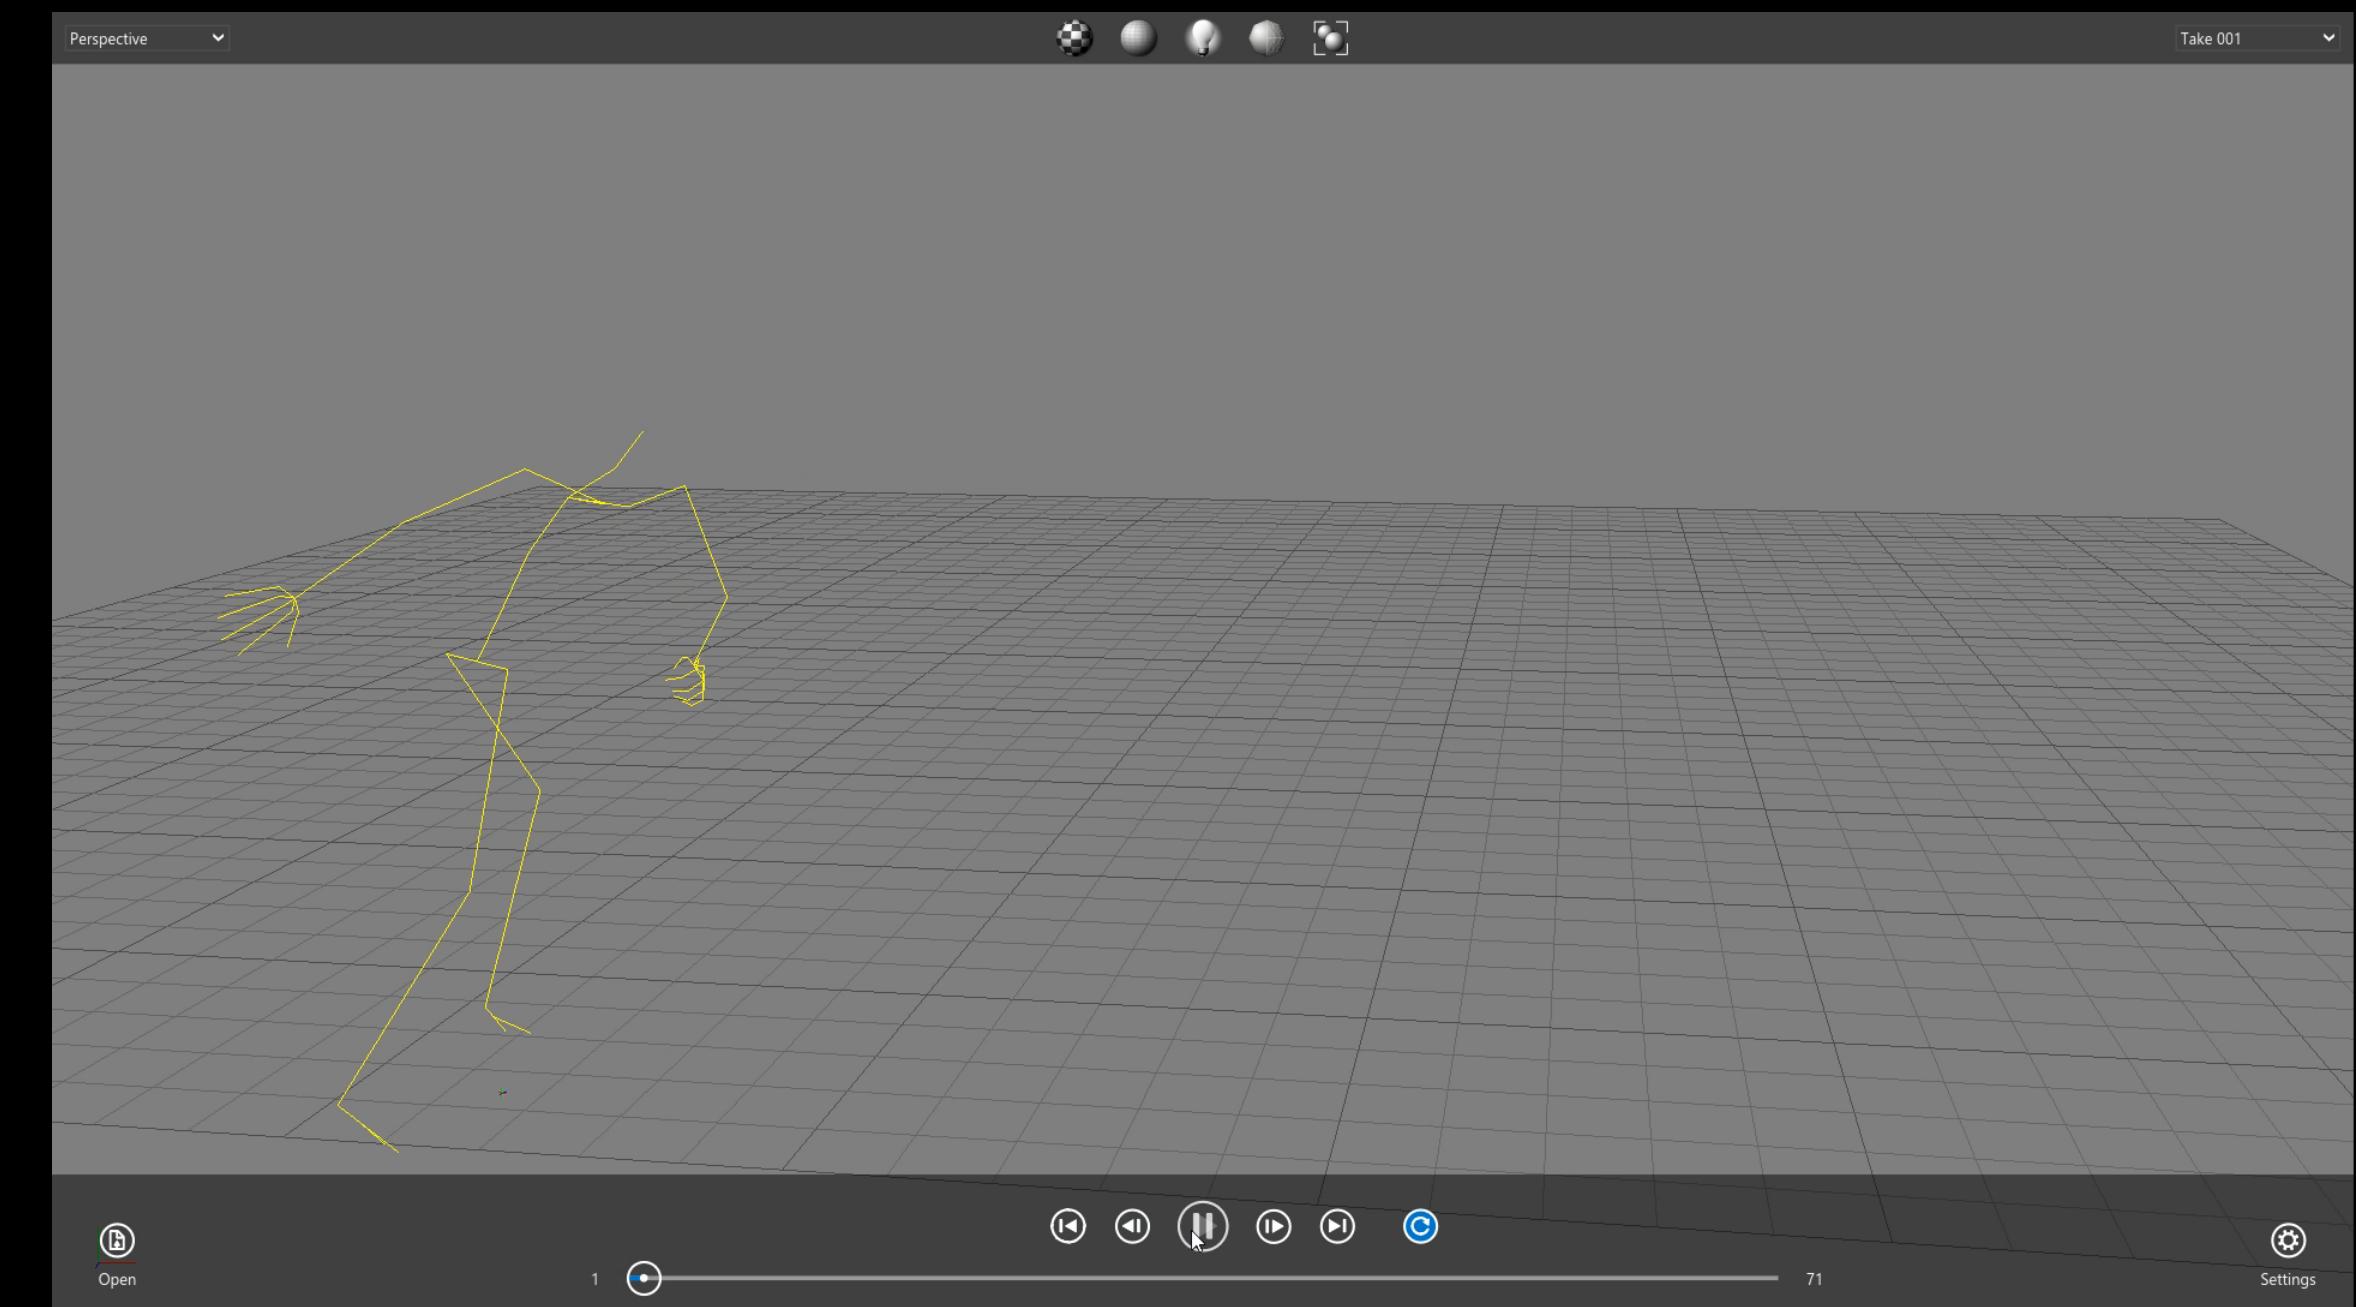
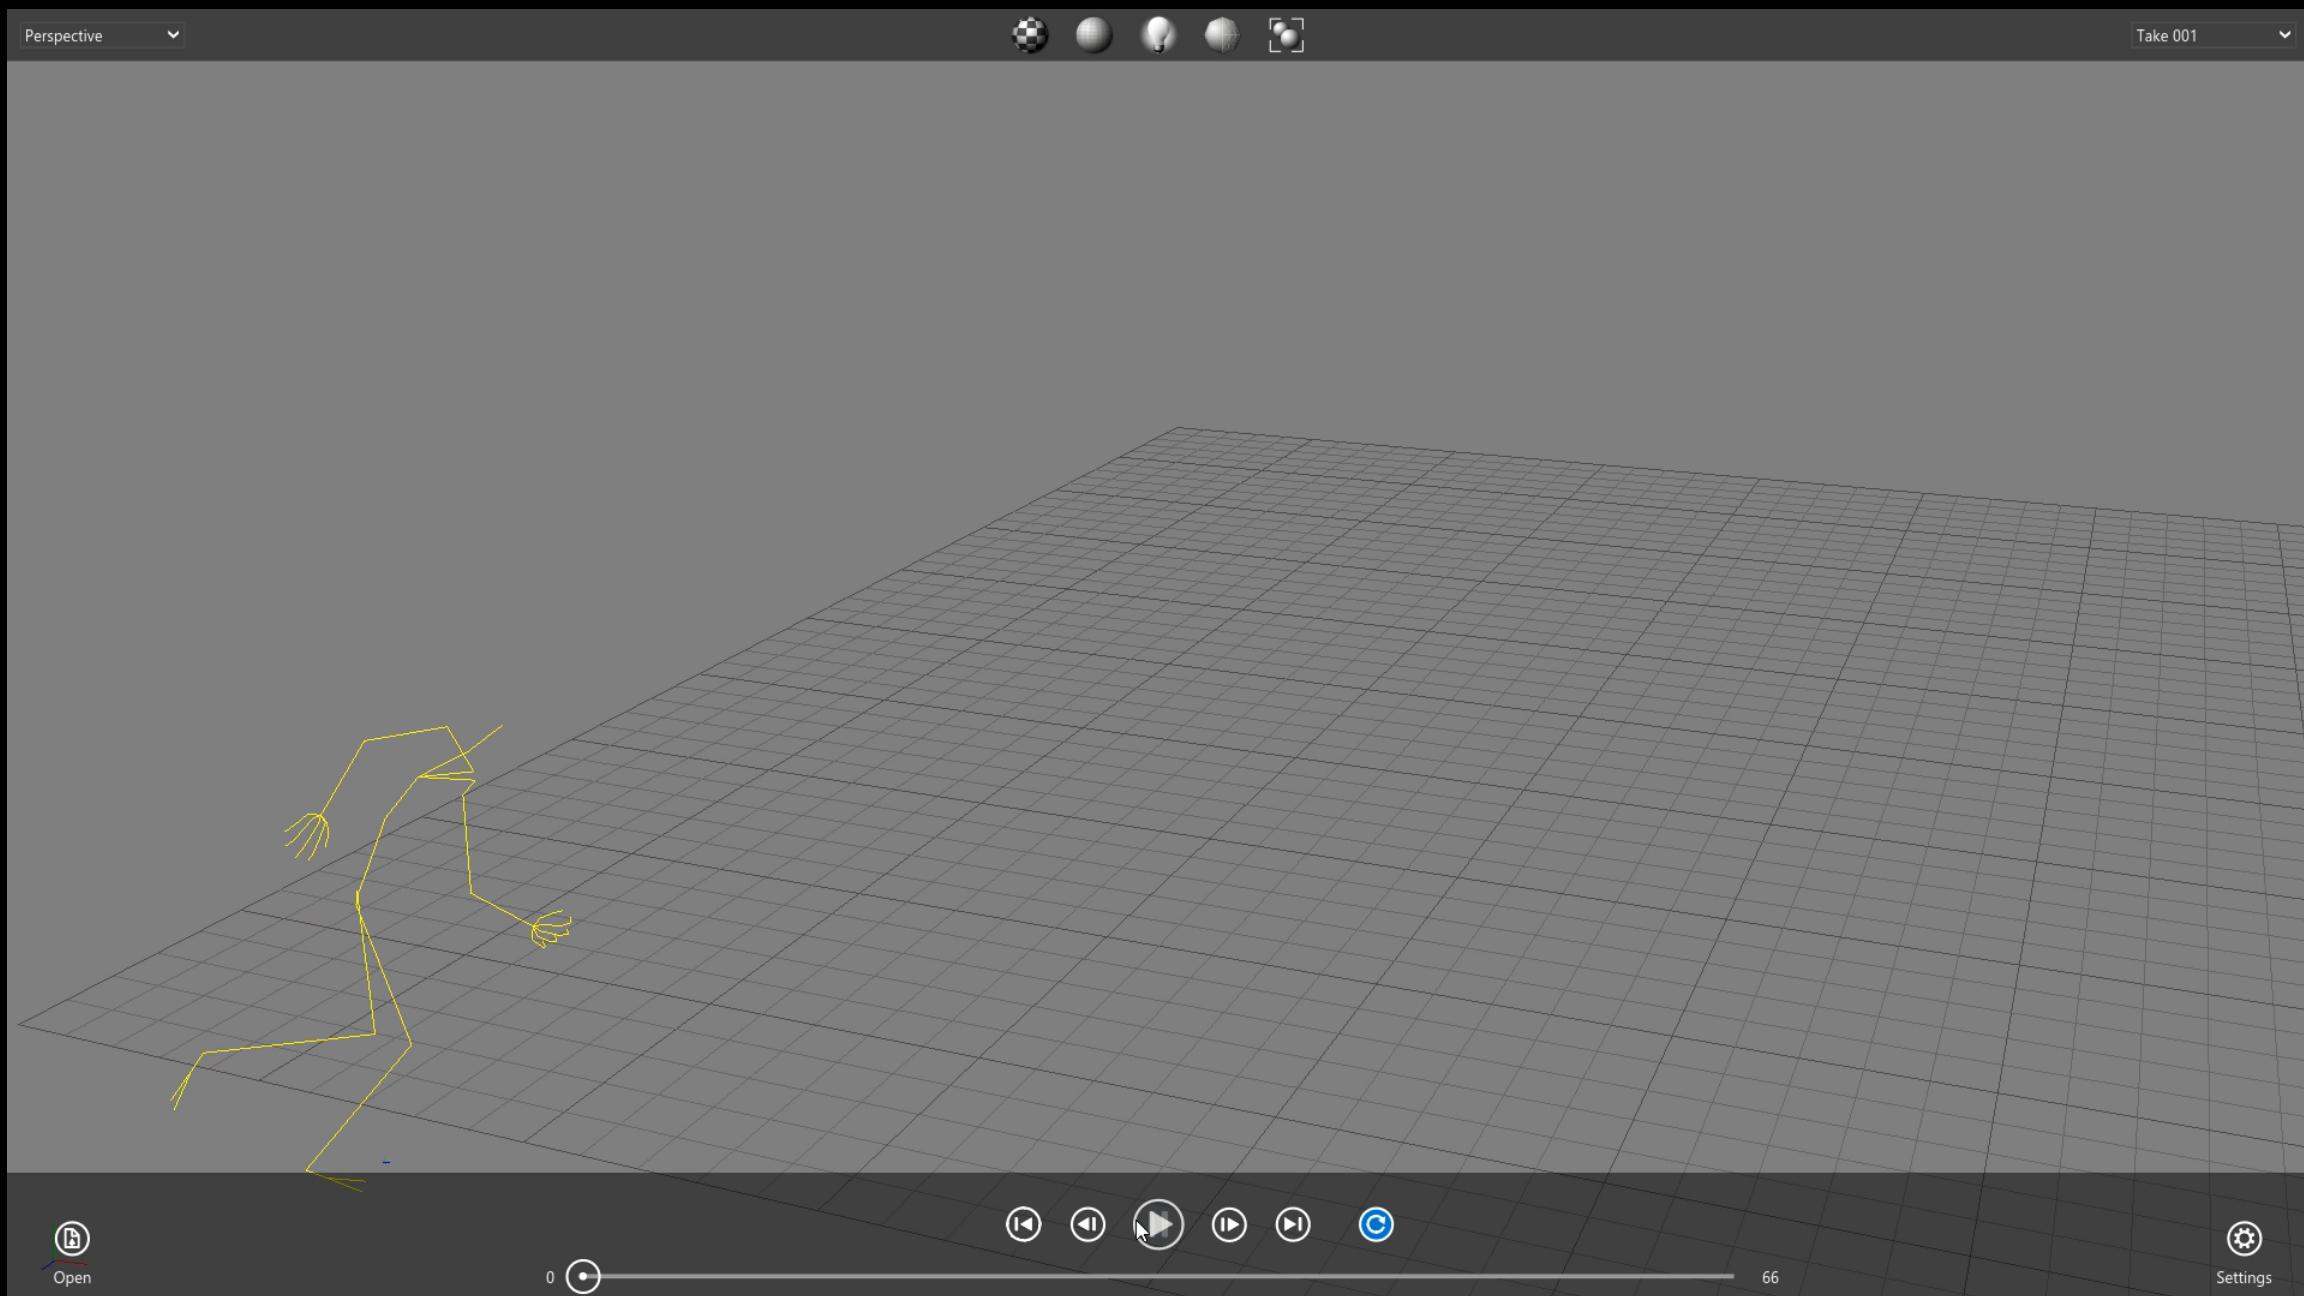
# The 4 No-No's

- Locomotion & Cyclic motions
- Phase as temporal progression
- Pose merging
- Fitting animation along predicted trajectory



# Locomotion & Cyclic motions

Most motion synthesis research emphasizes cyclic motions in general and locomotion in particular



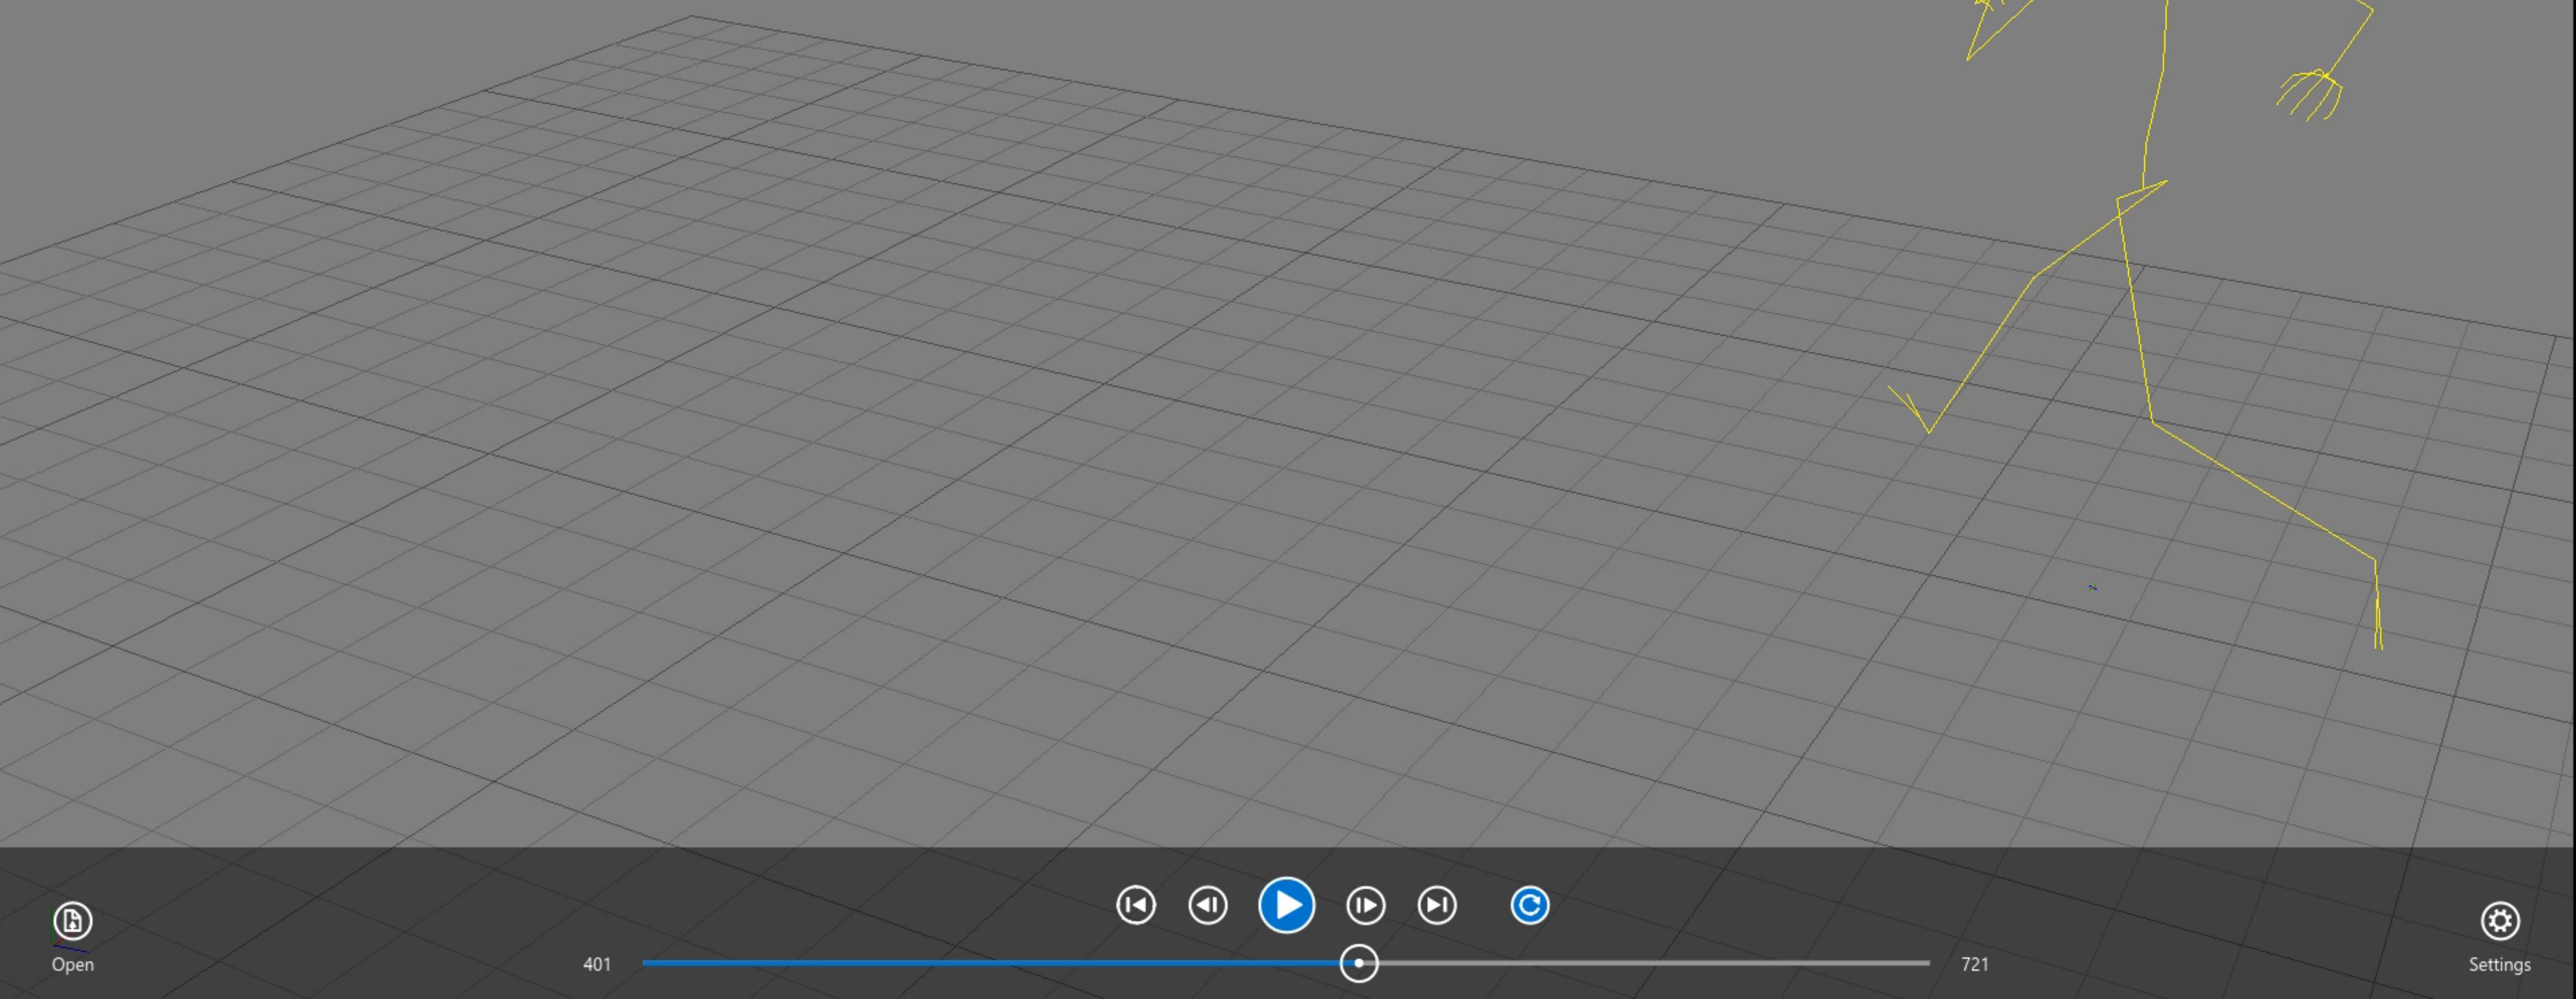
# Phase as temporal progression

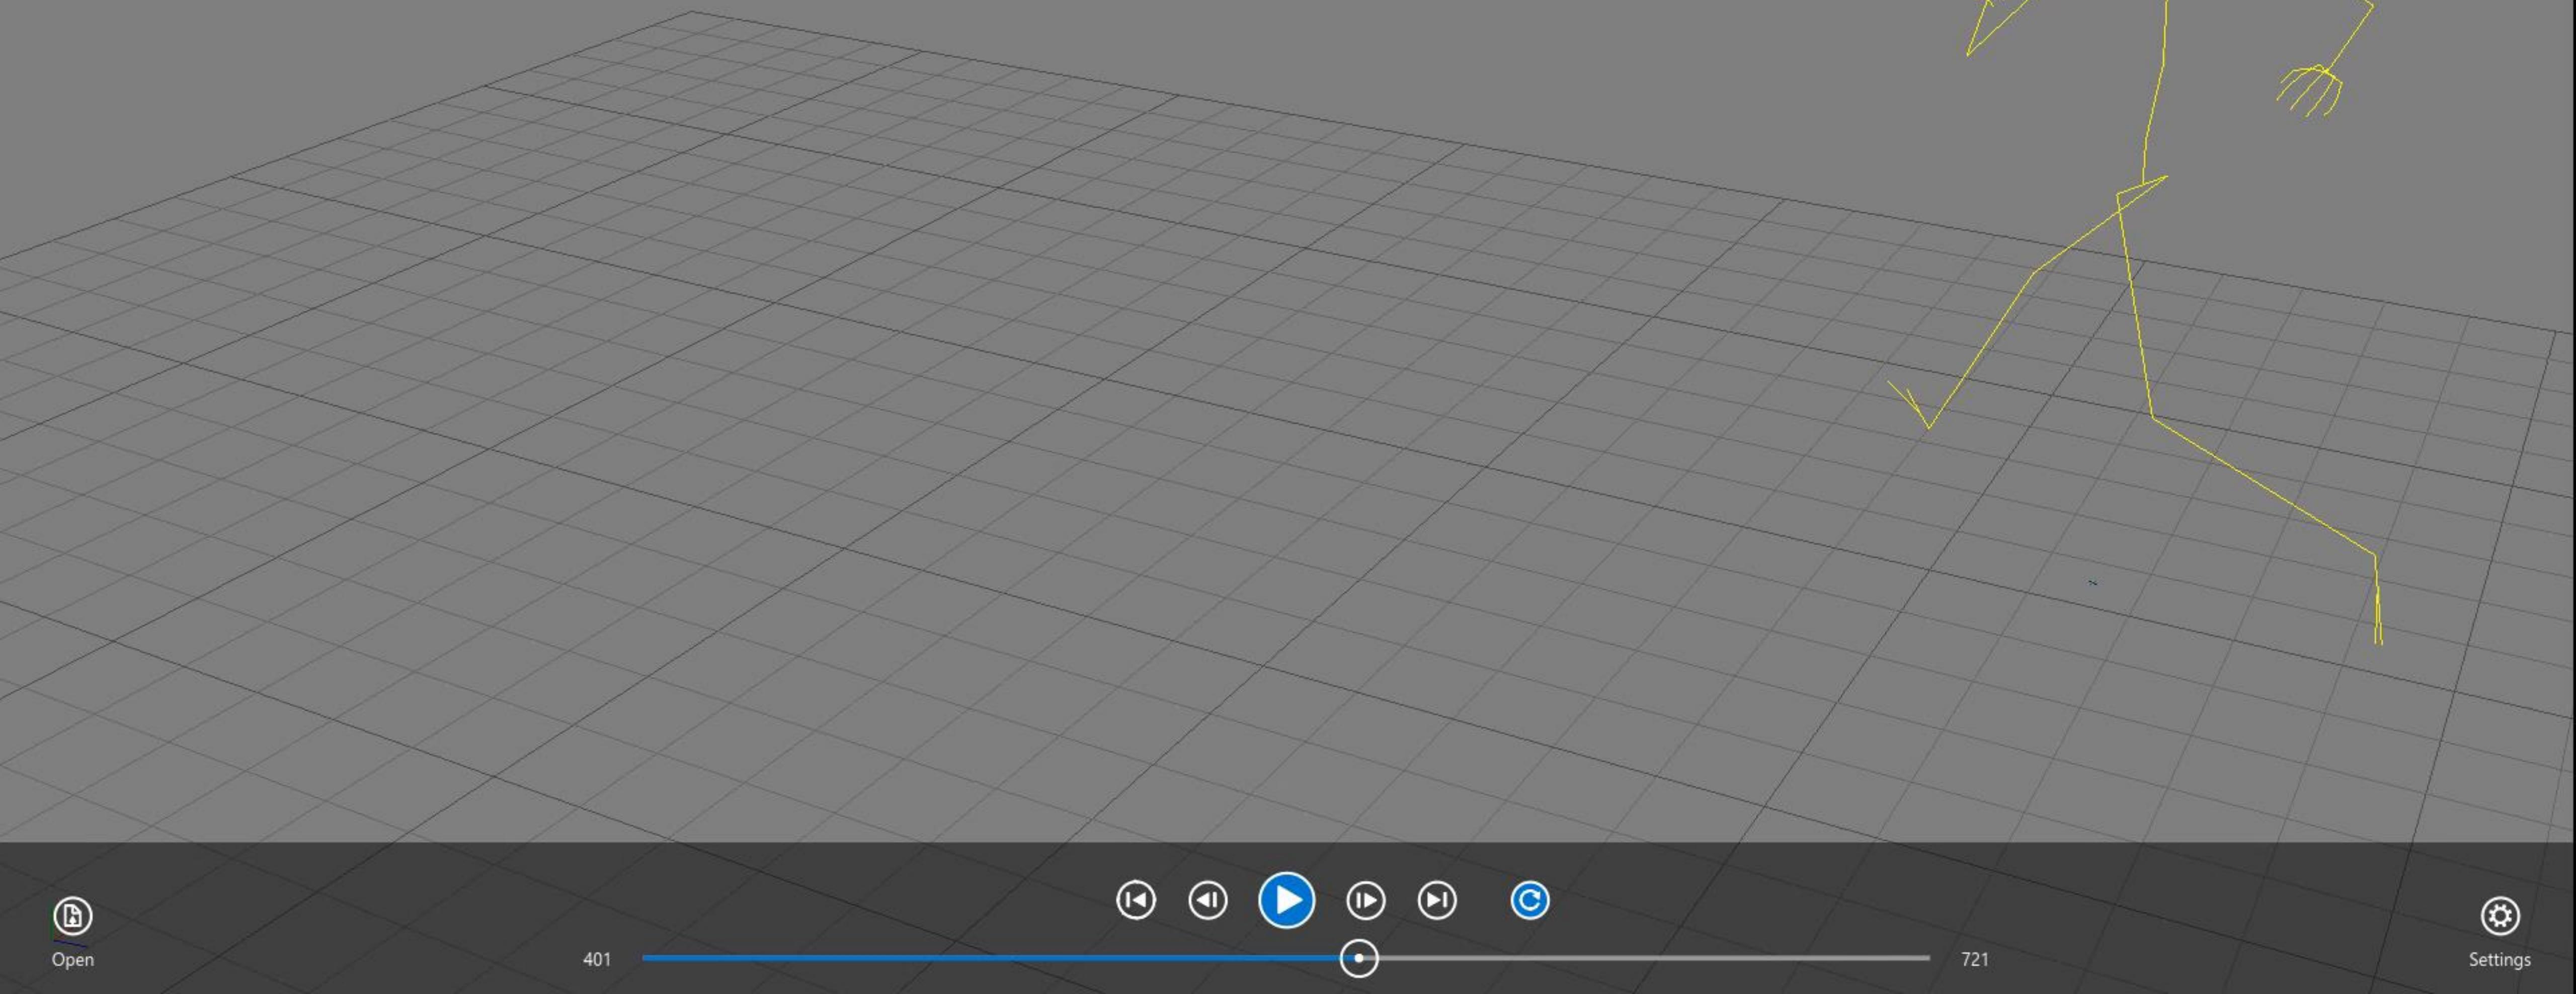
Most motion synthesis research uses the concept of a “phase”; scalar variable in the range  $0$  to  $2\pi$  representing the point in time of the current pose in the locomotion cycle

Not true → Given a pose that corresponds to  $\Theta$  all poses that corresponds to  $\Theta + \Delta t$  are “similar”

“Phase” can have an arbitrary meaning (footcontact, entire “action” like for example cartwheel) – not a general concept

```
while not stop:  
    direction = spline.interpolate(travel_distance, 'tangent')  
    facing_direction = spline.interpolate(travel_distance, 'direction')  
    avg_speed = spline.interpolate(travel_distance, 'amplitude').squeeze()  
    freq = spline.interpolate(travel_distance, 'frequency').squeeze()  
  
    rot, displacement, height = next_frame(current_phase, avg_speed, direction, facing_direction)  
  
    current_phase = (current_phase + freq) % (2*np.pi)  
    travel_distance += avg_speed  
    next_distance = travel_distance + displacement  
  
    if next_distance < spline.length():  
        rotations.append(rot)  
        positions.append((next_distance, height))  
    else:  
        # End of spline reached  
        stop = True
```

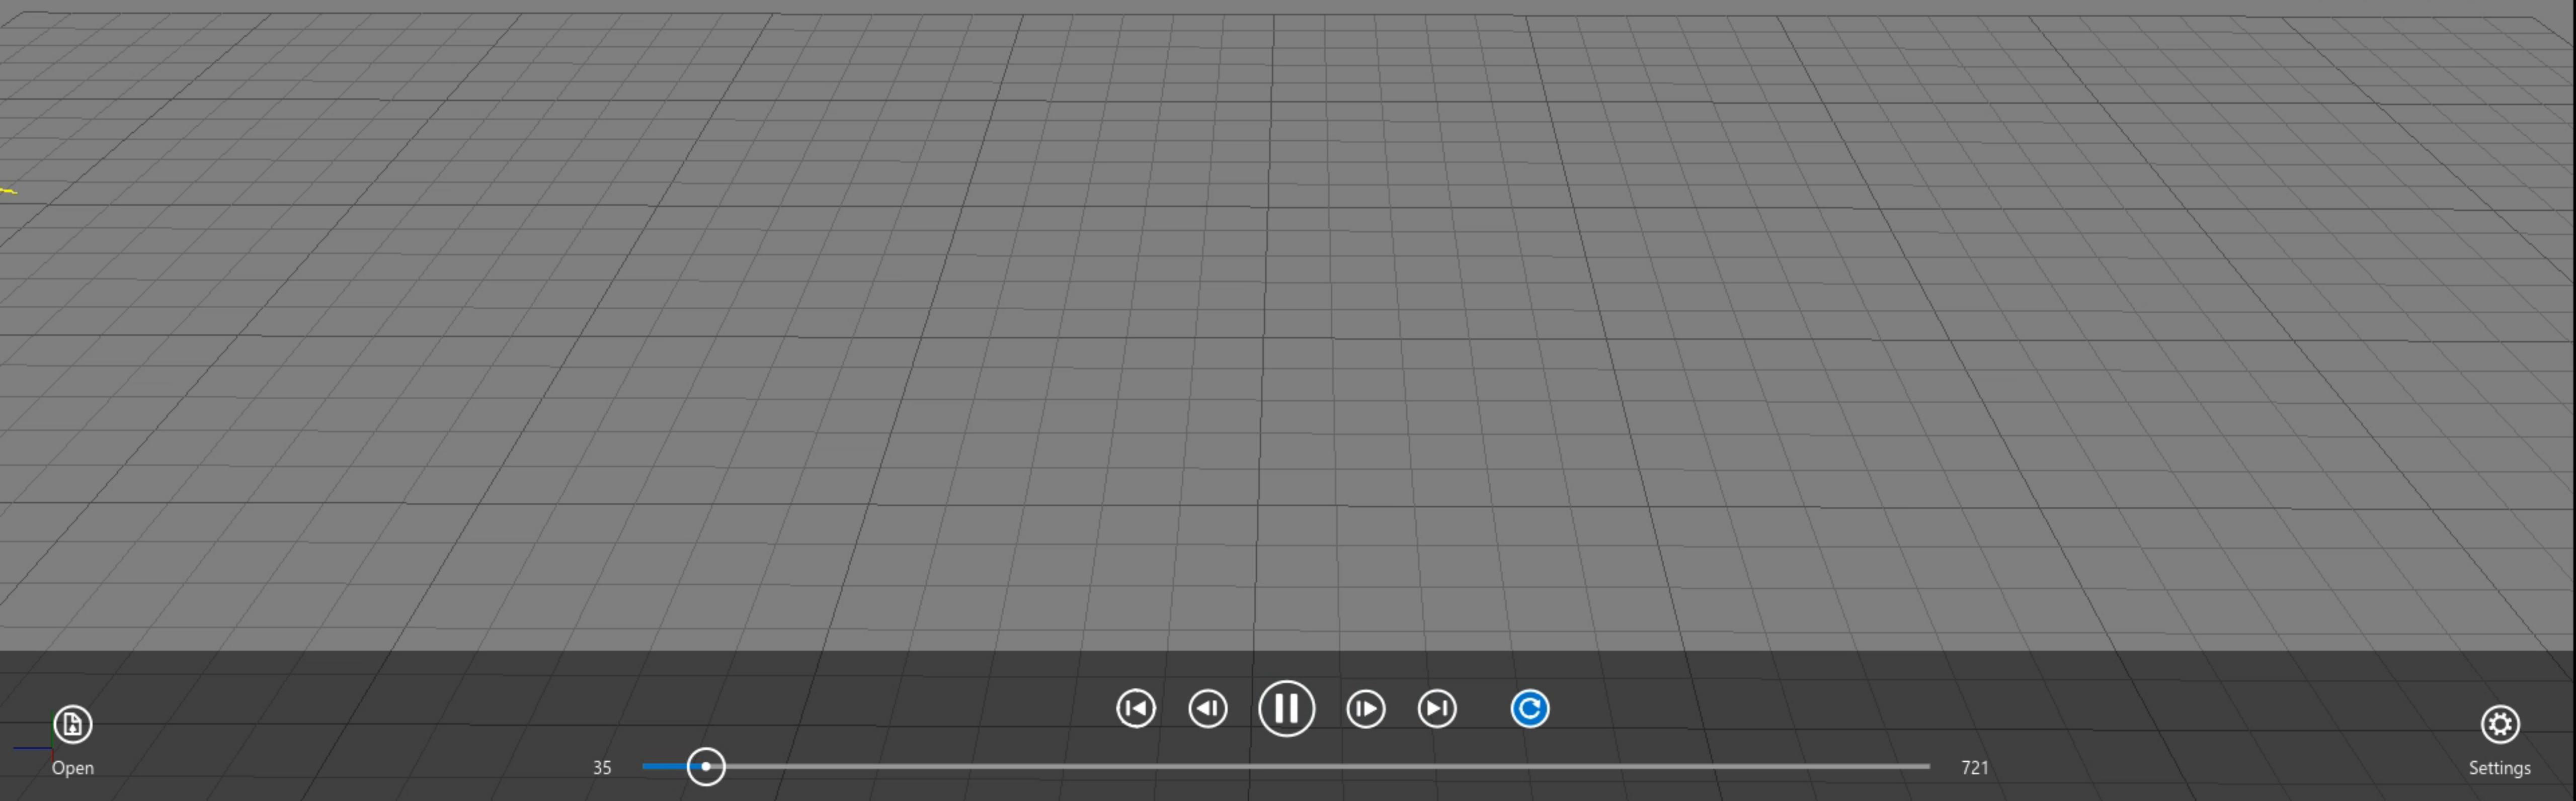




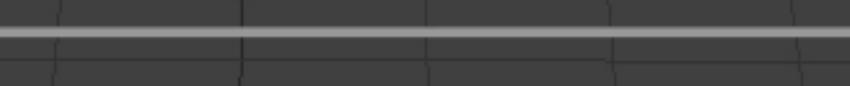
Perspective



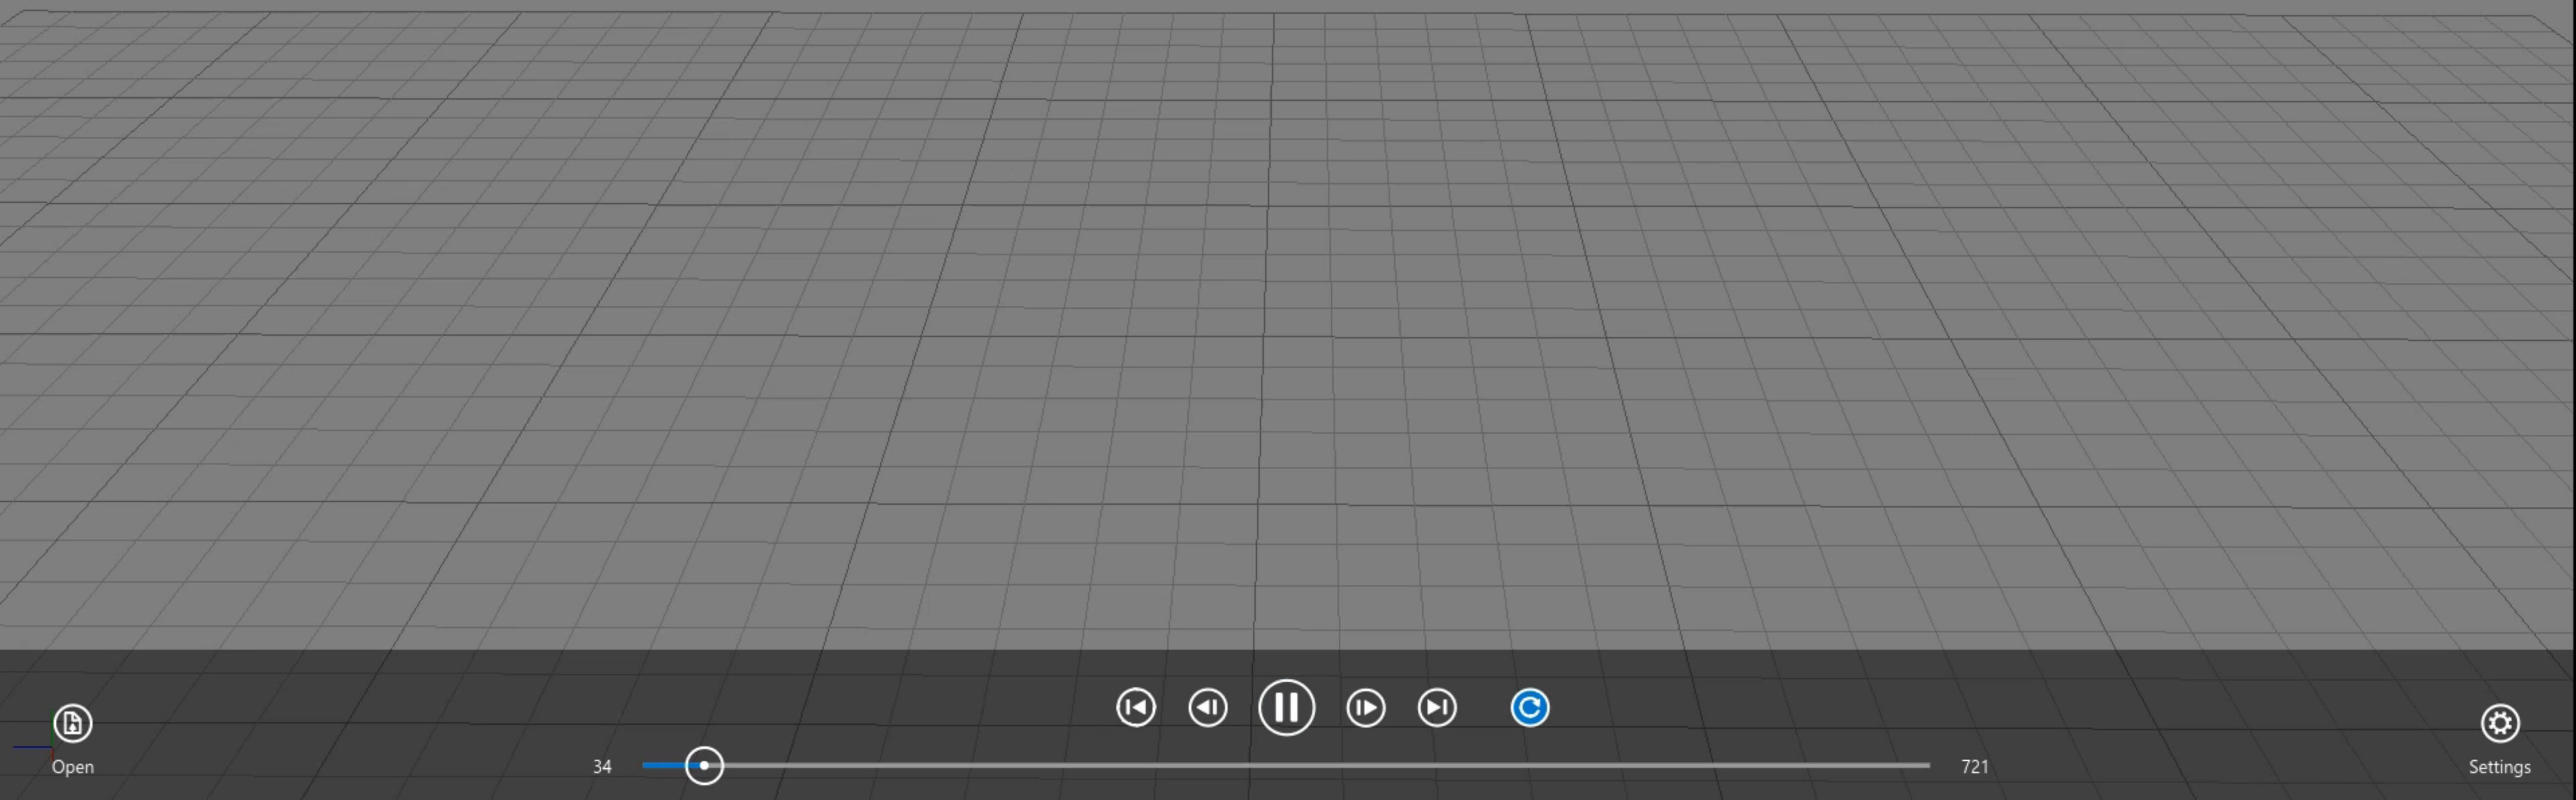
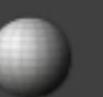
Take 001



Open



Settings



Open



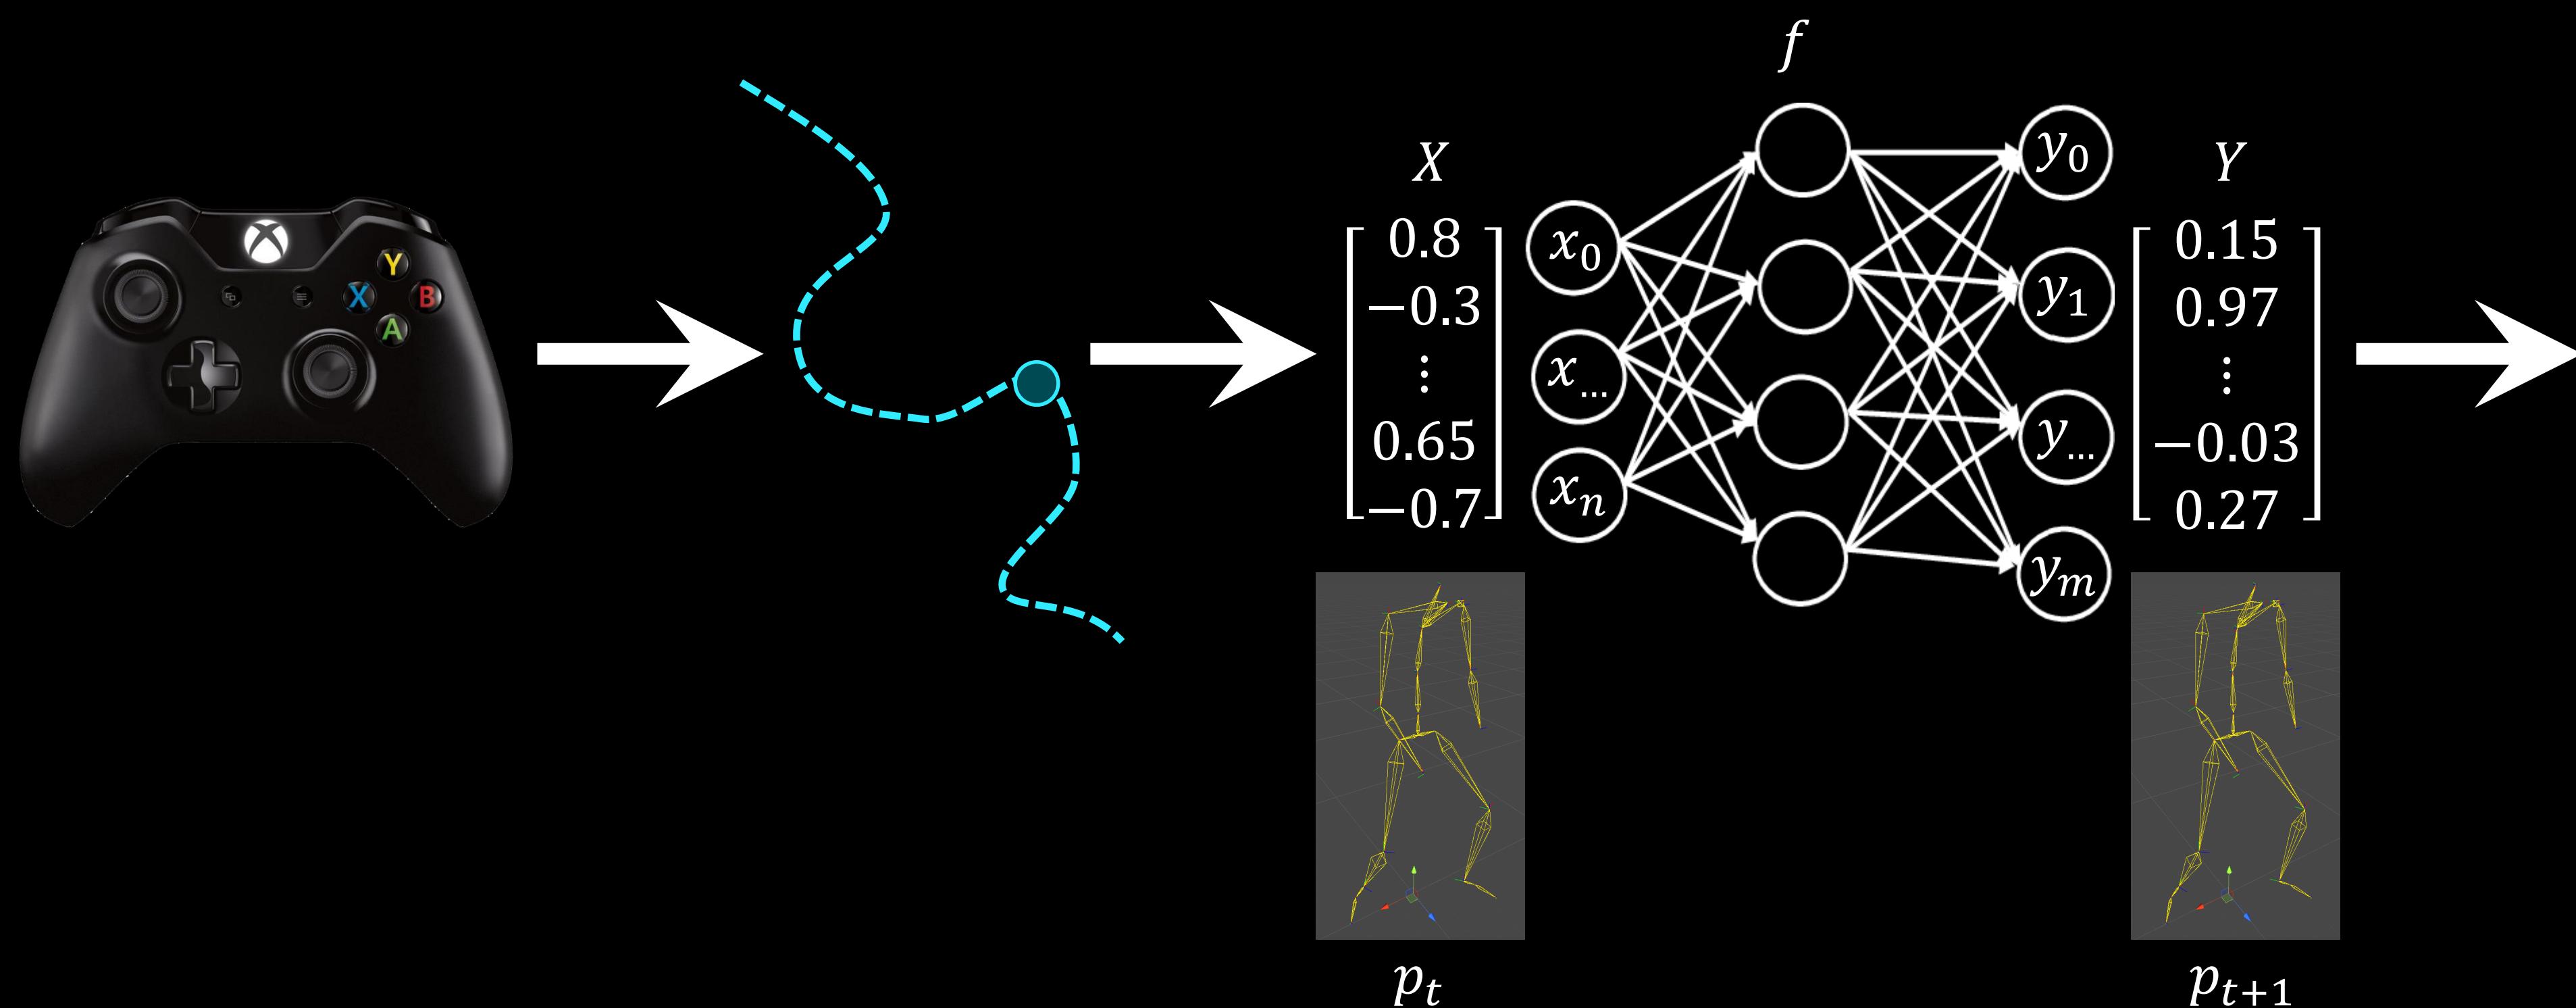
34

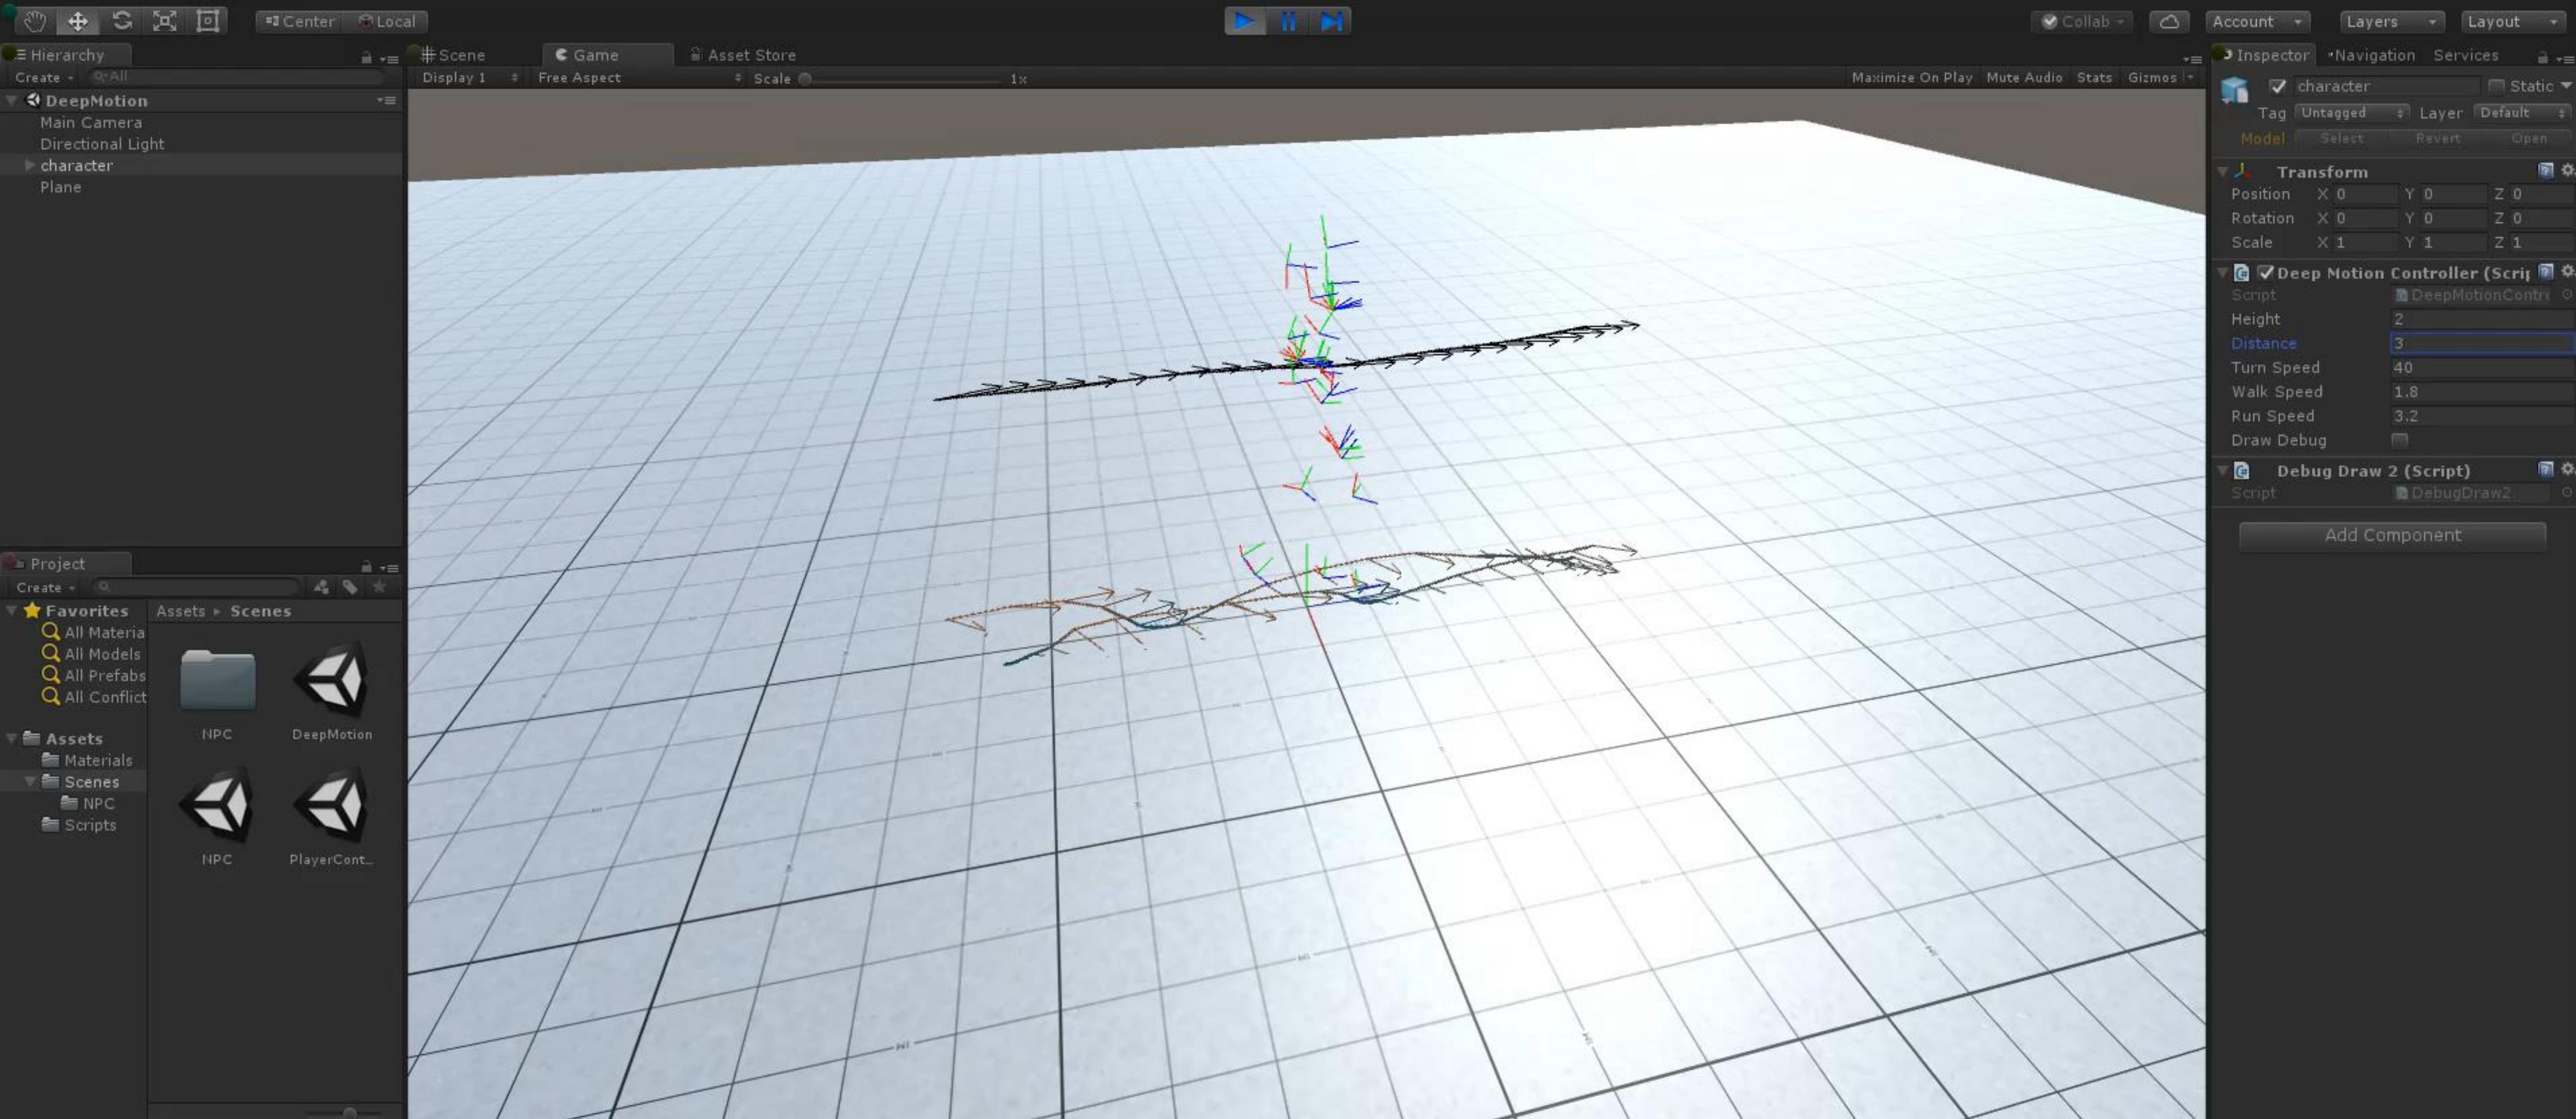
721



Settings

# Autoregressive methods





Console

Clear Collapse Clear on Play Error Pause Connected Plays + Stop for Assert Stop for Error

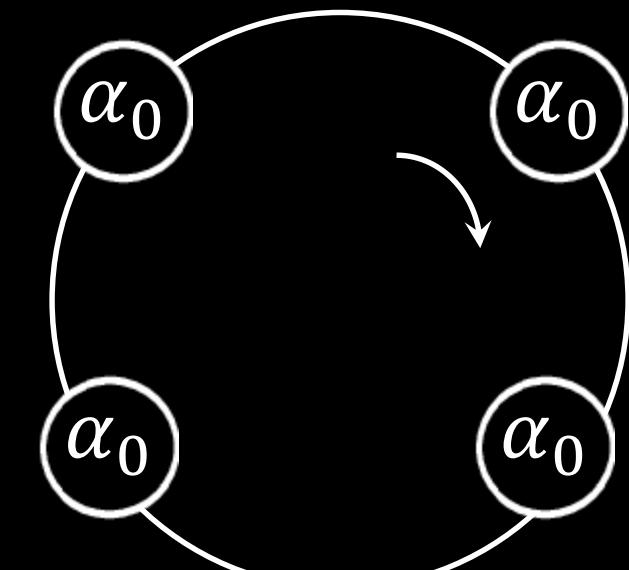
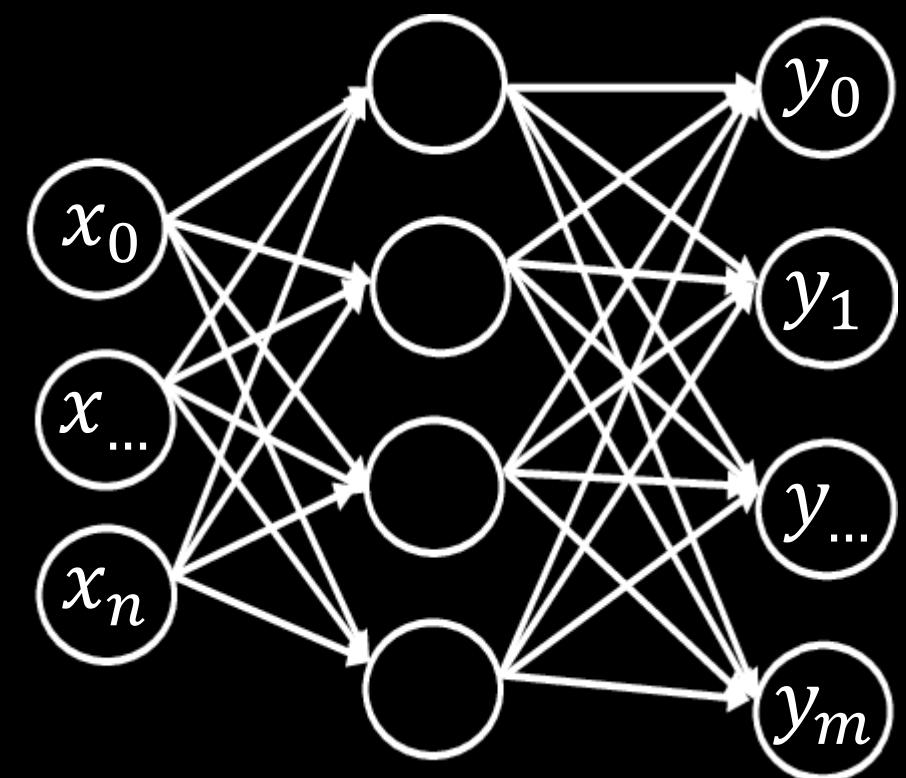
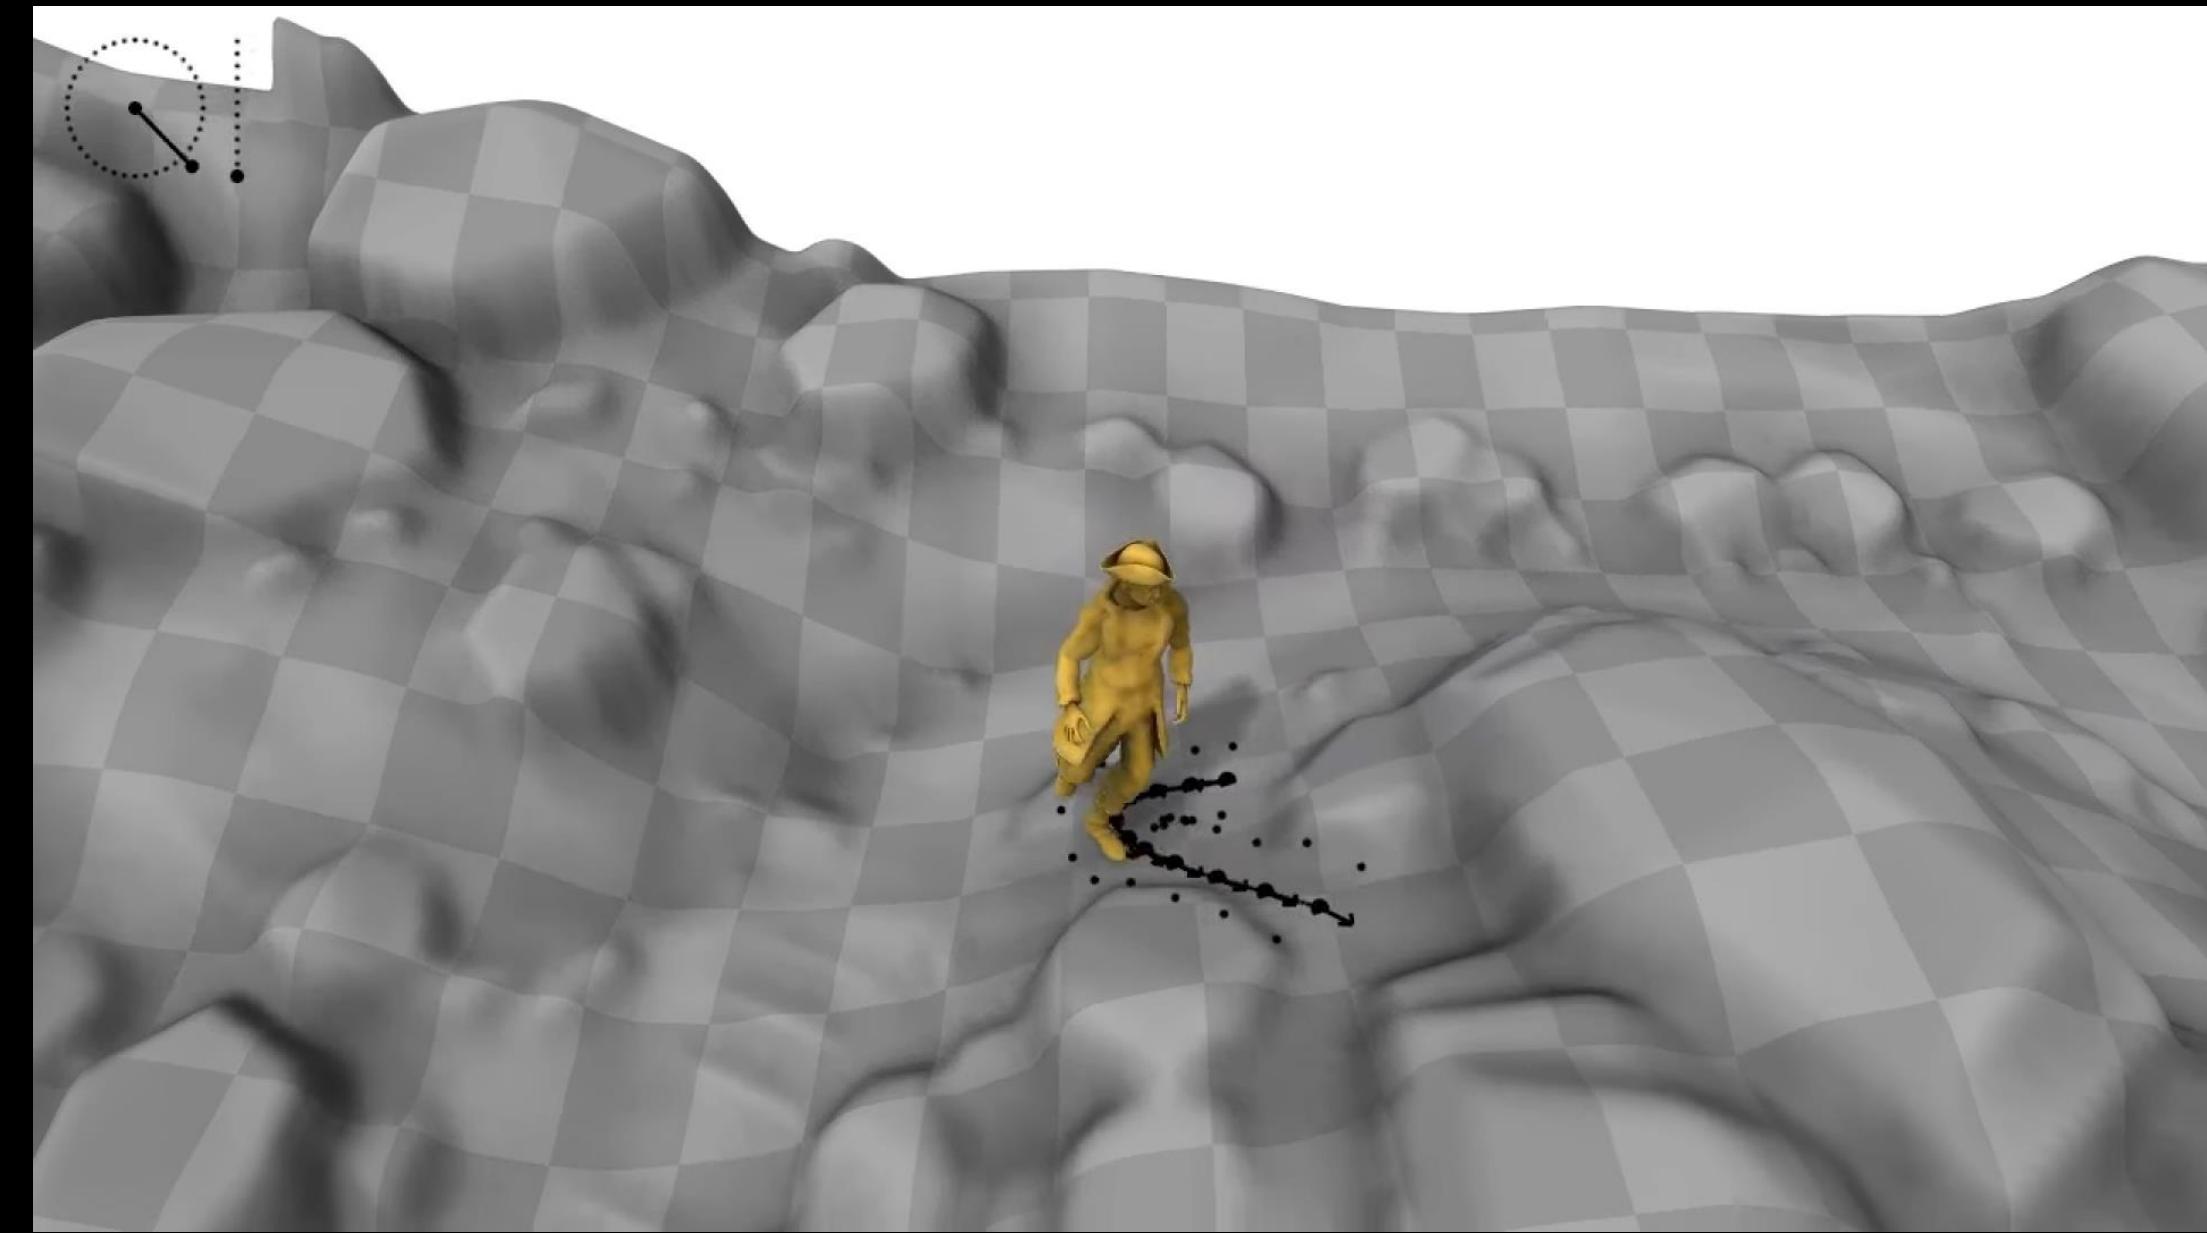
! Error loading launcher://unity/C:/Users/Mickey/AppData/Roaming/Unity/Packages/node\_modules/unity-editor-home/dist/index.html?code=jvIznuNpPAnsV6tAOJZwVw019f&locale=en&session\_state=e72a27f5f69cbbb1bb63d1d76d9fb41f77b41a44c0b08b4b7c563dc57c433df8.U1bbNWrH2JWM7cFB--C1DA003f#/login

! Error loading launcher://unity/C:/Users/Mickey/AppData/Roaming/Unity/Packages/node\_modules/unity-editor-home/dist/index.html?code=jvIznuNpPAnsV6tAOJZwVw019f&locale=en&session\_state=e72a27f5f69cbbb1bb63d1d76d9fb41f77b41a44c0b08b4b7c563dc57c433df8.U1bbNWrH2JWM7cFB--C1DA003f#/login Allocated: 53.0 MB Objects: 3170

# Phase Function Neural Networks

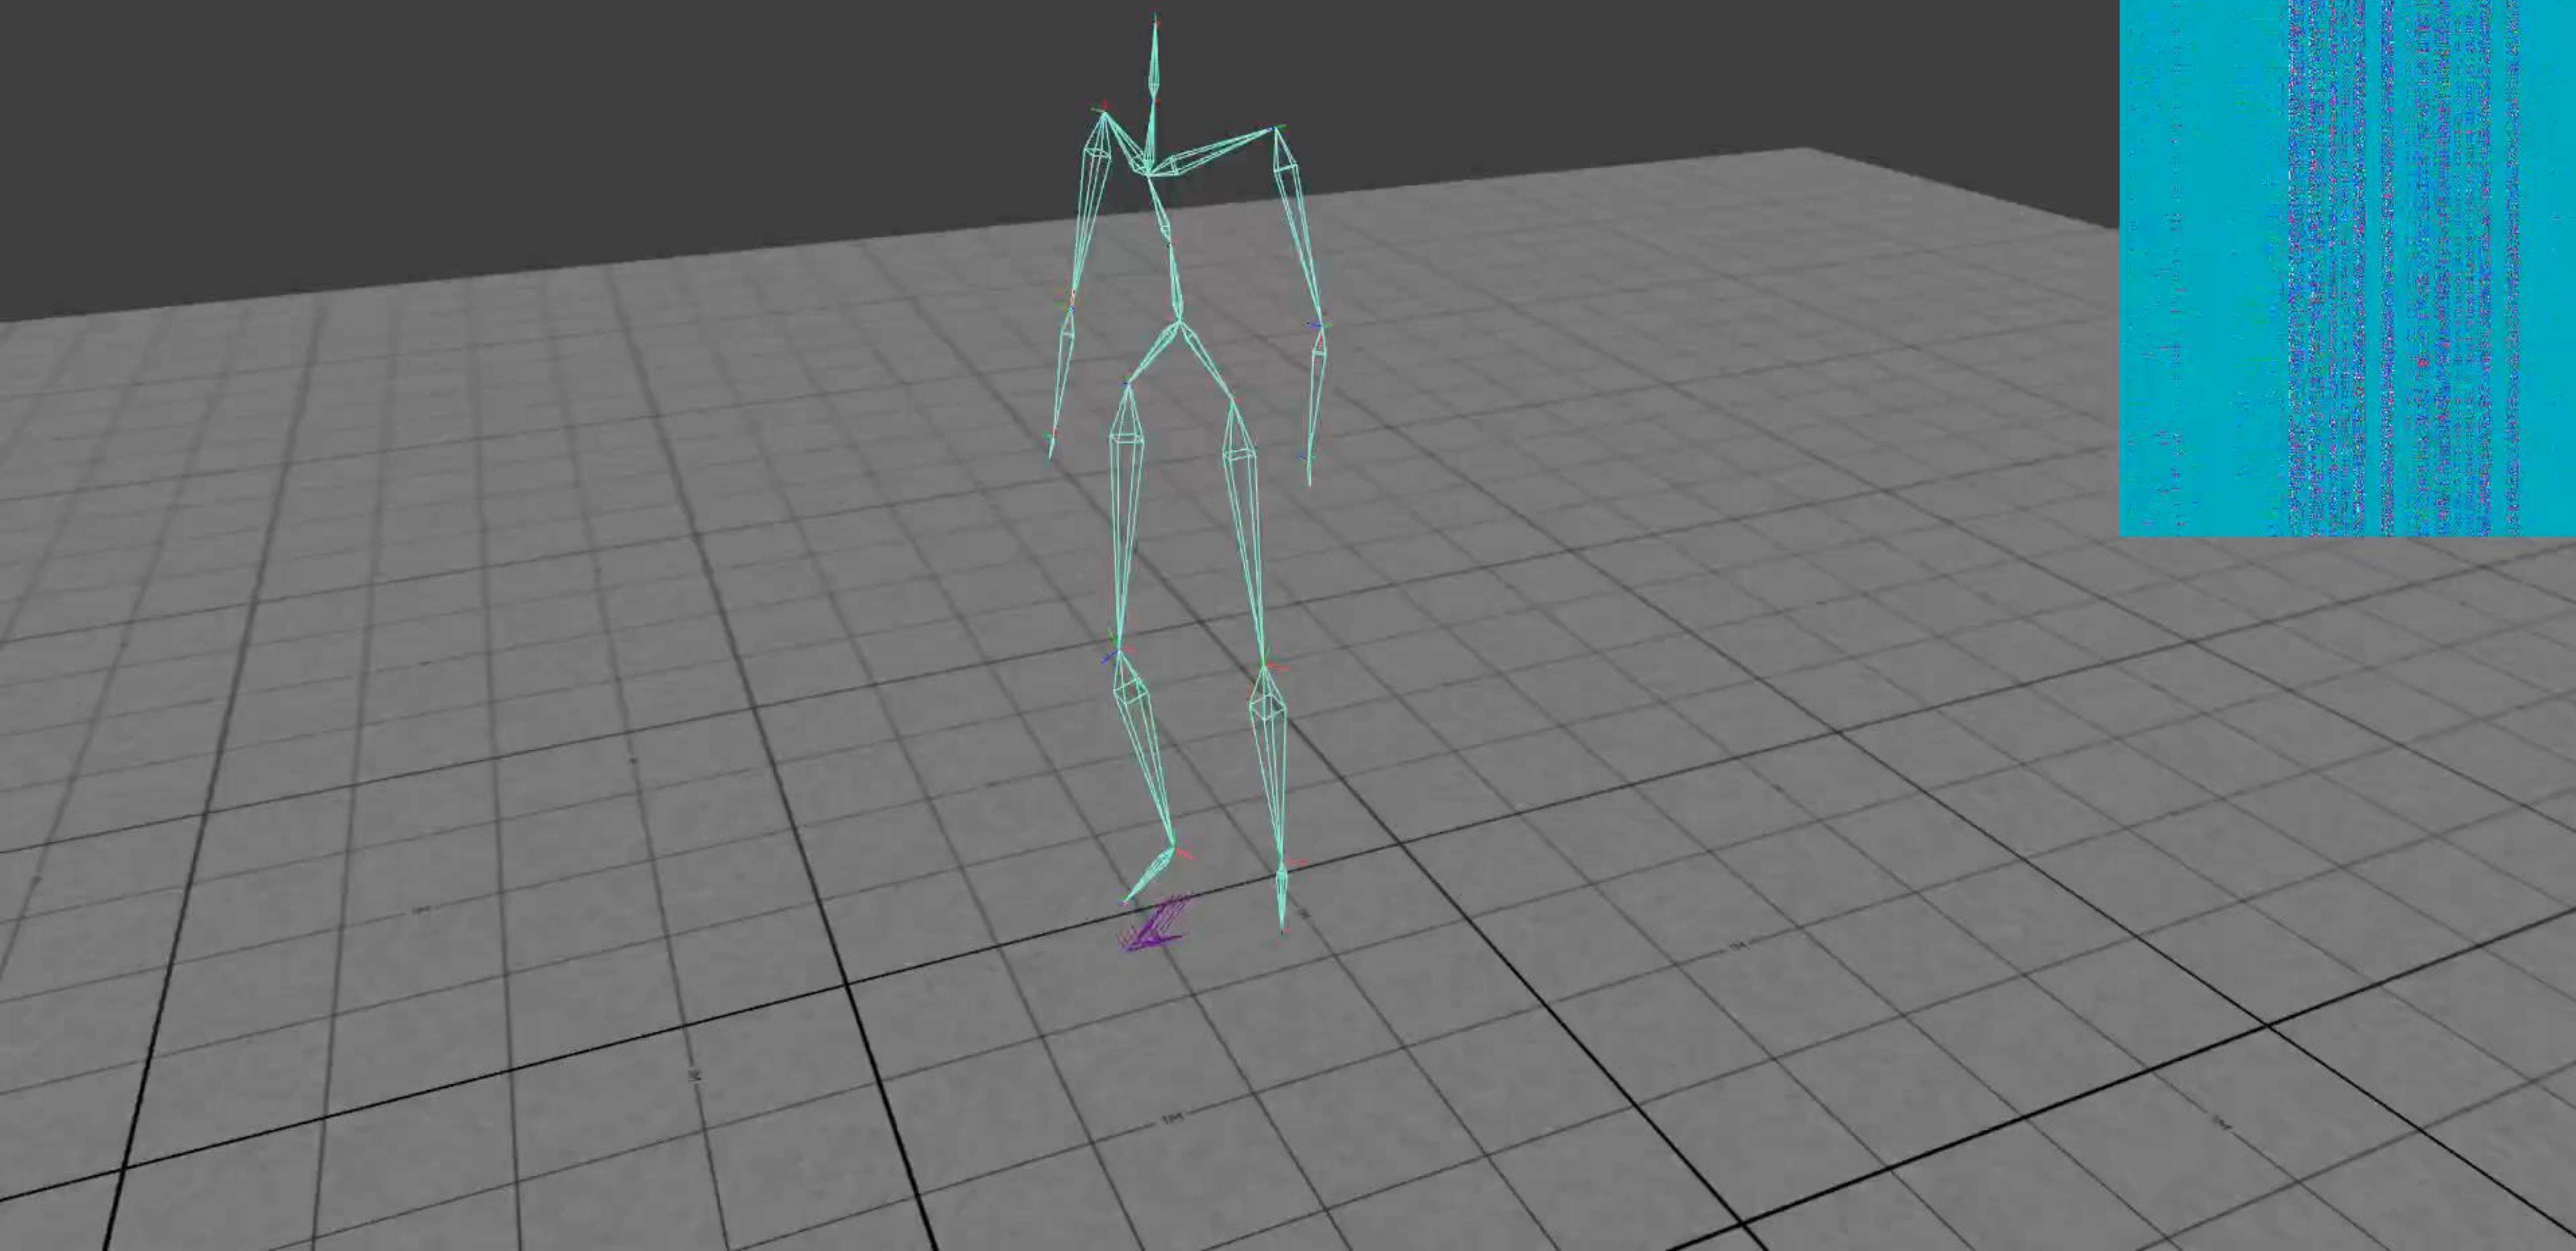
A neural network where the weights are generated as a function of the phase

The “phase” is the scalar variable in the range 0 to  $2\pi$  representing the point in time of the current pose in the locomotion cycle

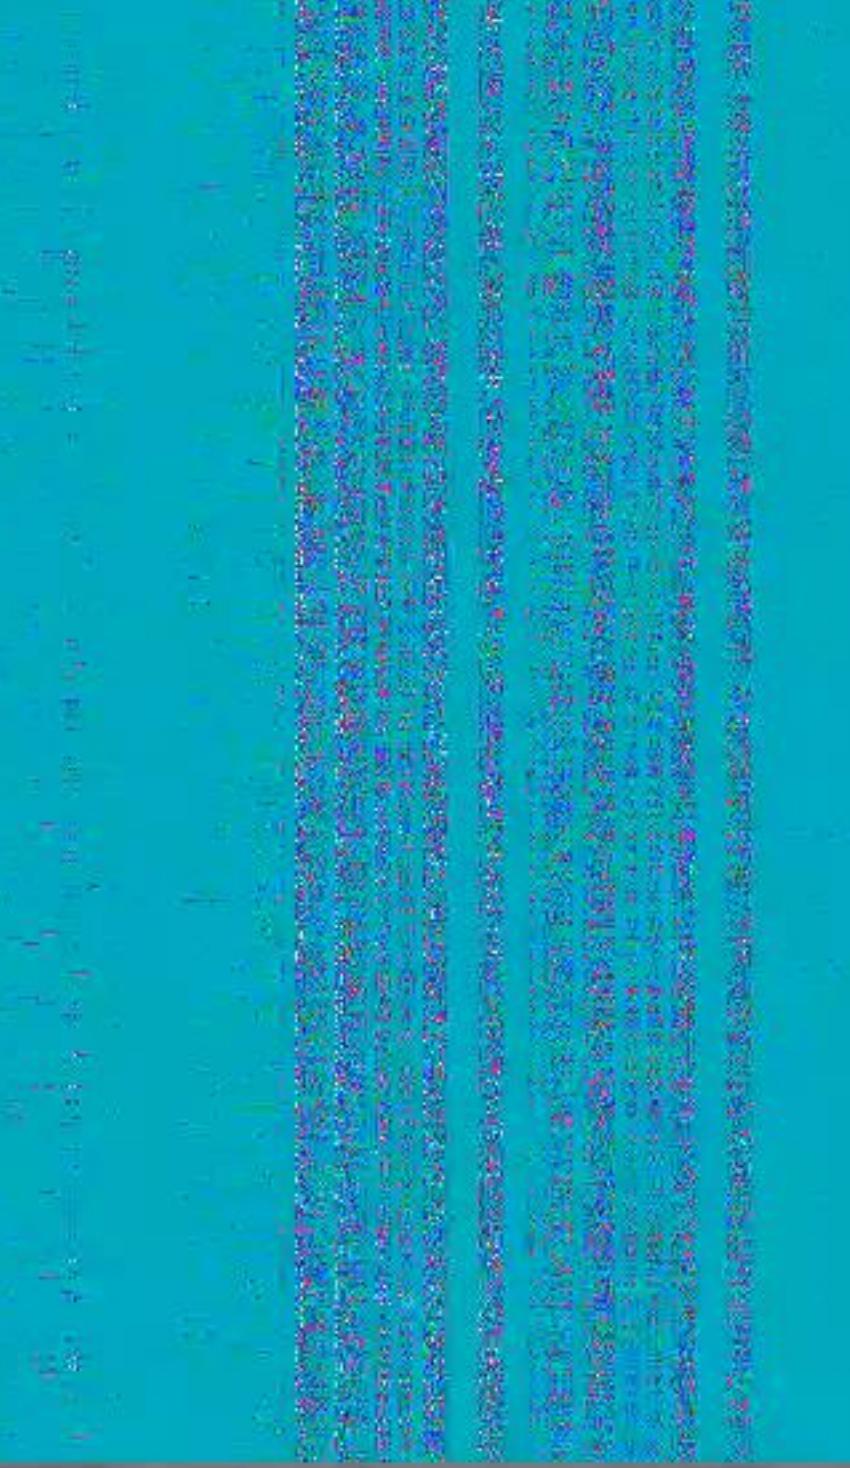
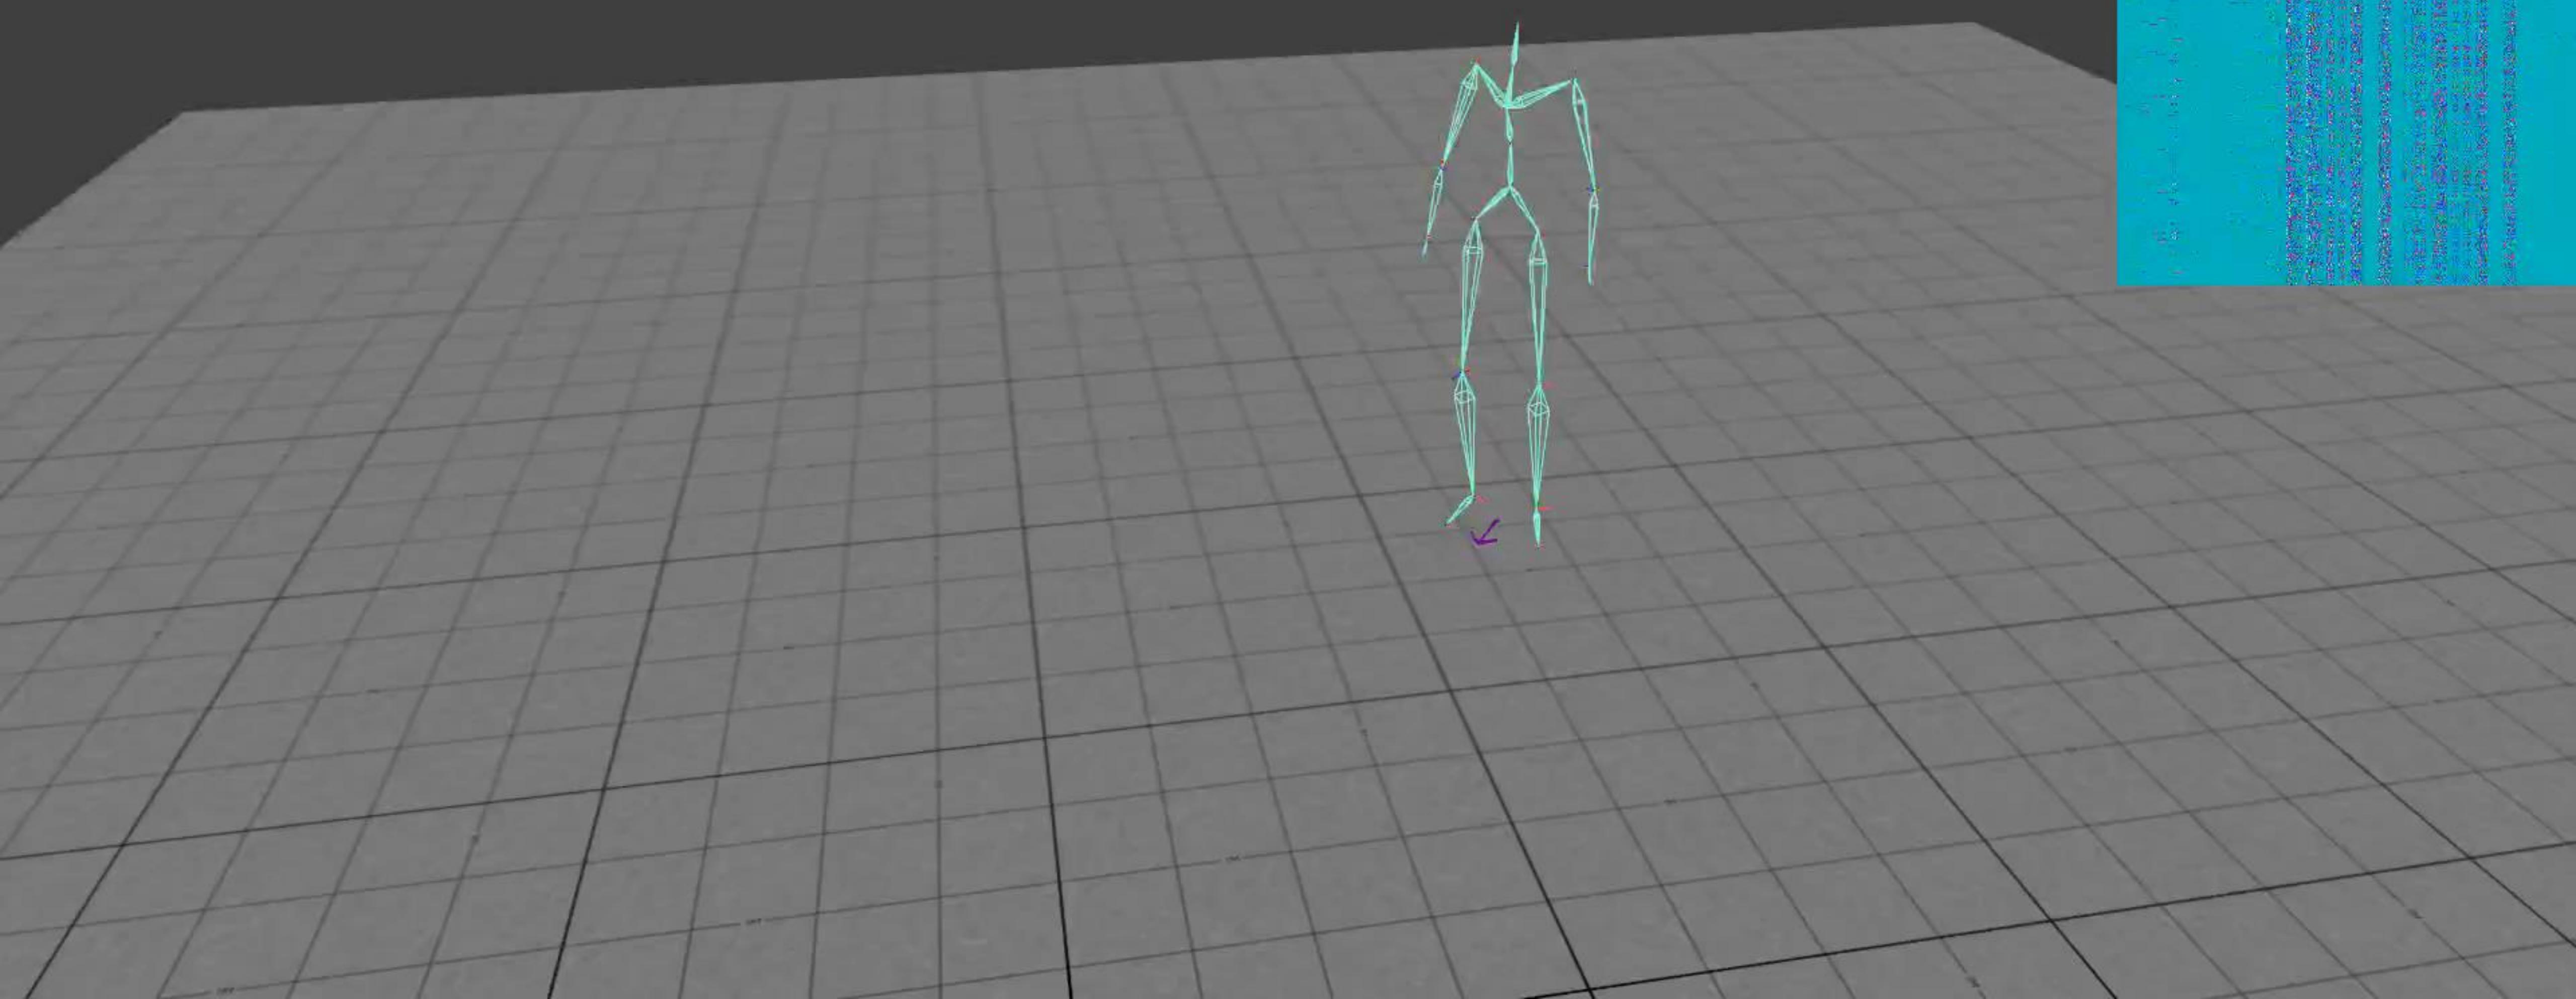


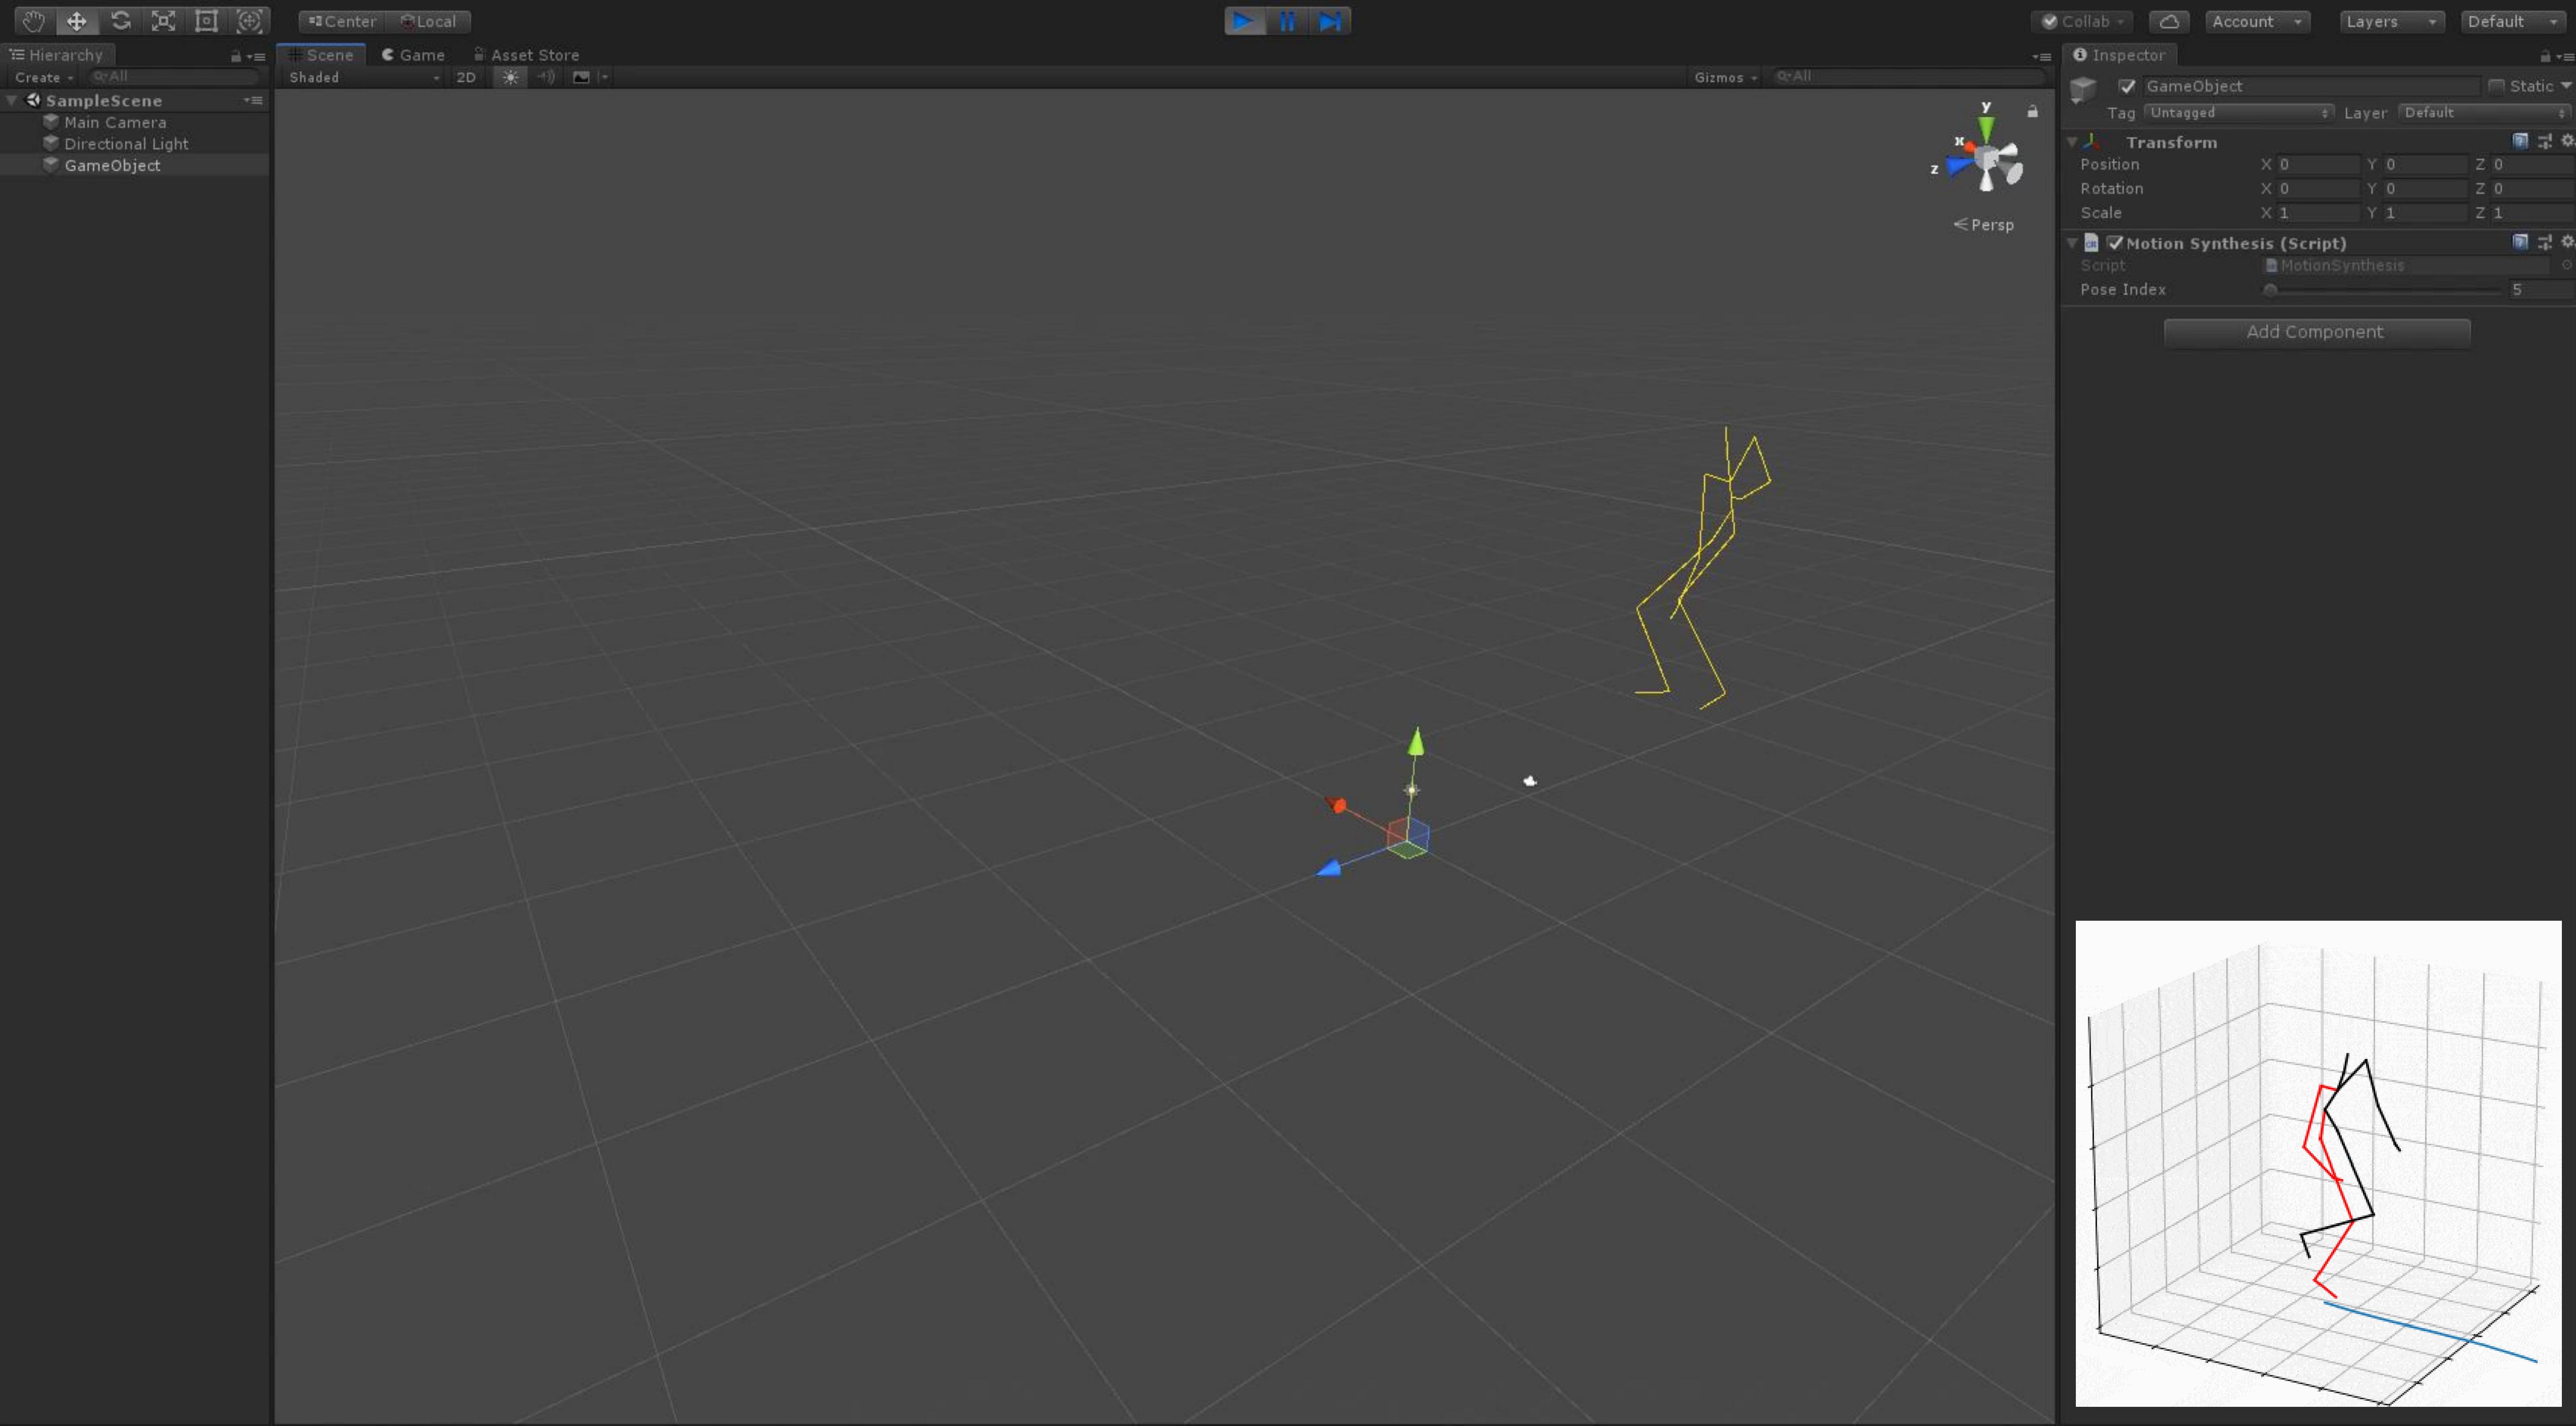
$$W_i = \Theta_i(p)$$

Ik OFF  
Prediction ON  
Target speed 5.50

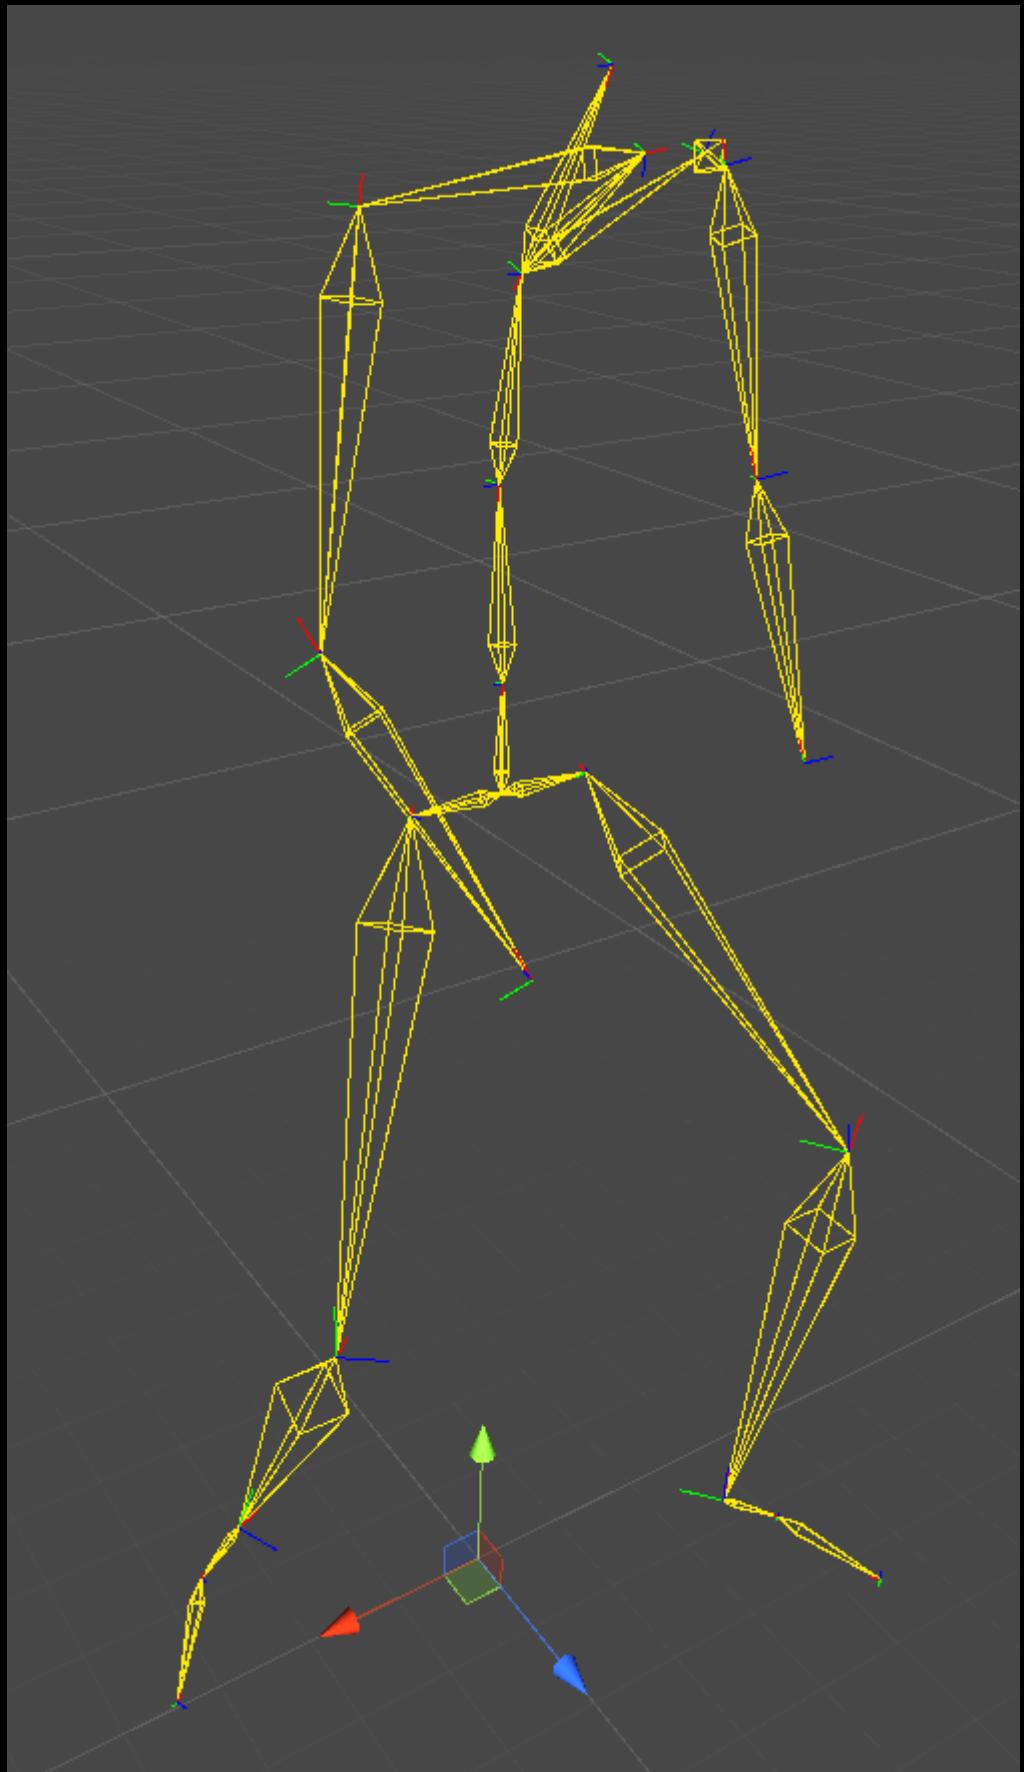
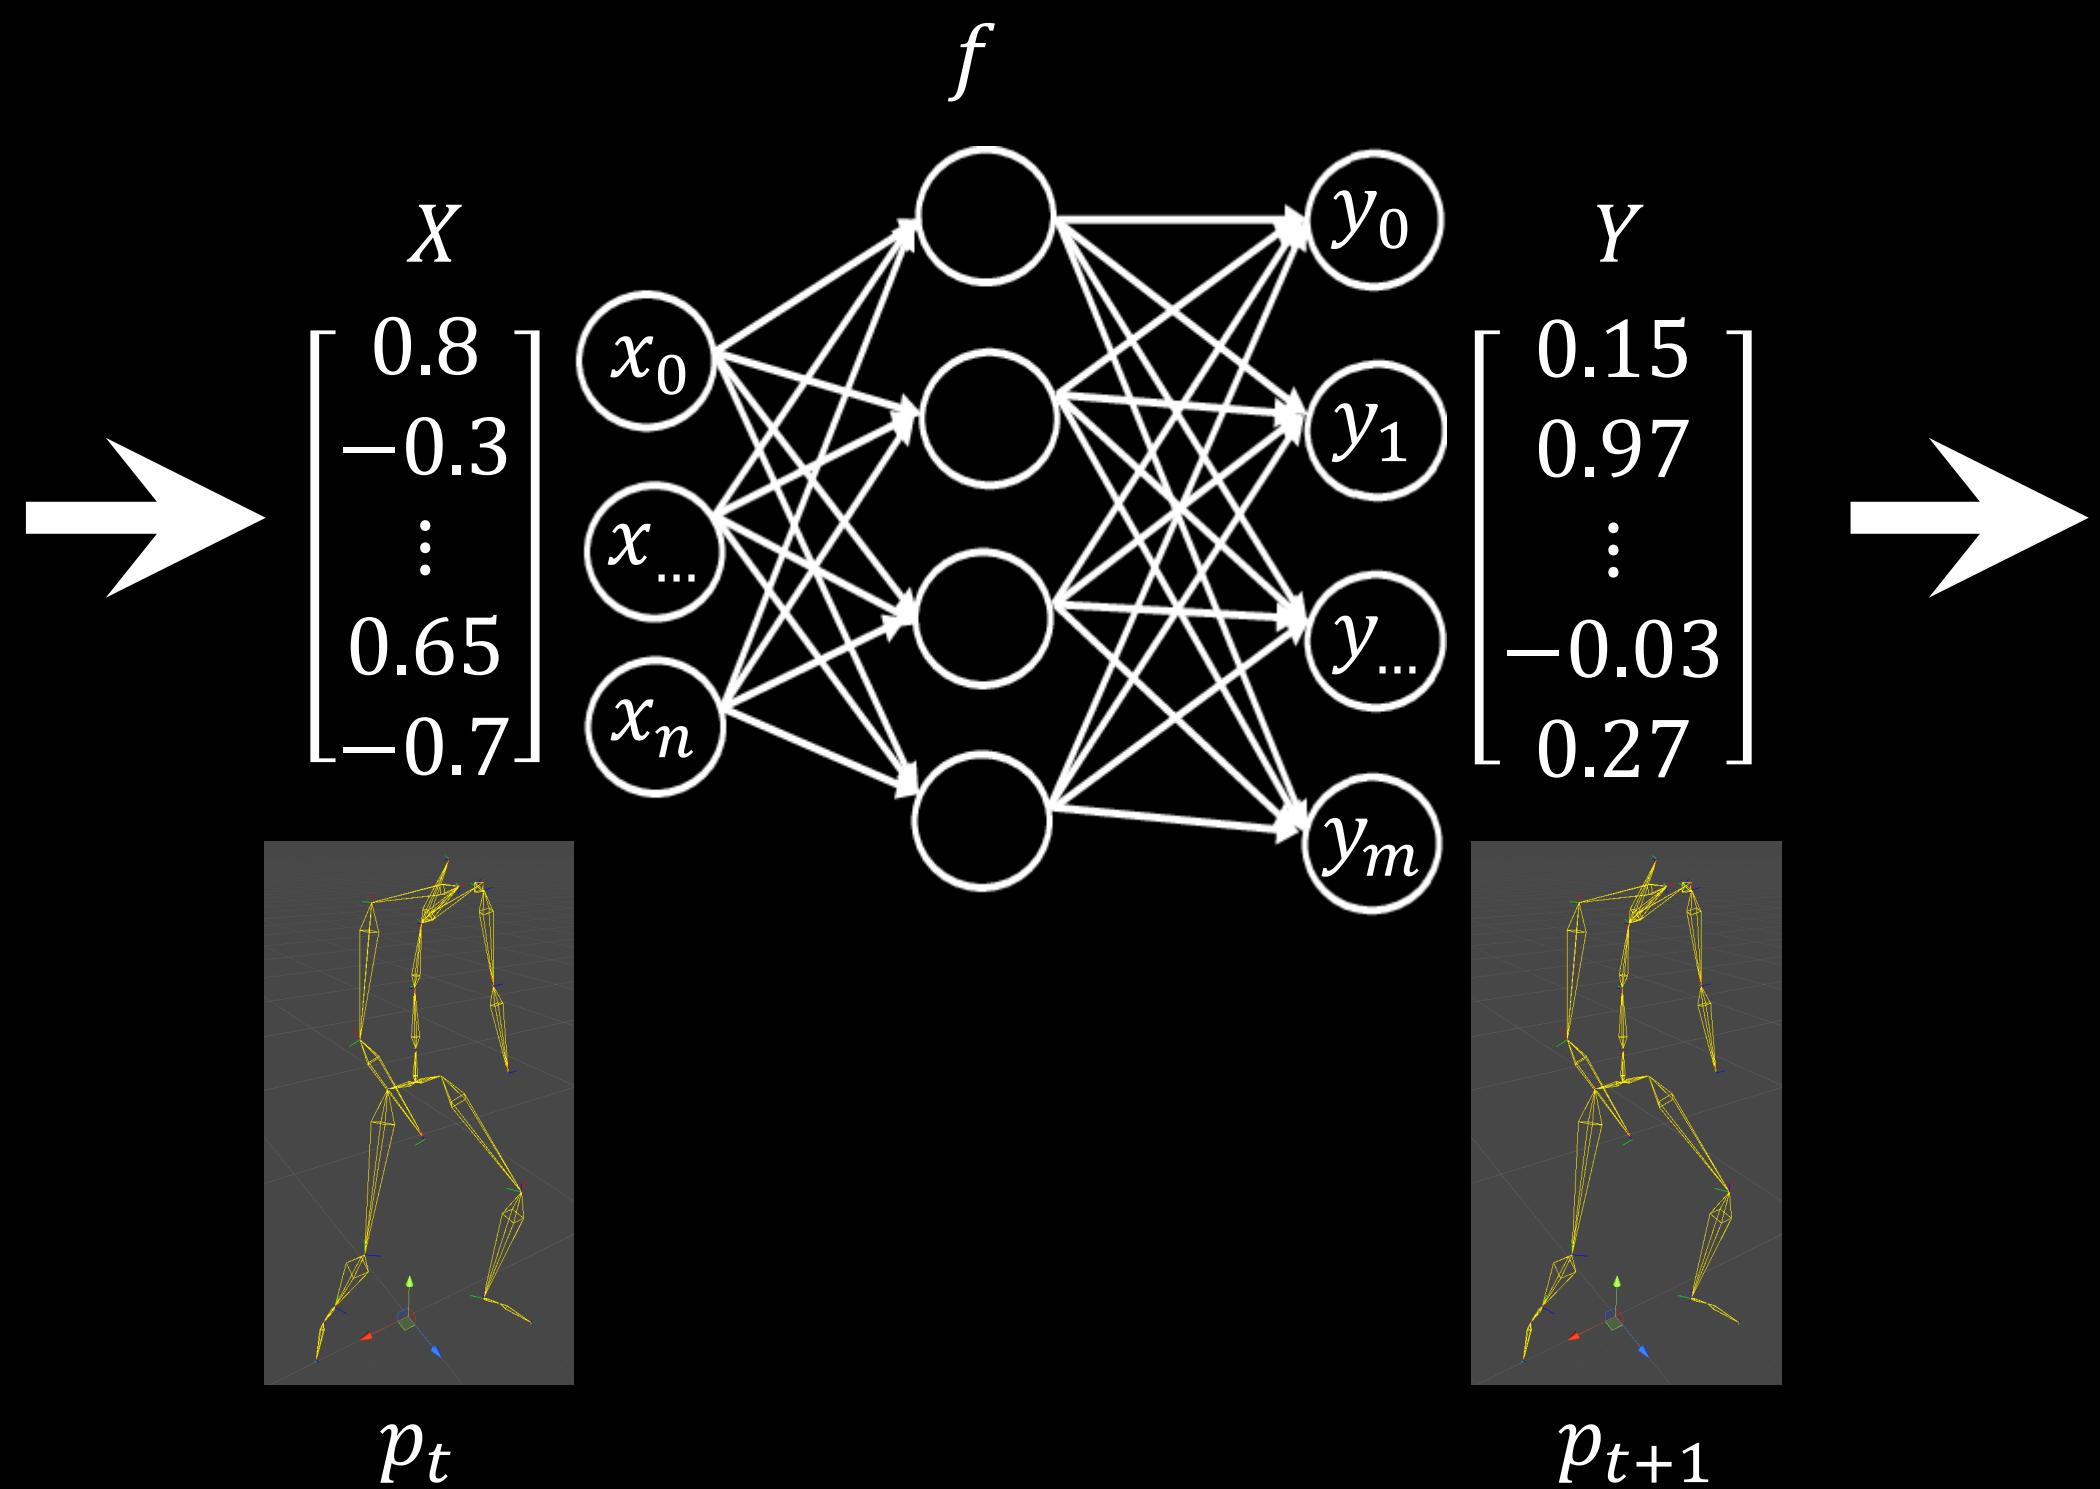
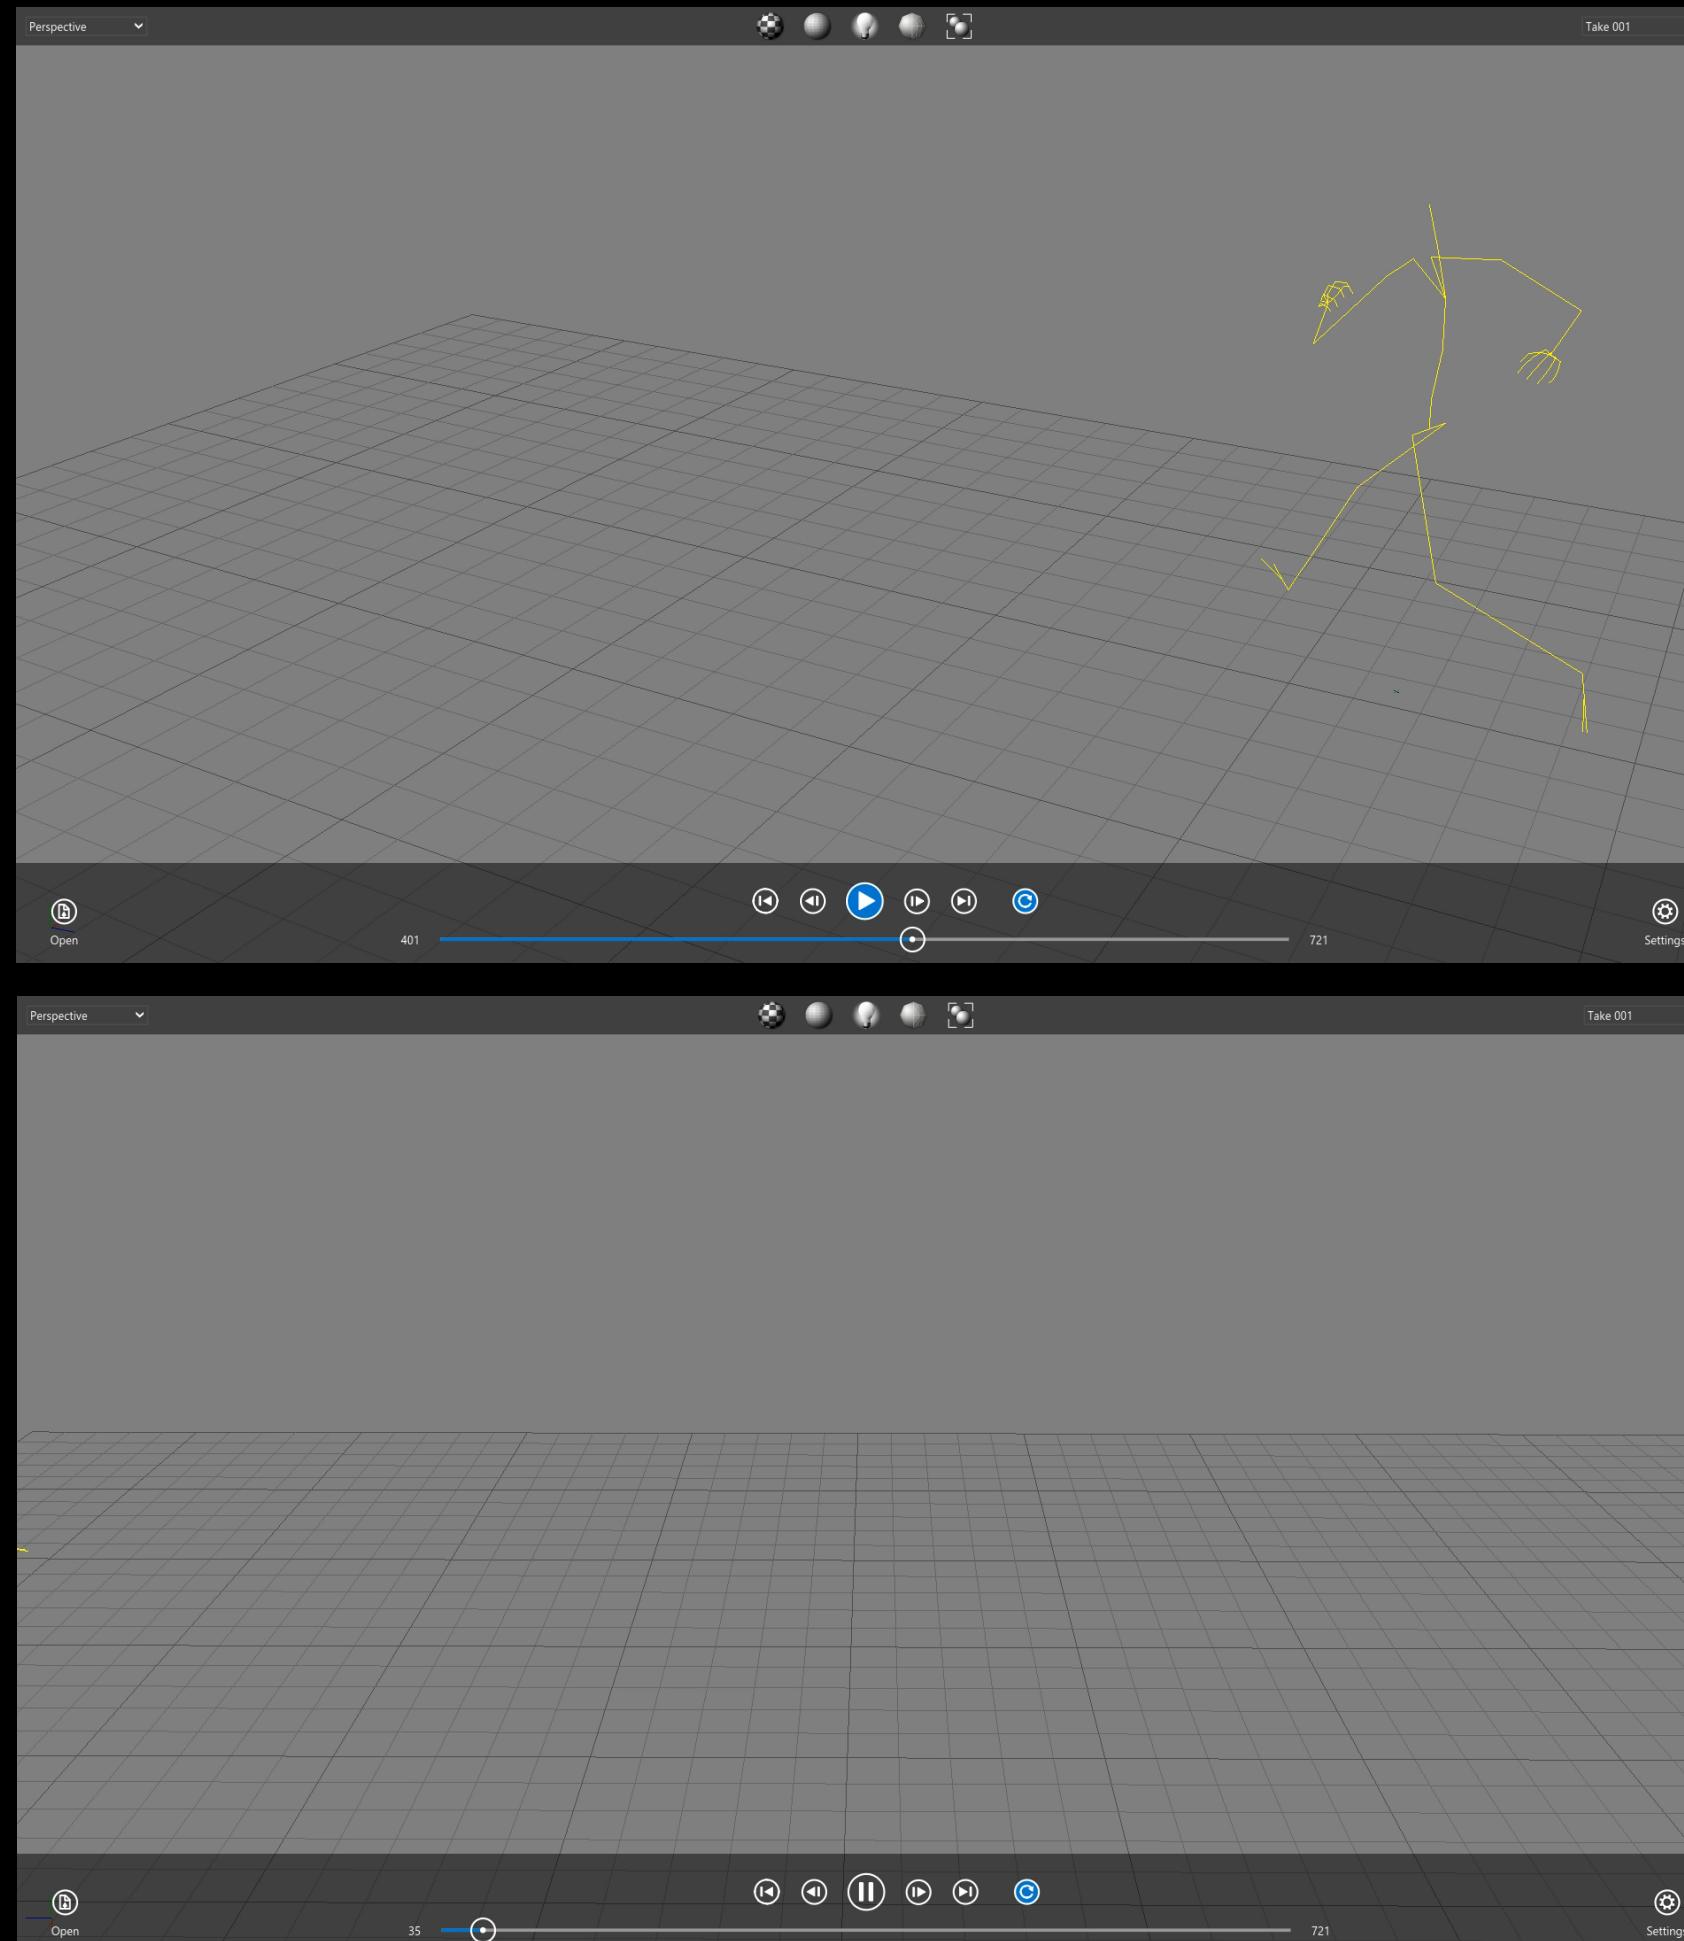


Ik OFF  
Prediction ON  
Target speed 5.50





# Autoregressive methods & Pose merging



# Autoregressive methods

Autoregressive network training averages the possible continuation candidates -> Loss of quality

"...NN is compact, requiring only a few megabytes of memory, even when trained on gigabytes of motion capture data..."

"...requires keeping all the motion data..."

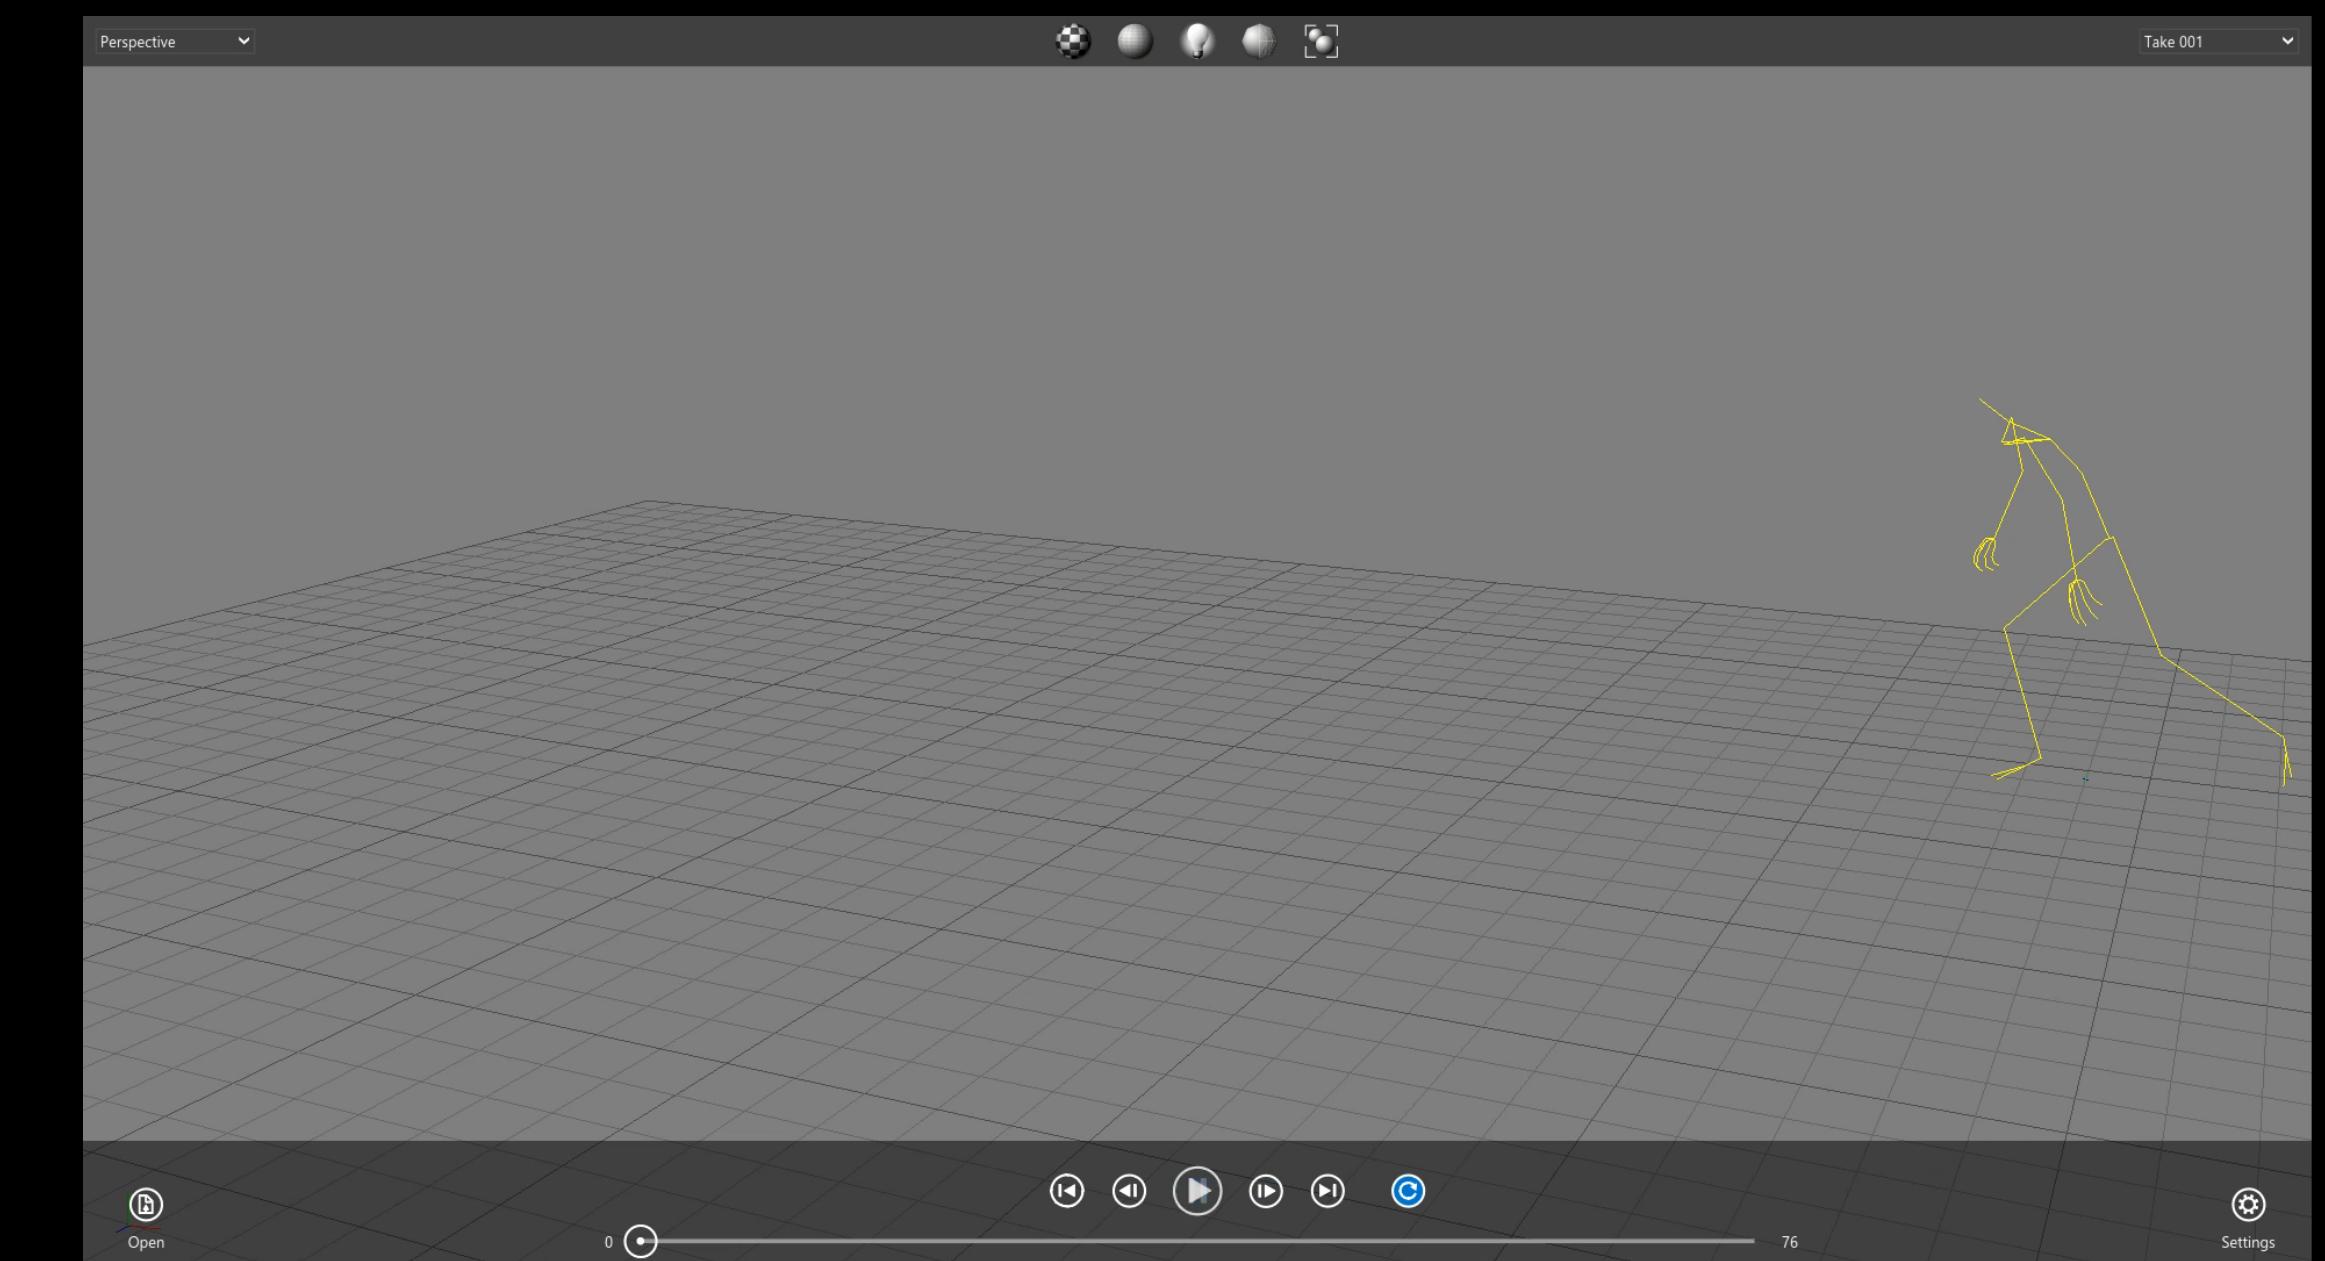
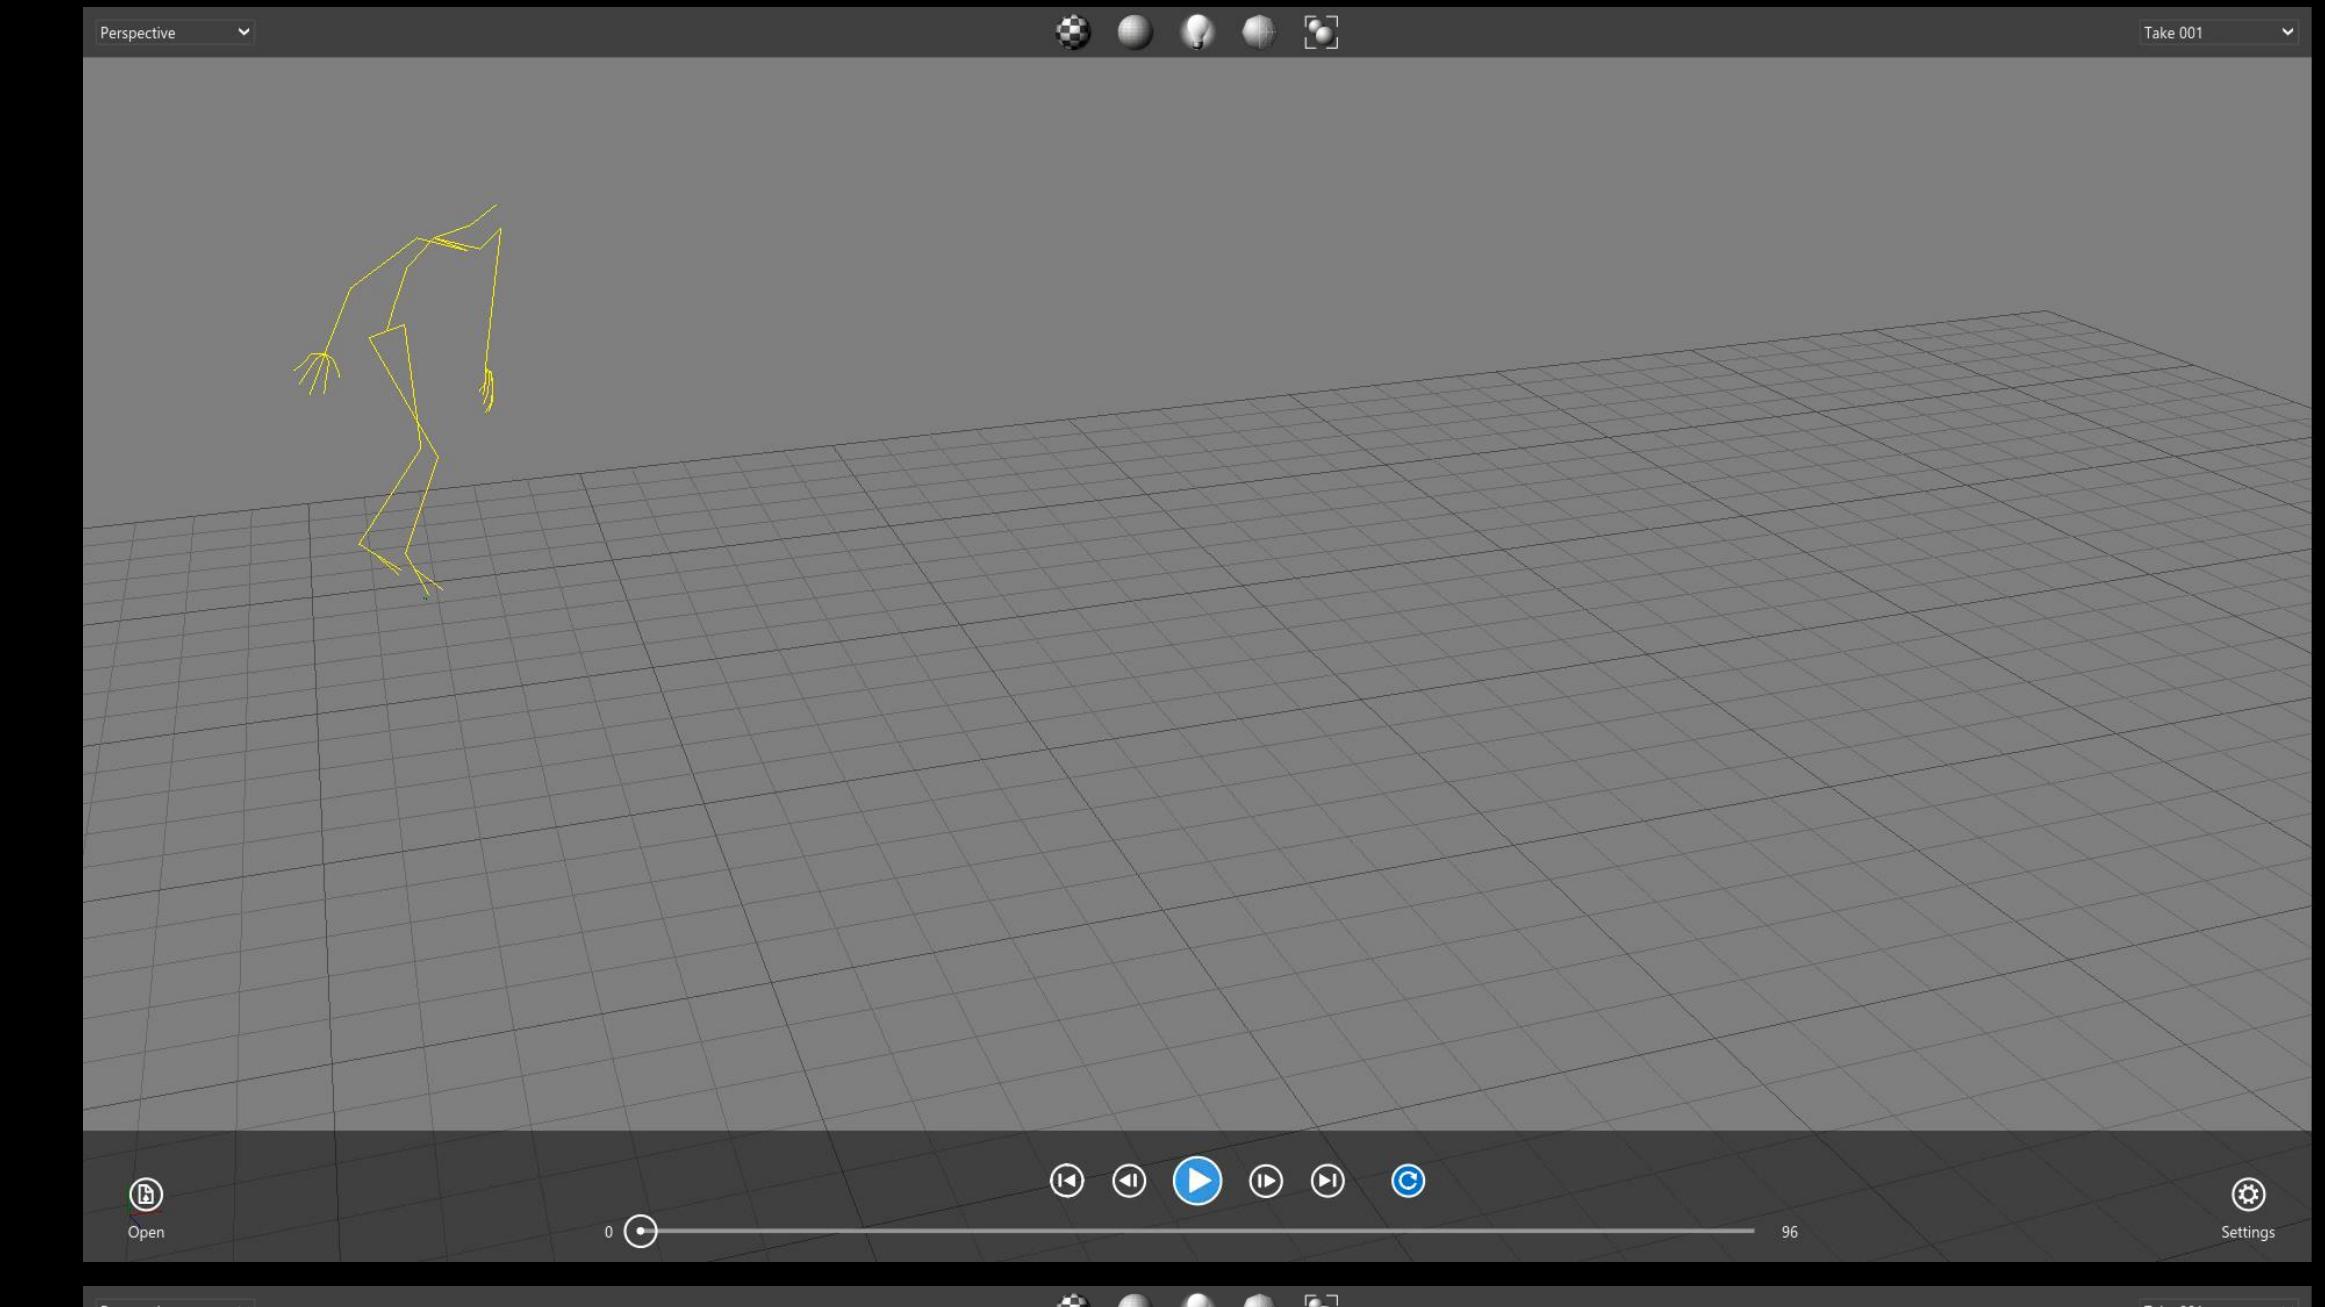
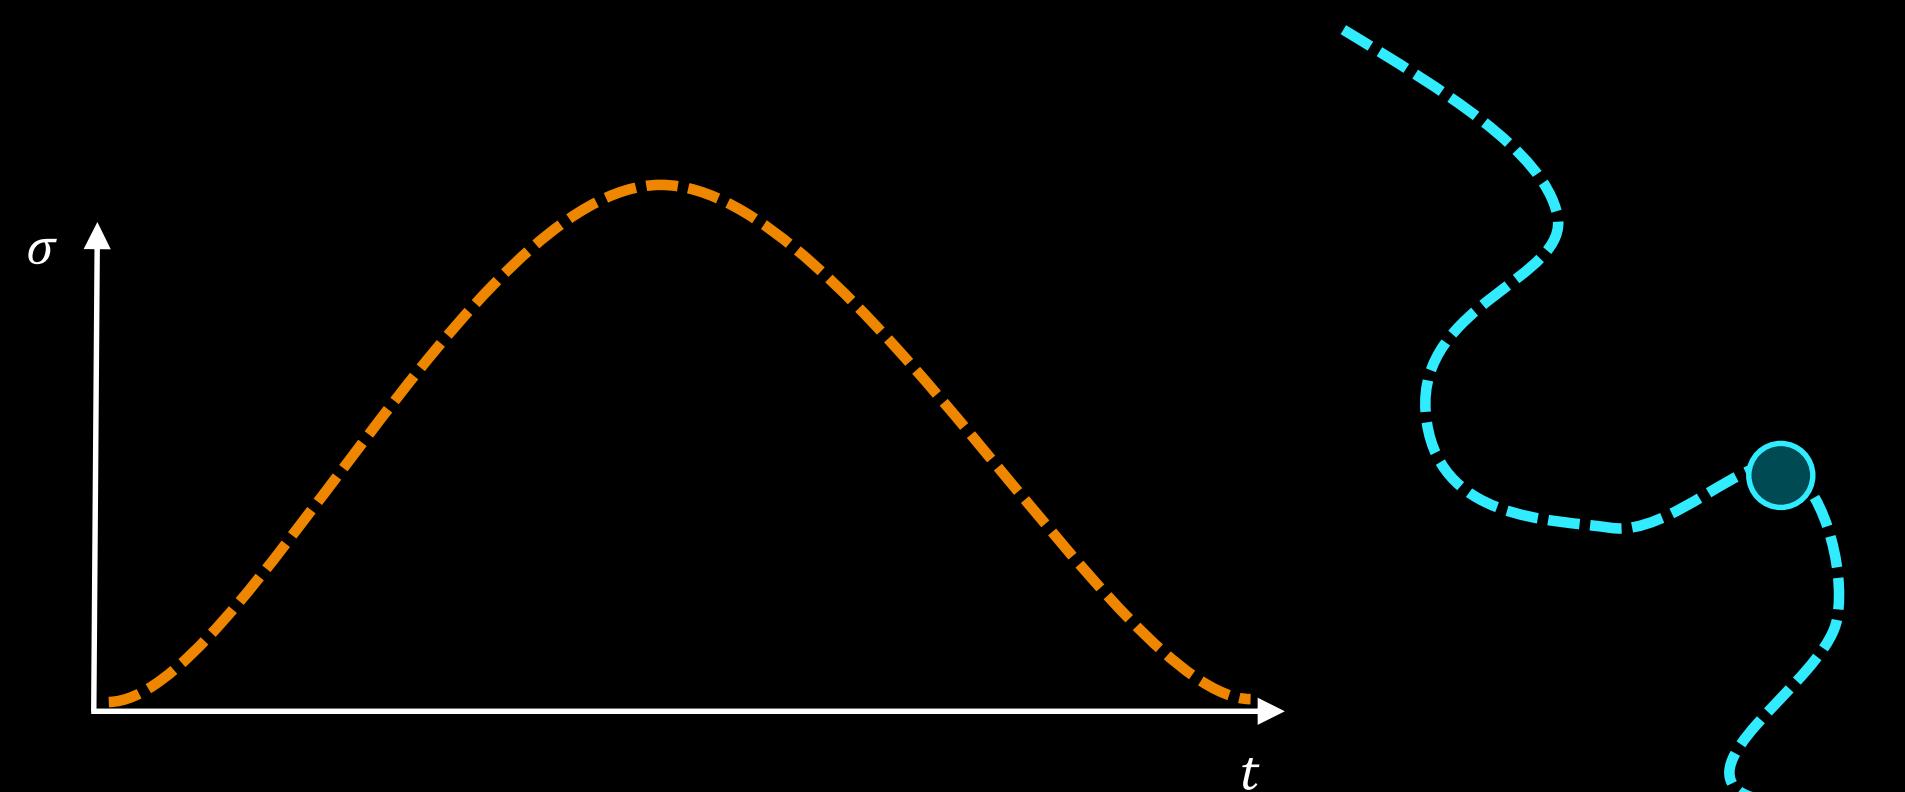


# Trajectory Control

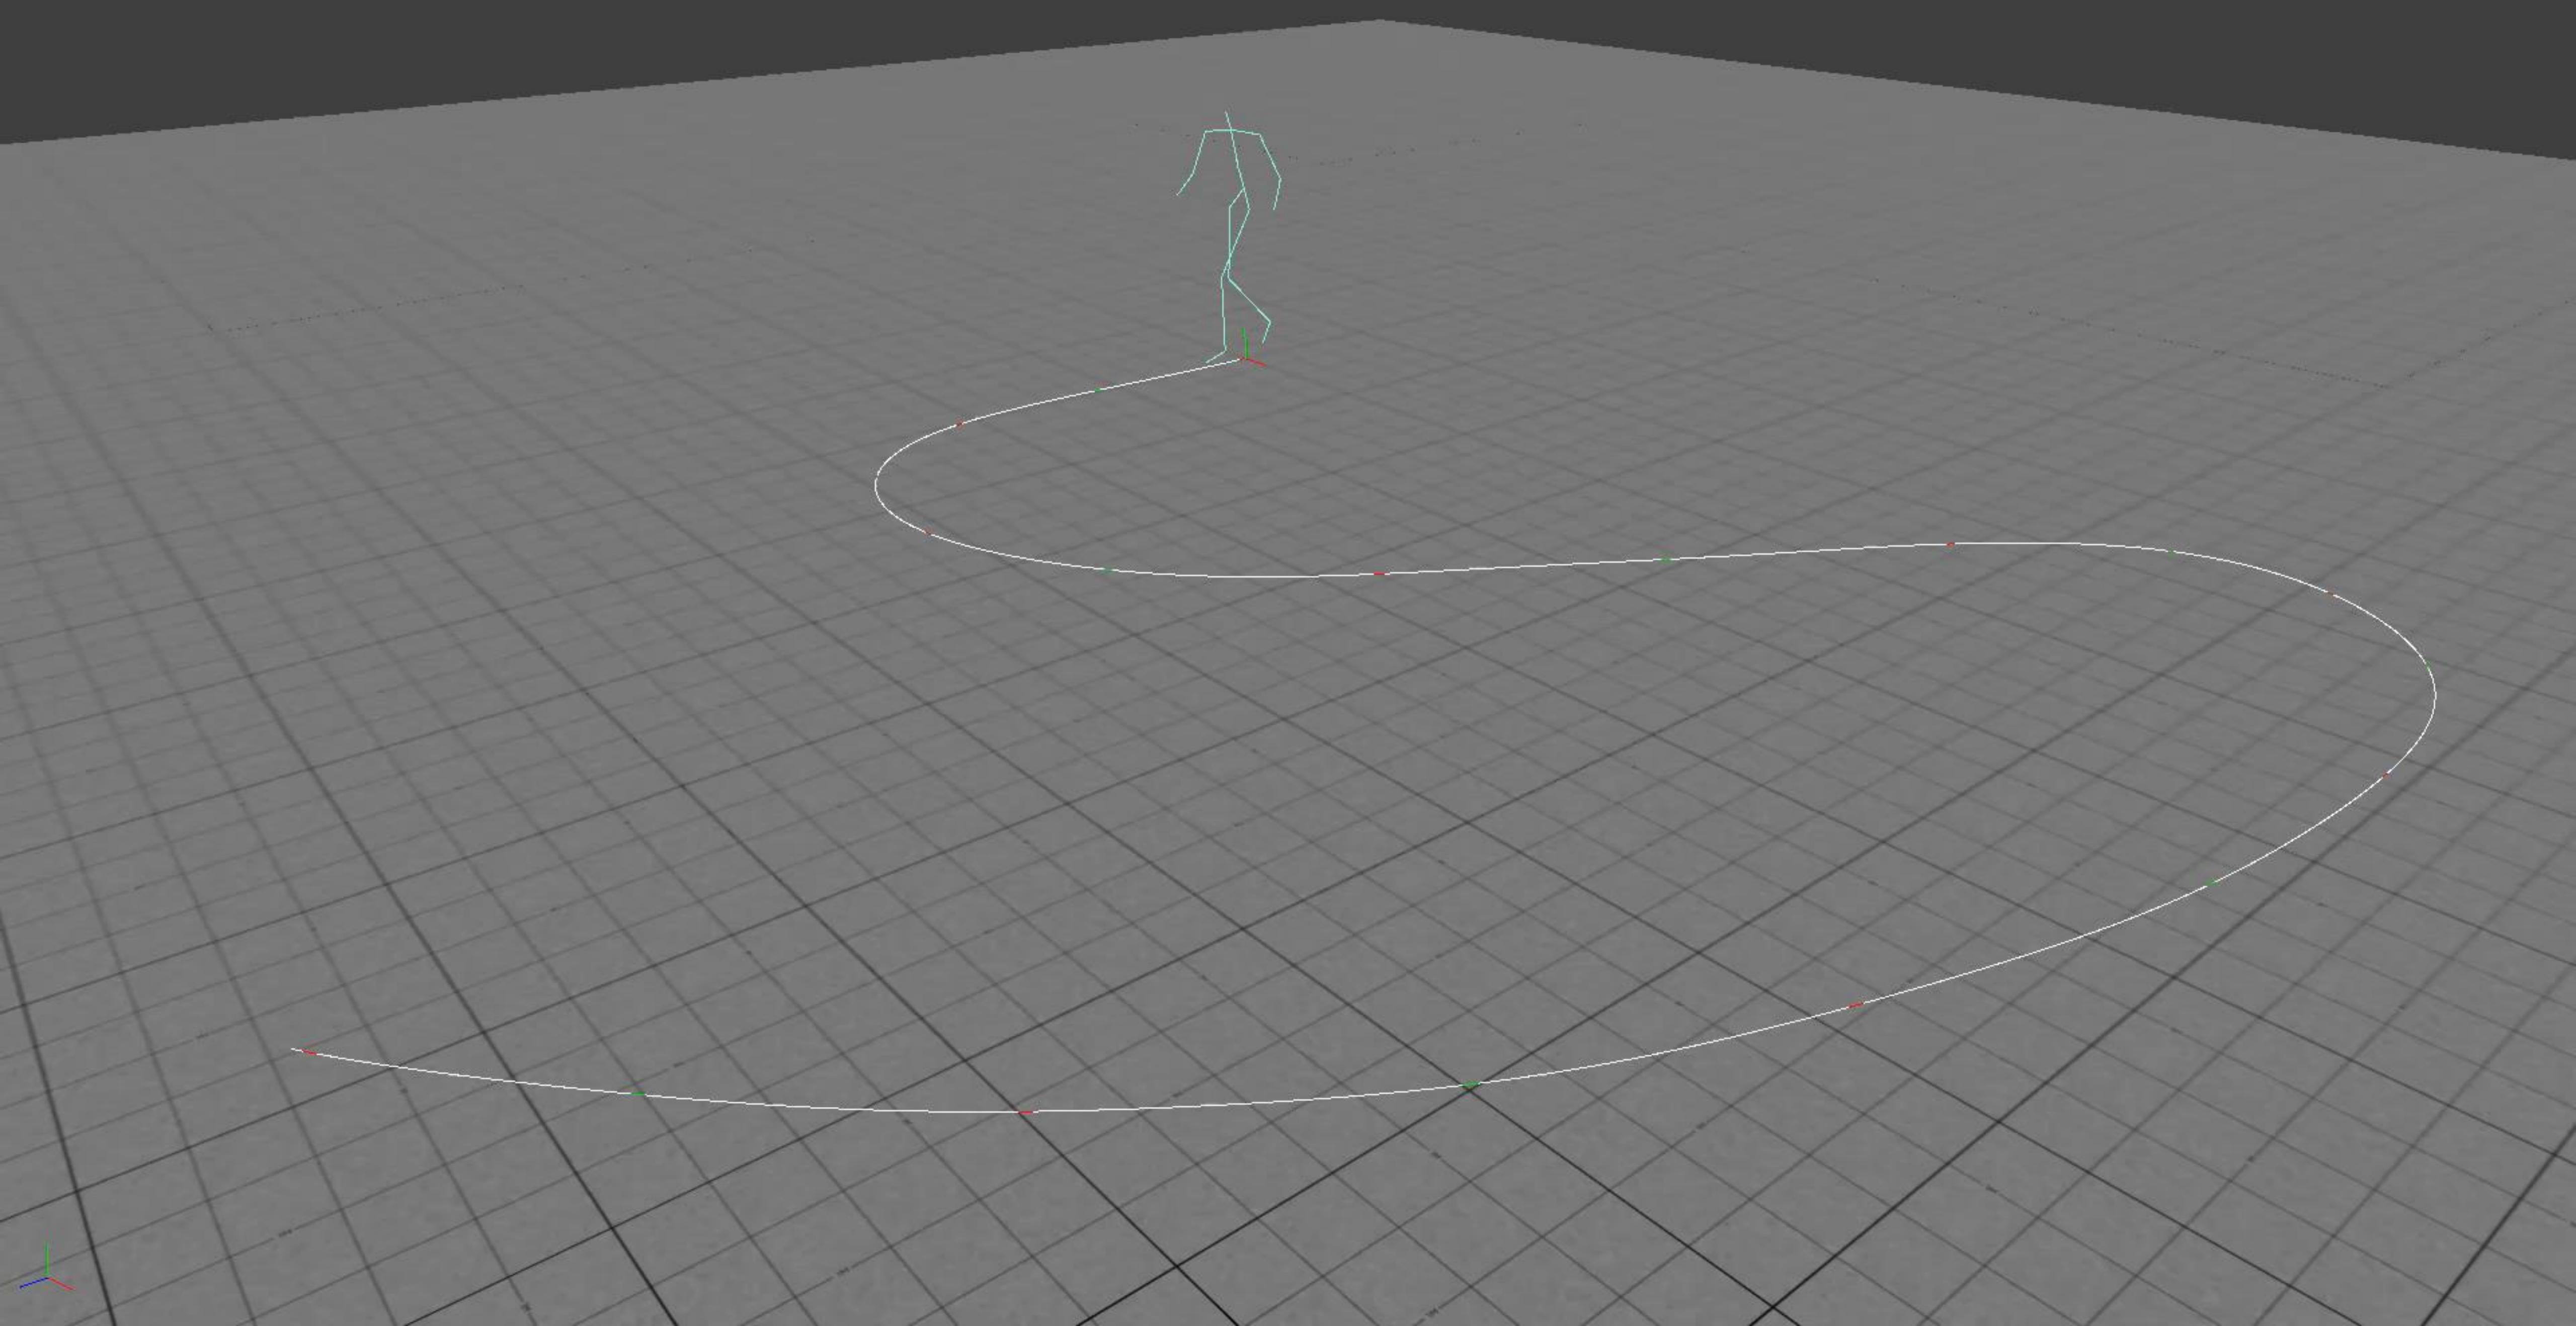
Trajectory annotations are used to guide the pose generation process

Expert gates can merge different movements  
(locomotion & jumping) -> Loss of quality

Trajectory dictates timing, but instead animation needs to dictate overall timing



Current time: 0.00  
Blend weight: 0.00



# Phase Function NN | Mode adaptive NN

Both approaches only work for locomotion

Climbing for example has no obvious phase

Basic assumption – any pose that corresponds to the same “phase” is similar

Poor quality

“Floating”

Don't use exponential maps in NN's

Variations won't be preserved but get averaged

Memory footprint is determined by number of weights – not the amount of animations used

Neural Networks do not memorize anything!

Slow runtime performance

my SSE implementation was ~0.9ms

4 samples of  $\Theta$  yields average result, more samples require a higher memory footprint

Phase mispredictions result in a tendency towards the mean pose

Extraordinary long training times (6+ hours)

Edit data, retrain, hope it appears

We can't predict the output  
...or ask why it was produced

Mode adaptive NN's don't work for bipeds



24  
20

10 Antidote

6080  
70



44

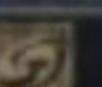


Lightning Blade



Estus Flask+10

15



0







MK|IceAndFire

# The holy grail

Fast turn-around times

Scaleable

Ground-truth motion synthesis

Controllable

Minimal memory footprint

Versatile

Fast runtime

Precise

Style





Kinematica

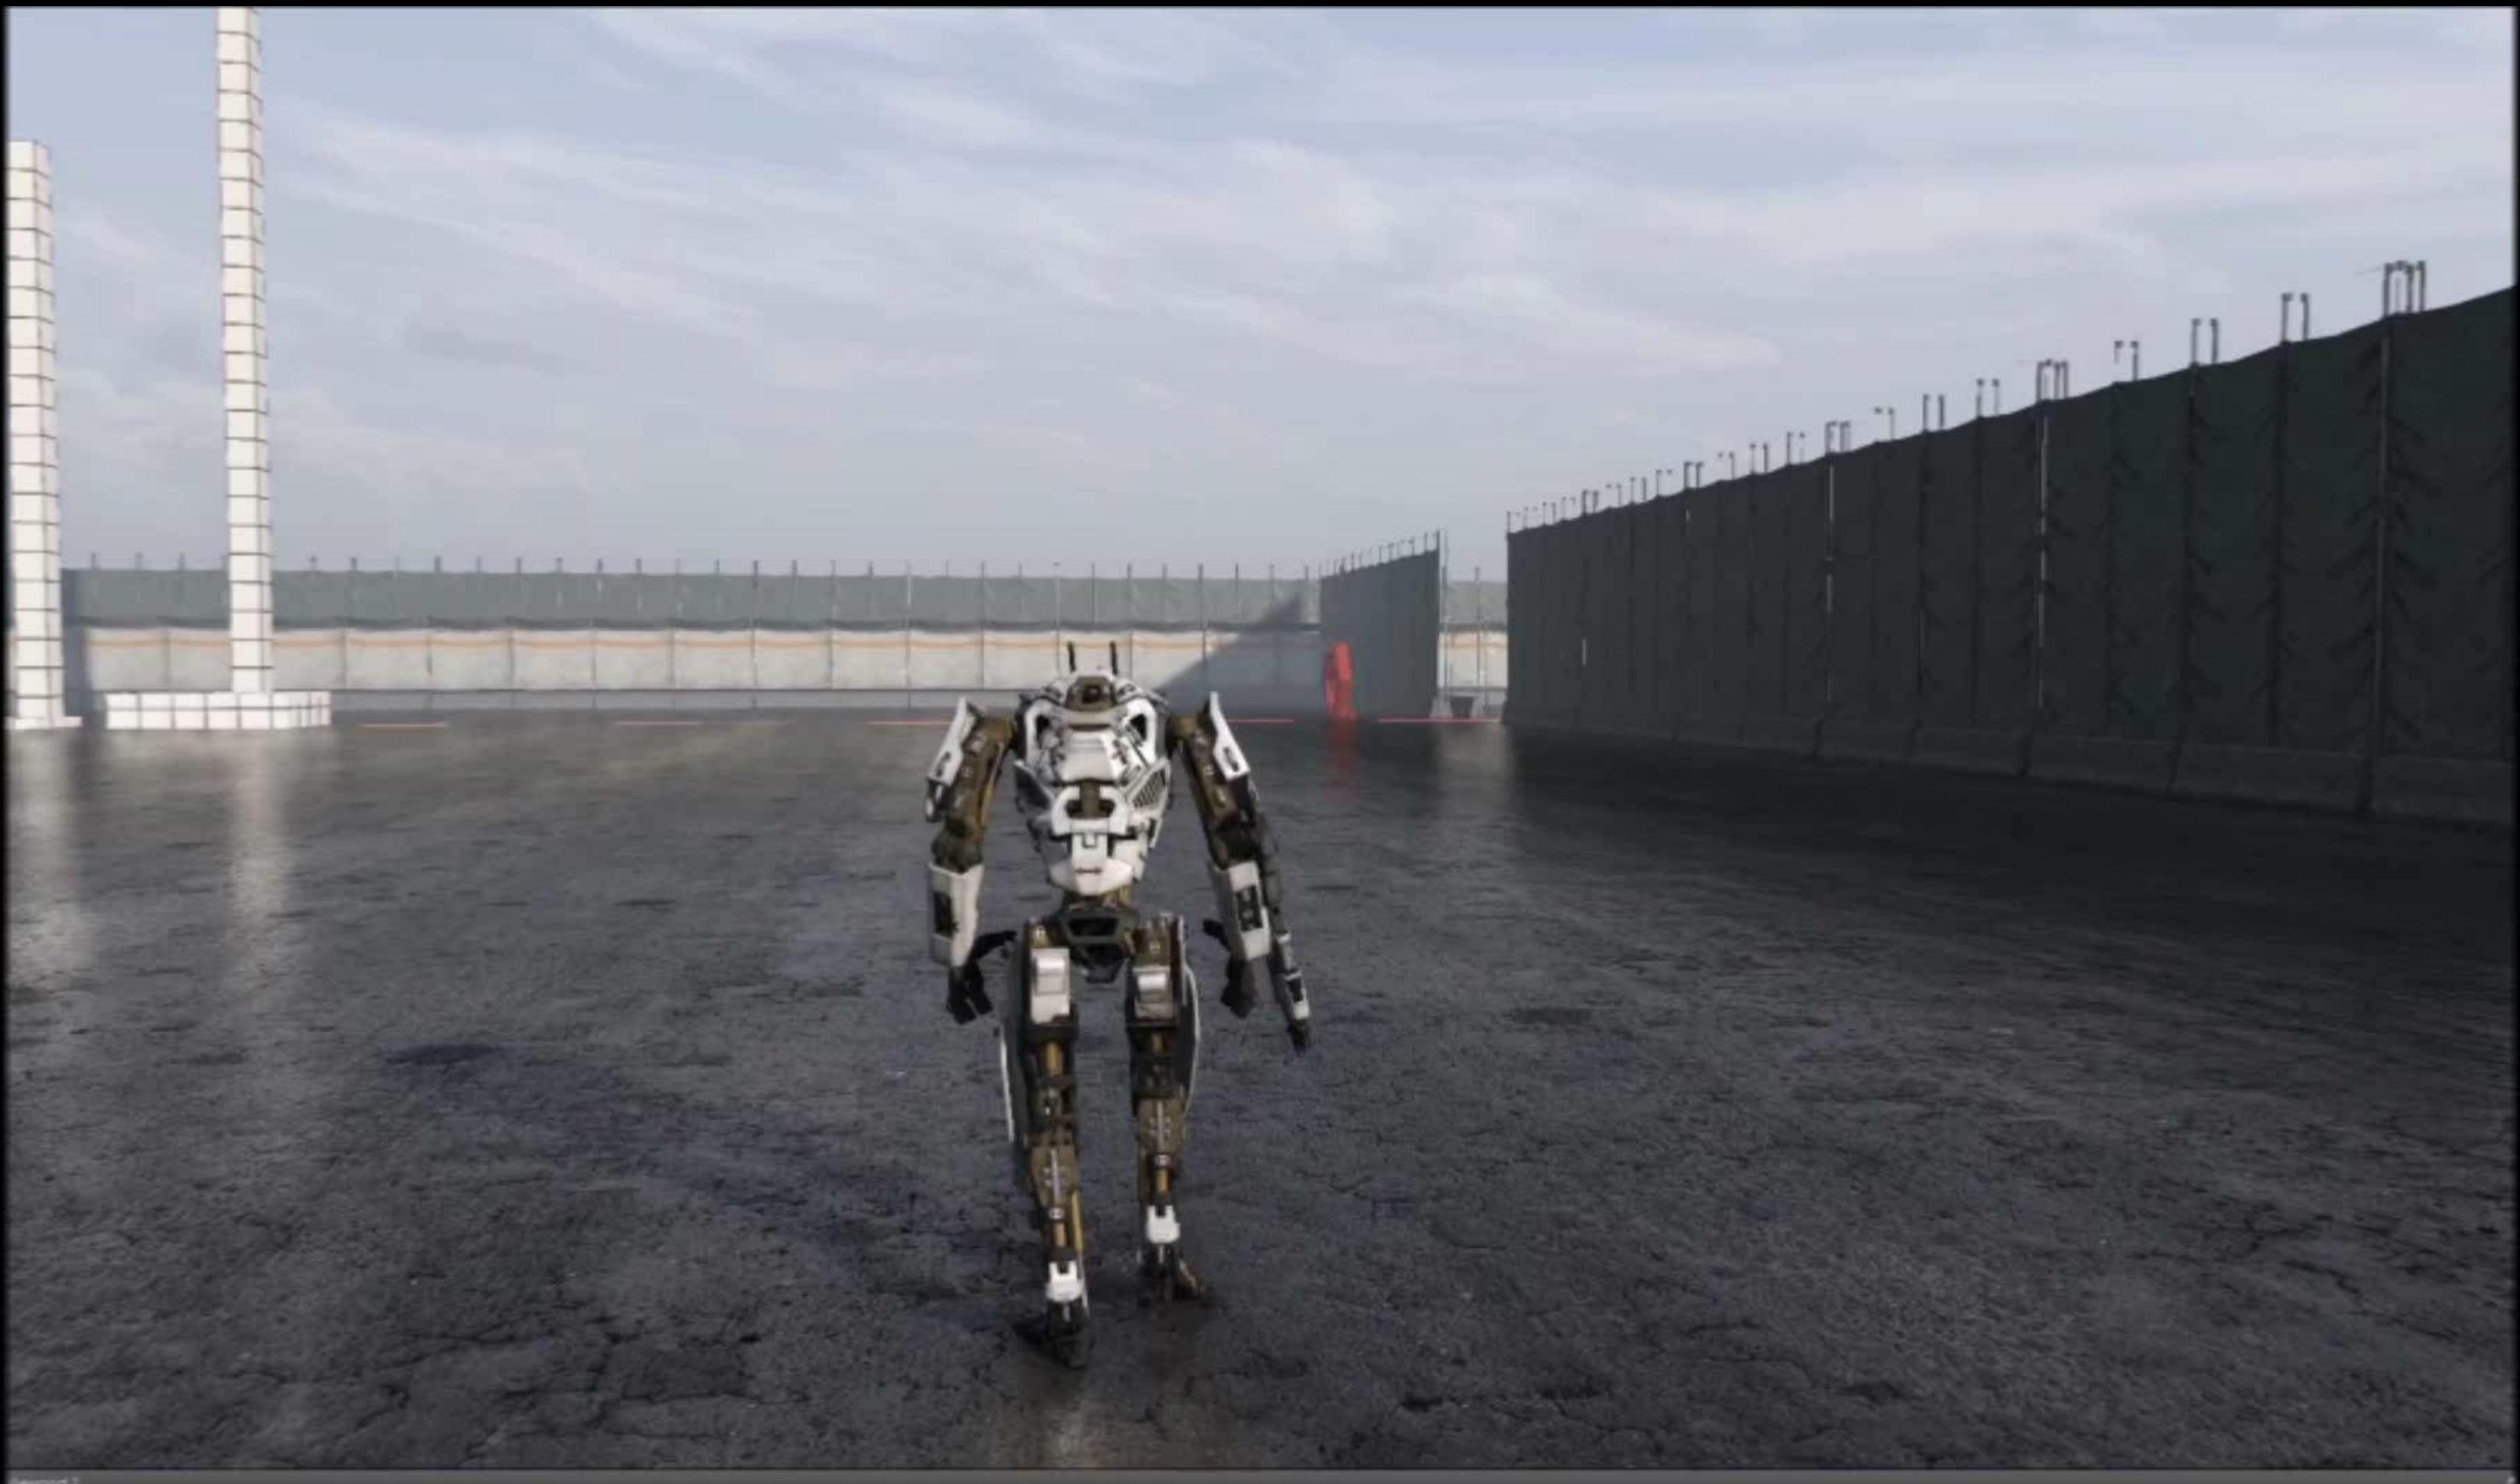
# Motion Matching

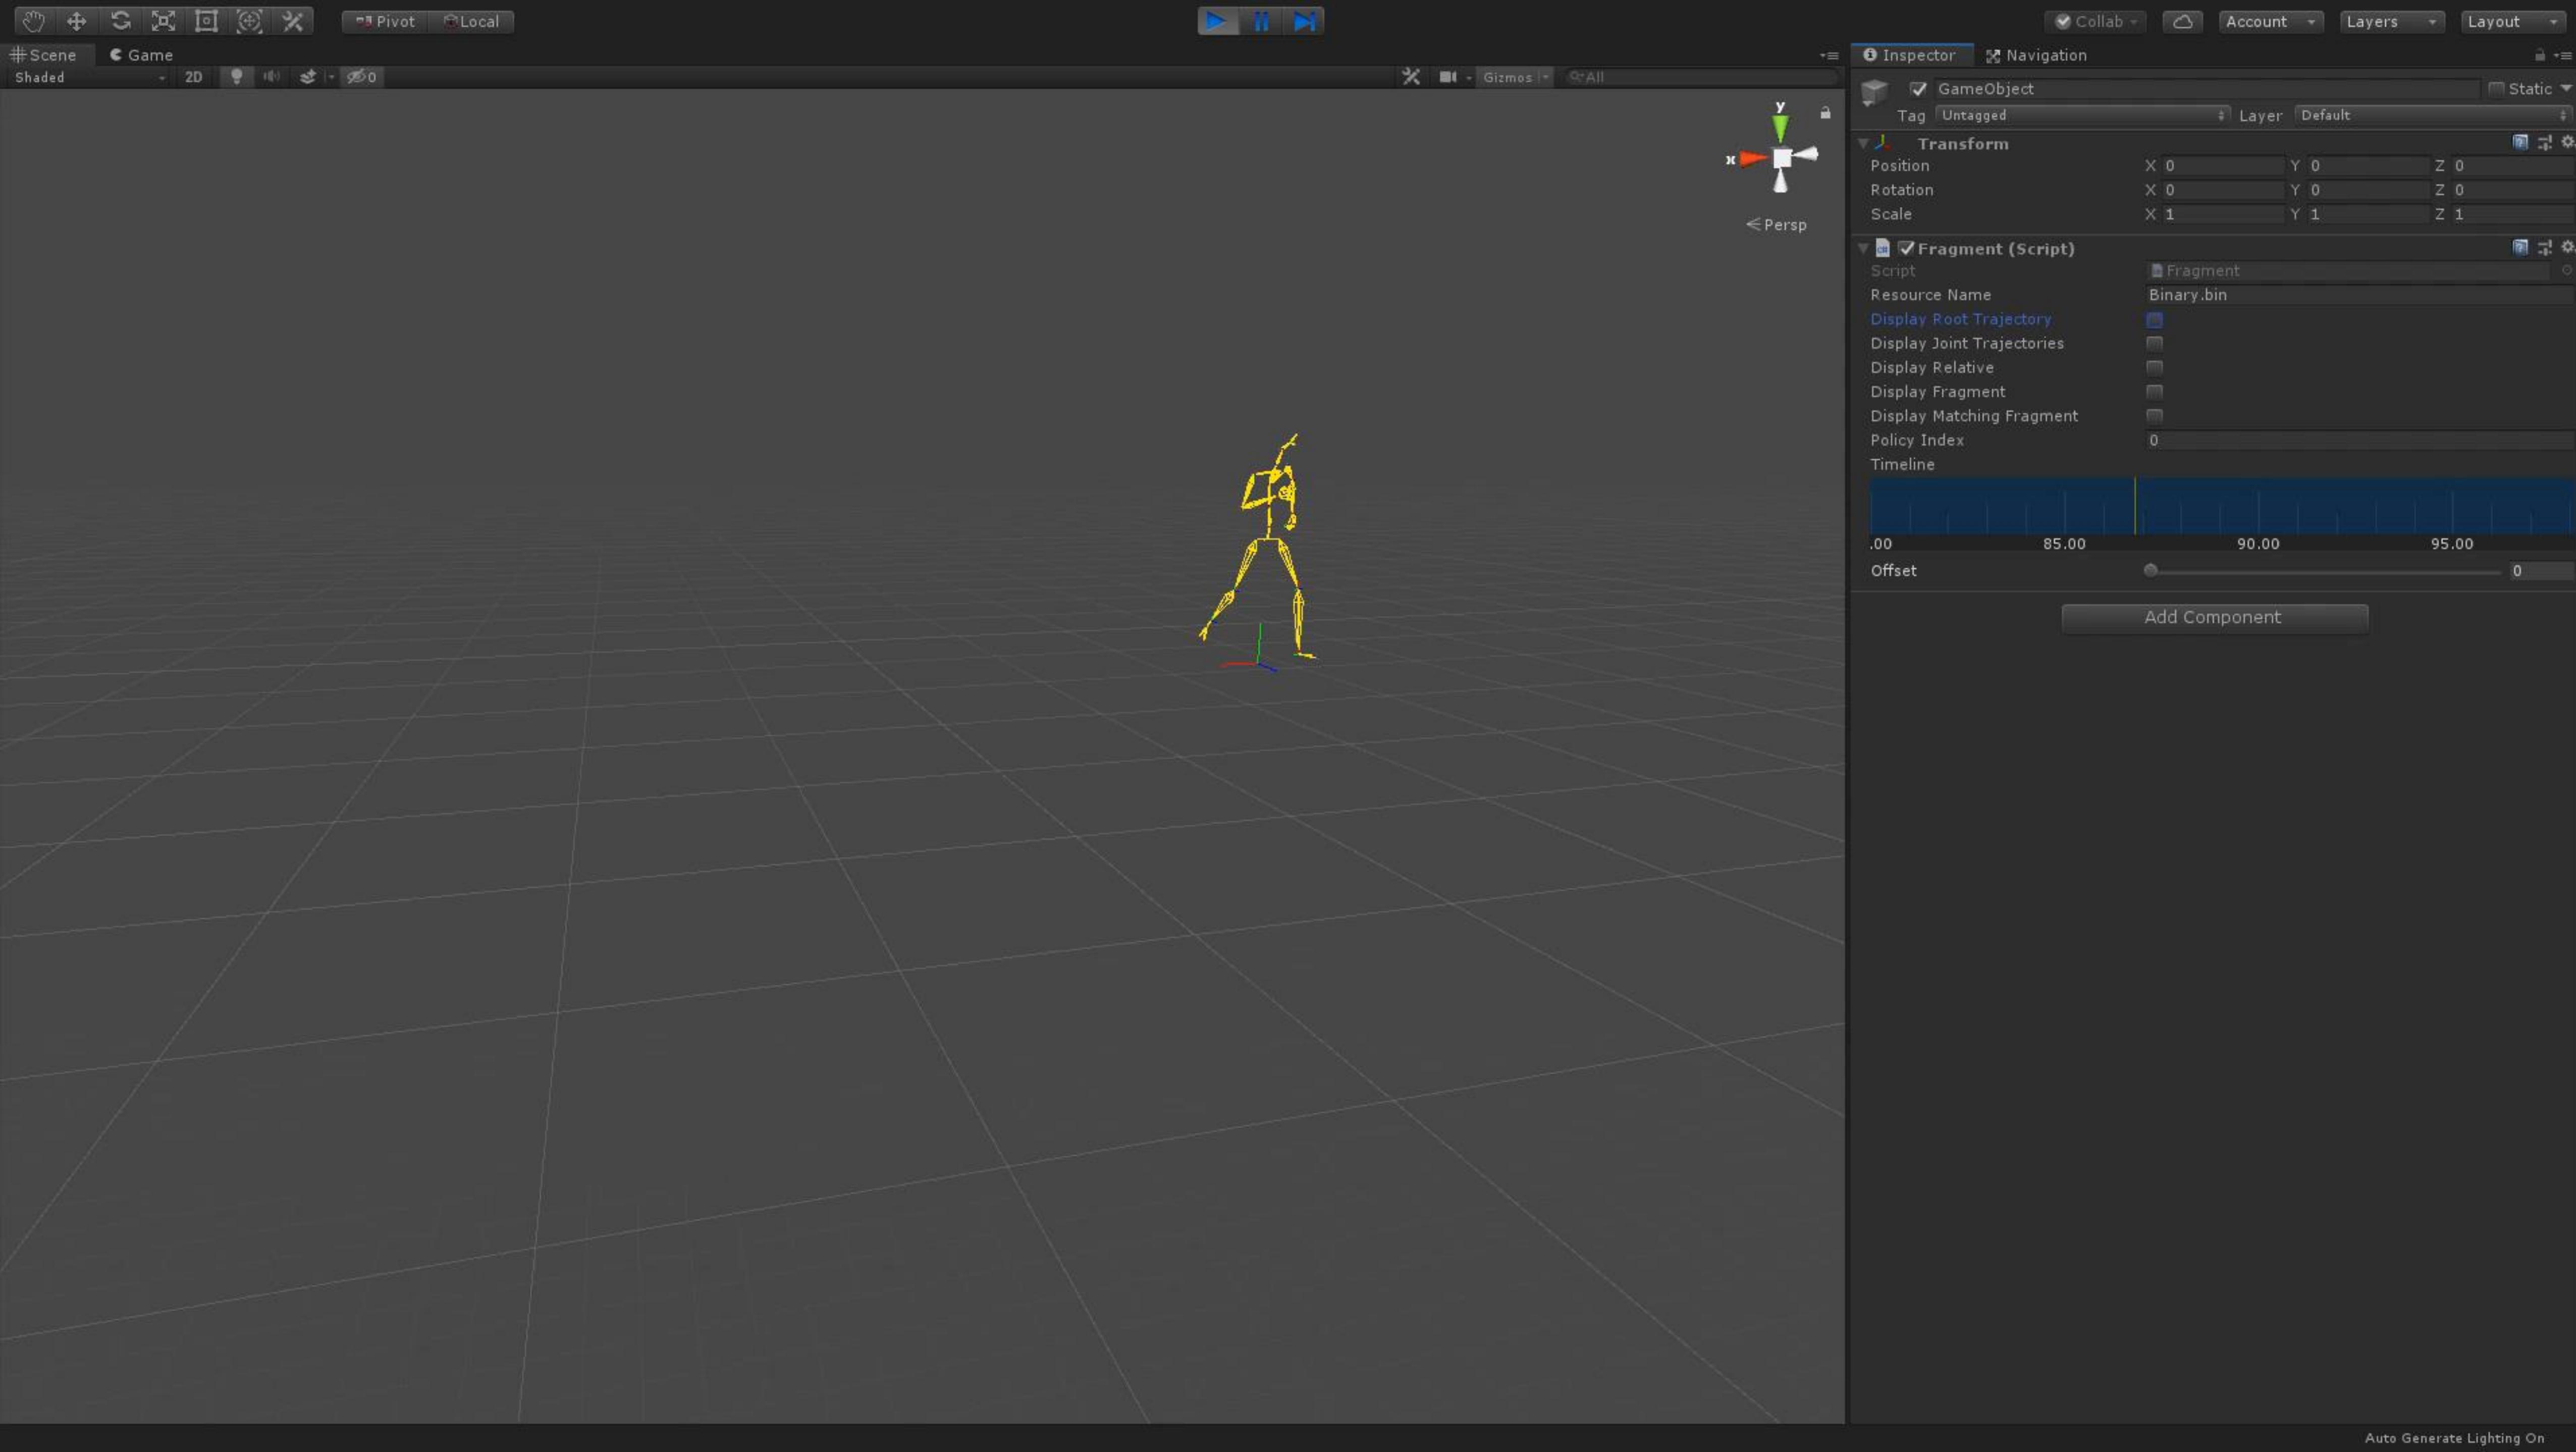
## Pros

- Preserves high quality result
- Does not rely on phases
- Relatively easy to implement

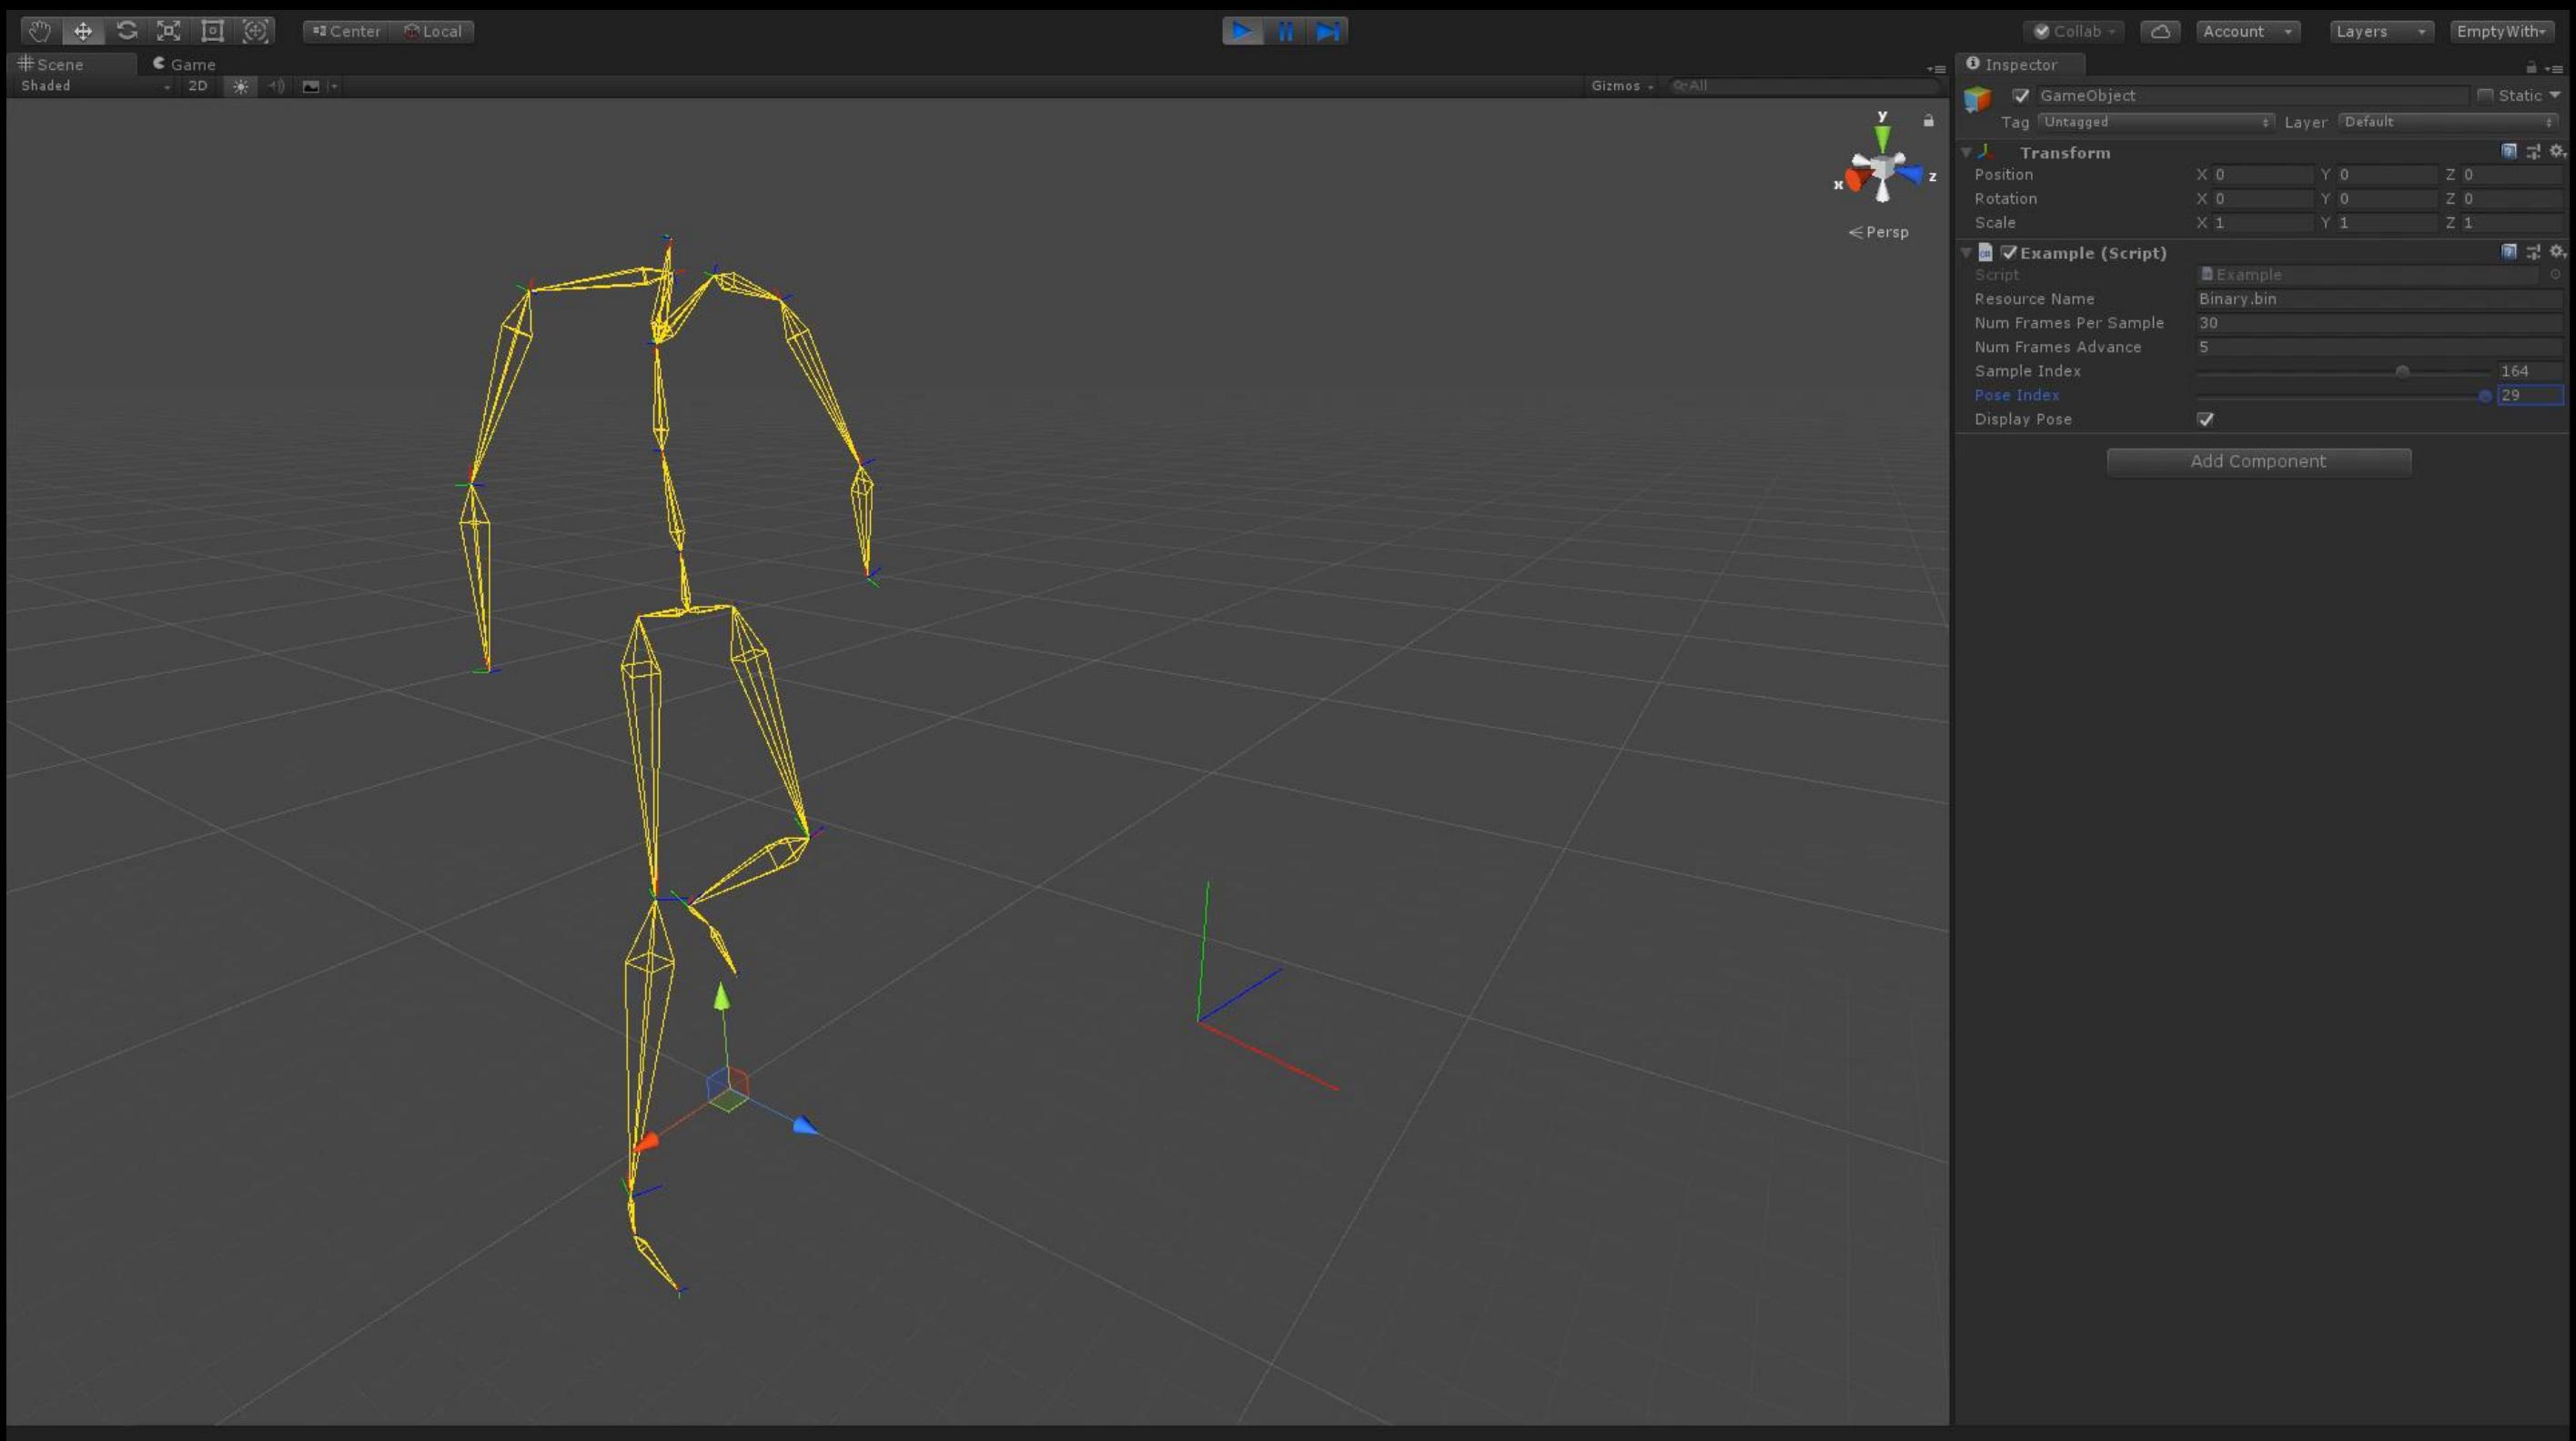
## Cons

- Prediction must match data
- Construction of cost functions
- Requires a lot of tweaking
- Doesn't scale well
- Duplicate data problem
- "Back-in-time" problem





# Motion Fragments



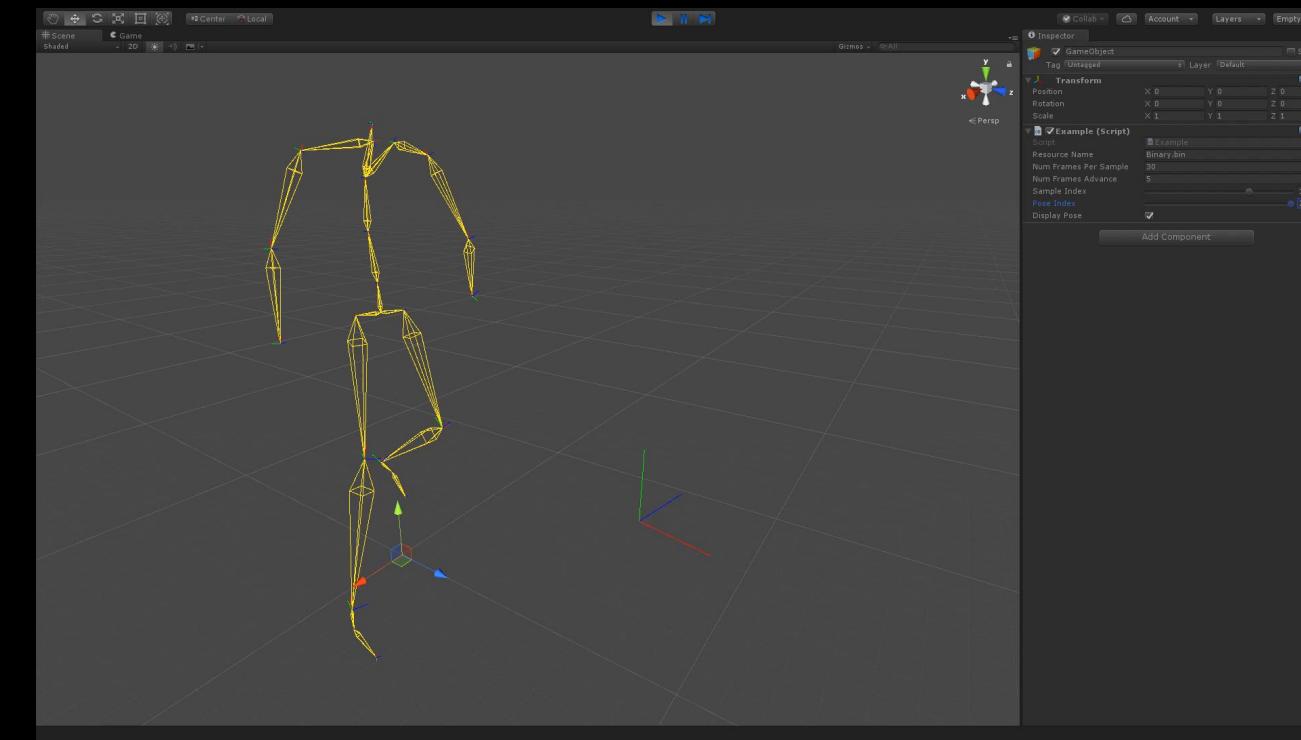
A motion fragment is represented as a matrix  $F_i \in M^{\tau \times (m+1)}(\mathbb{R})$  where each entry  $v_j^t$  contains the velocity of joint  $j$  at some time  $t \in \{i - \tau, \dots, i + \tau\}$ , and  $j = 0$  represents the root transform

$$F_i = \begin{bmatrix} v_0^i & v_0^{i+1} & \dots & v_0^{i+\tau} \\ v_1^{i-\tau} & v_1^{i-\tau+1} & \dots & v_1^i \\ v_2^{i-\tau} & v_2^{i-\tau+1} & \dots & v_2^i \\ \vdots & \vdots & \vdots & \vdots \\ v_m^{i-\tau} & v_m^{i-\tau+1} & \dots & v_m^i \end{bmatrix}$$

$$v_j^t = J_j^t - J_j^{t+1} (T_r^{t+1-1} T_r^t)$$

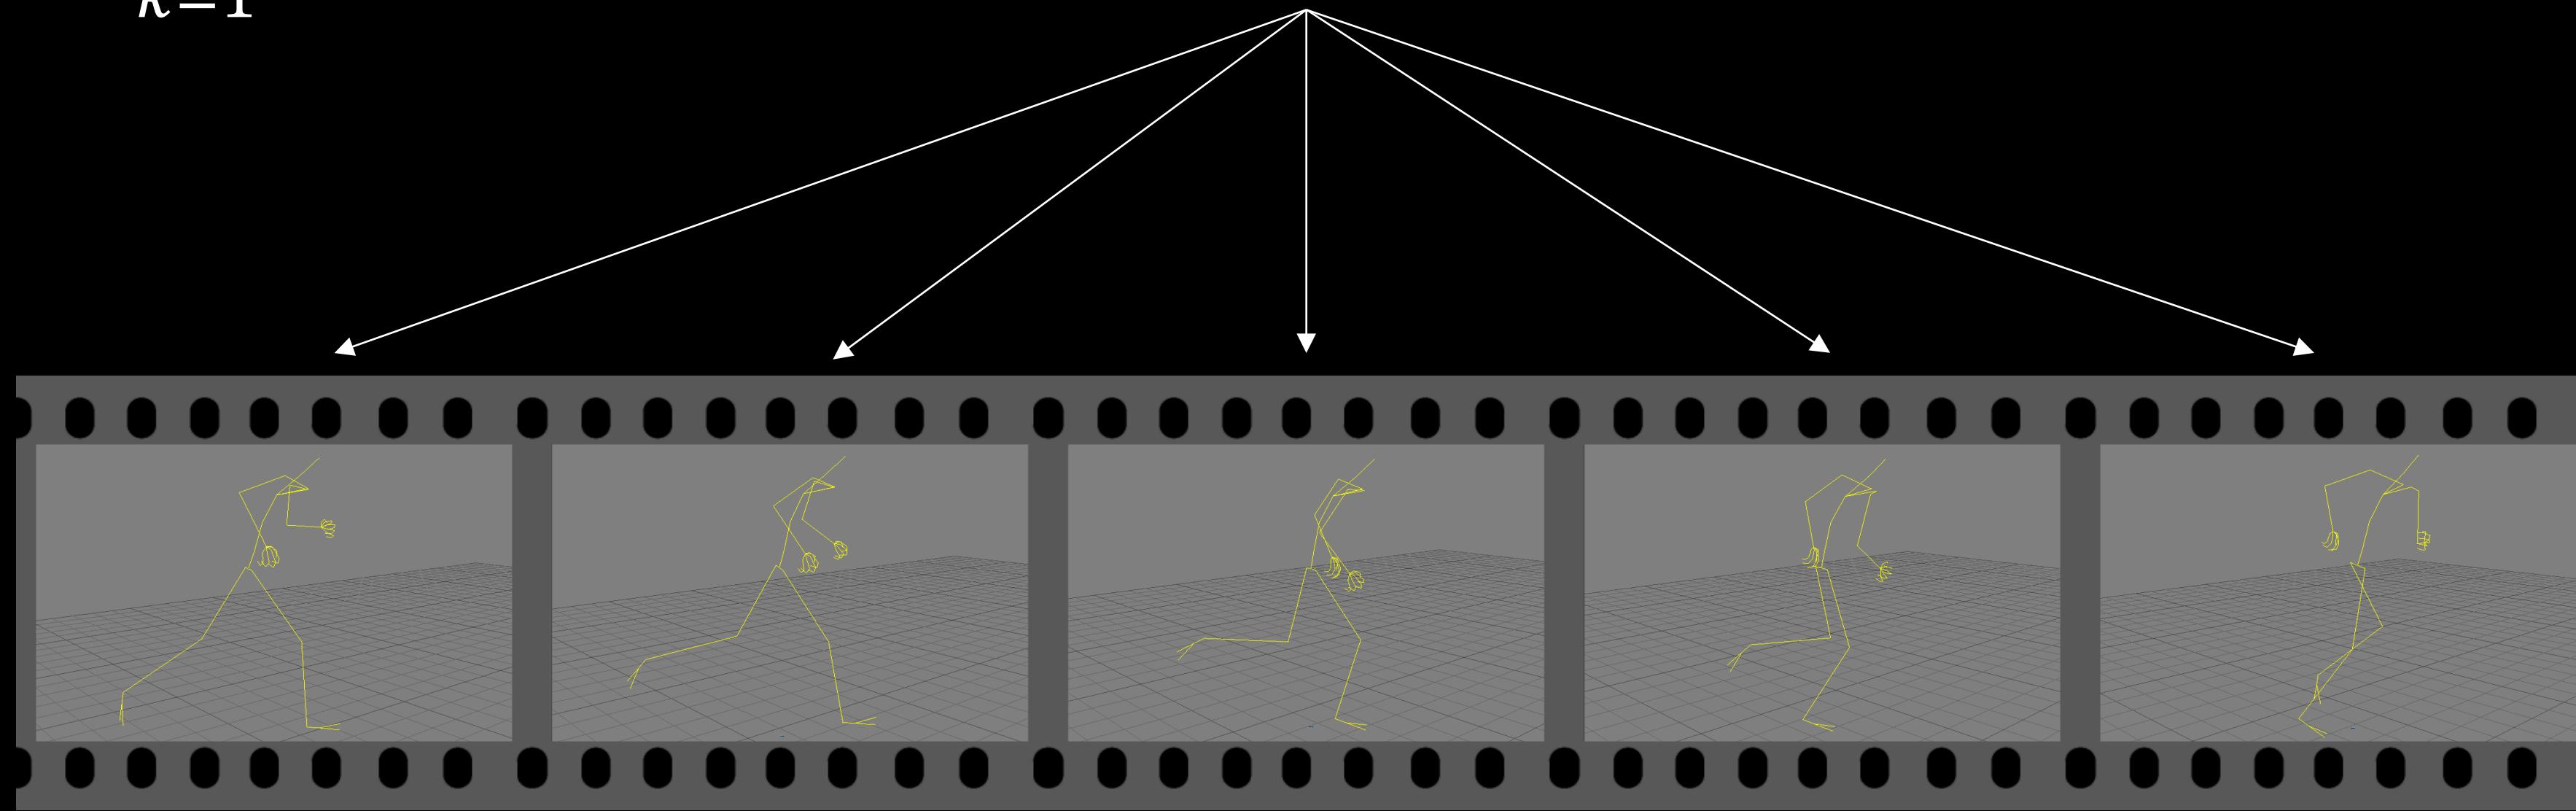
# Kinematic Algorithm

$$s(F, F') = \sum_{k=1}^n |F_k - F'|_k^2$$



Query  $\rightarrow F(y)$

$$\arg \min_x s(F(x), F(y))$$



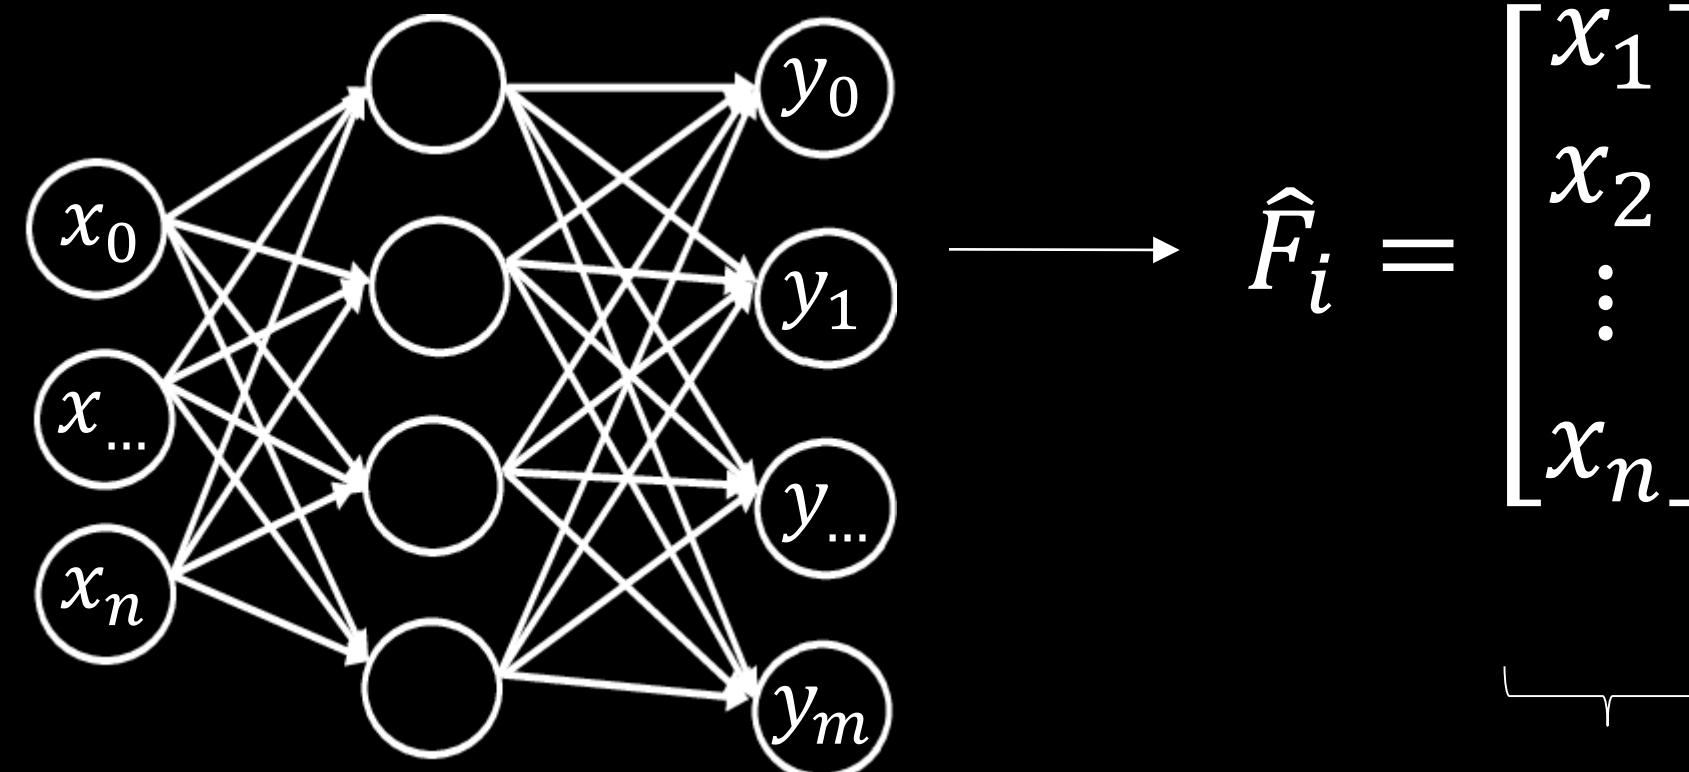
Motion Database

# Nearest Neighbor Search

$$F_i = \begin{bmatrix} v_0^i & v_0^{i+1} & \dots & v_0^{i+\tau} \\ v_1^{i-\tau} & v_1^{i-\tau+1} & \dots & v_1^i \\ v_2^{i-\tau} & v_2^{i-\tau+1} & \dots & v_2^i \\ \vdots & \vdots & \vdots & \vdots \\ v_m^{i-\tau} & v_m^{i-\tau+1} & \dots & v_m^i \end{bmatrix}$$

300+ Scalar Values  
1200 bytes per fragment  
For 70.000 poses > 82 Mb

Product Quantization for Nearest Neighbor Search  
[Jegou, Douze, Schmid, 2011]

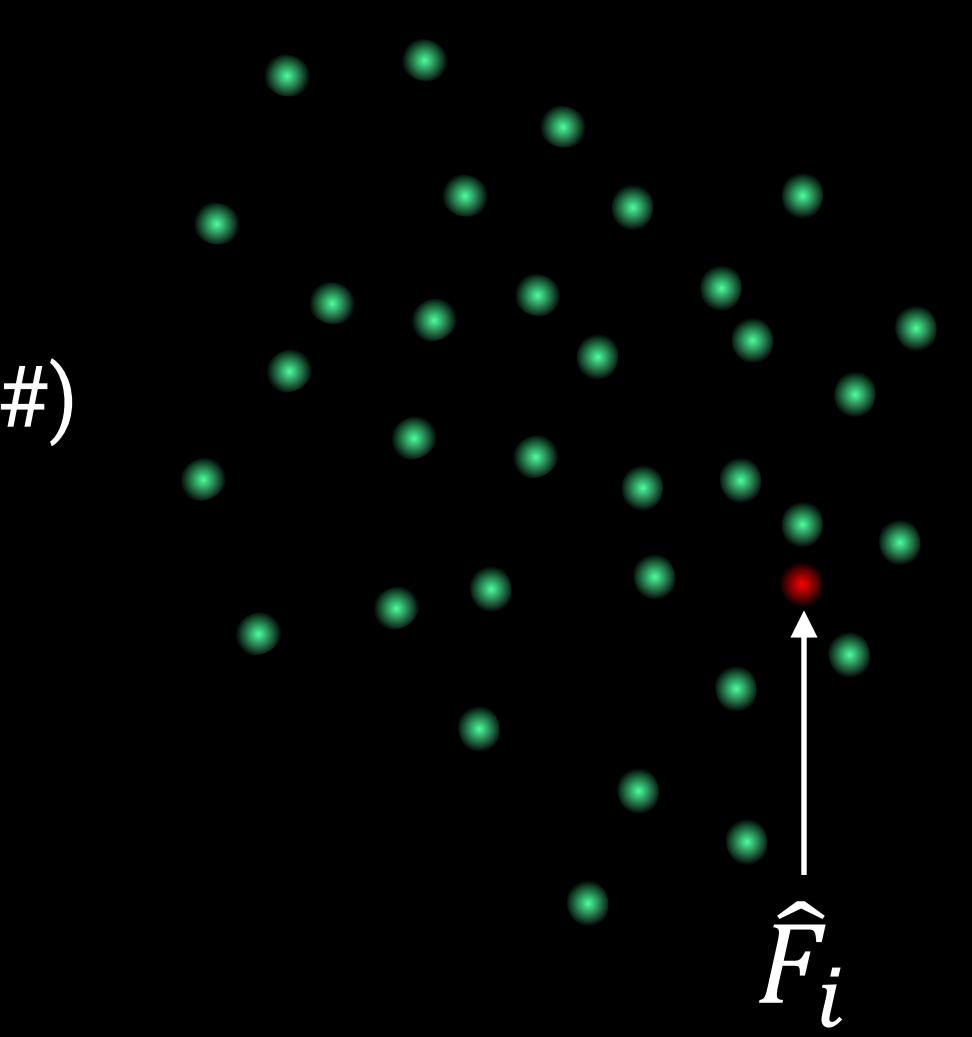


< 64 bytes per fragment  
For 70.000 poses -> ~3 Mb

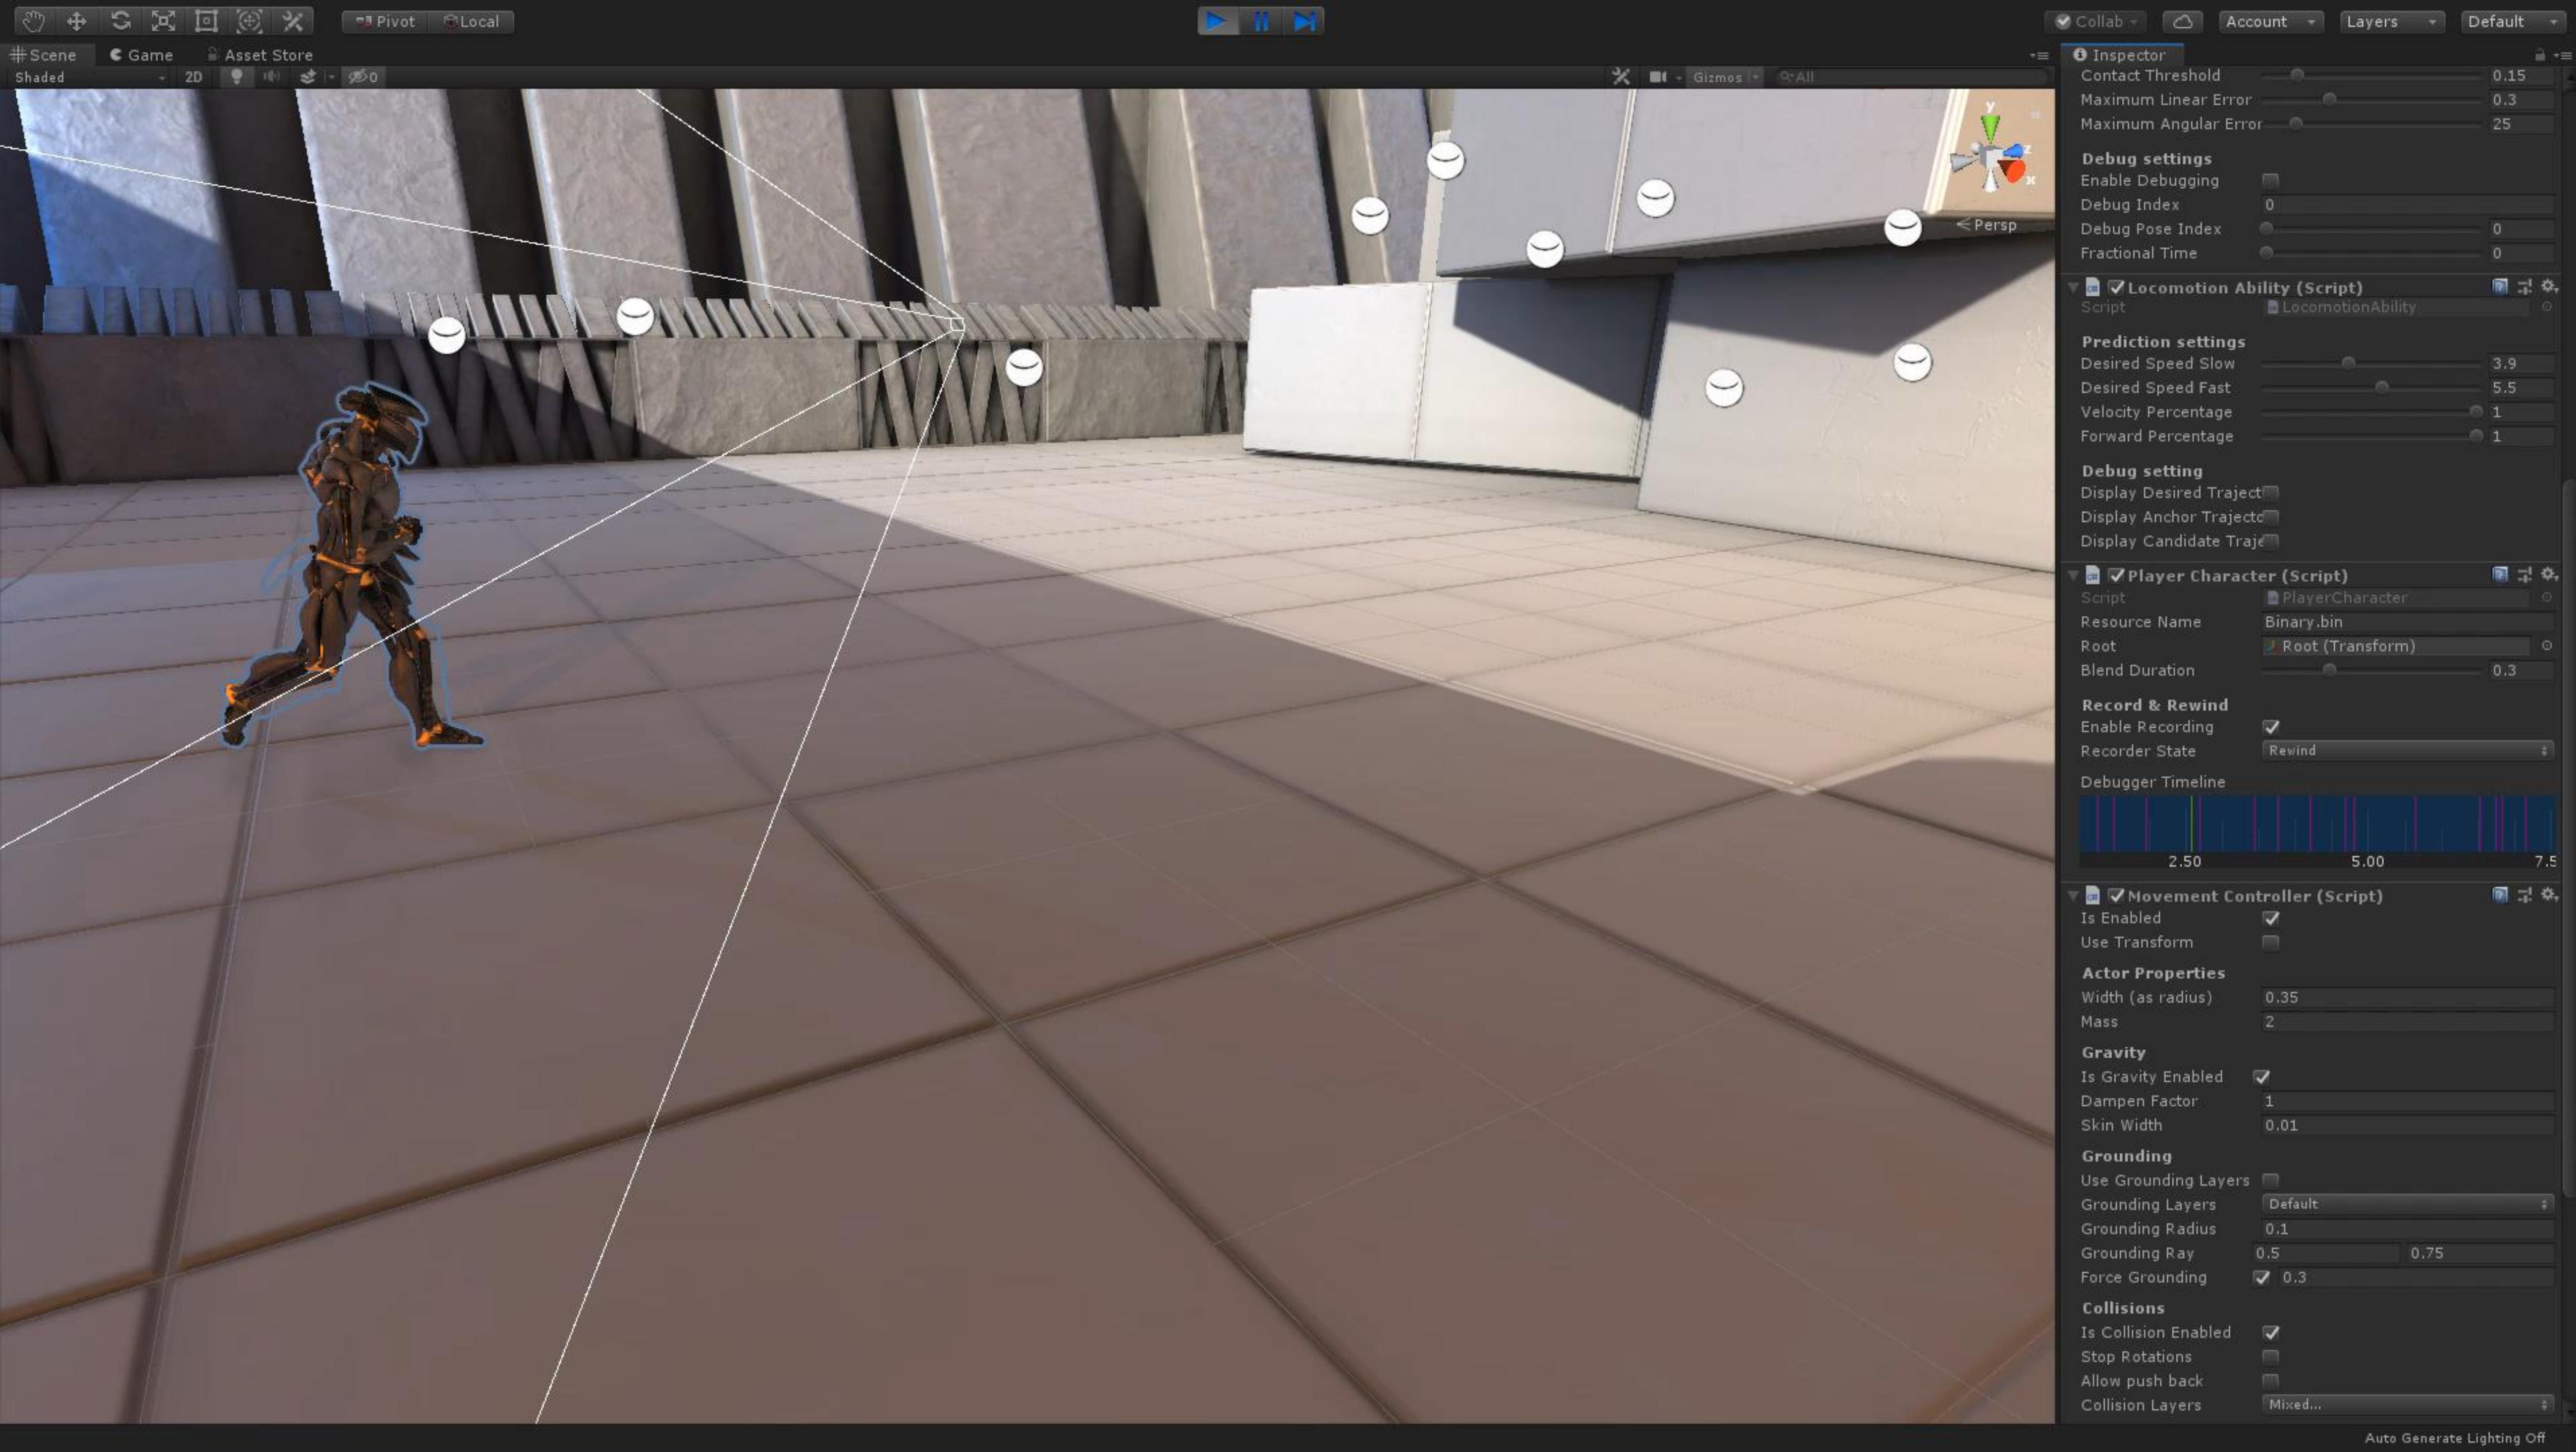
Nearest Neighbor Search time < 0.05 ms (HP C#)

Sub-Linear Nearest Neighbor Search -  $k = 1$

Short training time (< 5 minutes)

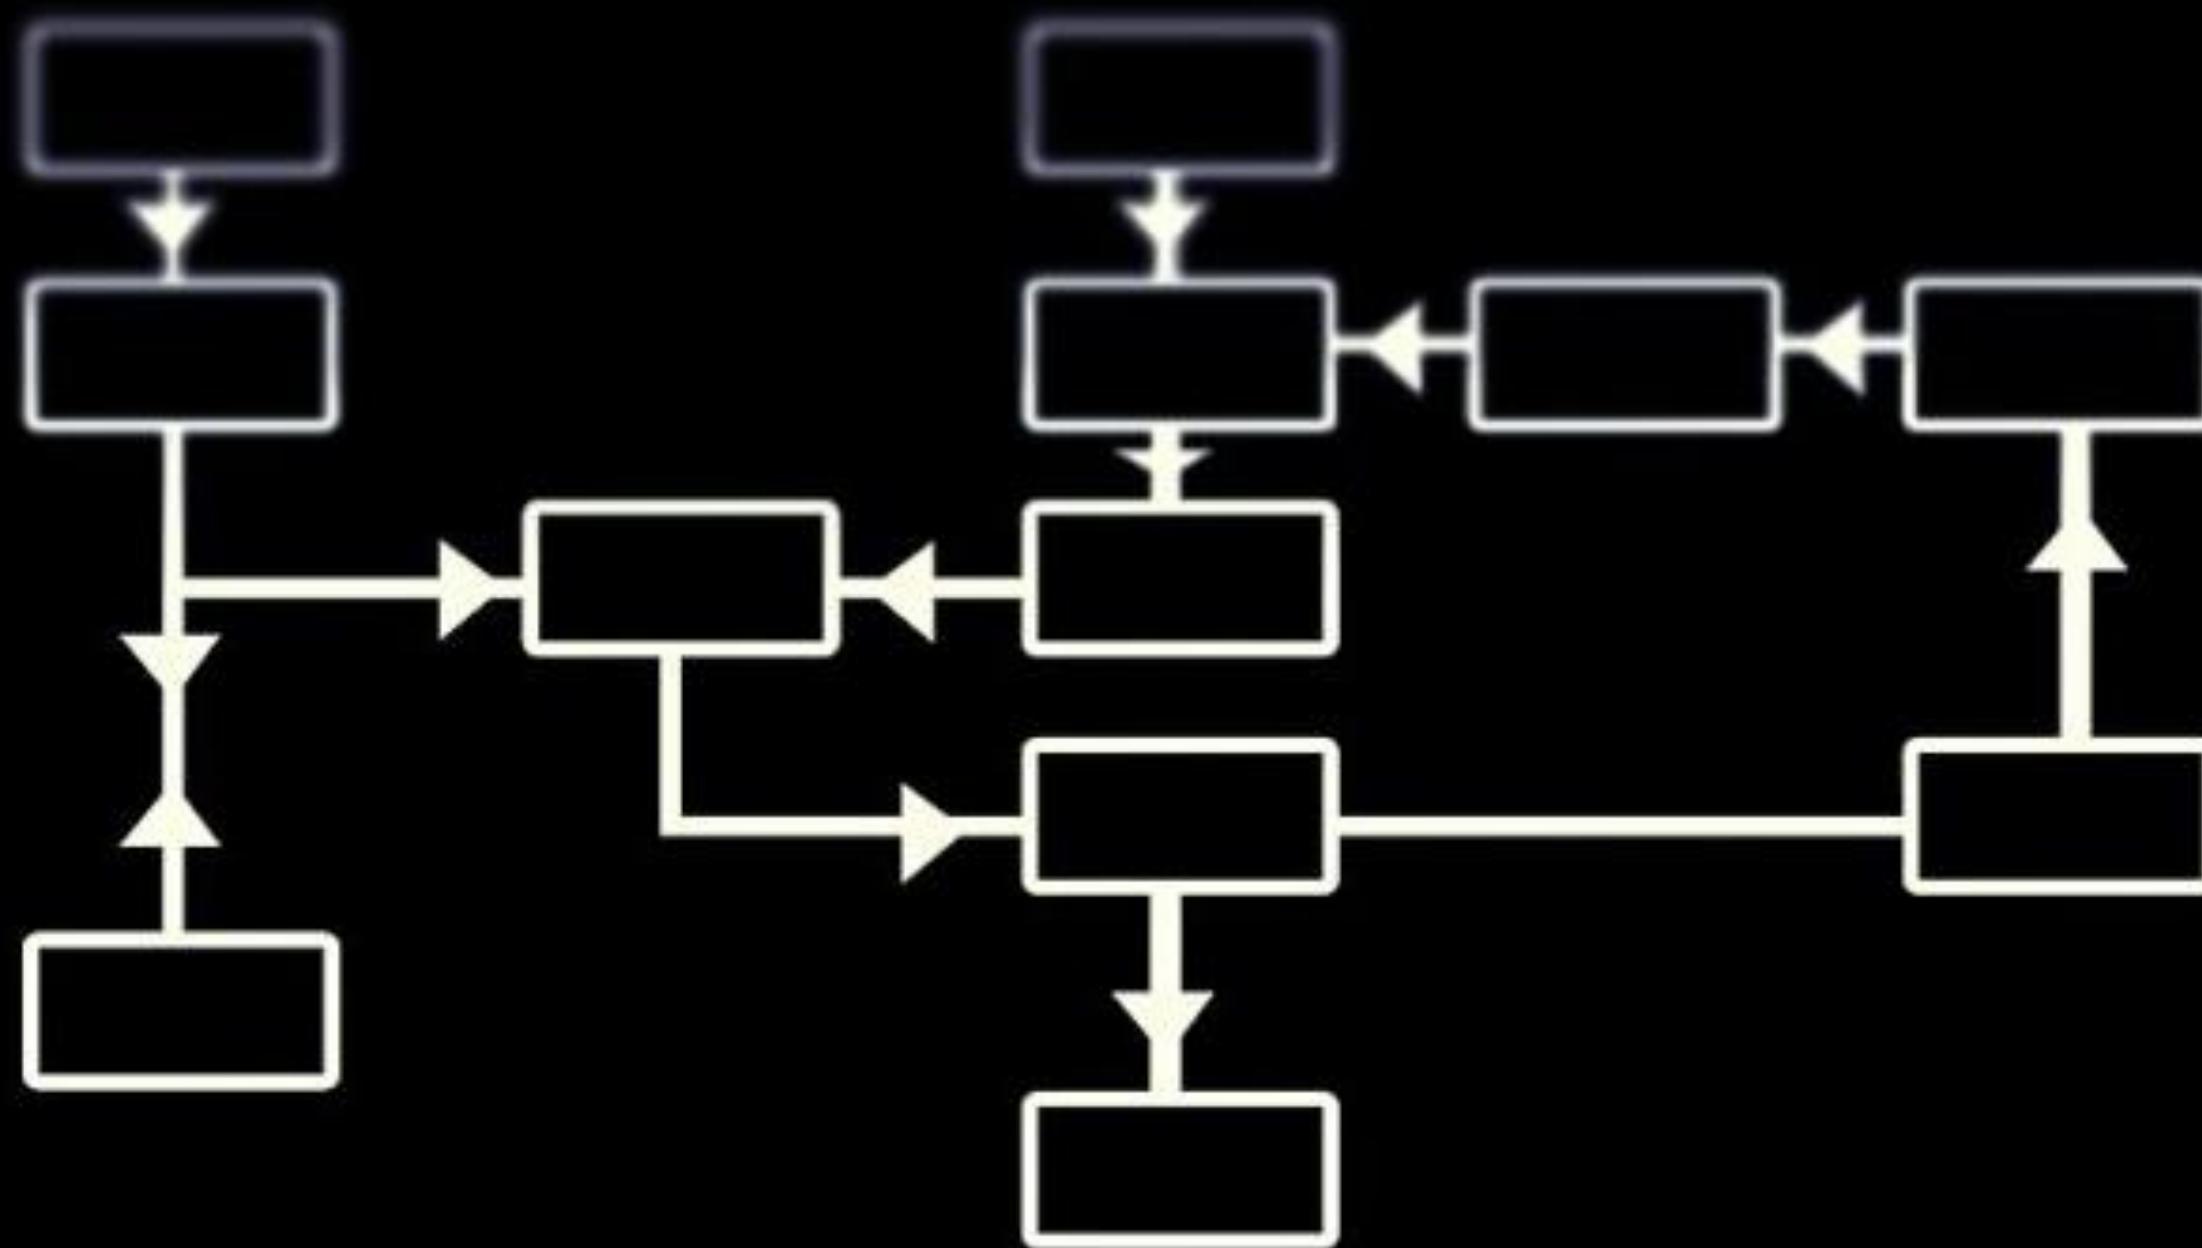




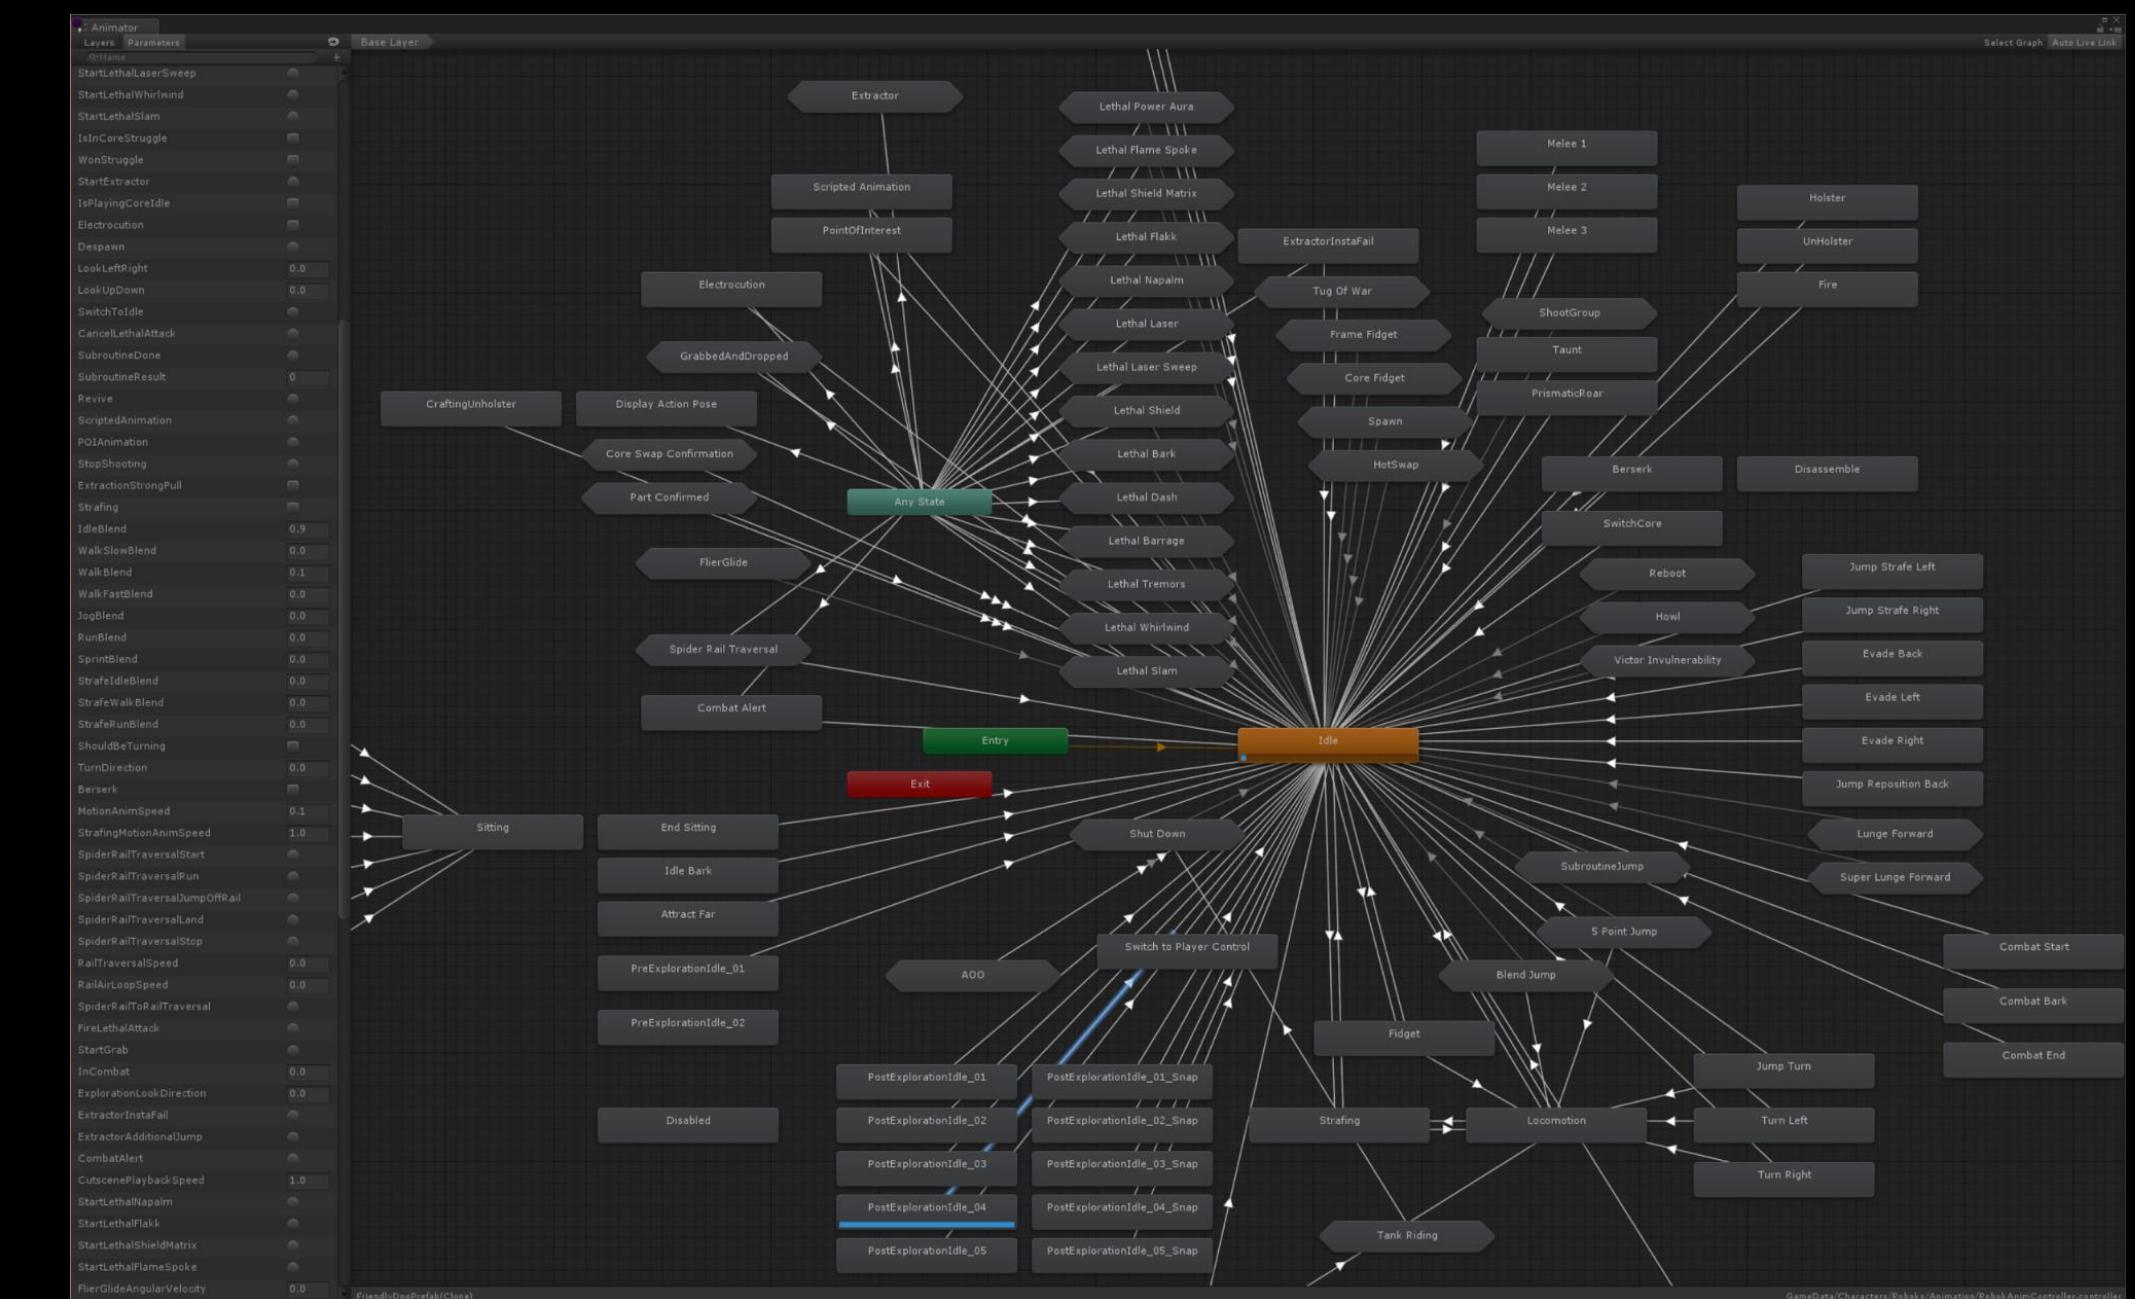


# Game State | Animation Graphs

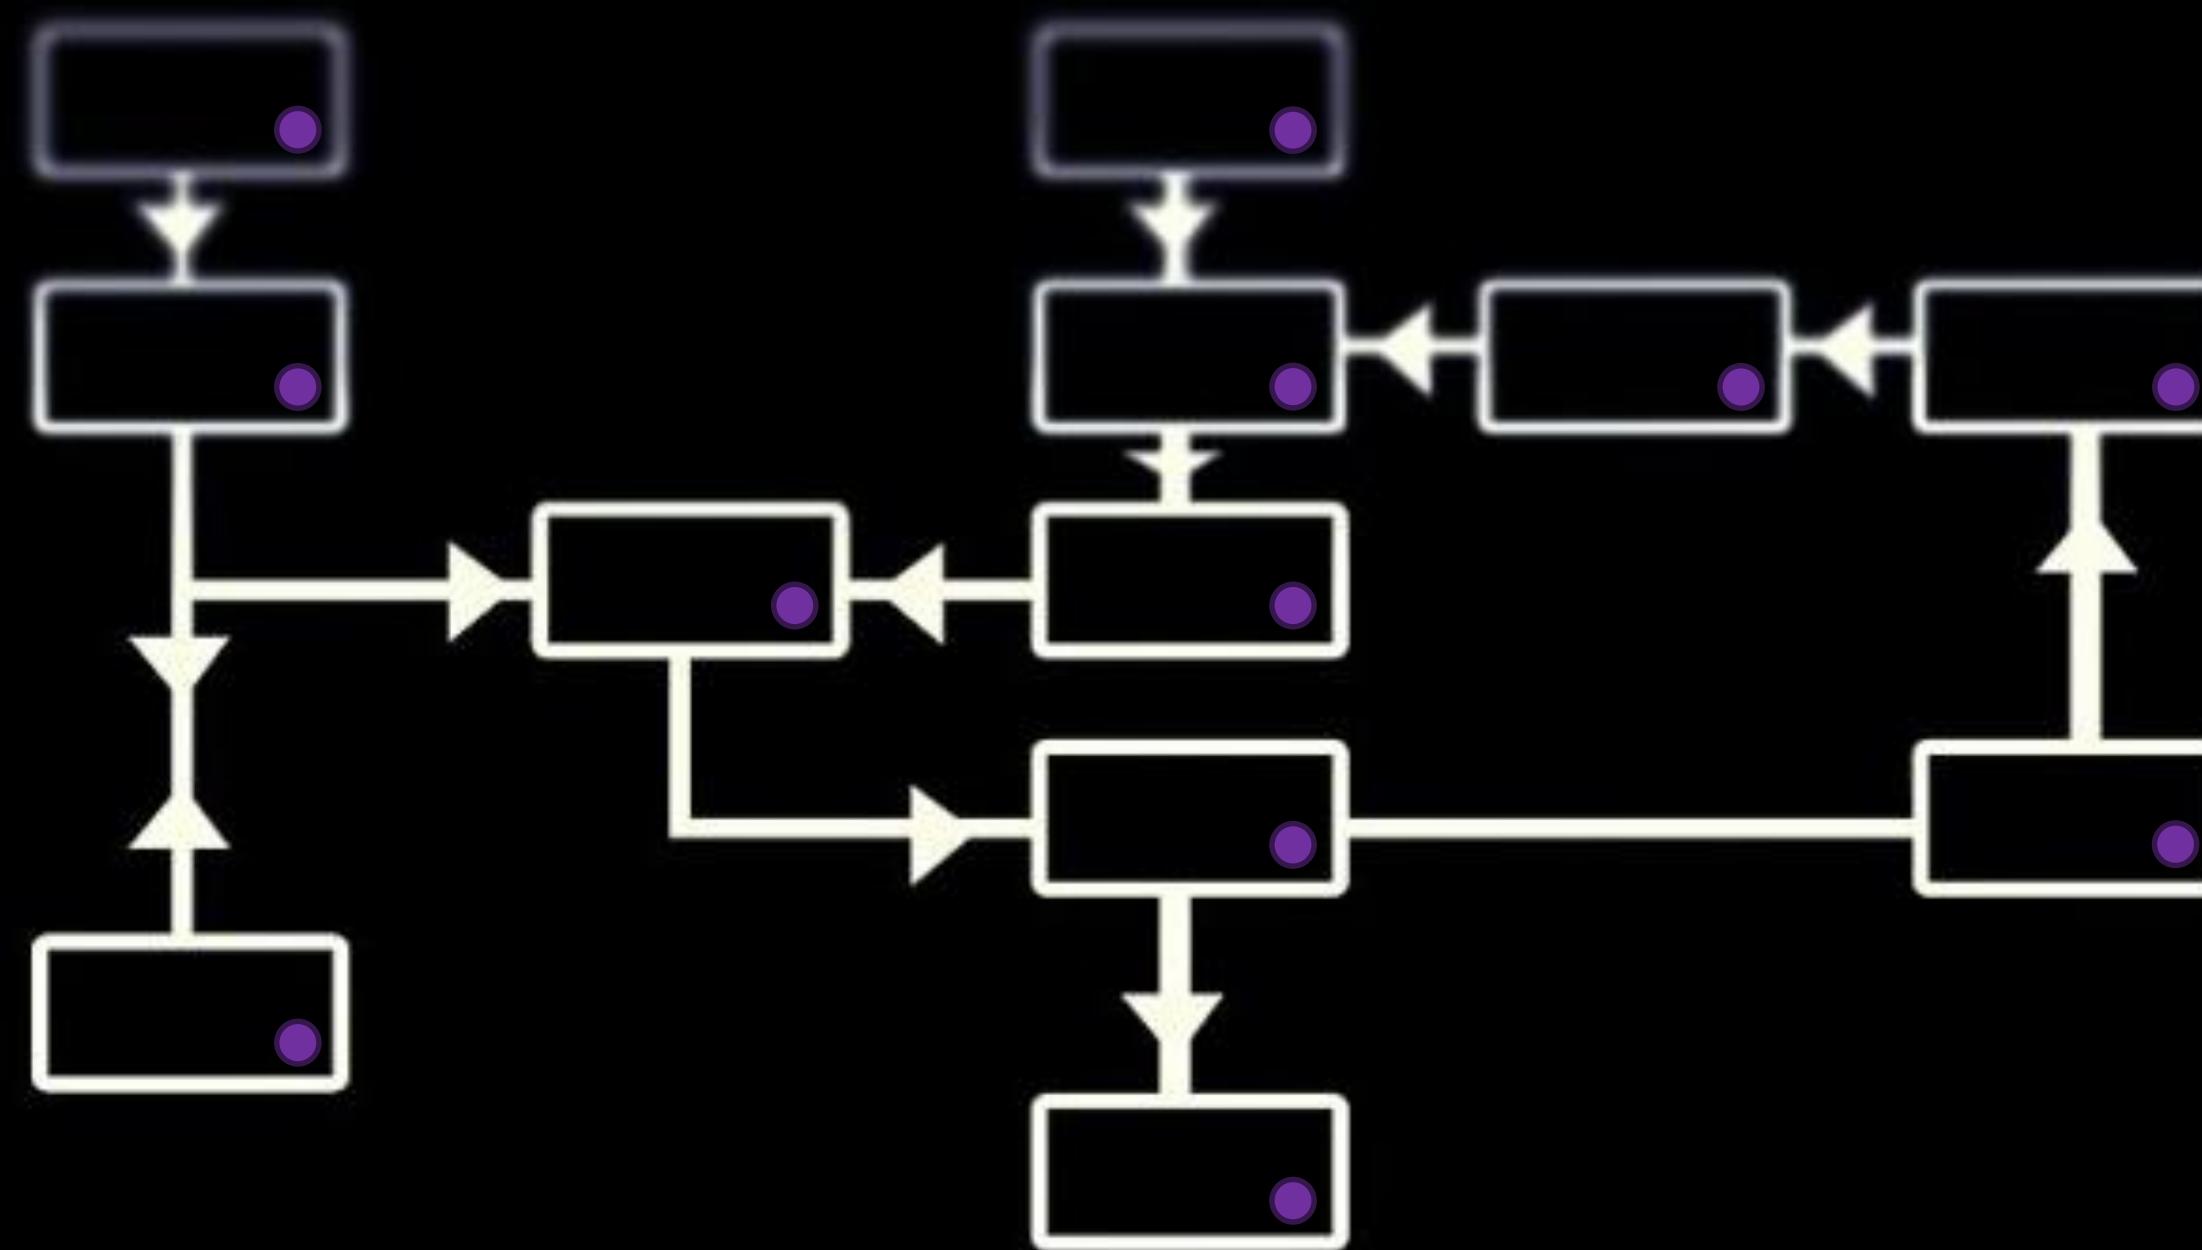
Game Logic



Animation Logic



# Game State



# Abilities



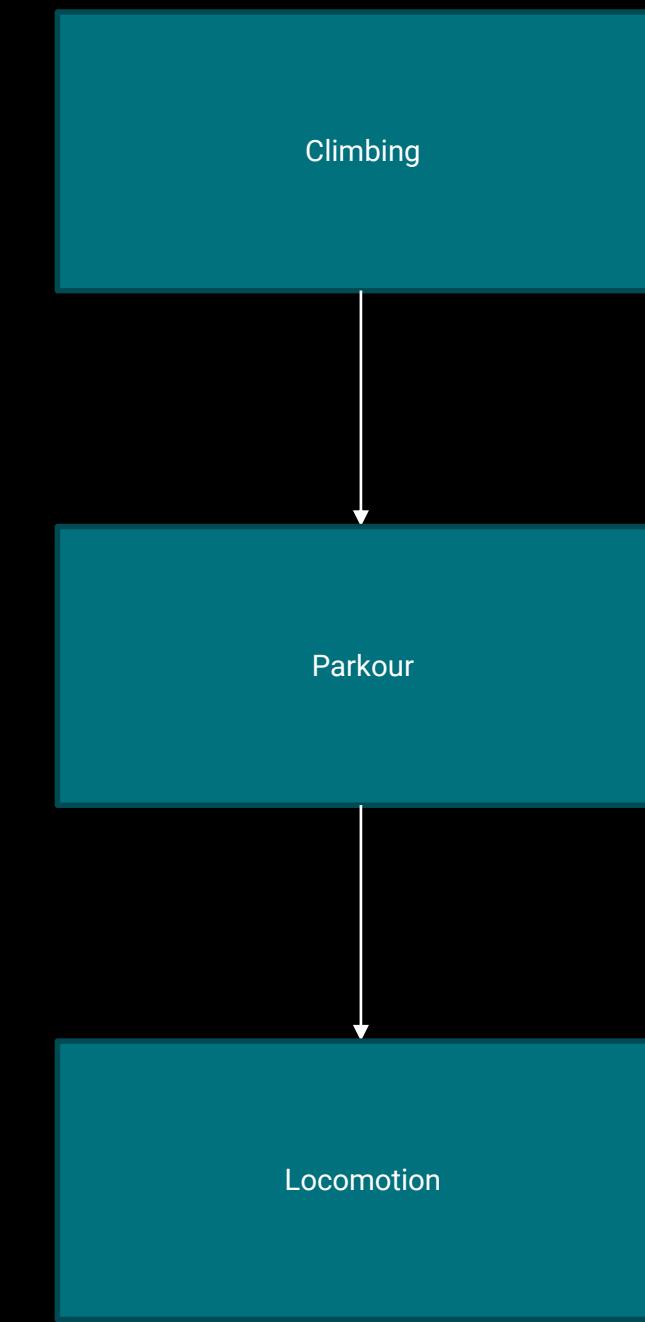
Kinematica's goal is to provide a **complete** alternative to animation graphs

Parkour, Climbing, Melee Combat,  
Synchronized movements, One-off actions, etc...

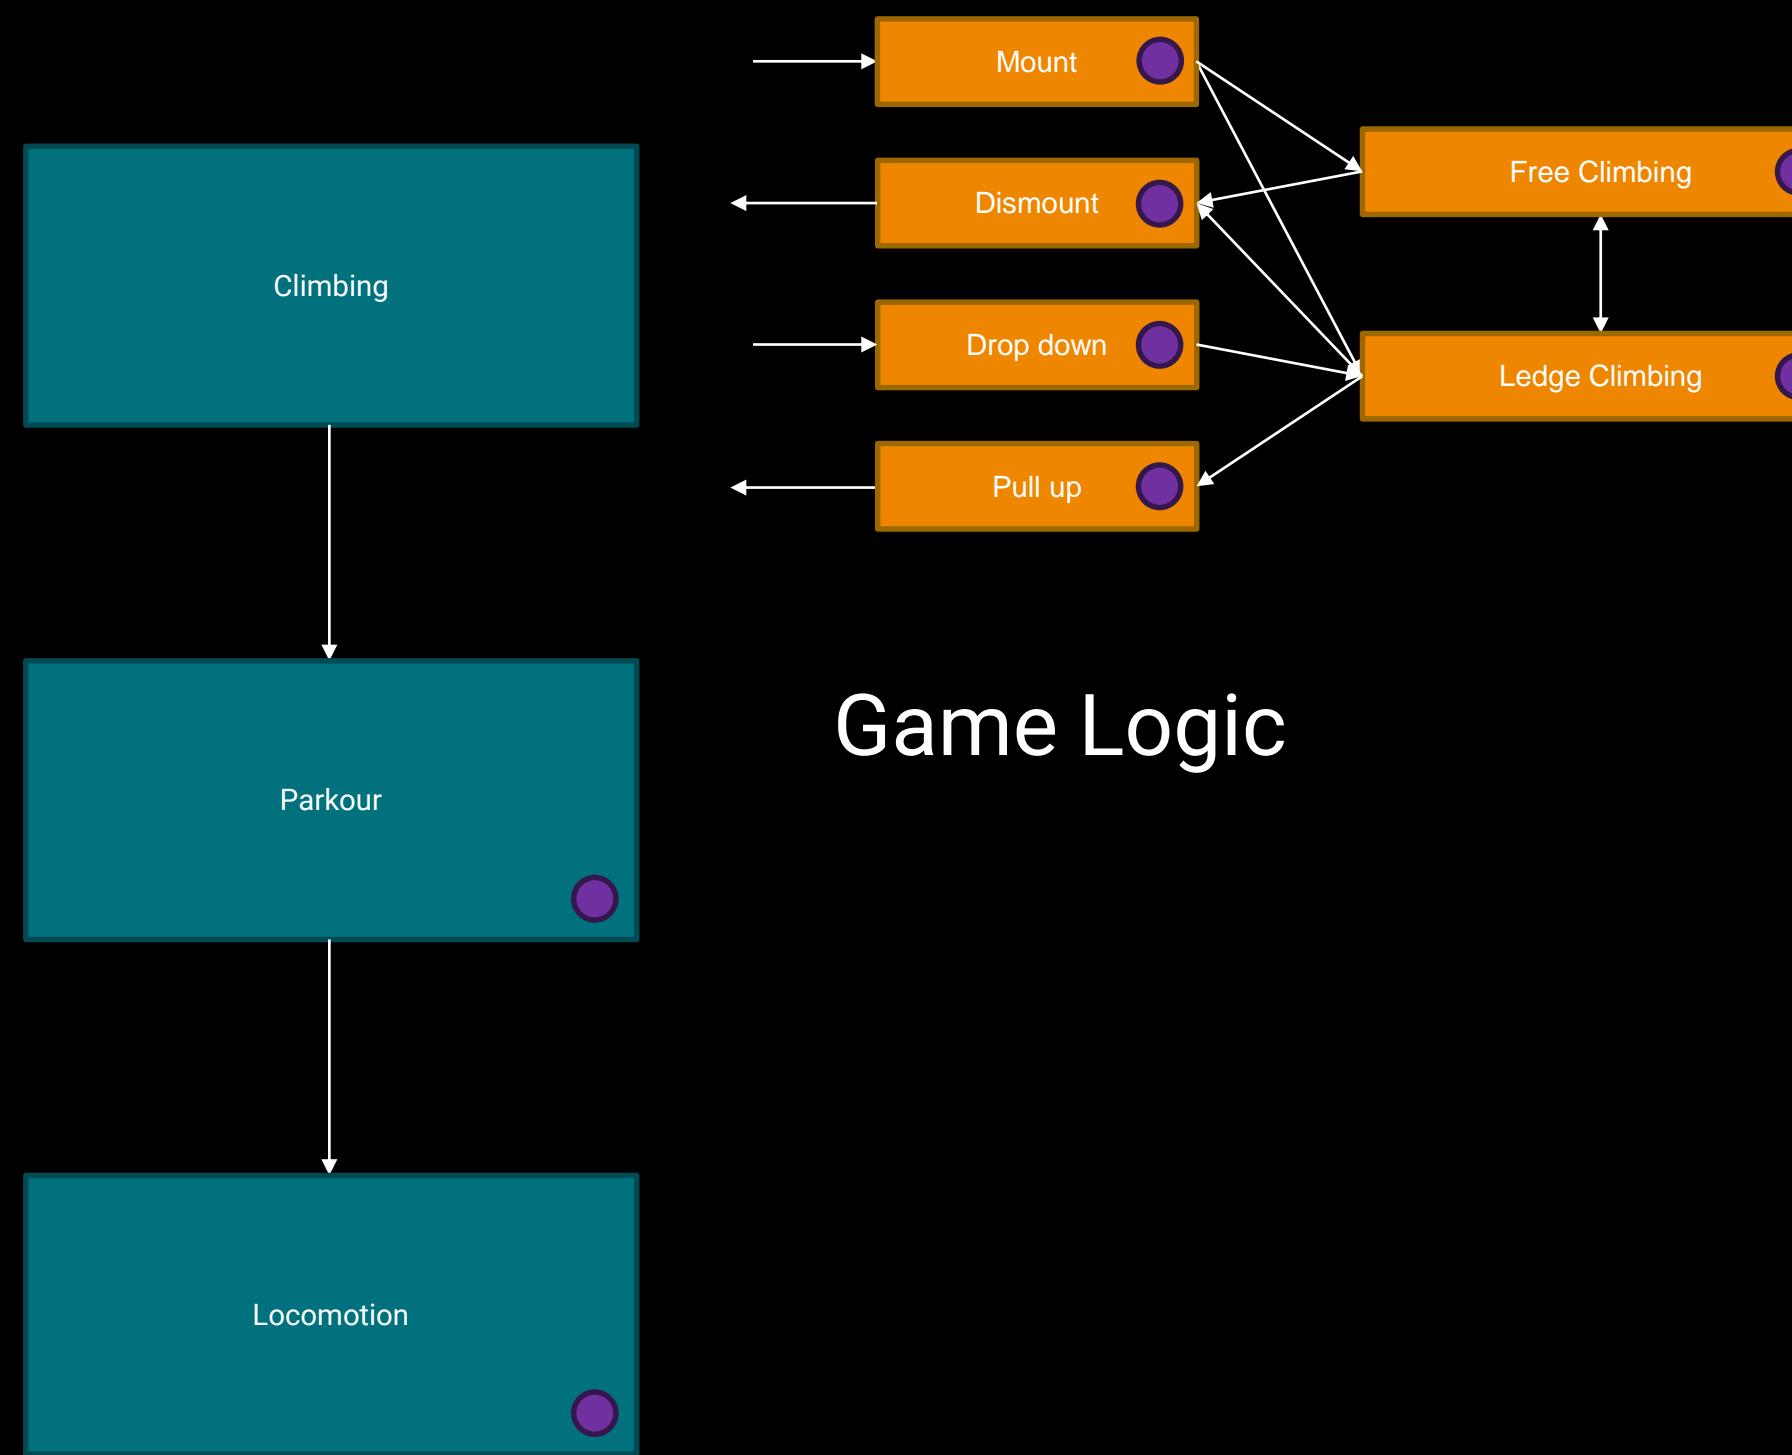
Prioritized list of abilities

Executed in order

This is not a super-imposed concept, i.e.  
Kinematica does not call into abilities



# Kinematica Components



Policies execute as part of the game code (ideally as C# jobs)

Policies execute variations of similarity searches depending on game logic

Similarity searches can be based on 1:1 or n:m fragments



# Motion Library

All poses are arranged into a large matrix  $D \in M^{(m+1) \times D}(\mathbb{R})$  where each column corresponds to a pose  $\mathcal{J} = \{J_i; i =$

$$D = \begin{bmatrix} T_r^1 & T_r^2 & \cdots & T_r^D \\ J_1^1 & J_1^2 & \cdots & J_1^D \\ \vdots & \vdots & \ddots & \vdots \\ J_m^1 & J_m^2 & \cdots & J_m^D \end{bmatrix}$$

Tagging segregates the motion library into addressable islands

Markers carry an arbitrary user-defined payload and are associated with discrete frames

Policies utilize tags and markers in user-defined similarity searches

```
namespace Parkour
{
    public struct Contact
    {
        public struct Anchor
        {
            // Frame index relative to the tag start
            // frame at which the anchor has been placed
            public int atFrame;

            // Transform relative to the trajectory transform
            // of the frame at which the anchor has been placed
            public AffineTransform transform;
        }

        public Type type;
        public int firstFrame;
        public int numFrames;
        public int firstContact;
        public int numContacts;
        public int layerMask;
        public Anchor anchor;
    }
}
```

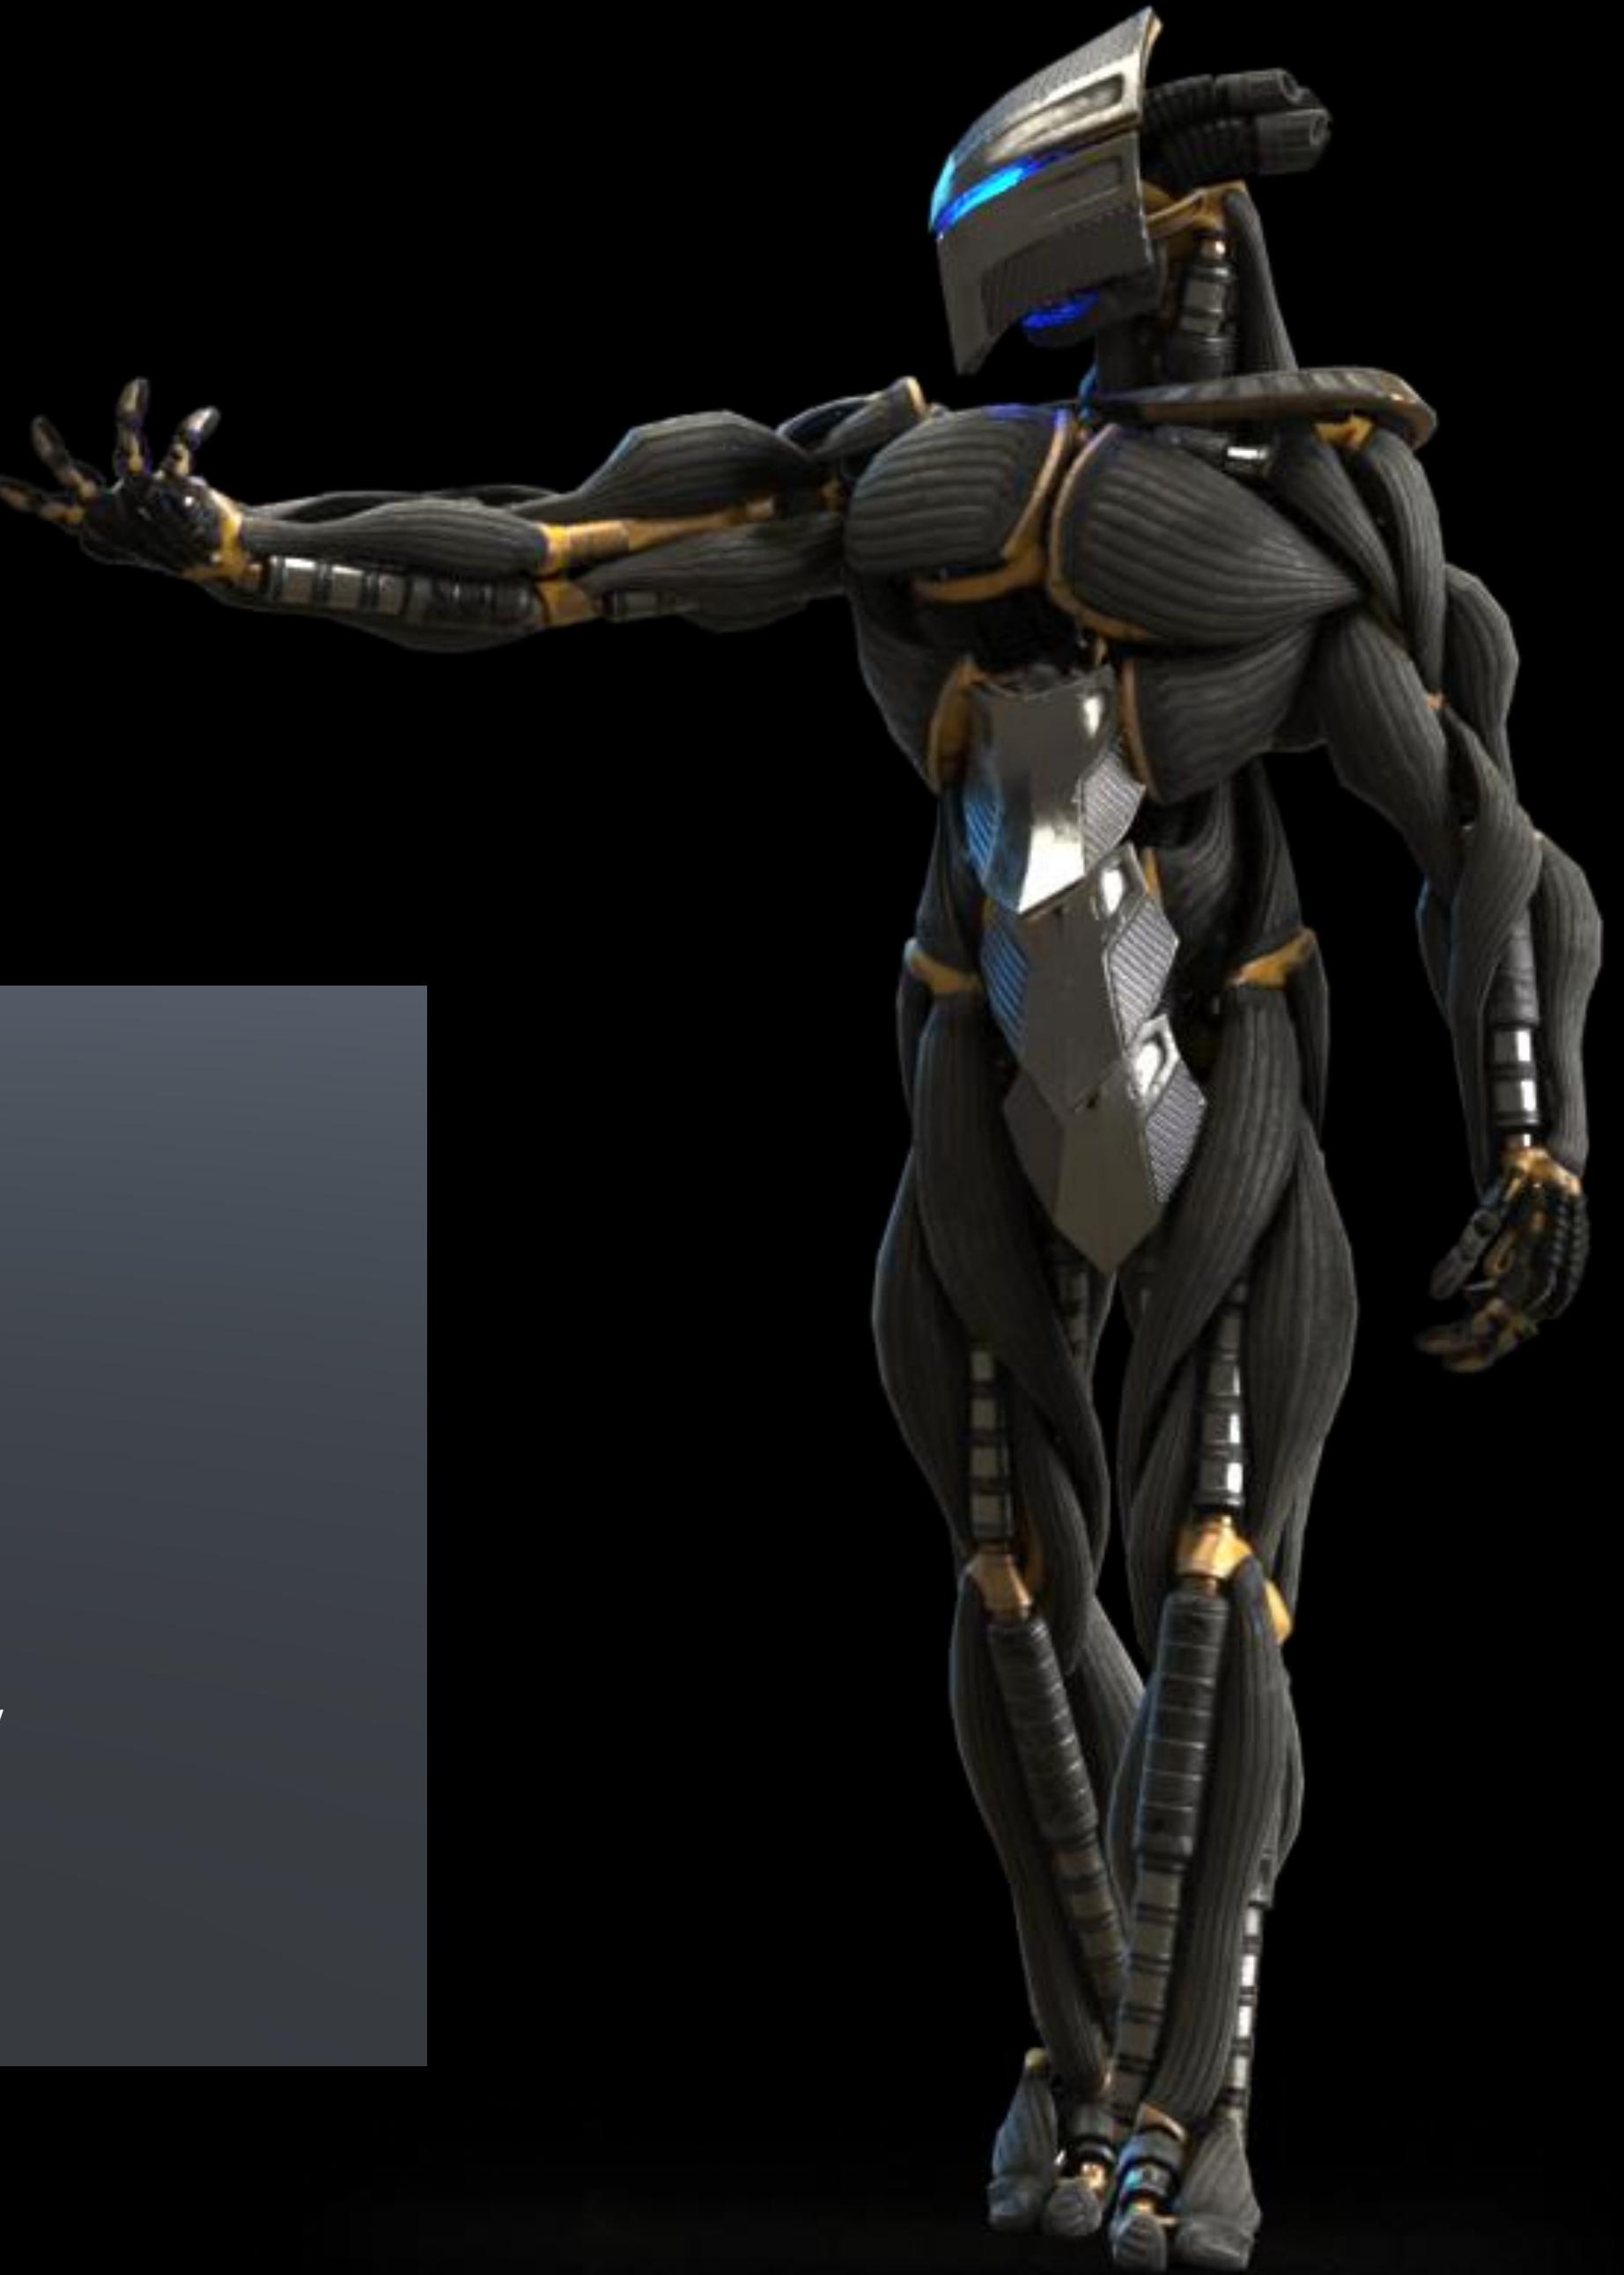
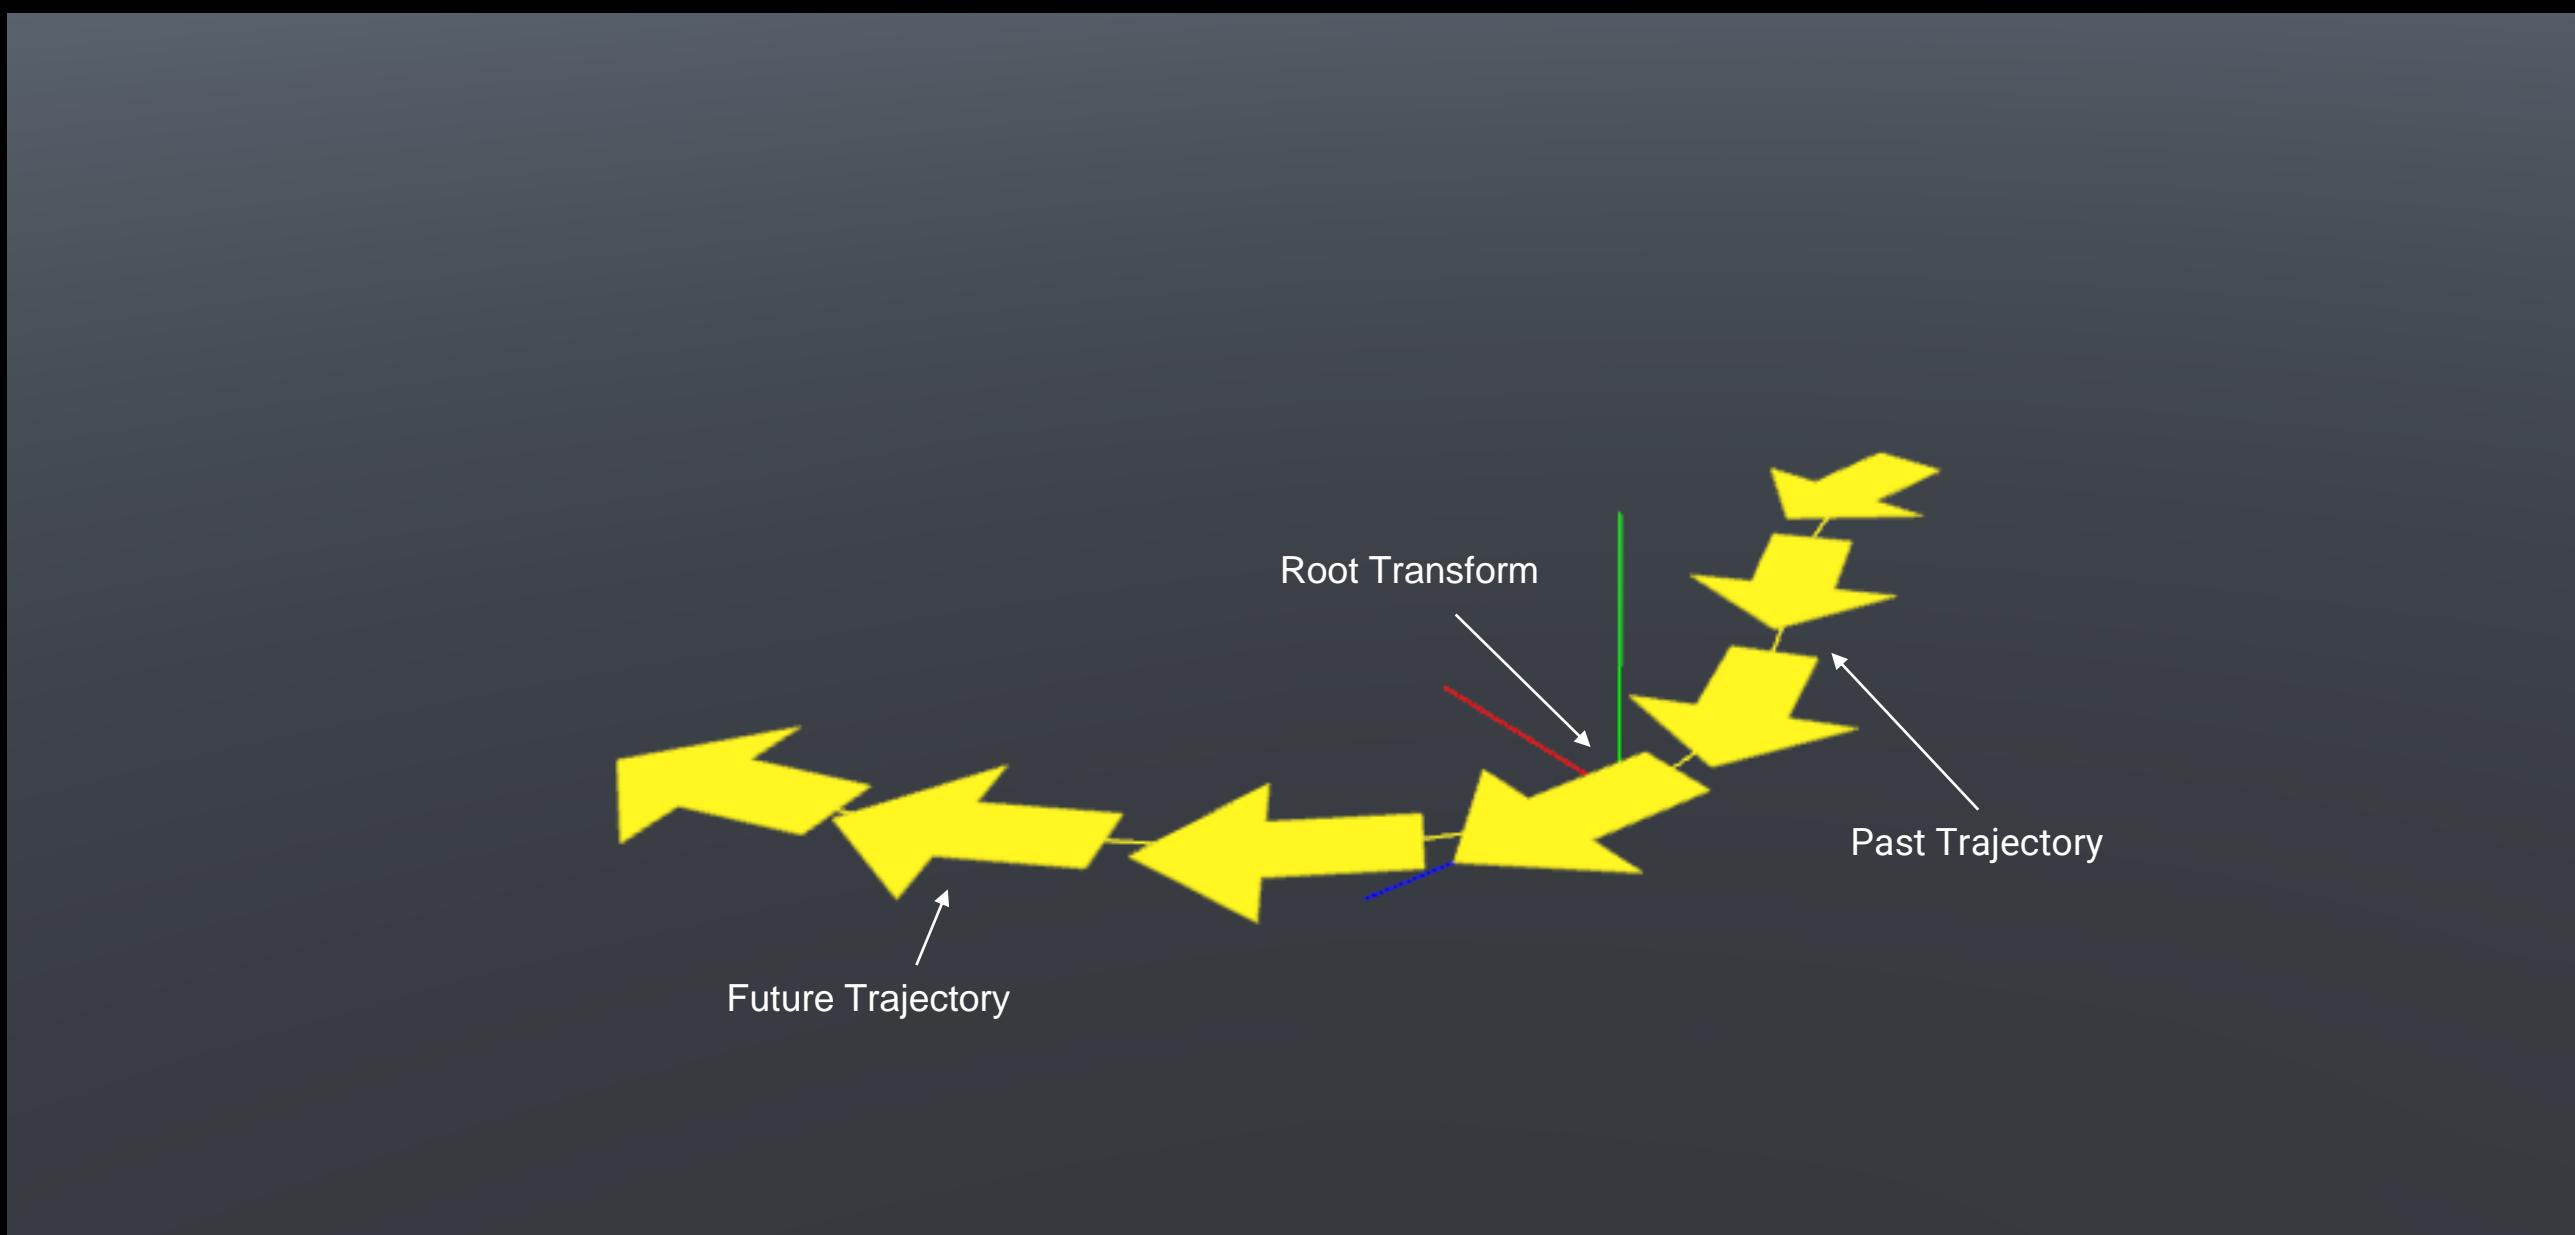


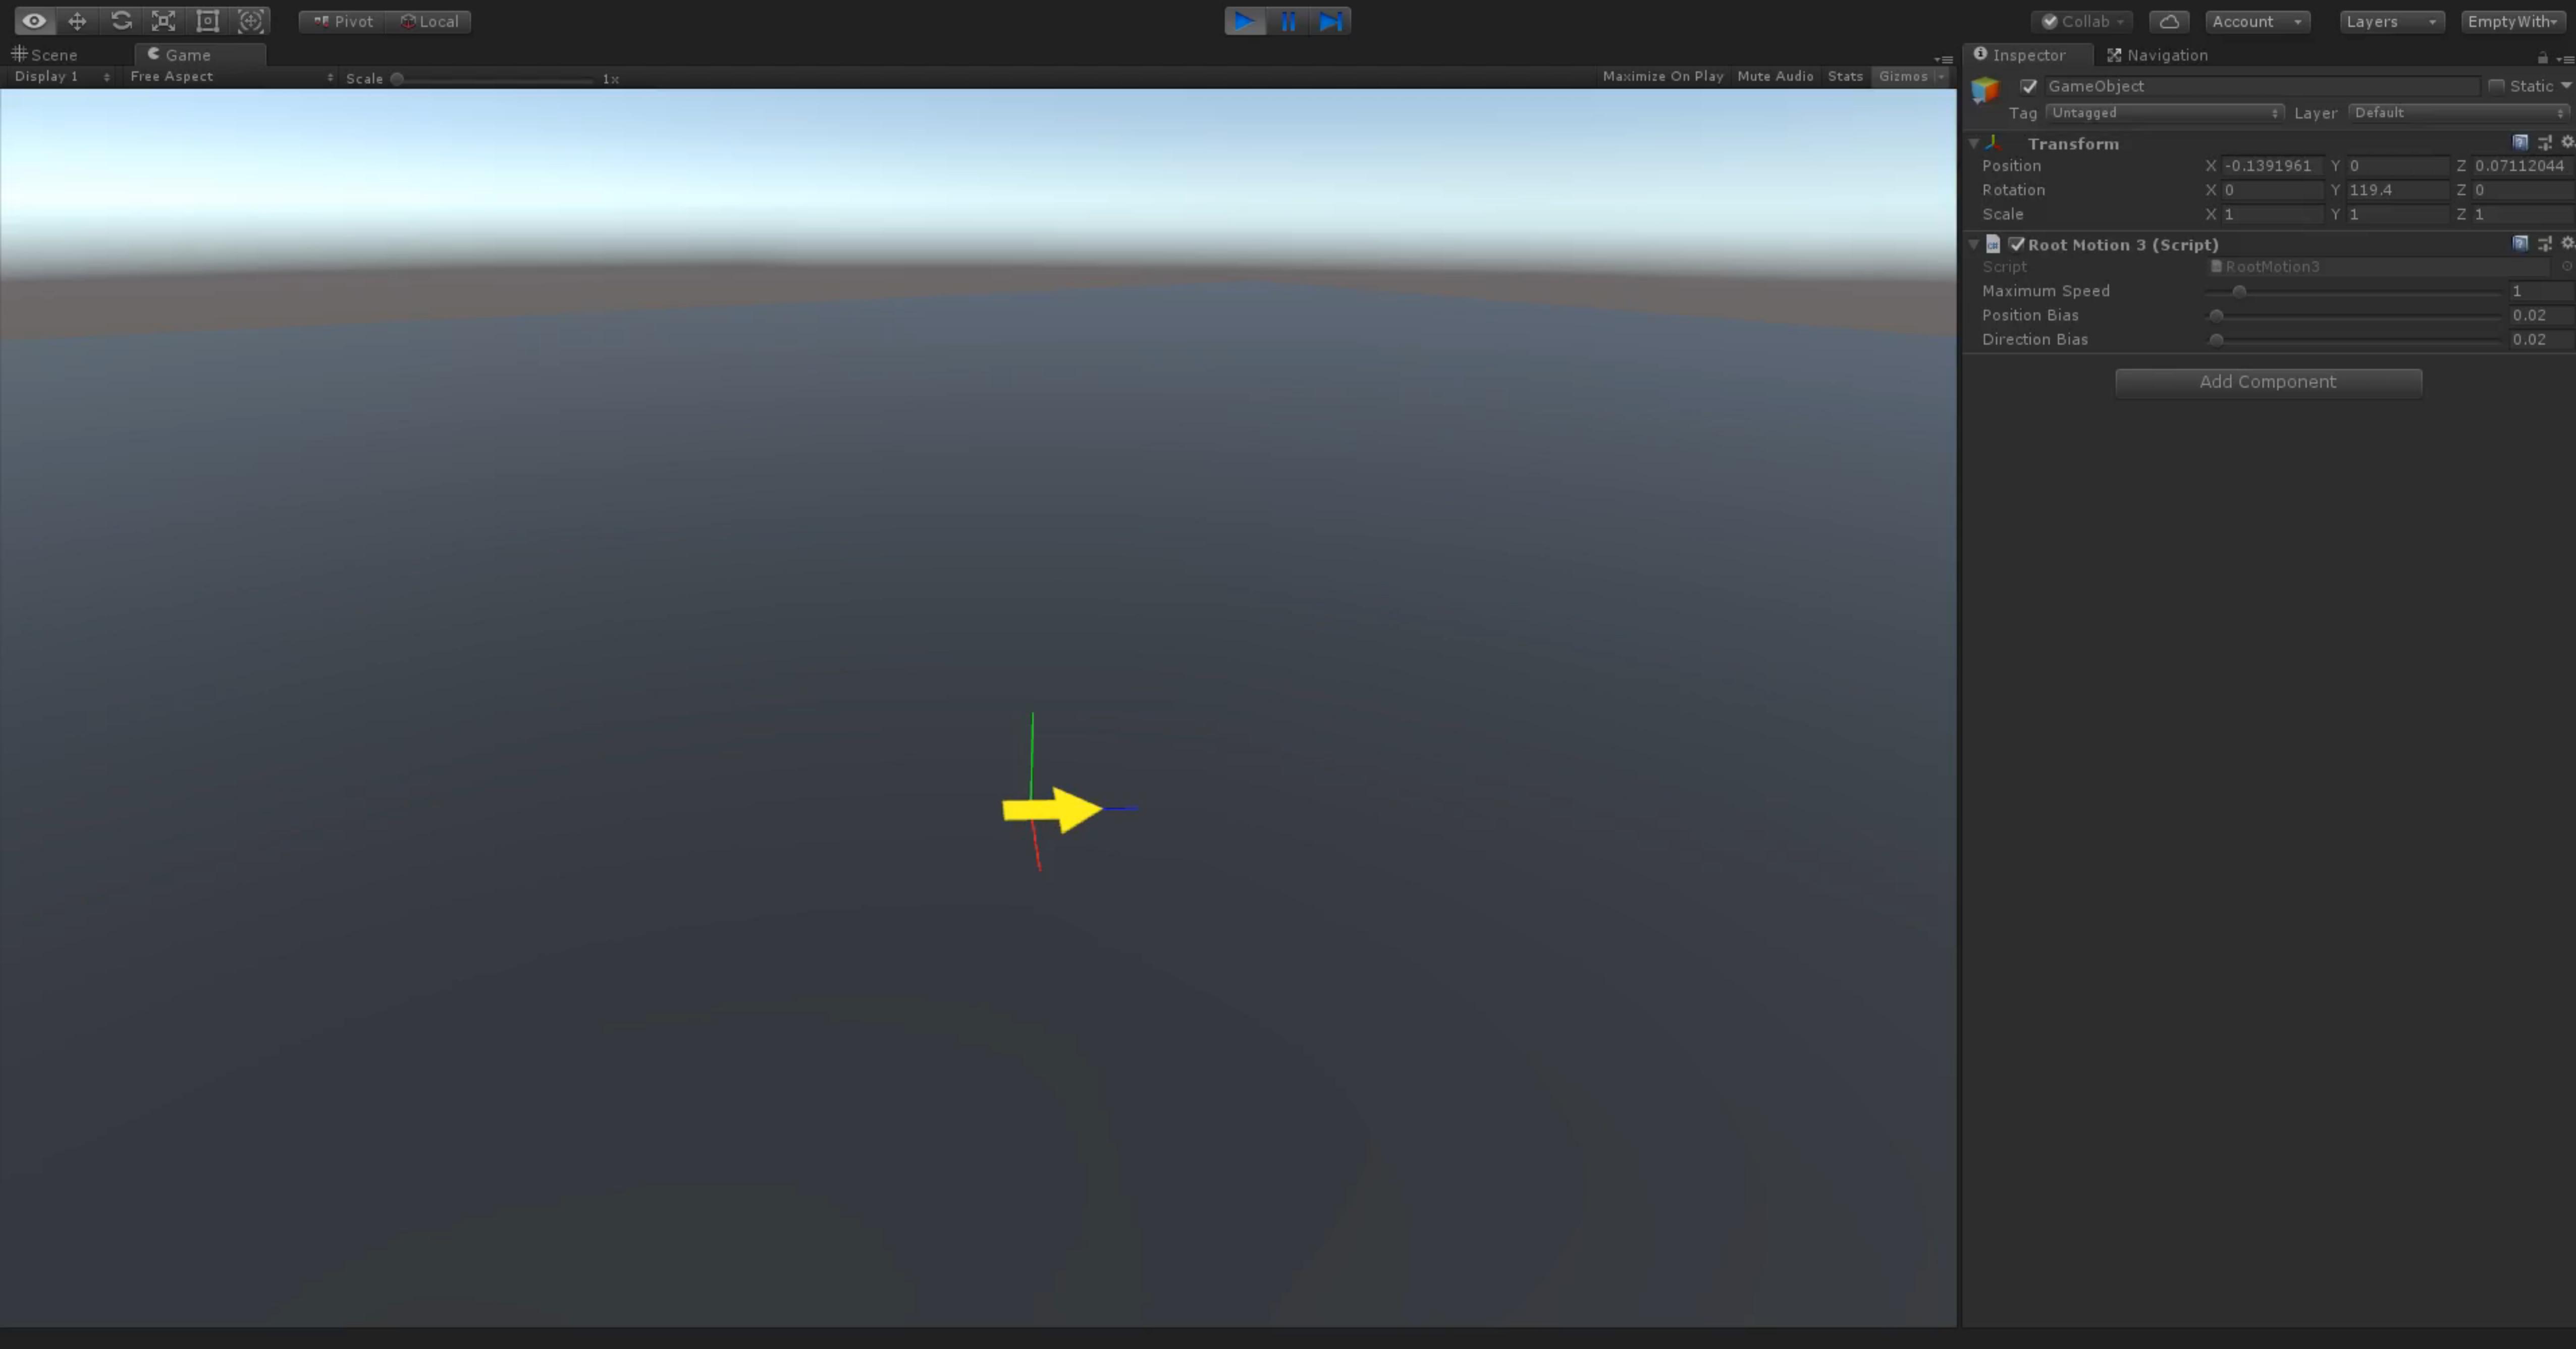
# Locomotion



# Trajectory Prediction

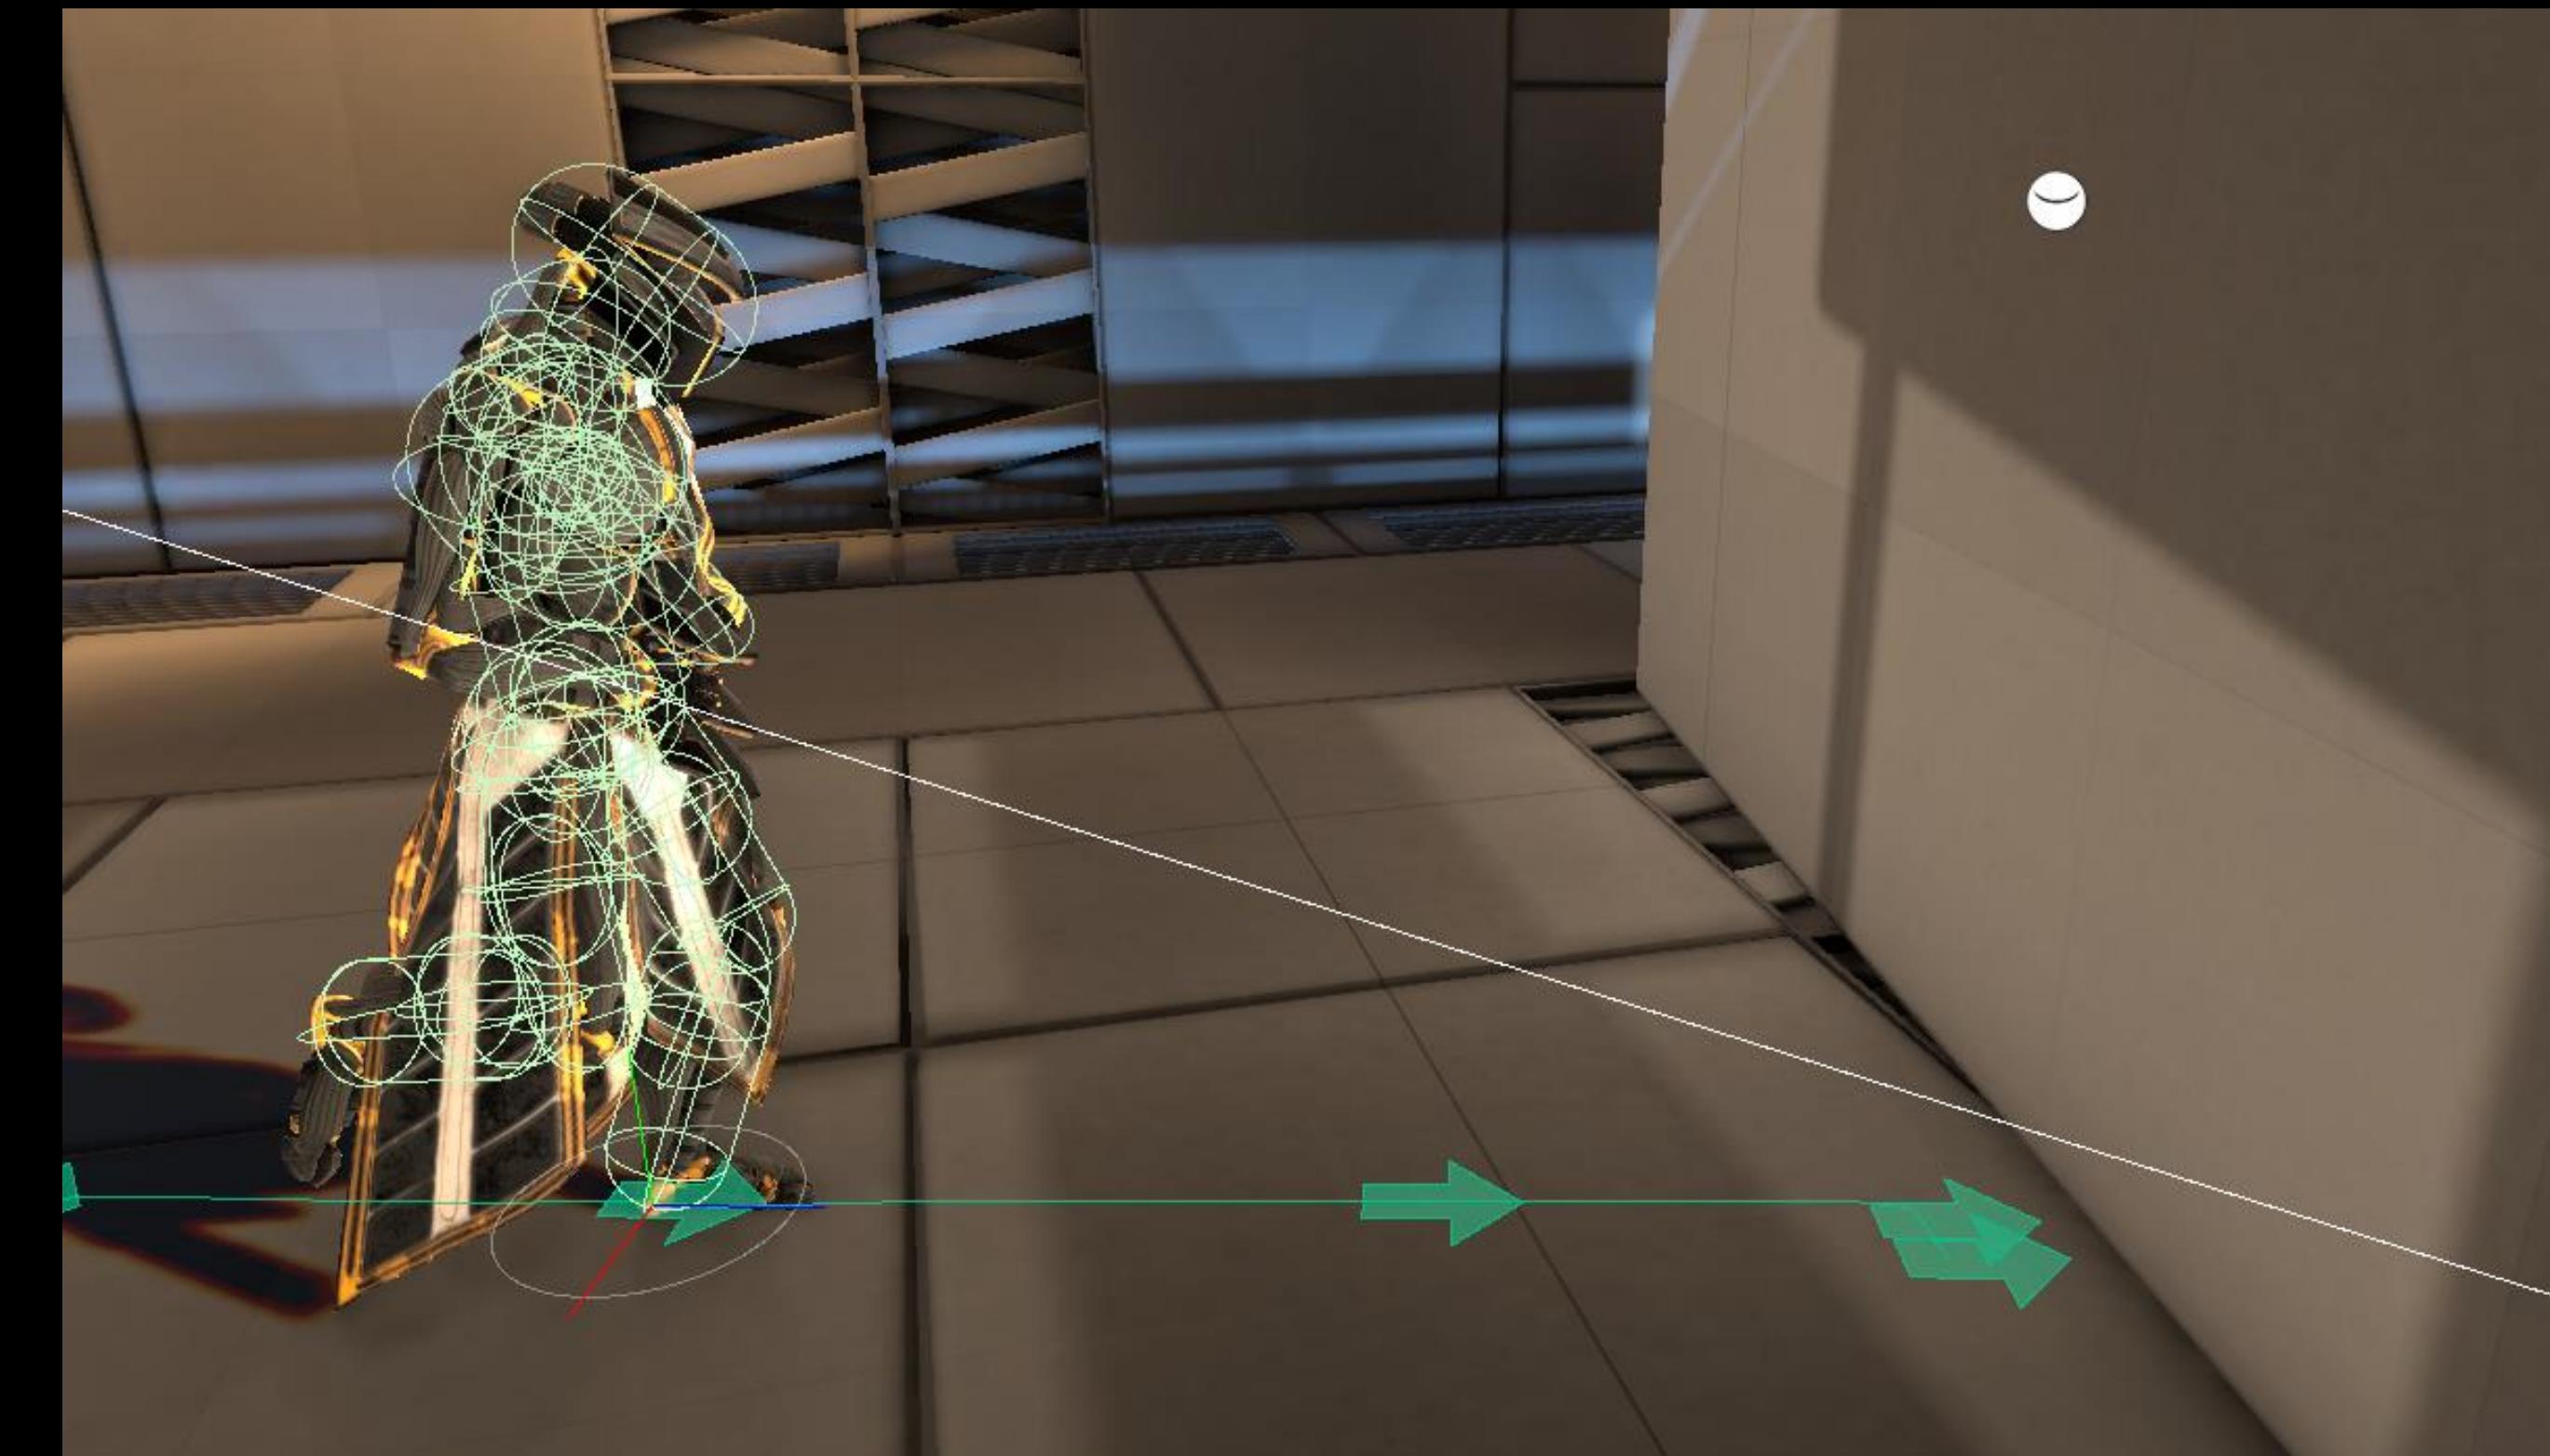
Using our knowledge about the intended target location (NPC) or desired velocity (PC) we can generate a predicted path over the time horizon

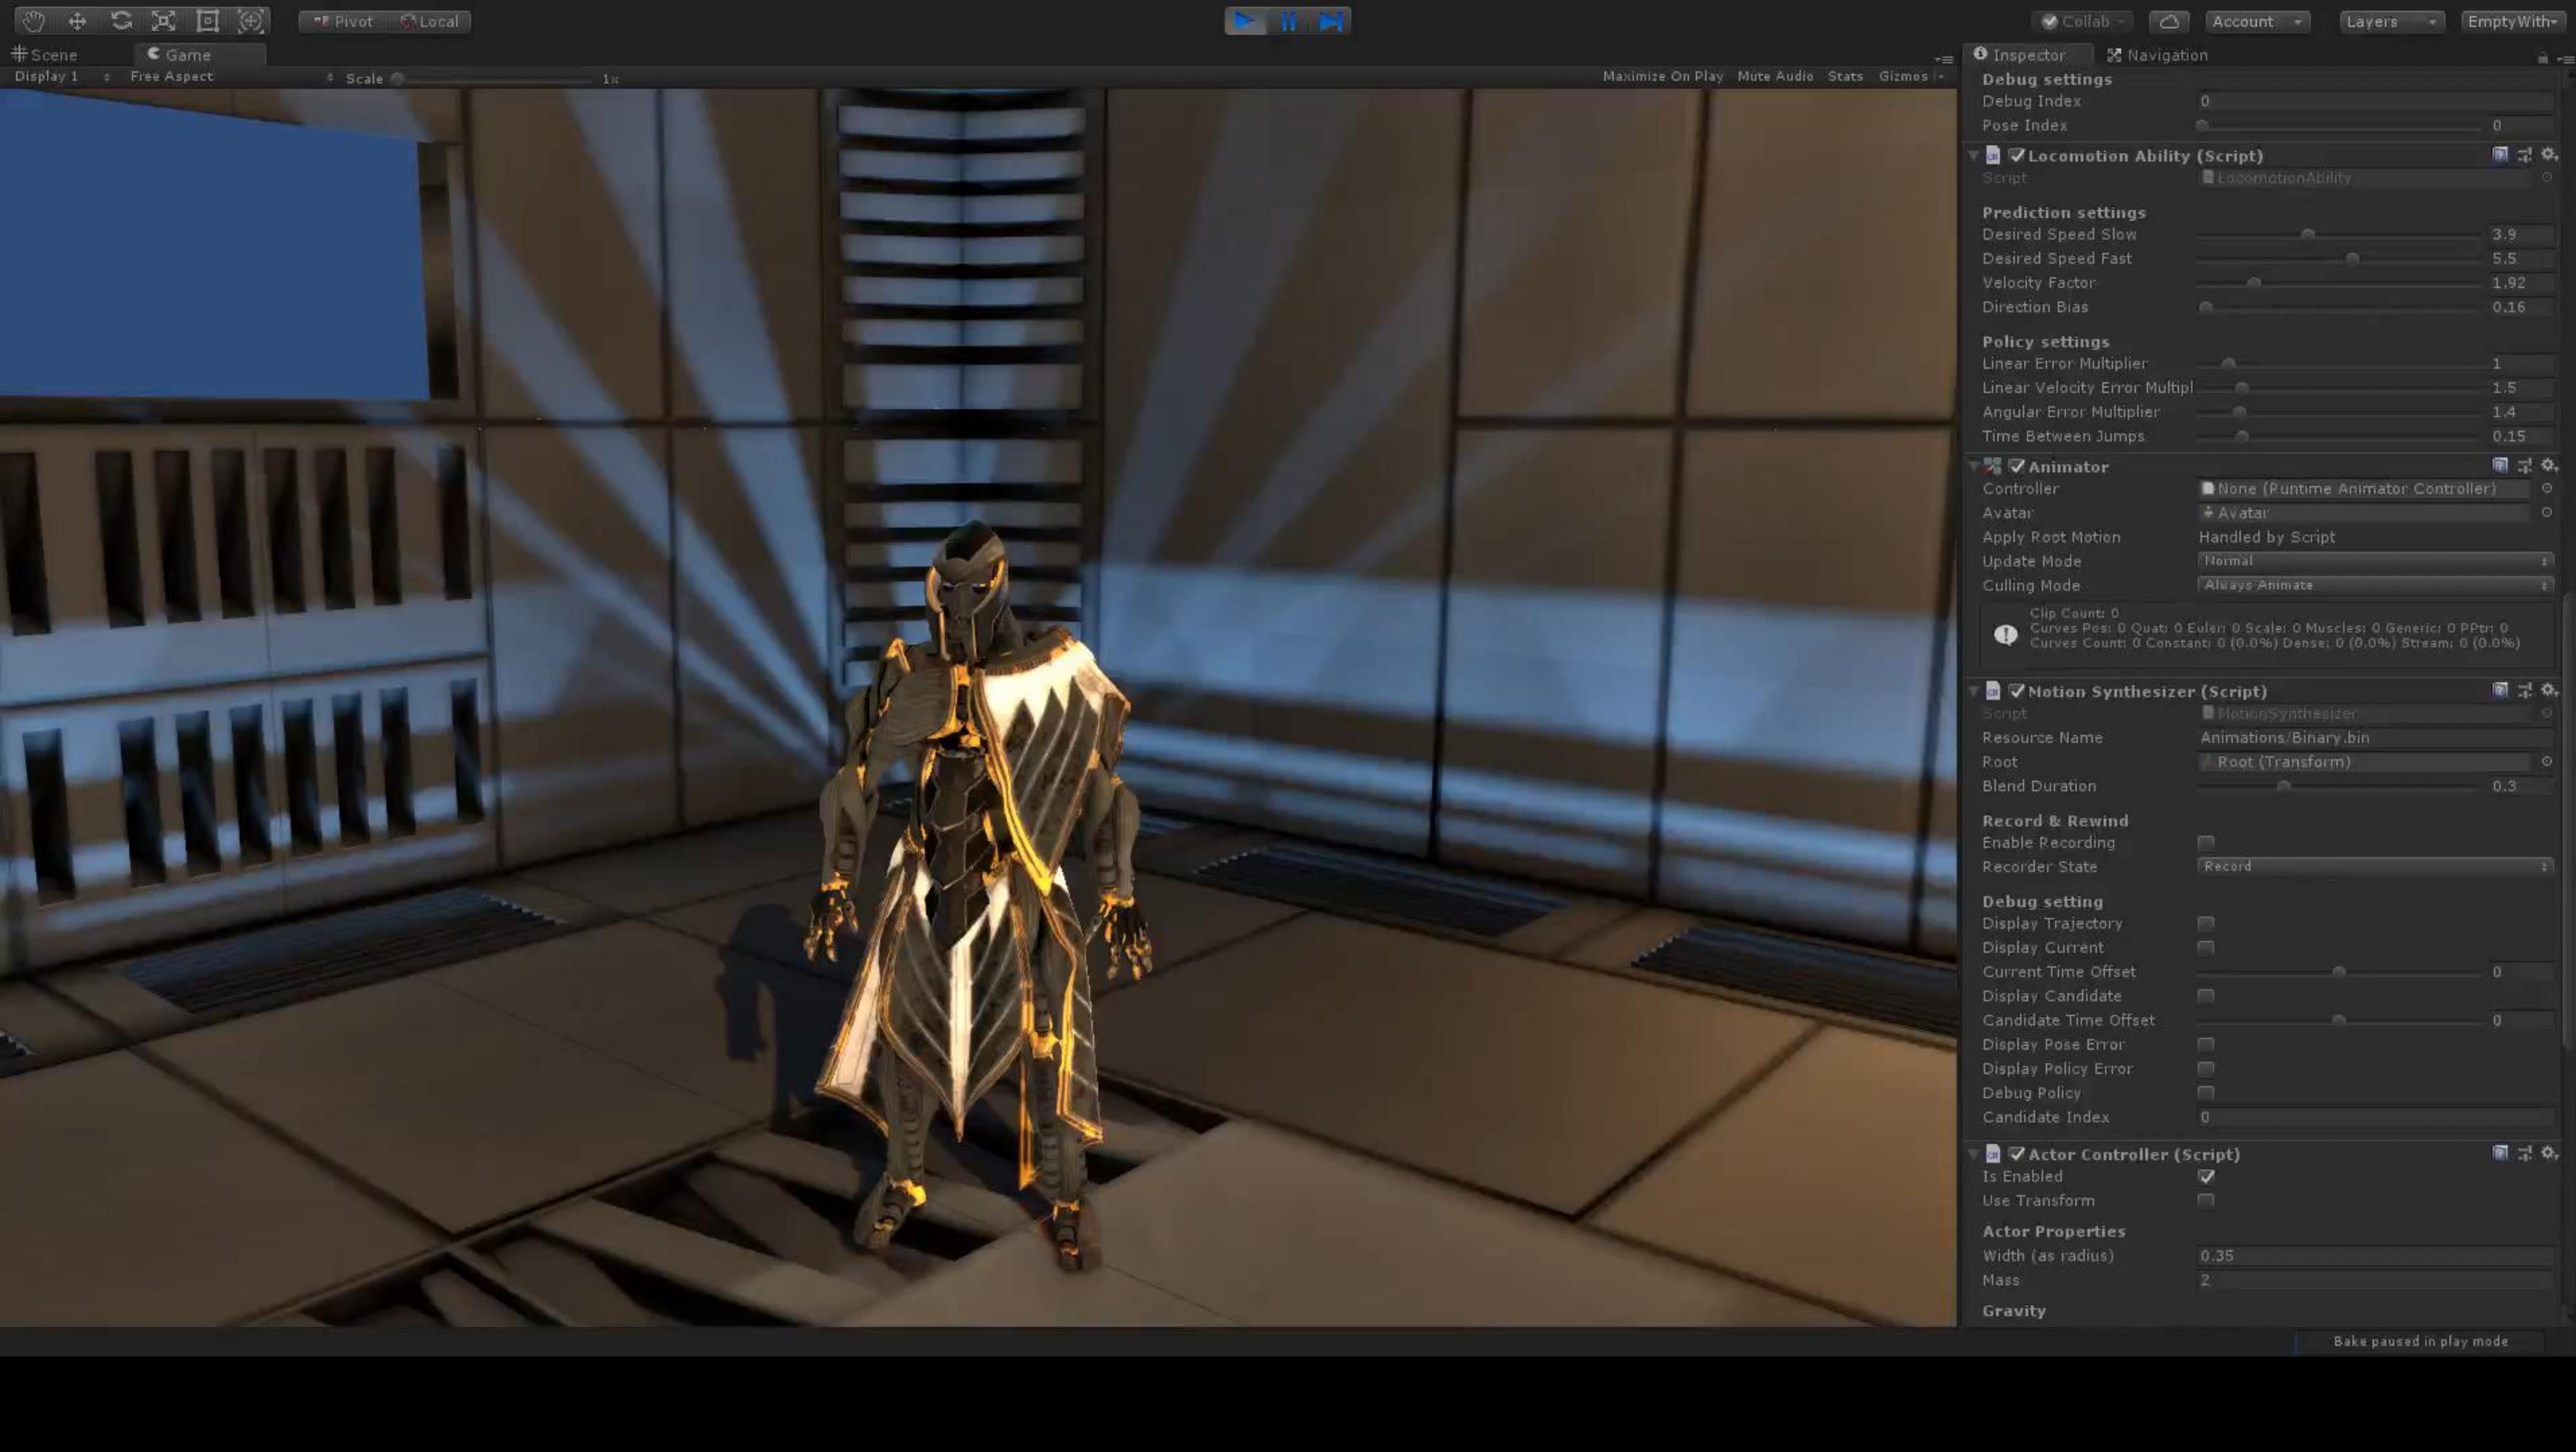




# Collision detection

During the generation of the predicted future trajectory we can perform collision detection with the environment and other characters



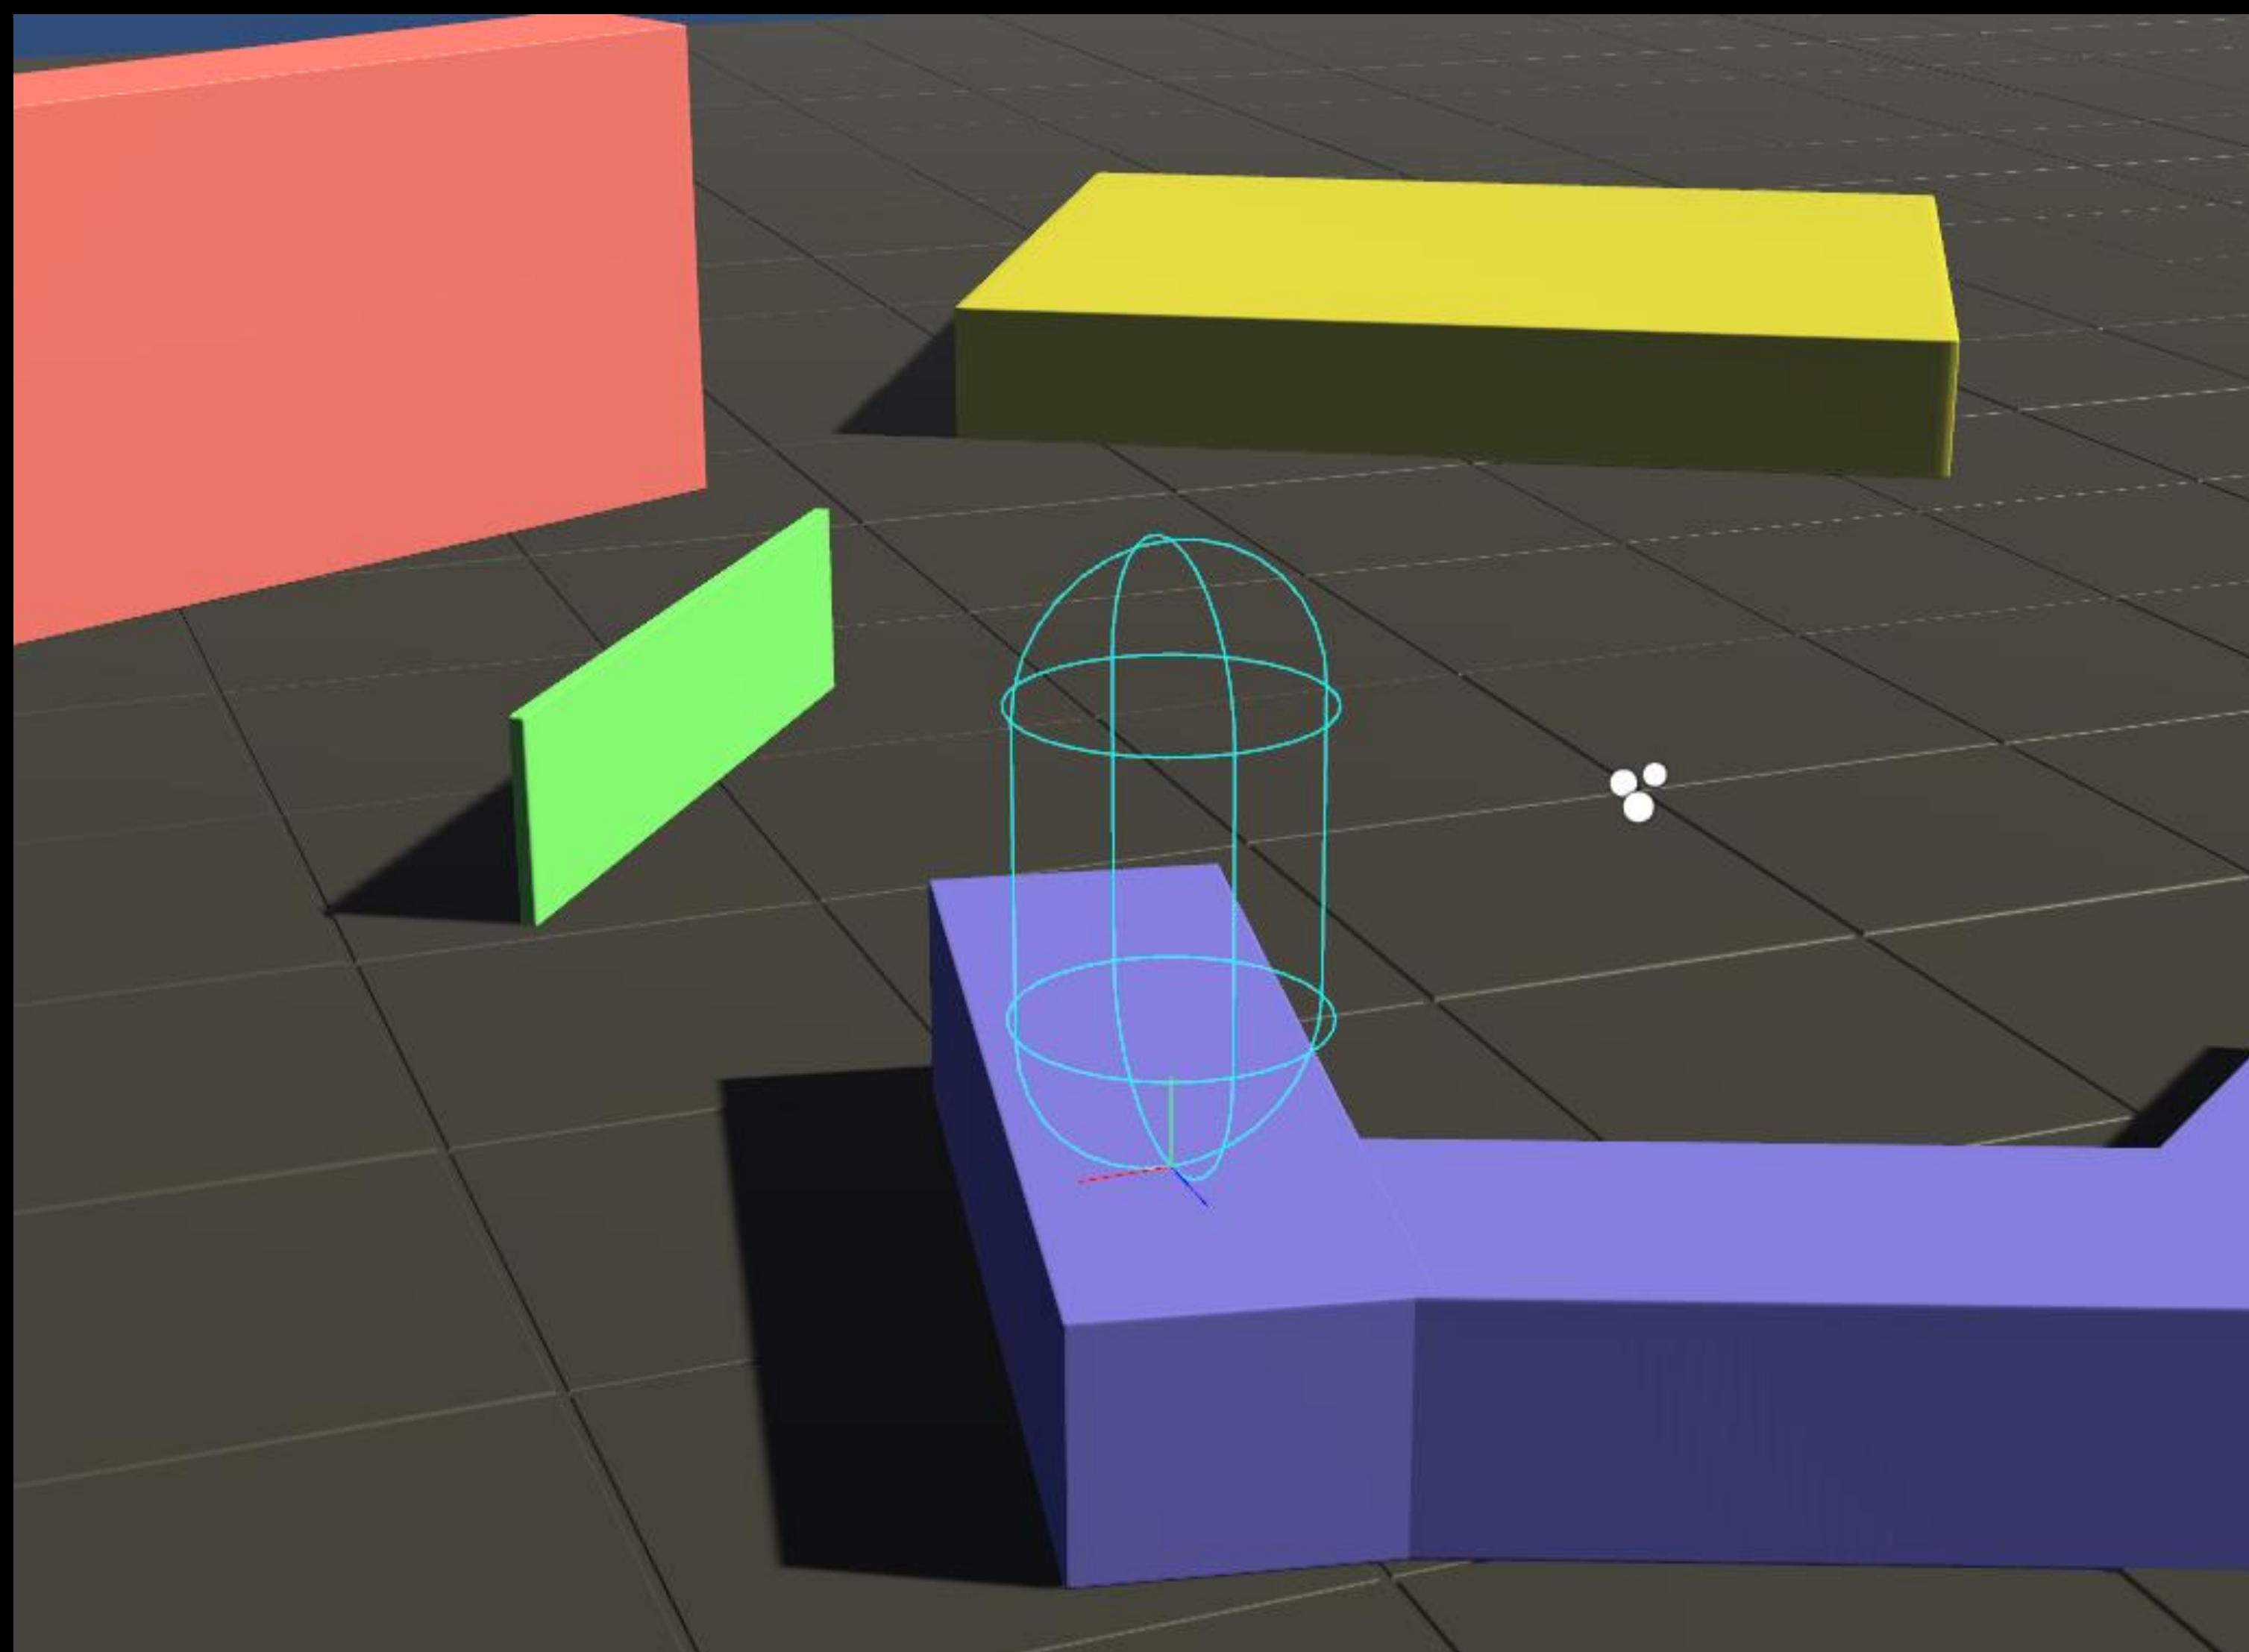


# Trajectory Prediction

During trajectory prediction we use a character controller with the ability to forward simulate the collision world

Allows us to detect collisions in advance and plan accordingly

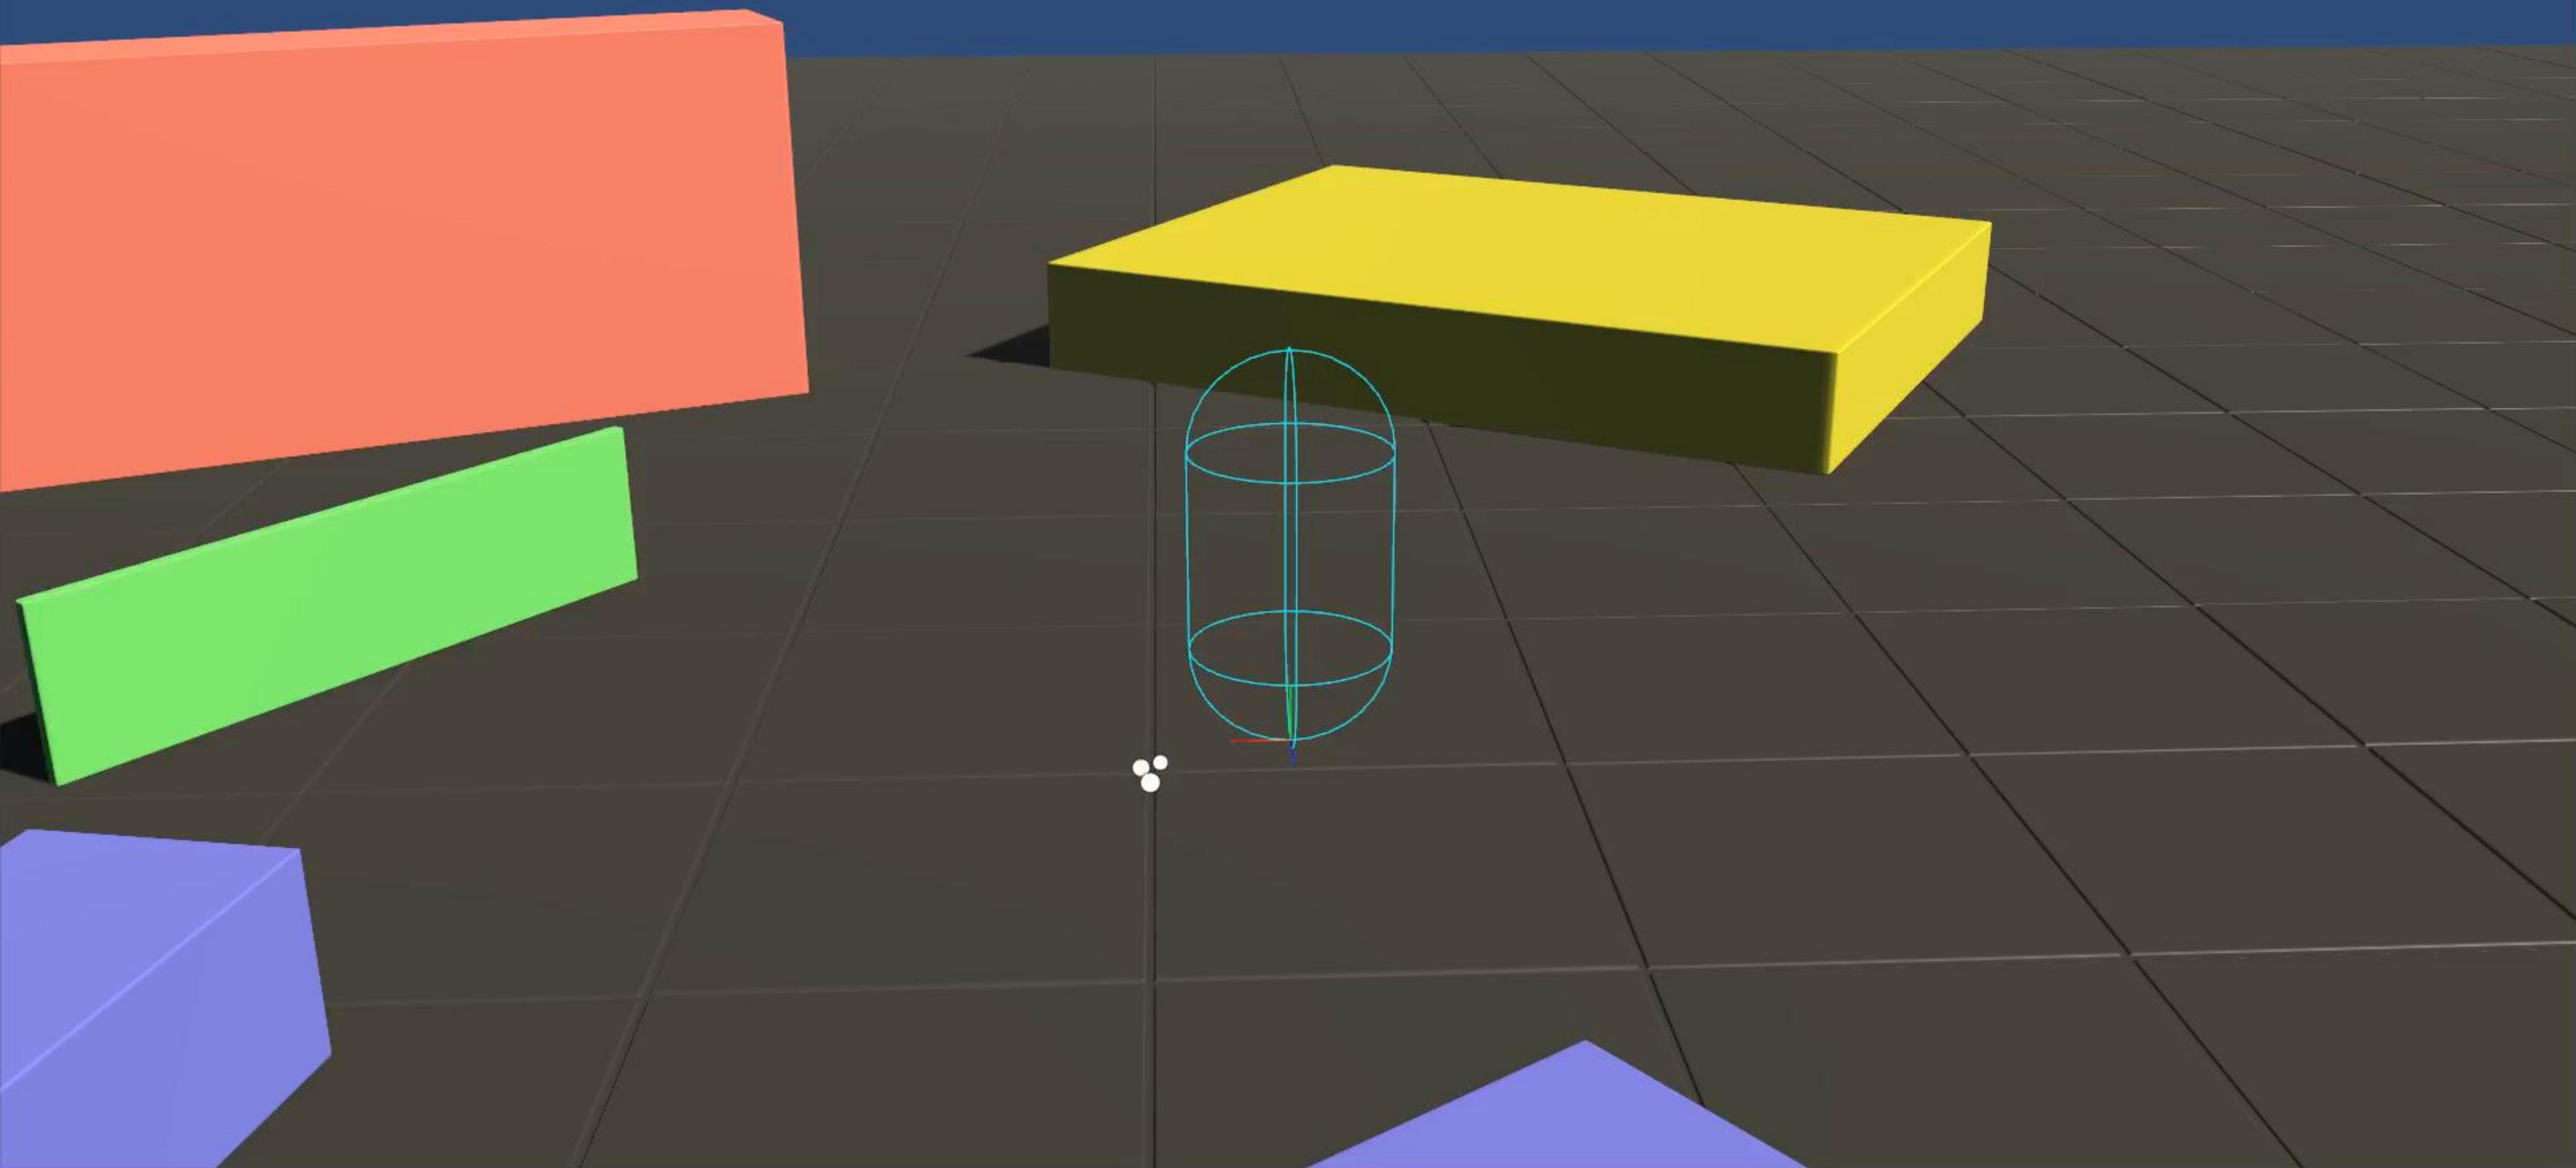
Controller has full knowledge of which objects it collides with during normal frame-by-frame processing as well as during the prediction phase and can be safely rolled back in time to return to a previous simulation step



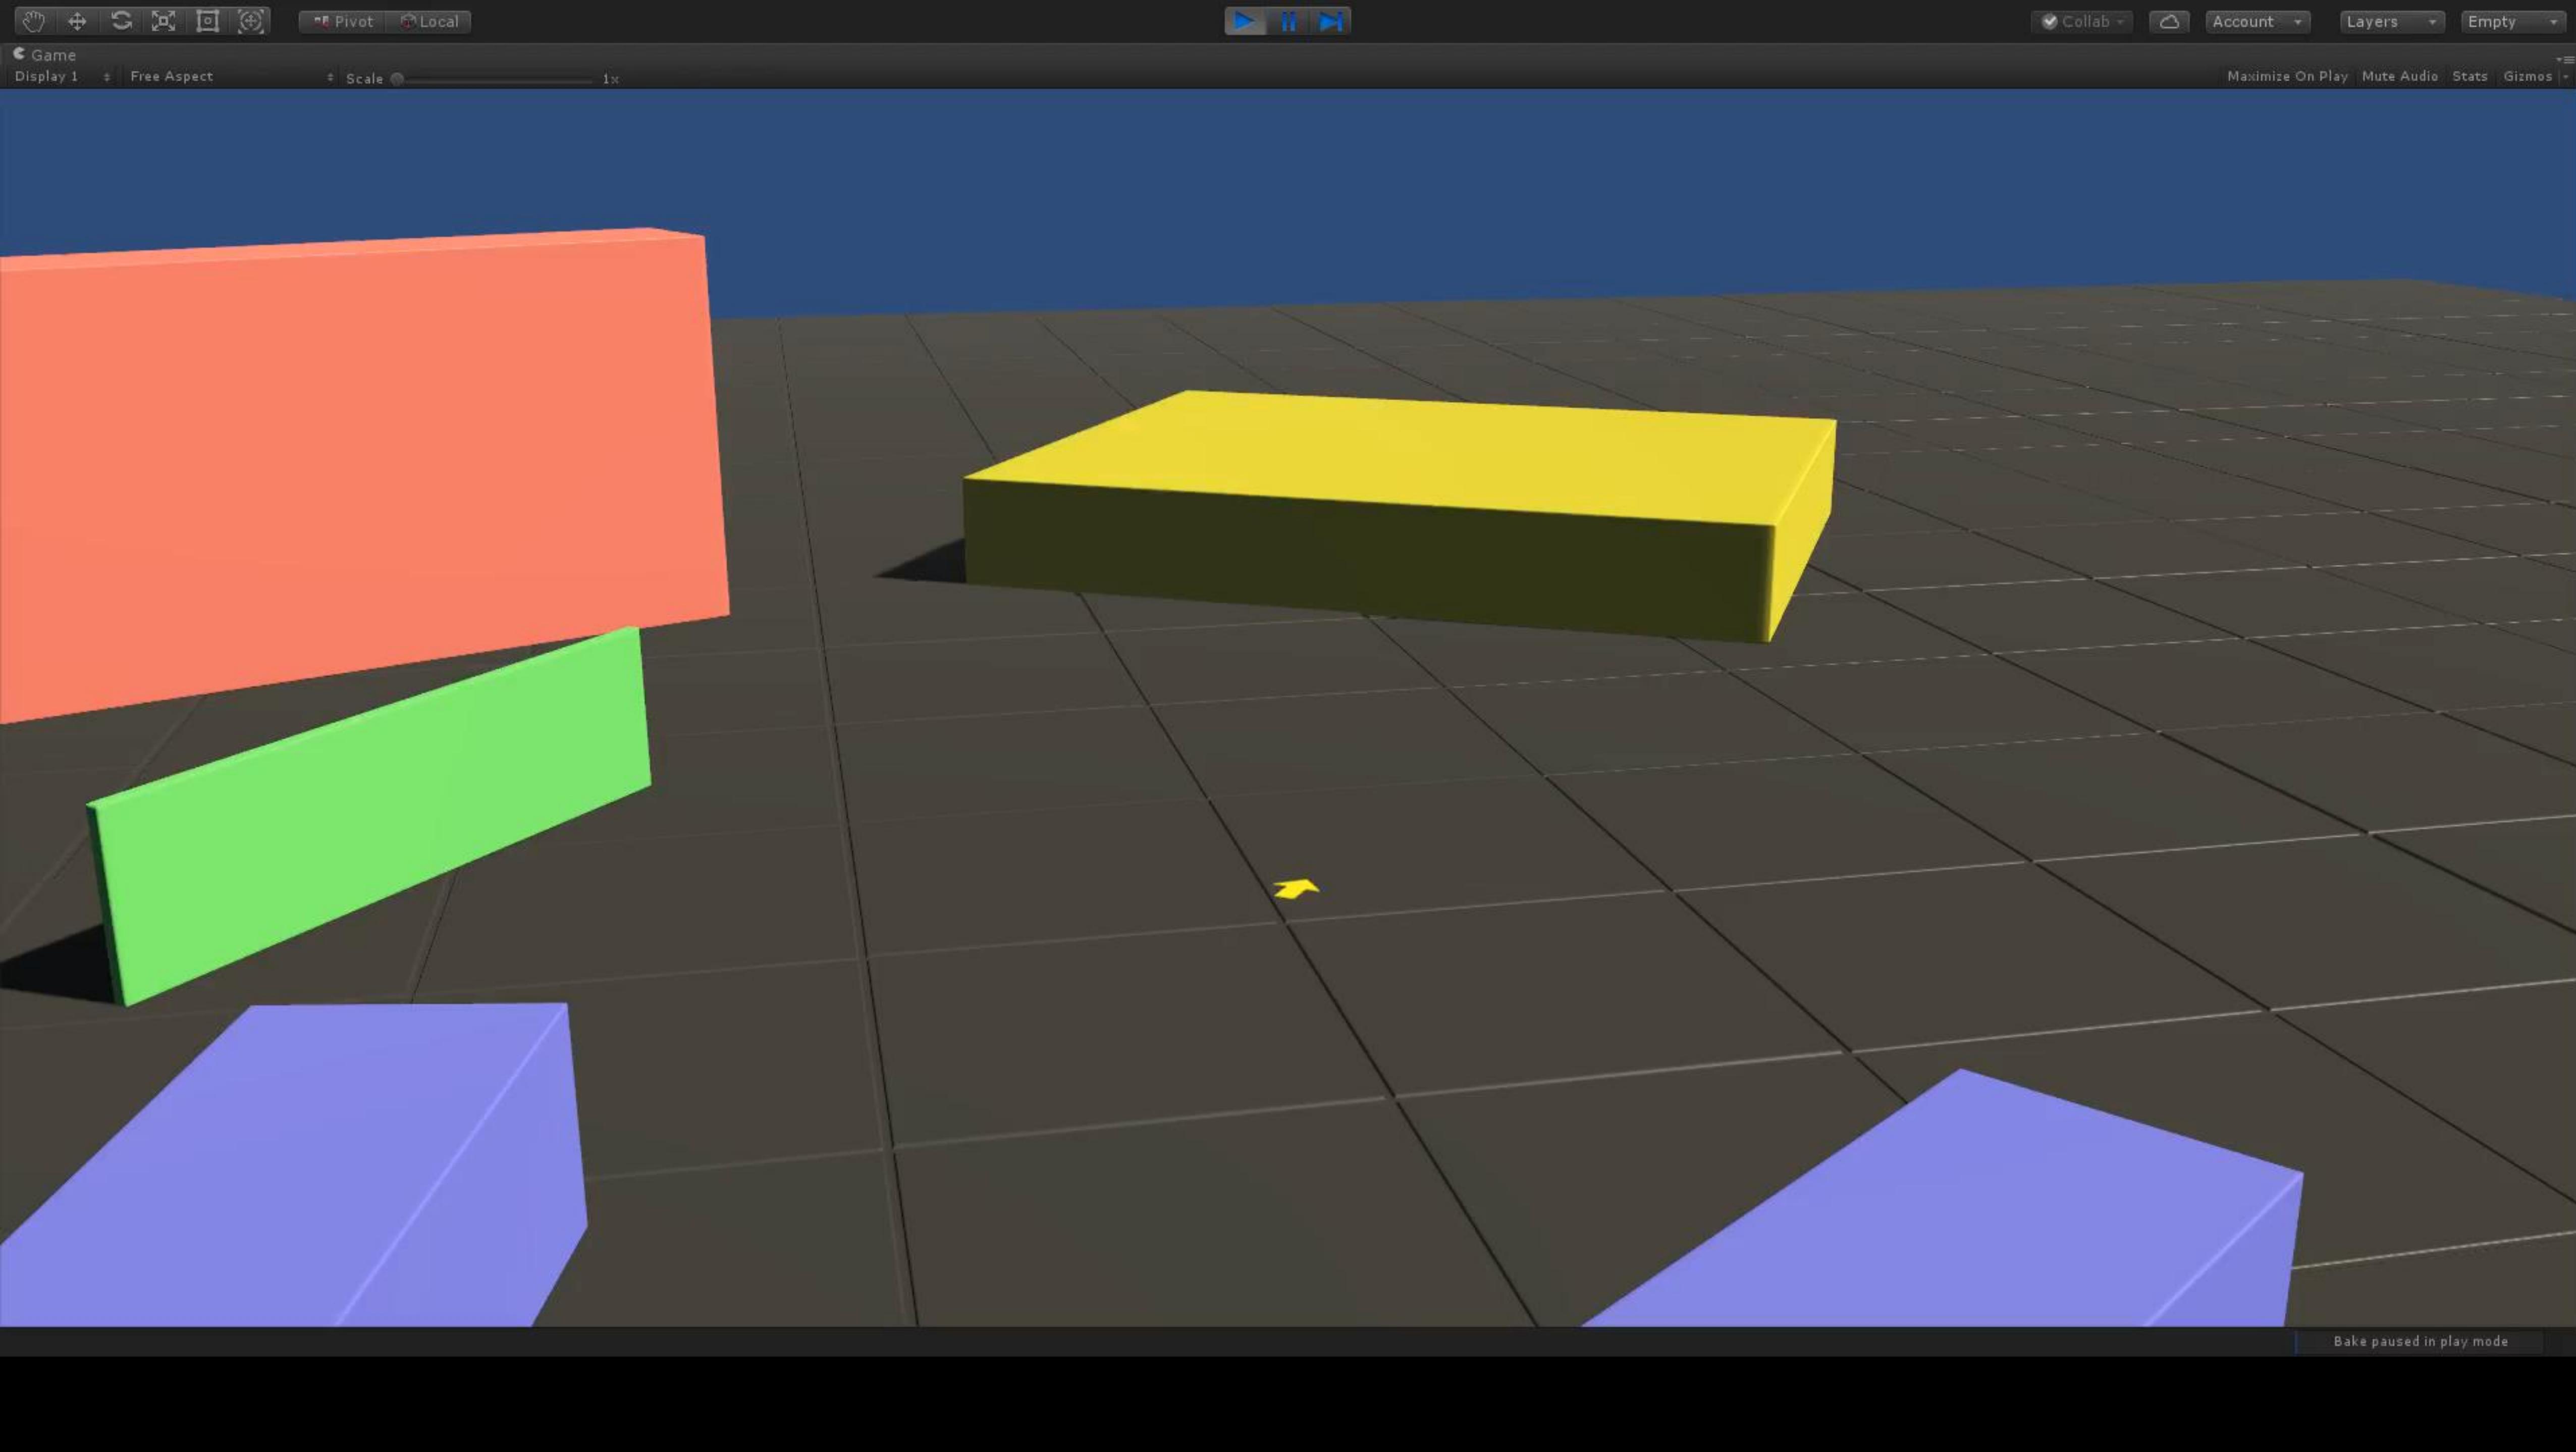


Game  
Display 1 Free Aspect Scale 1x

Maximize On Play Mute Audio Stats Gizmos

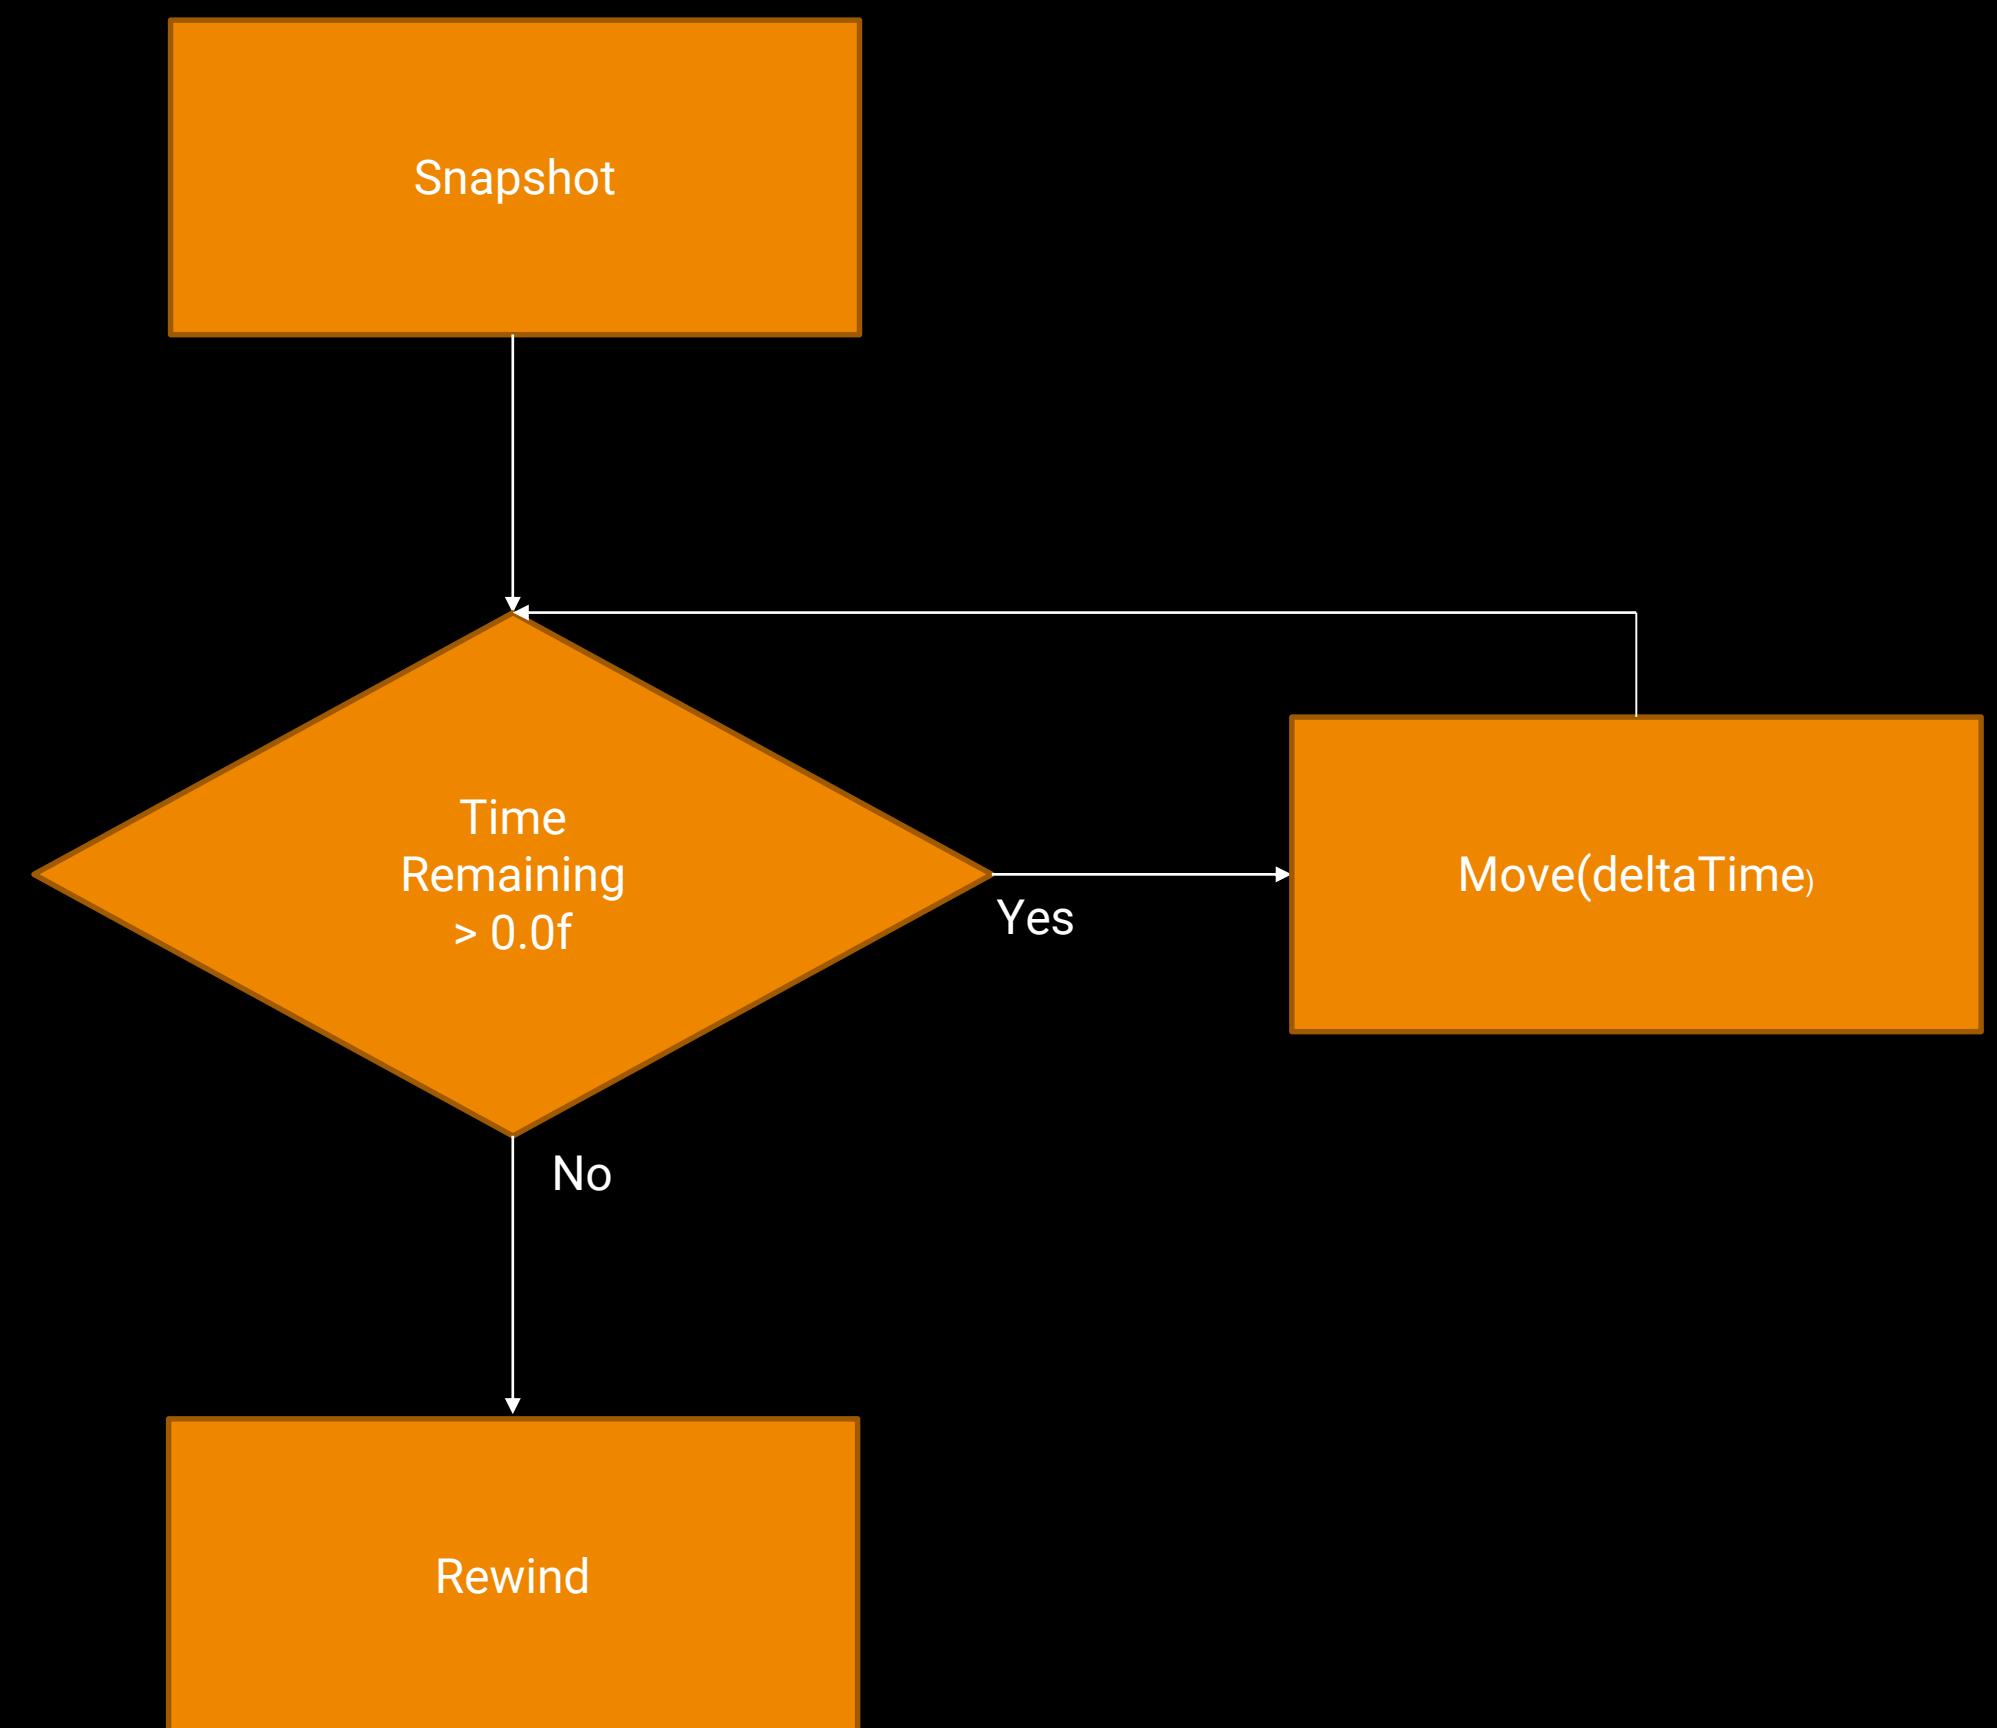


Bake paused in play mode



Bake paused in play mode

# Forward Prediction





# Parkour

# Anchors & Contacts

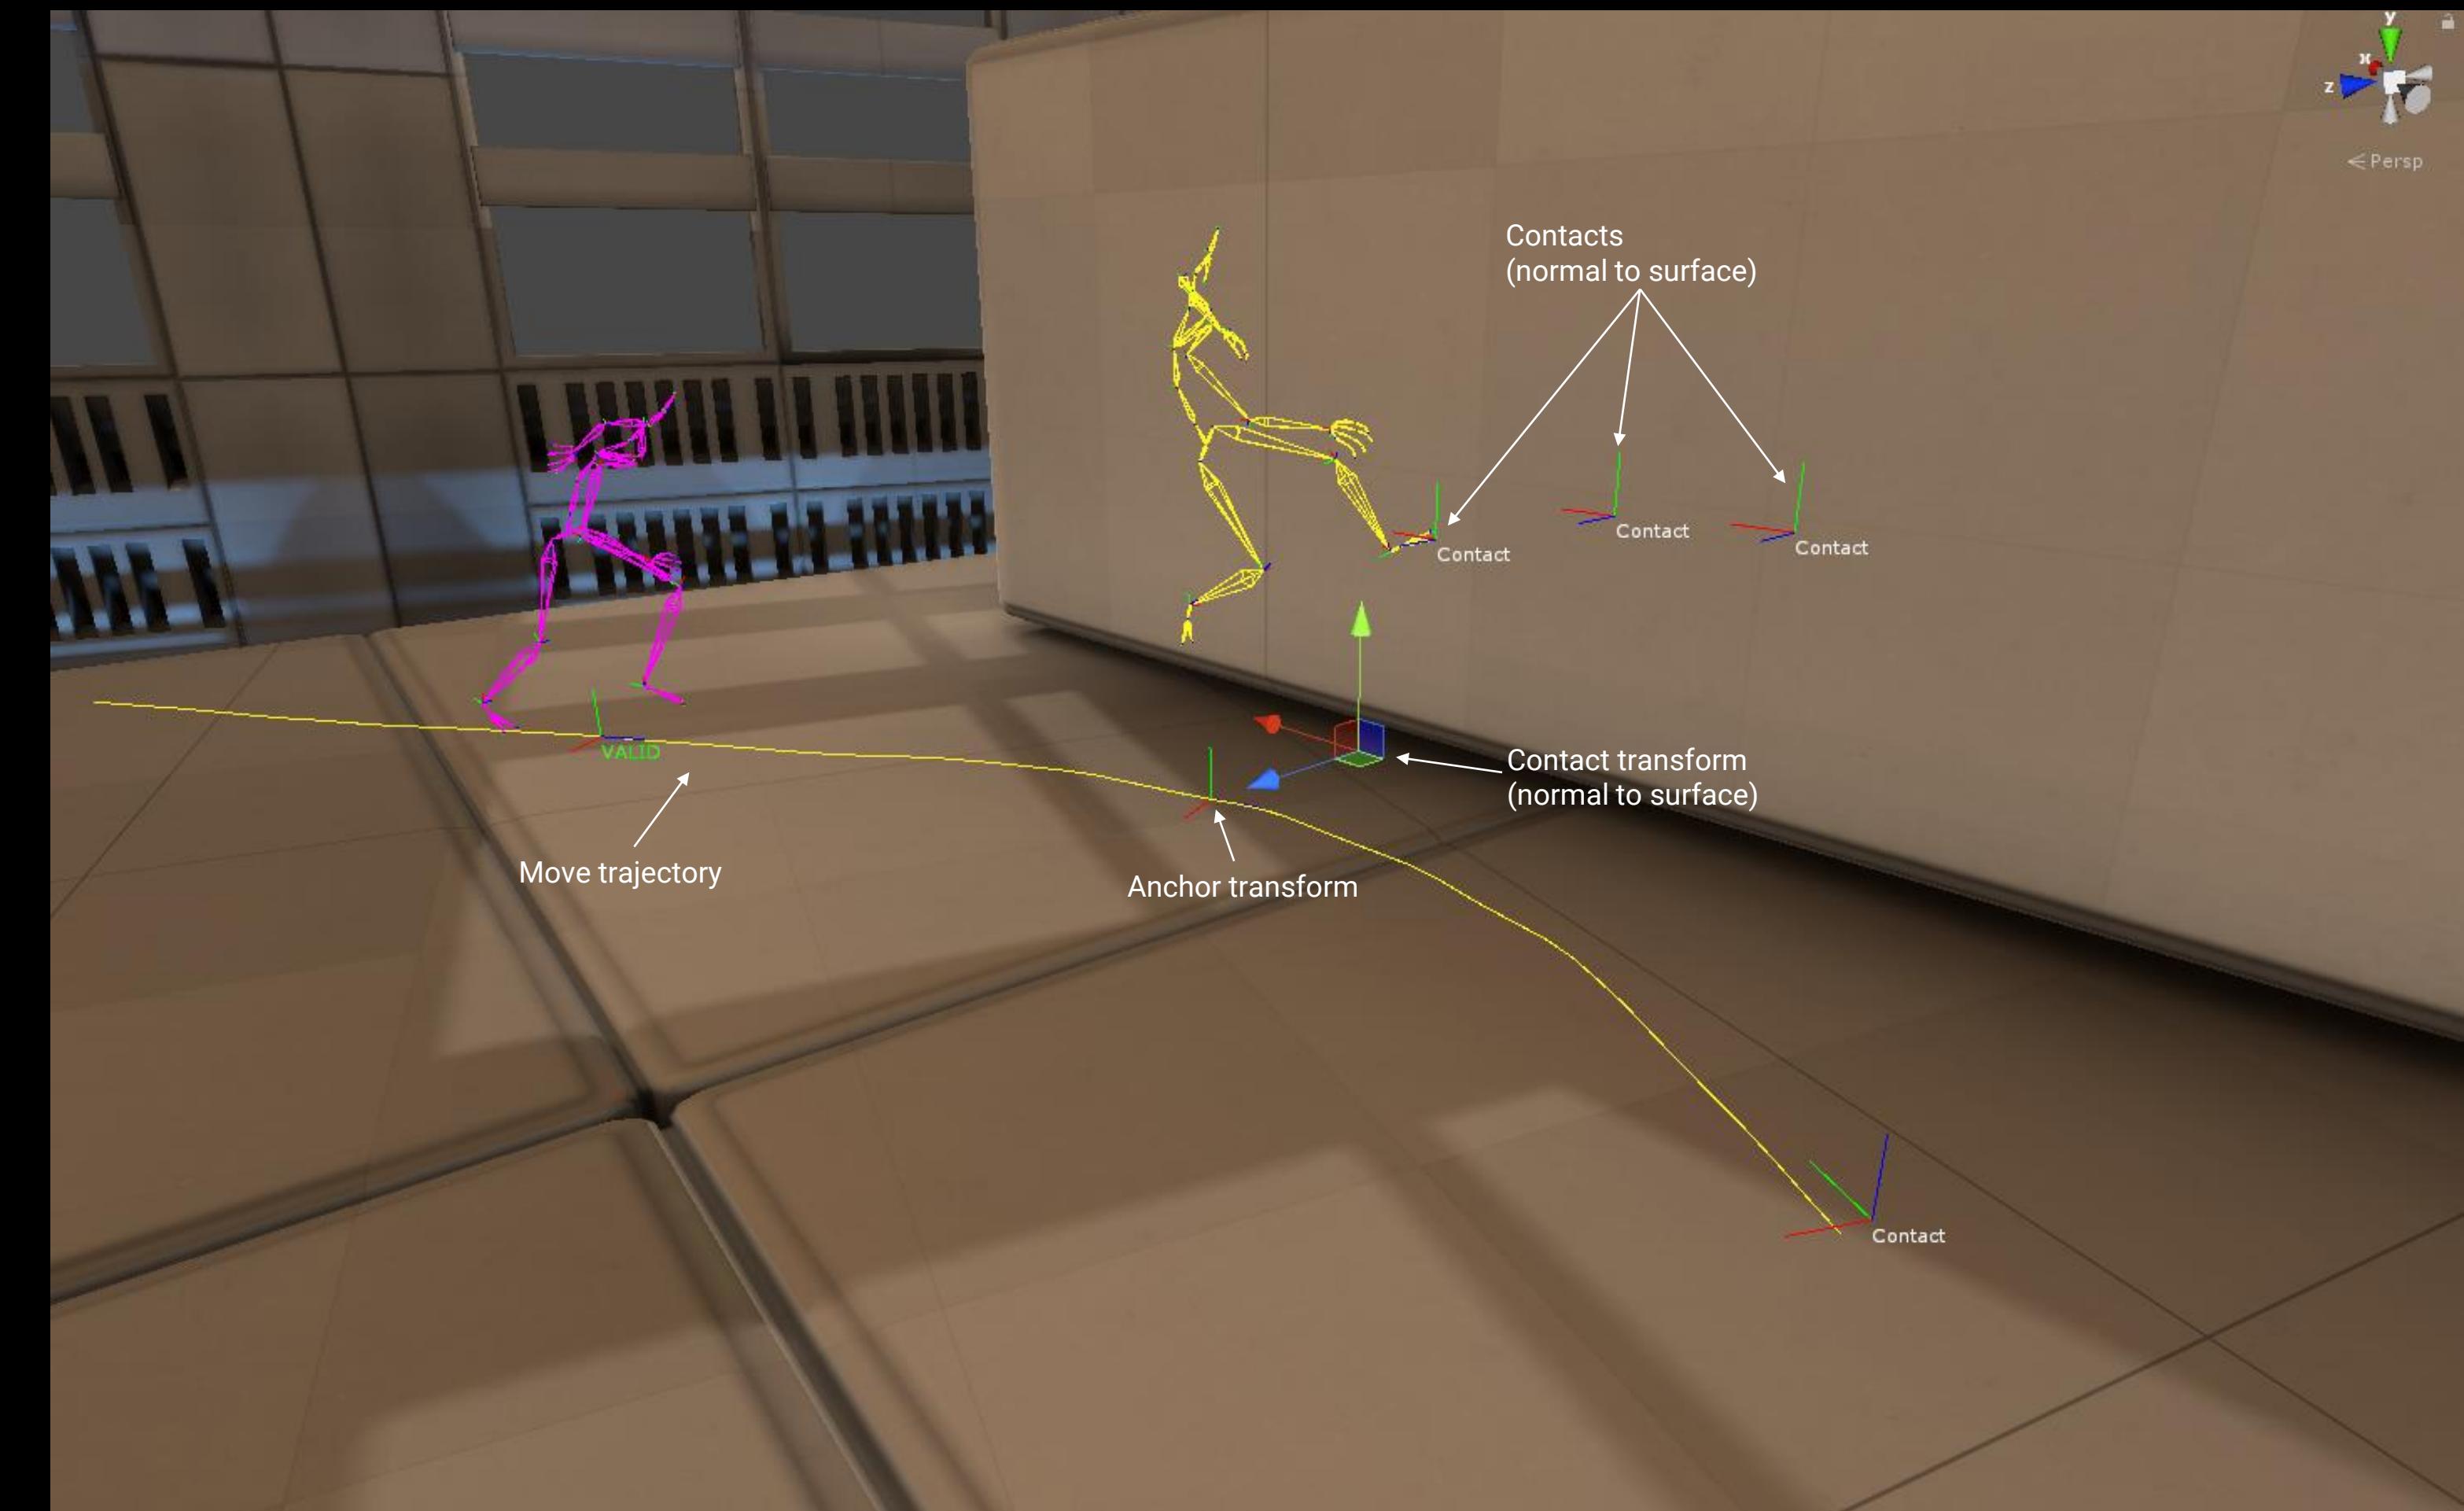
Parkour moves are designed to make precise contacts with the environment

The goal is to generate a predicted future trajectory for a specific parkour move

We use pose annotations to indicate which joint makes contact including the corresponding surface normal

We denote the transform between the contact transform and the root transform of the first contact point as “anchor transform”

Given contact transform  $\rightarrow$  Move trajectory





# Anchors & Contacts

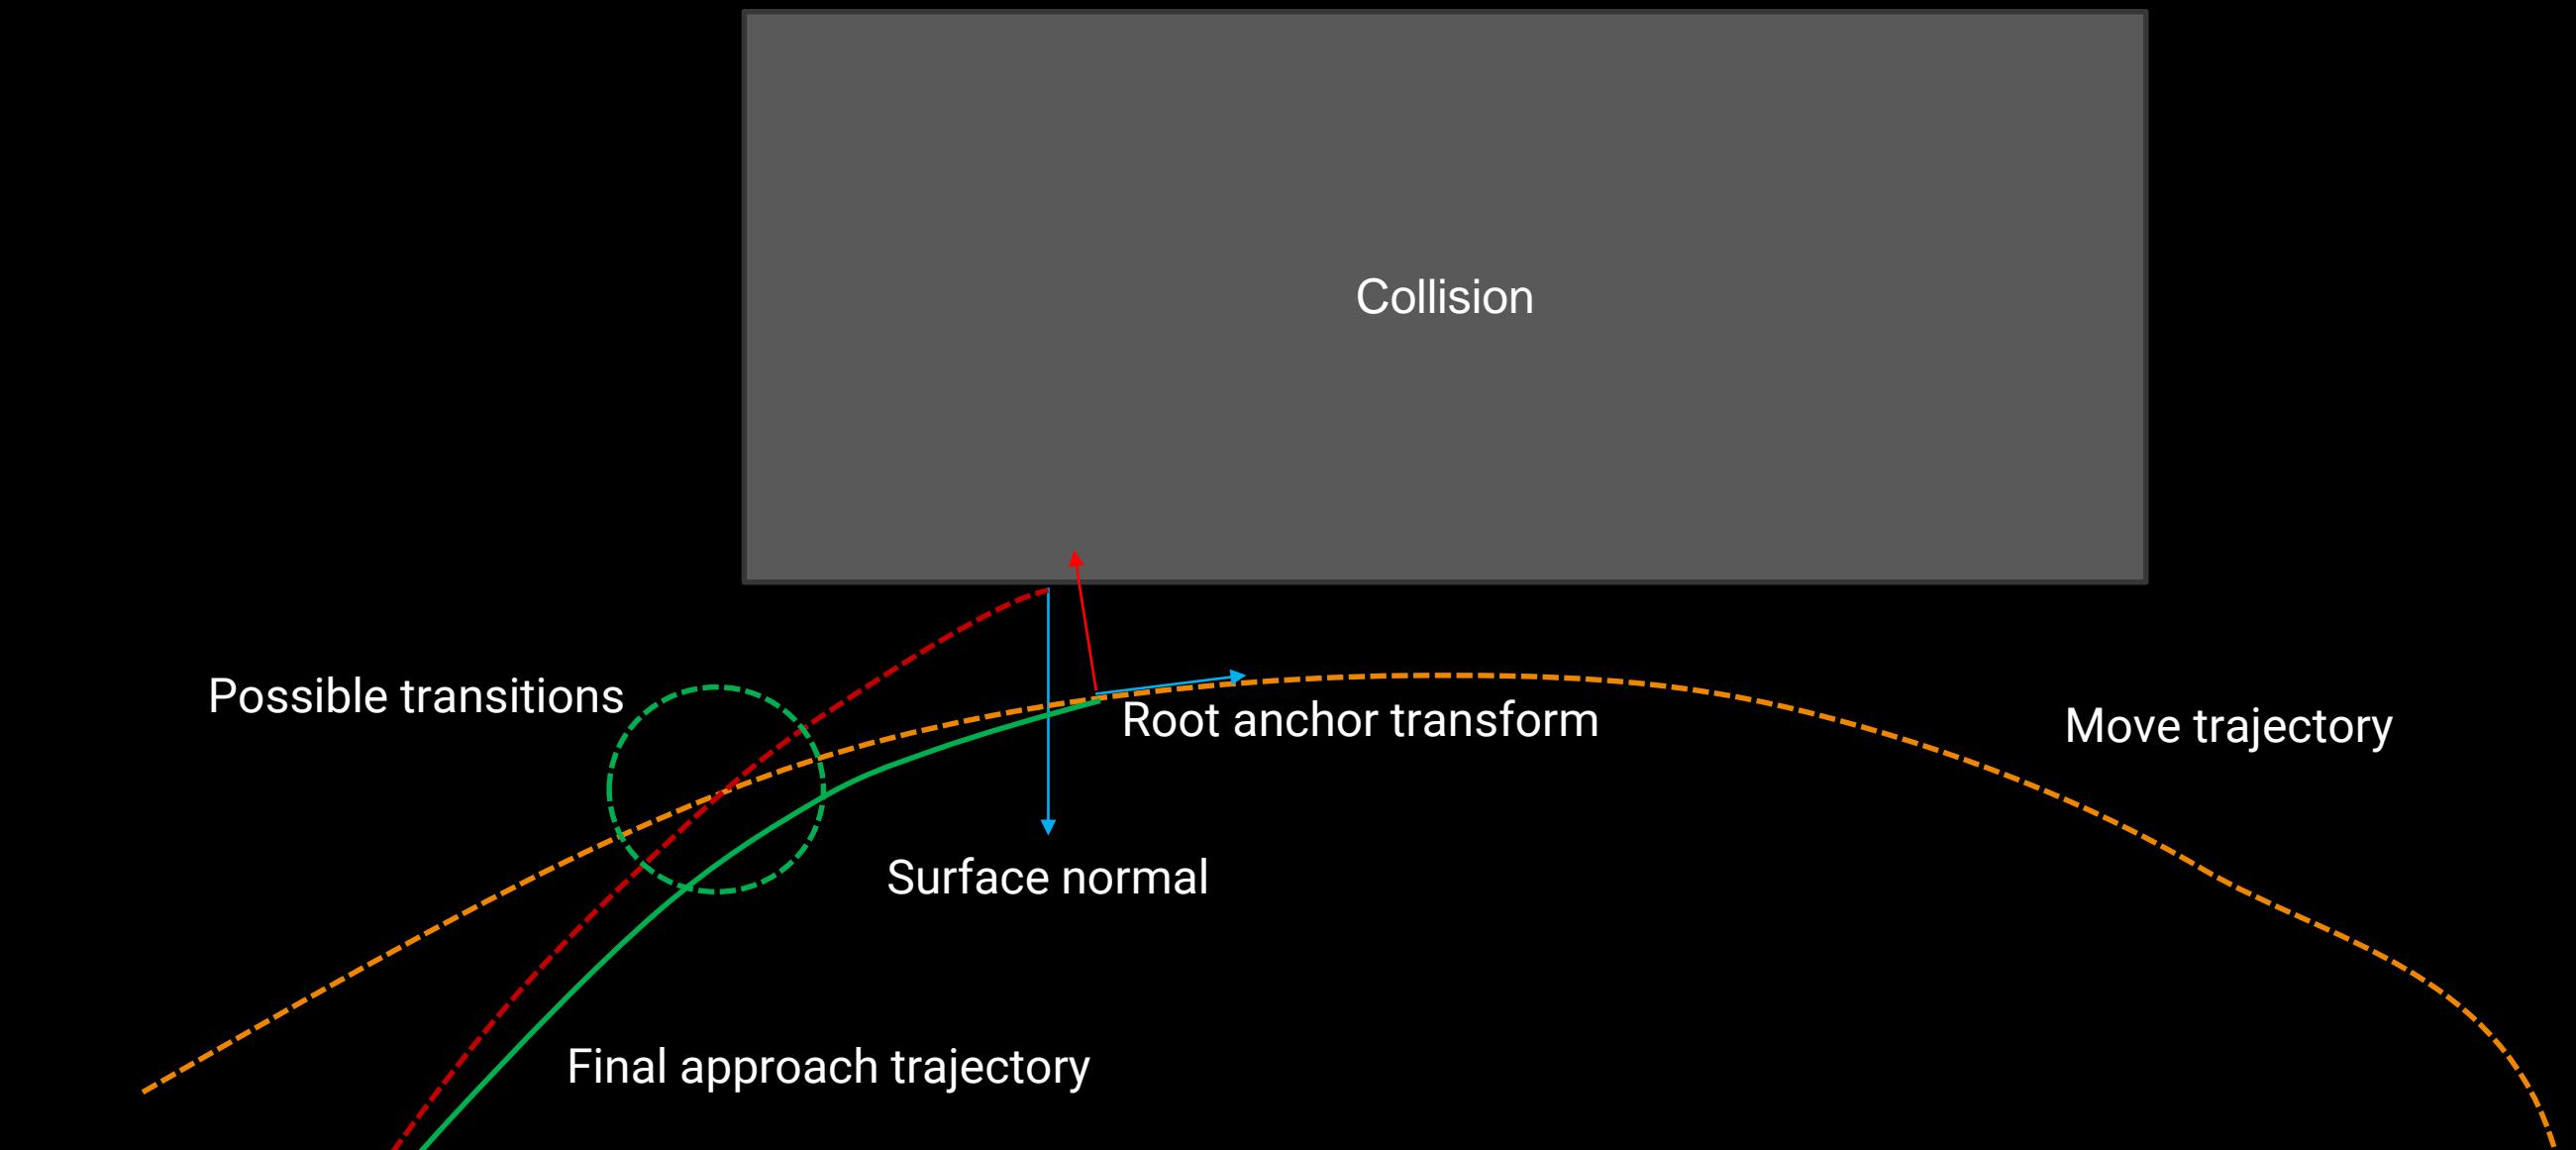
The goal is to generate a predicted future trajectory that “leads into” a specific parkour move

In case the predicted future trajectory detects a collision during the prediction phase...

We generate a “contact transform” which in turn allows us to anchor the entire move in world space

Now we can find possible transitions between the predicted future trajectory and the move trajectory

We generate a new predicted future trajectory based on the result

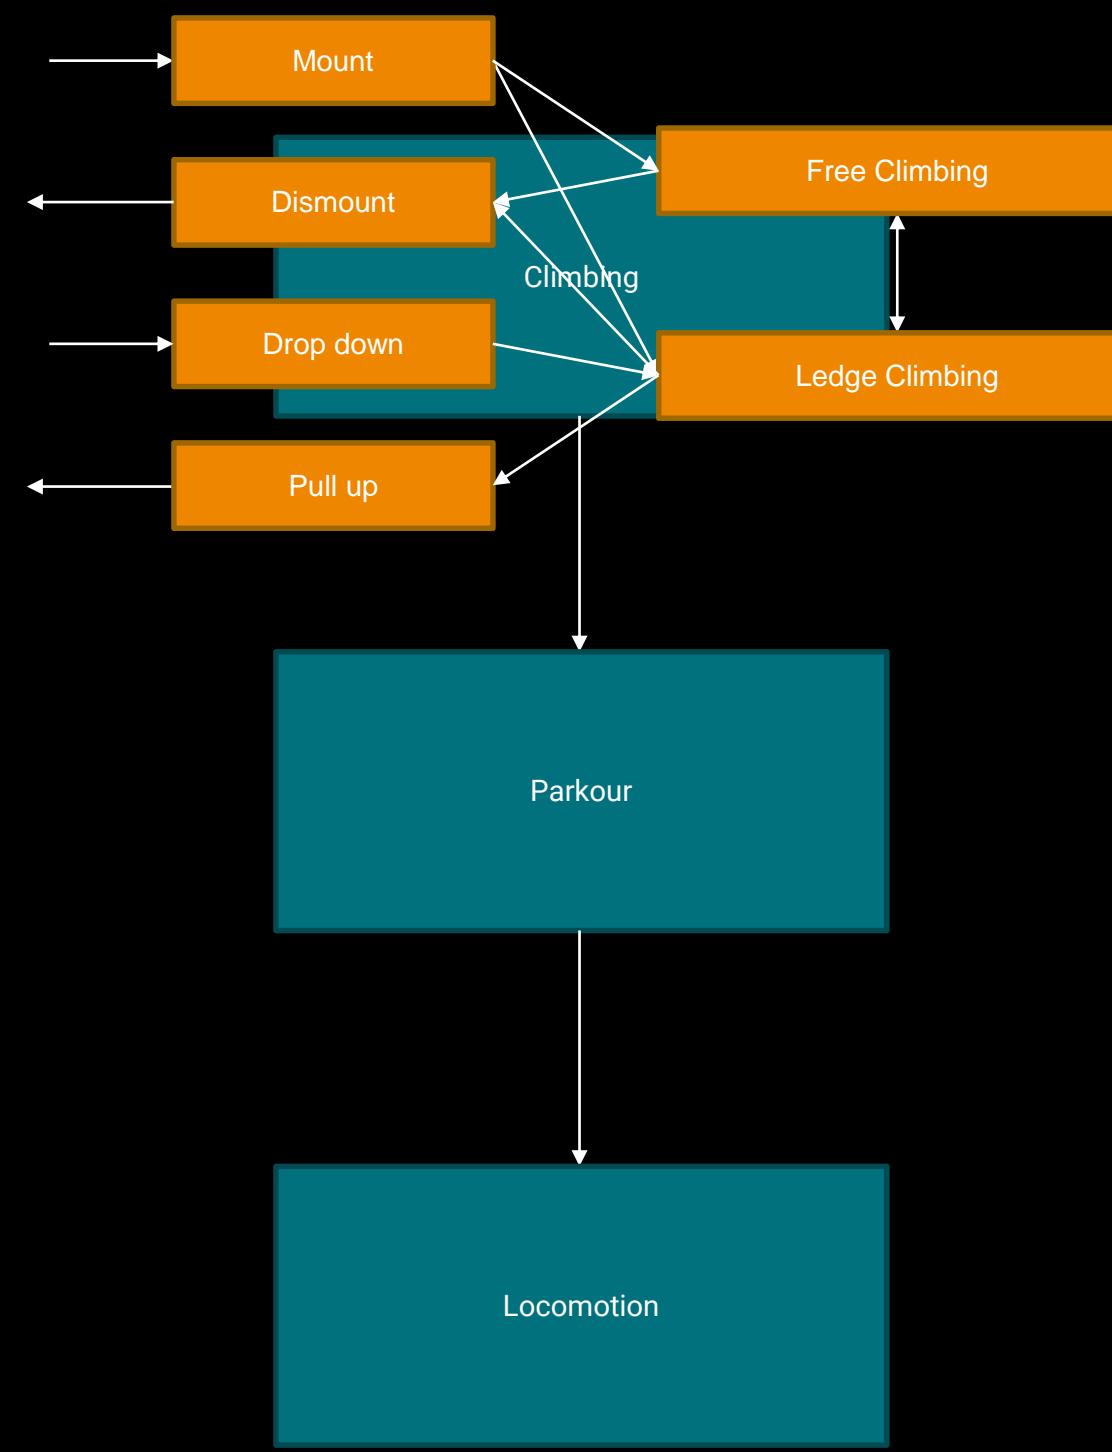






# Climbing

# Climbing States



Climbing is the most complex state

Several internal states

Transitions

Multiple movement types

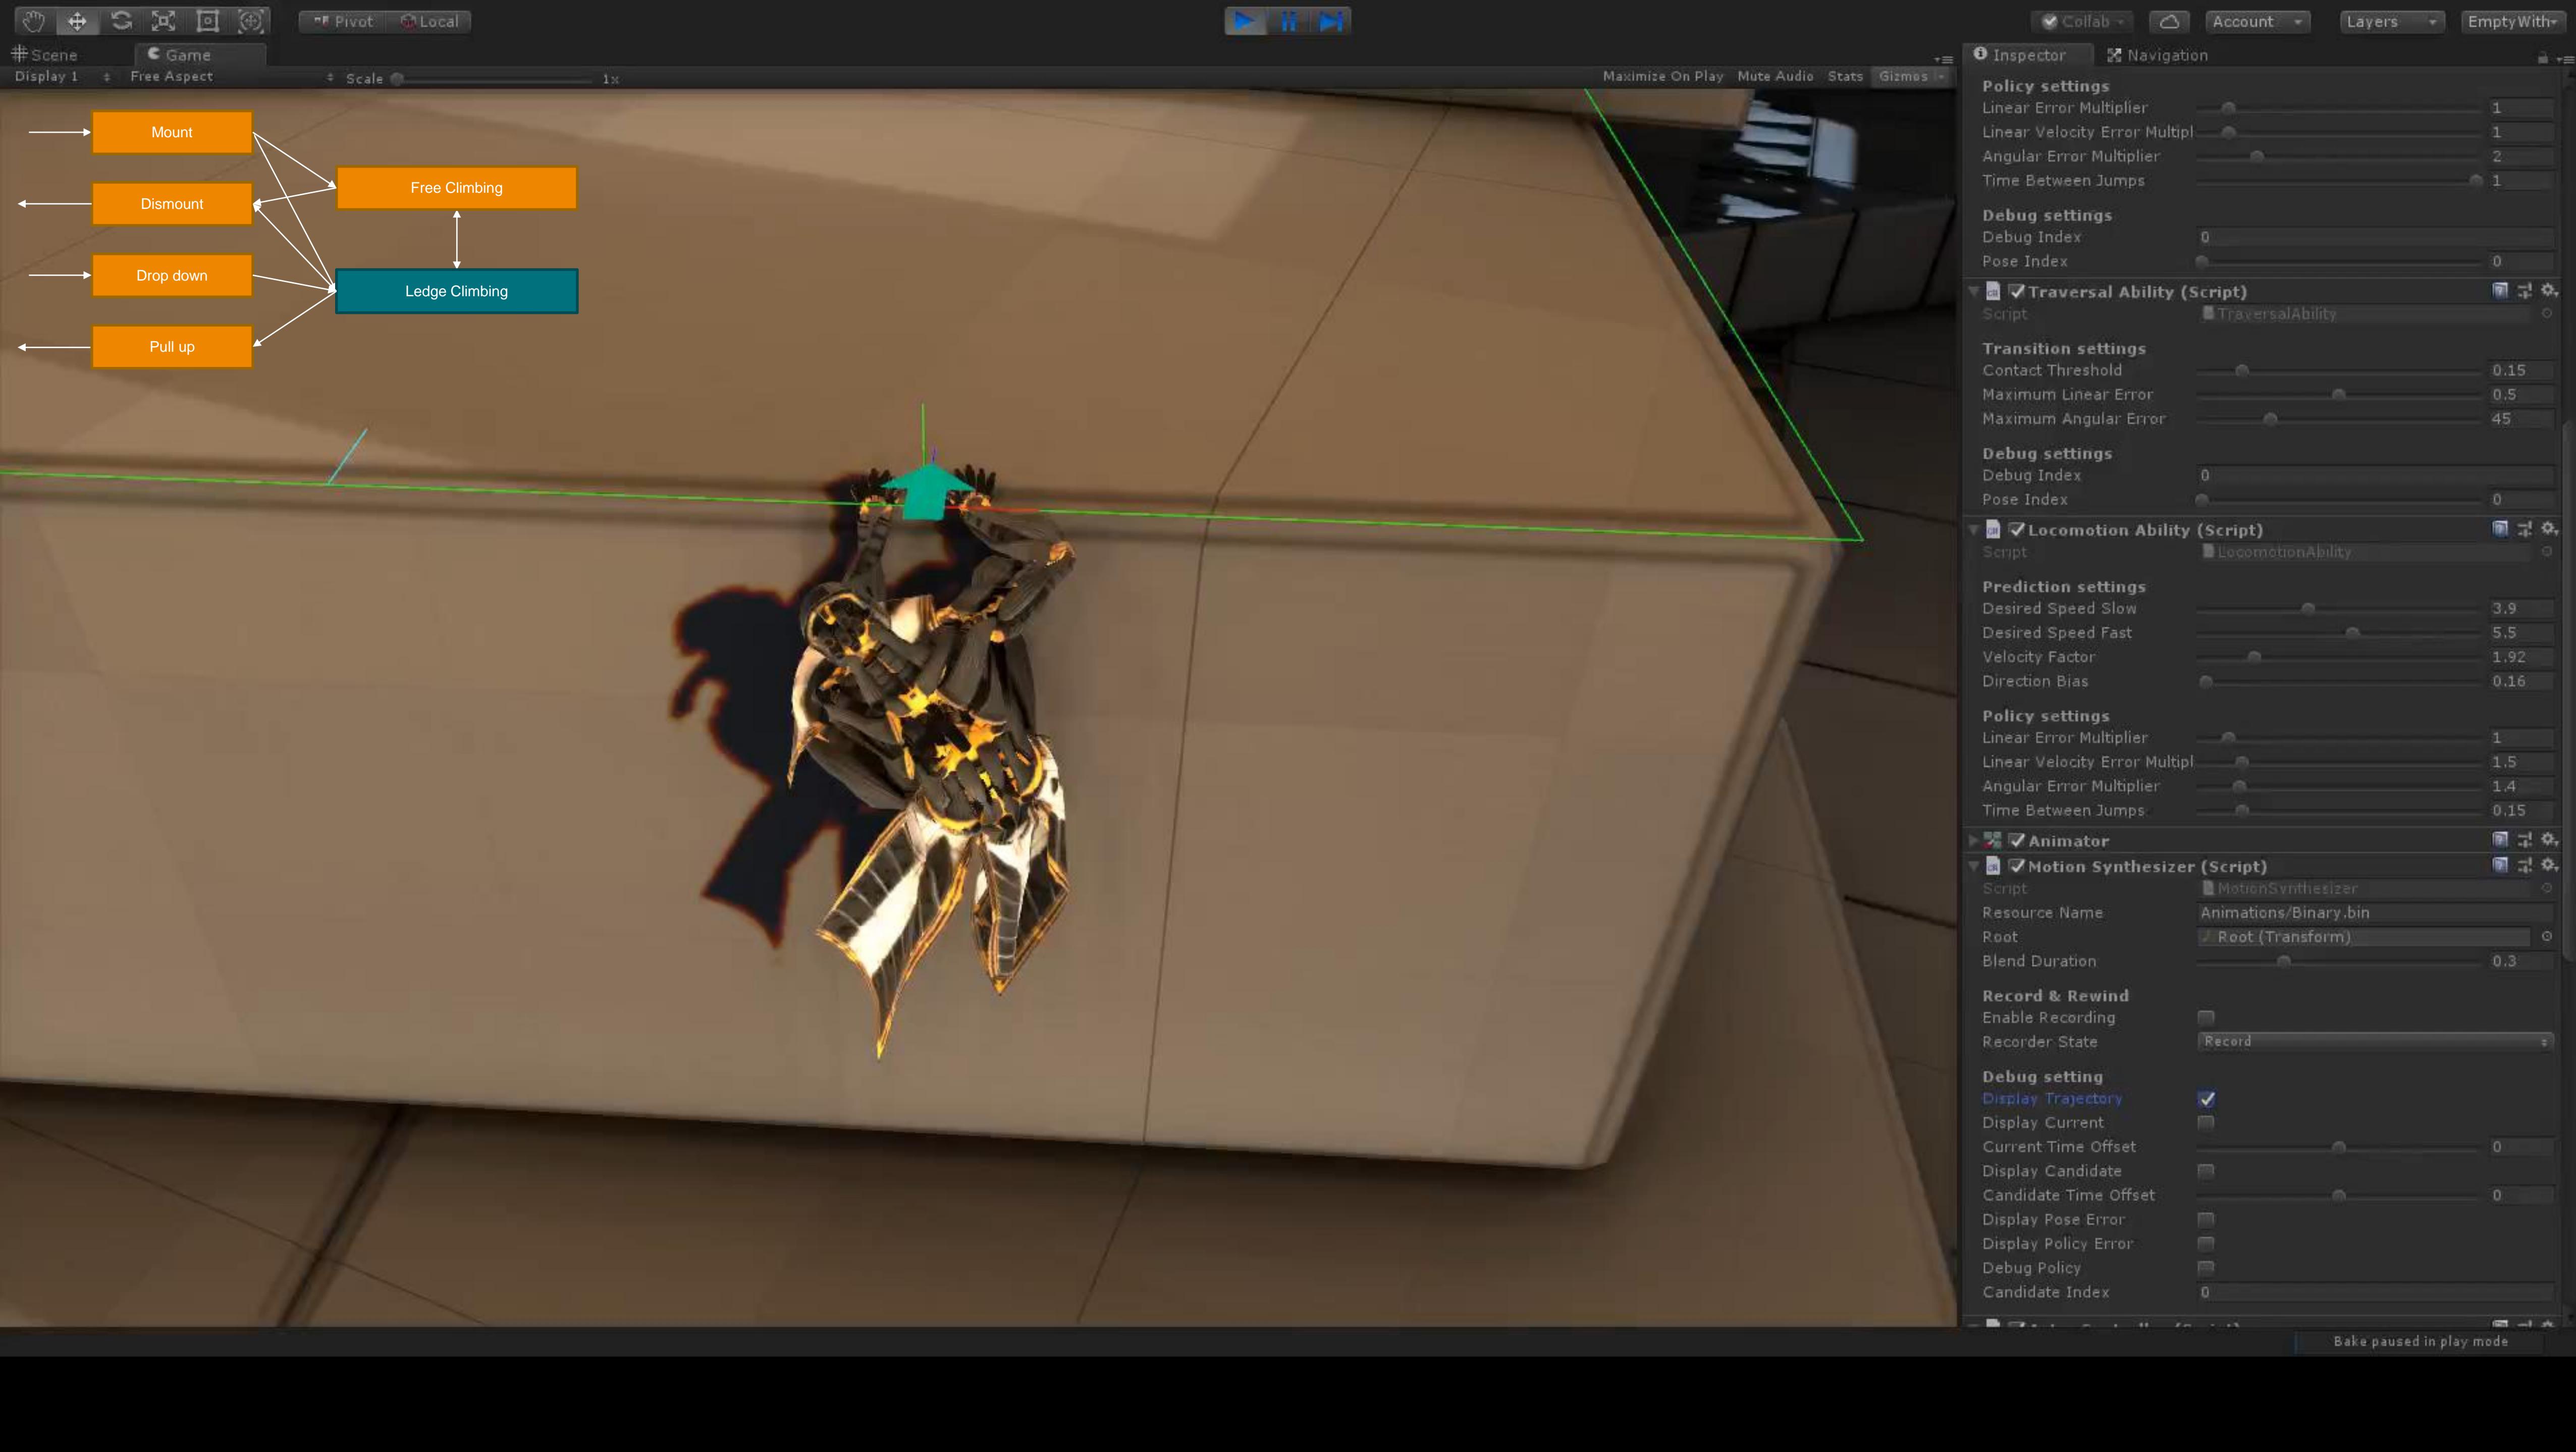














# Ledge Climbing Model

Ledge geometry gets constructed on-the-fly

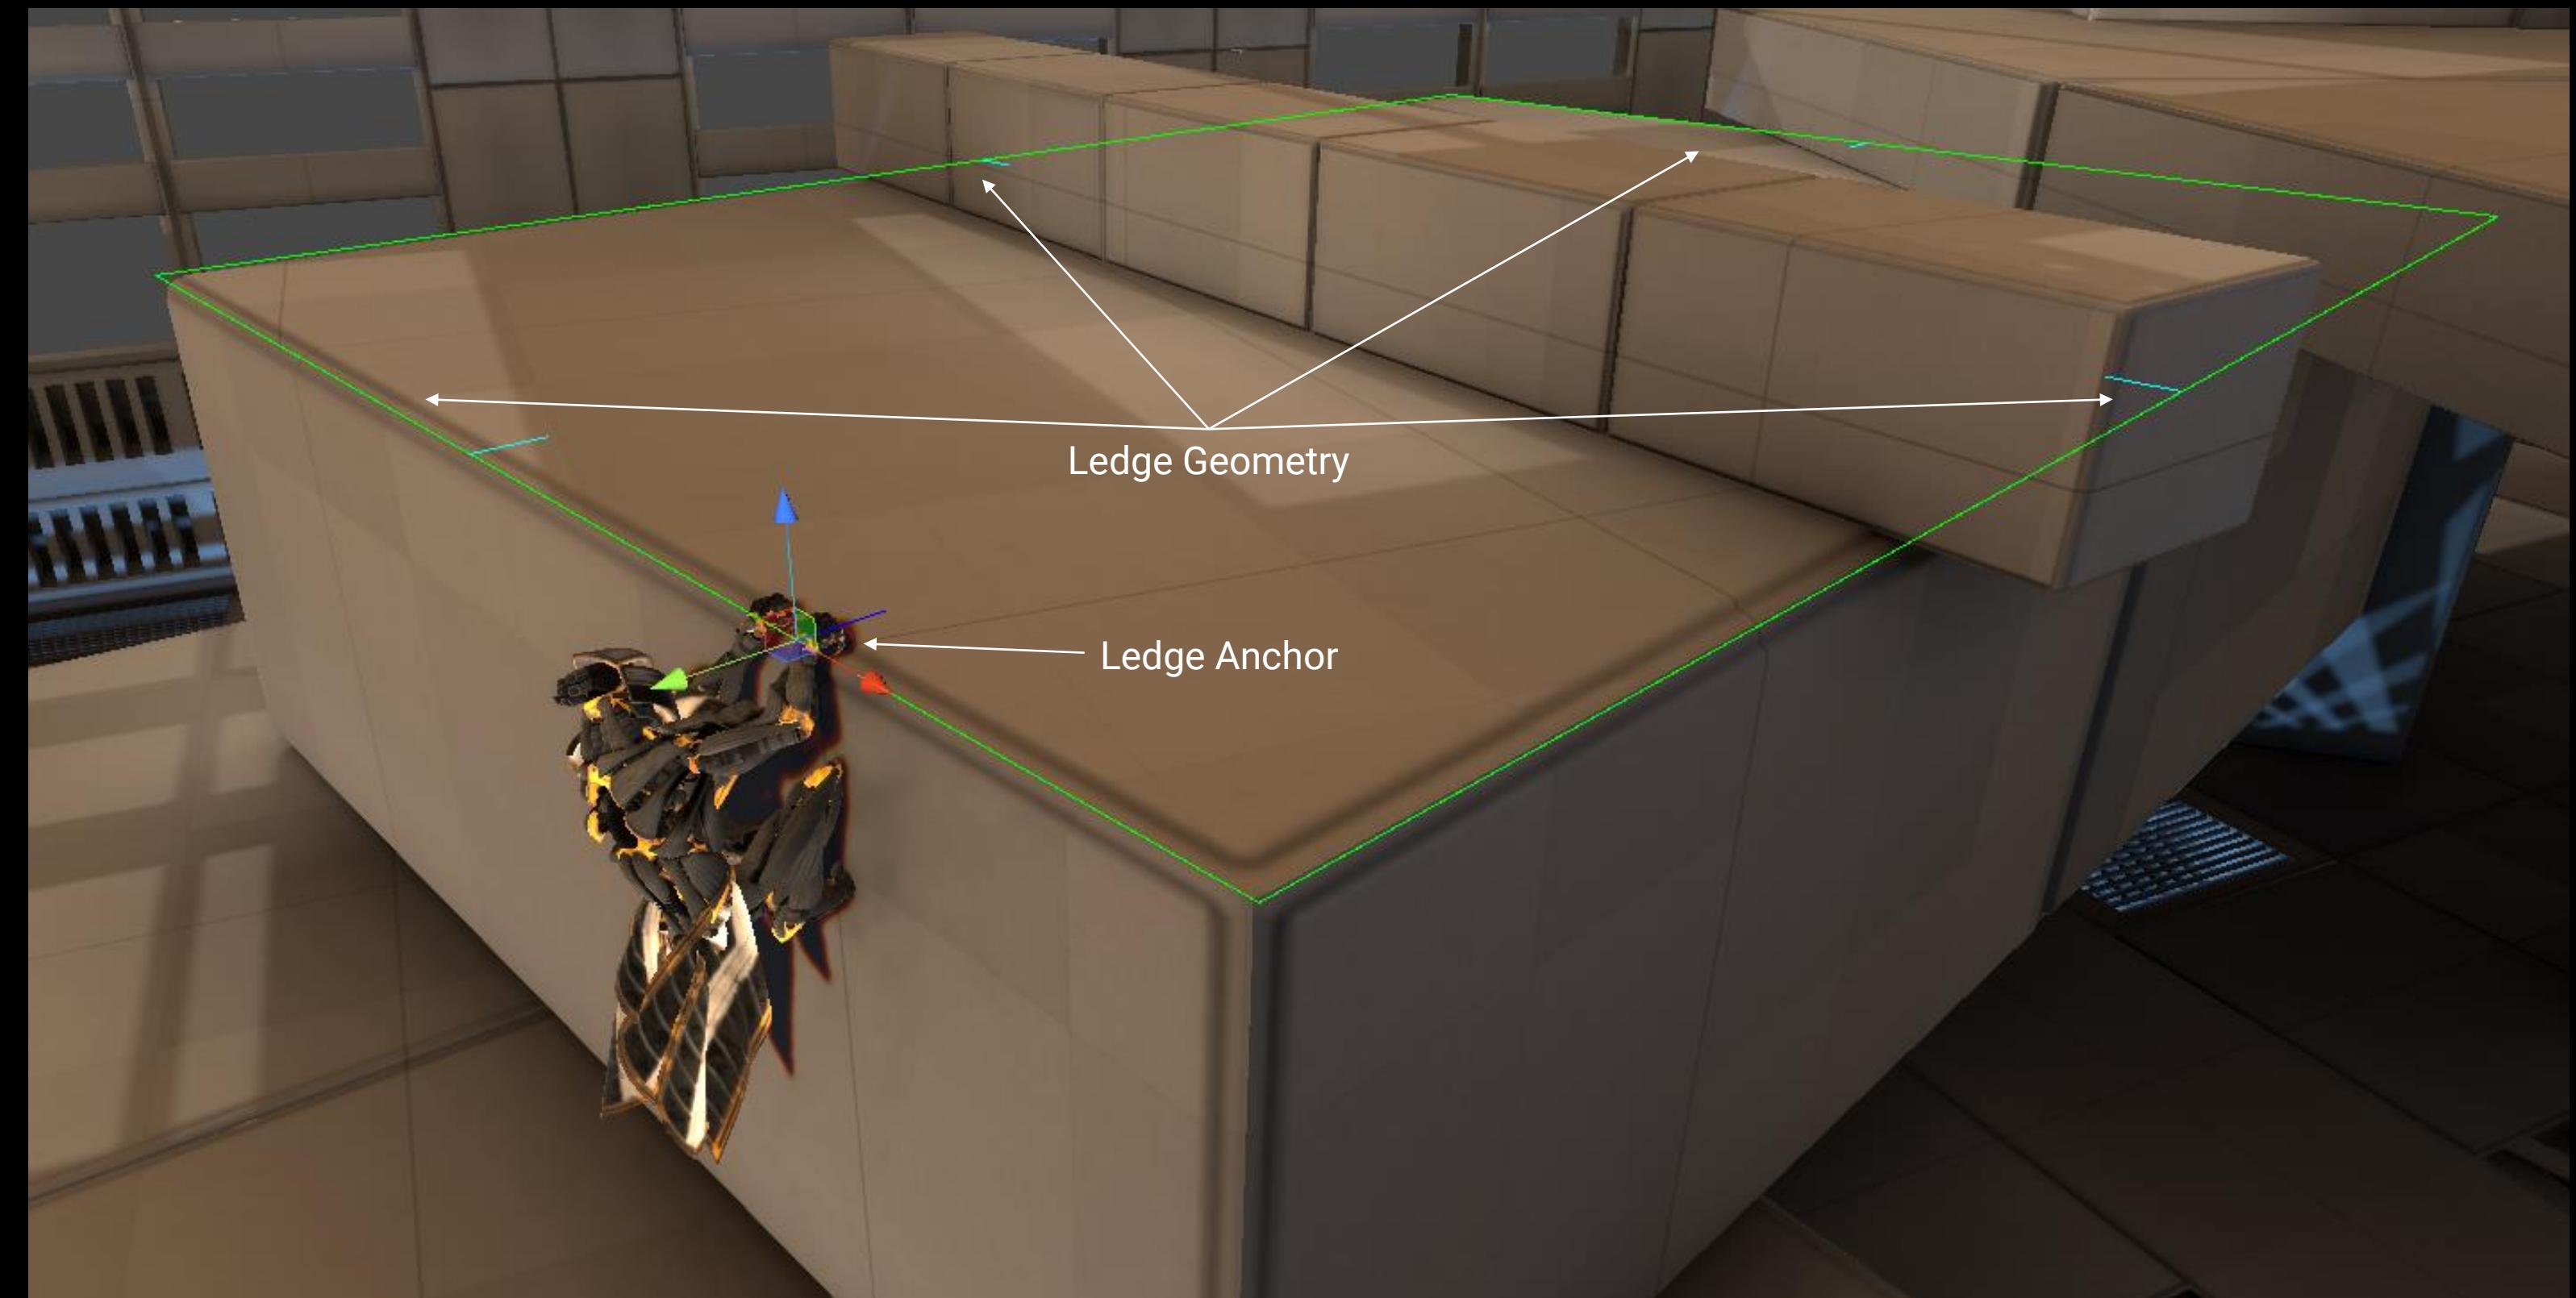
Ledge anchor is (line index, fraction)

Ledge anchor can be advanced based on distance

Ledge anchor can be constructed from world space position

Full predictive model

Snapshot() / Move() / Rewind()



# Free Climbing Model

Wall geometry gets constructed on-the-fly

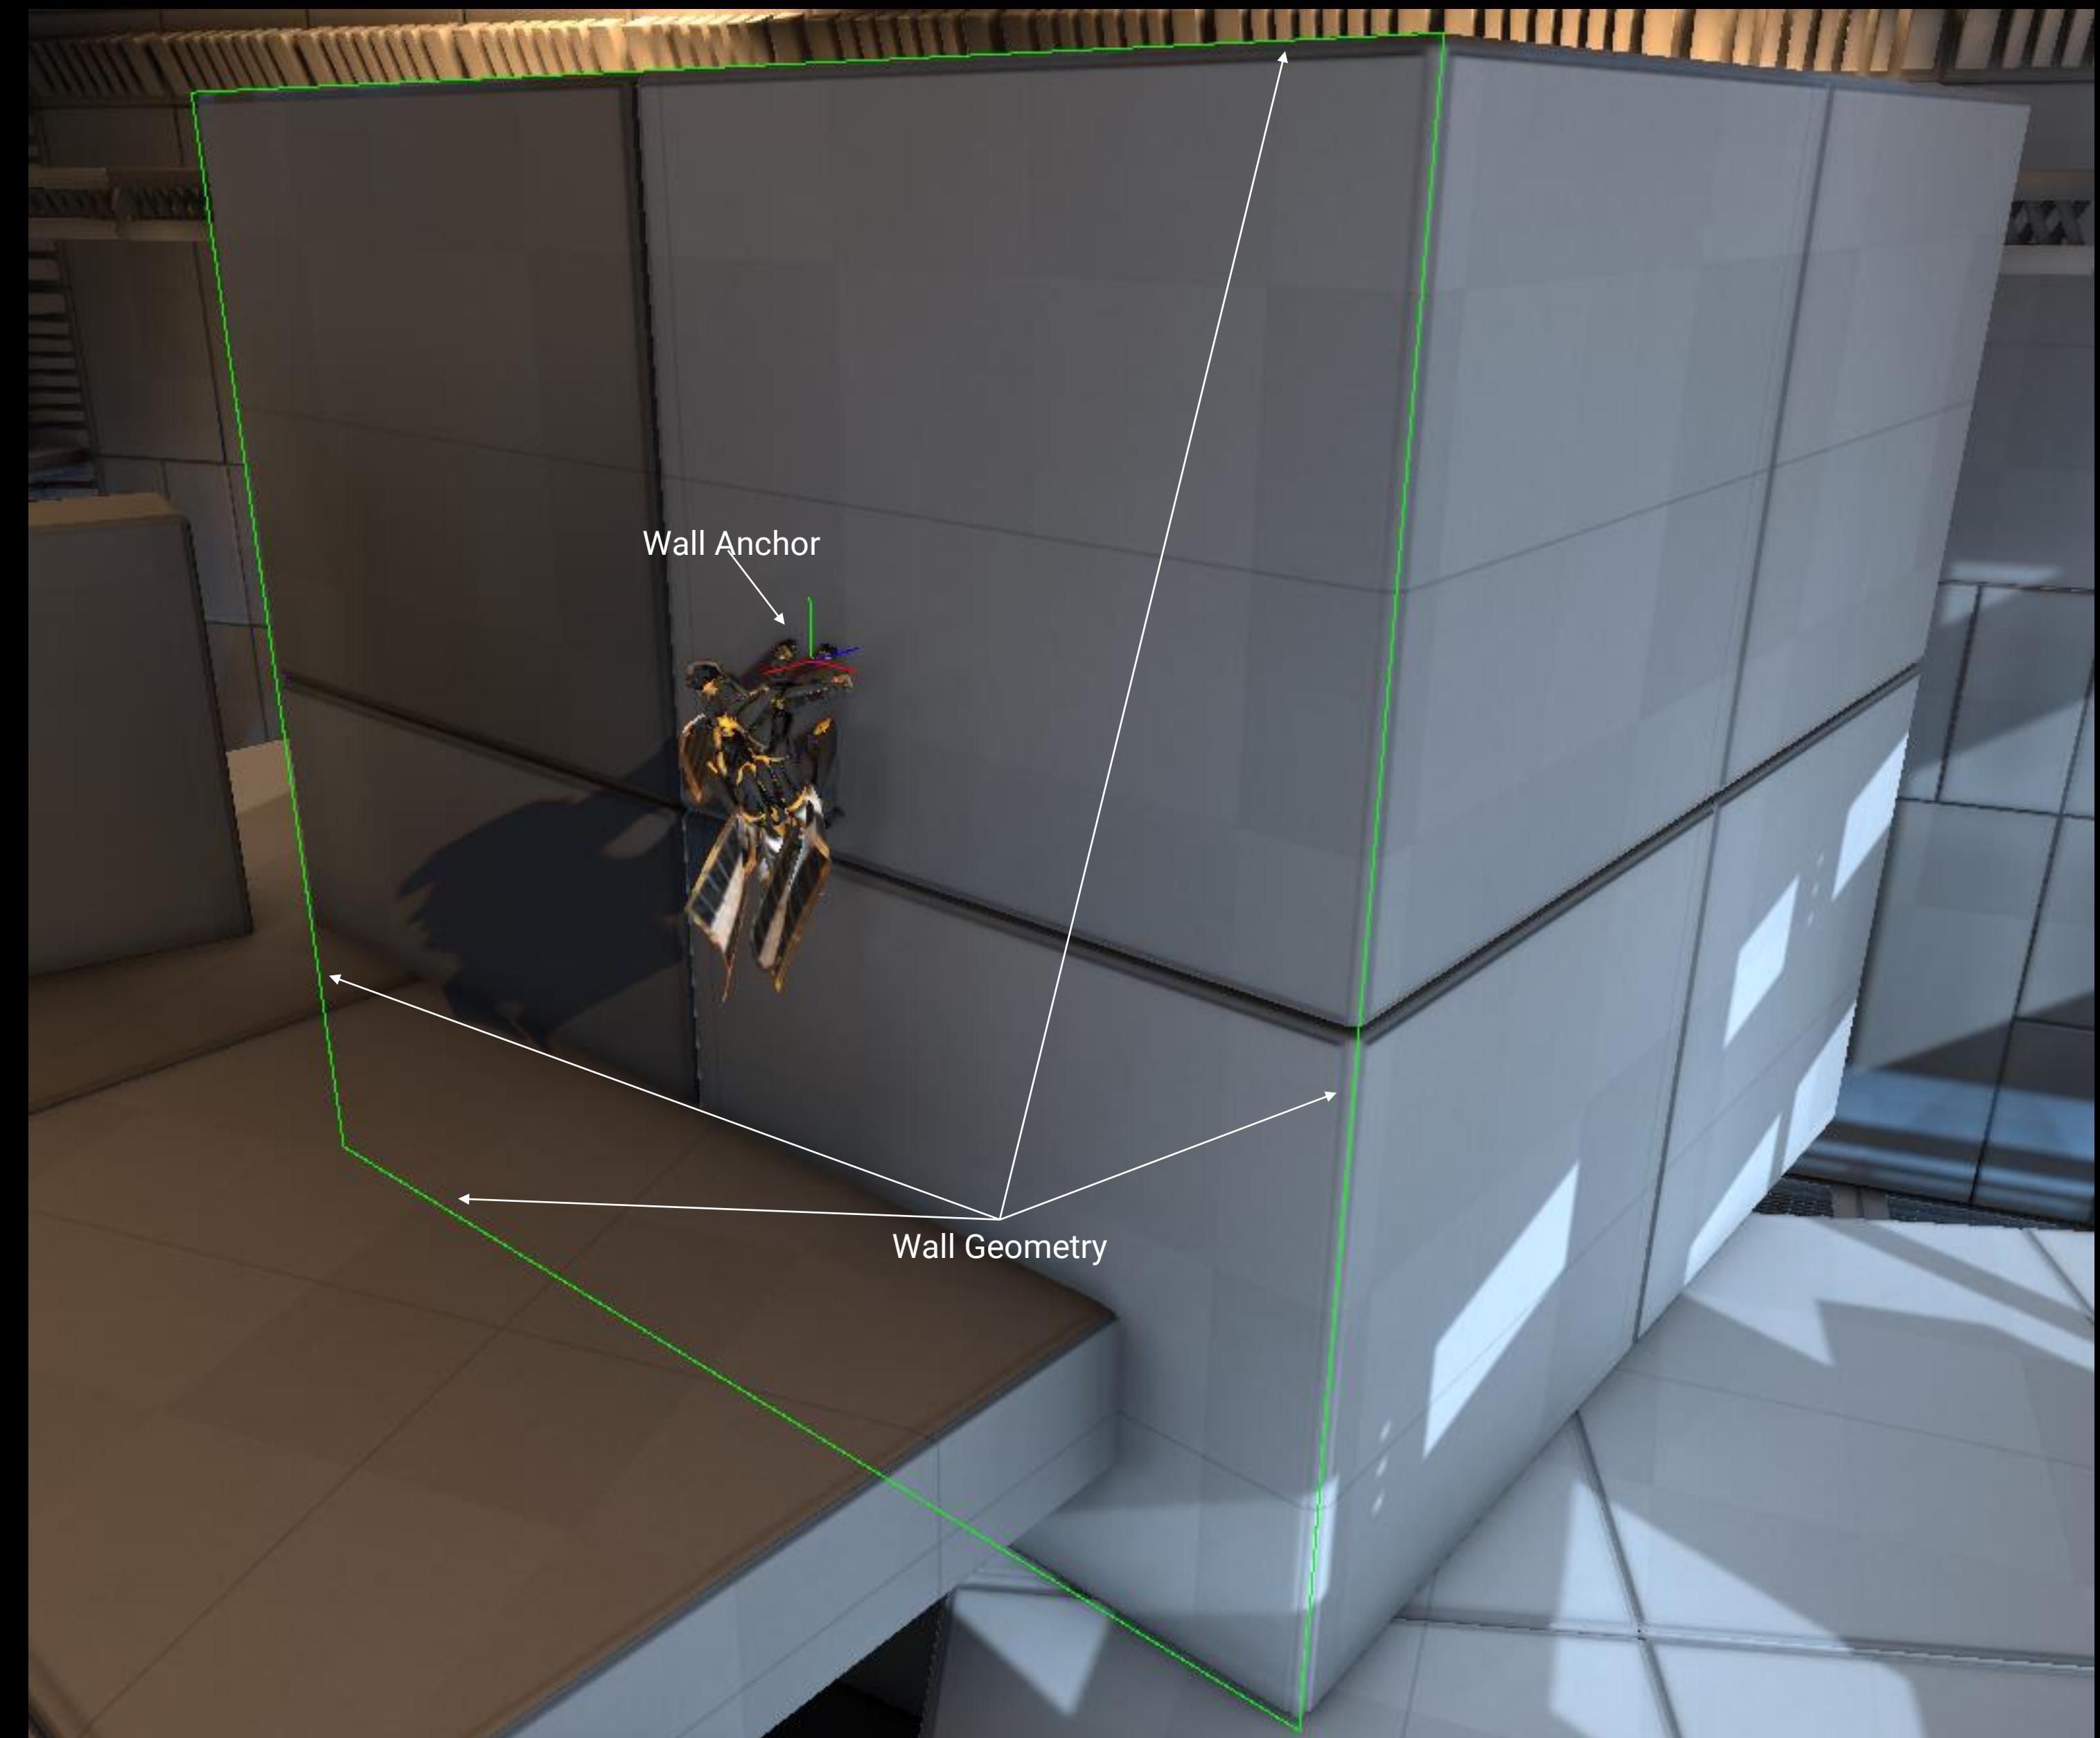
Wall anchor is (Normalized UV coordinate)

Wall anchor can be moved inside geometry bounds (2d displacement vector)

Wall anchor can be constructed from world space position

Full predictive model

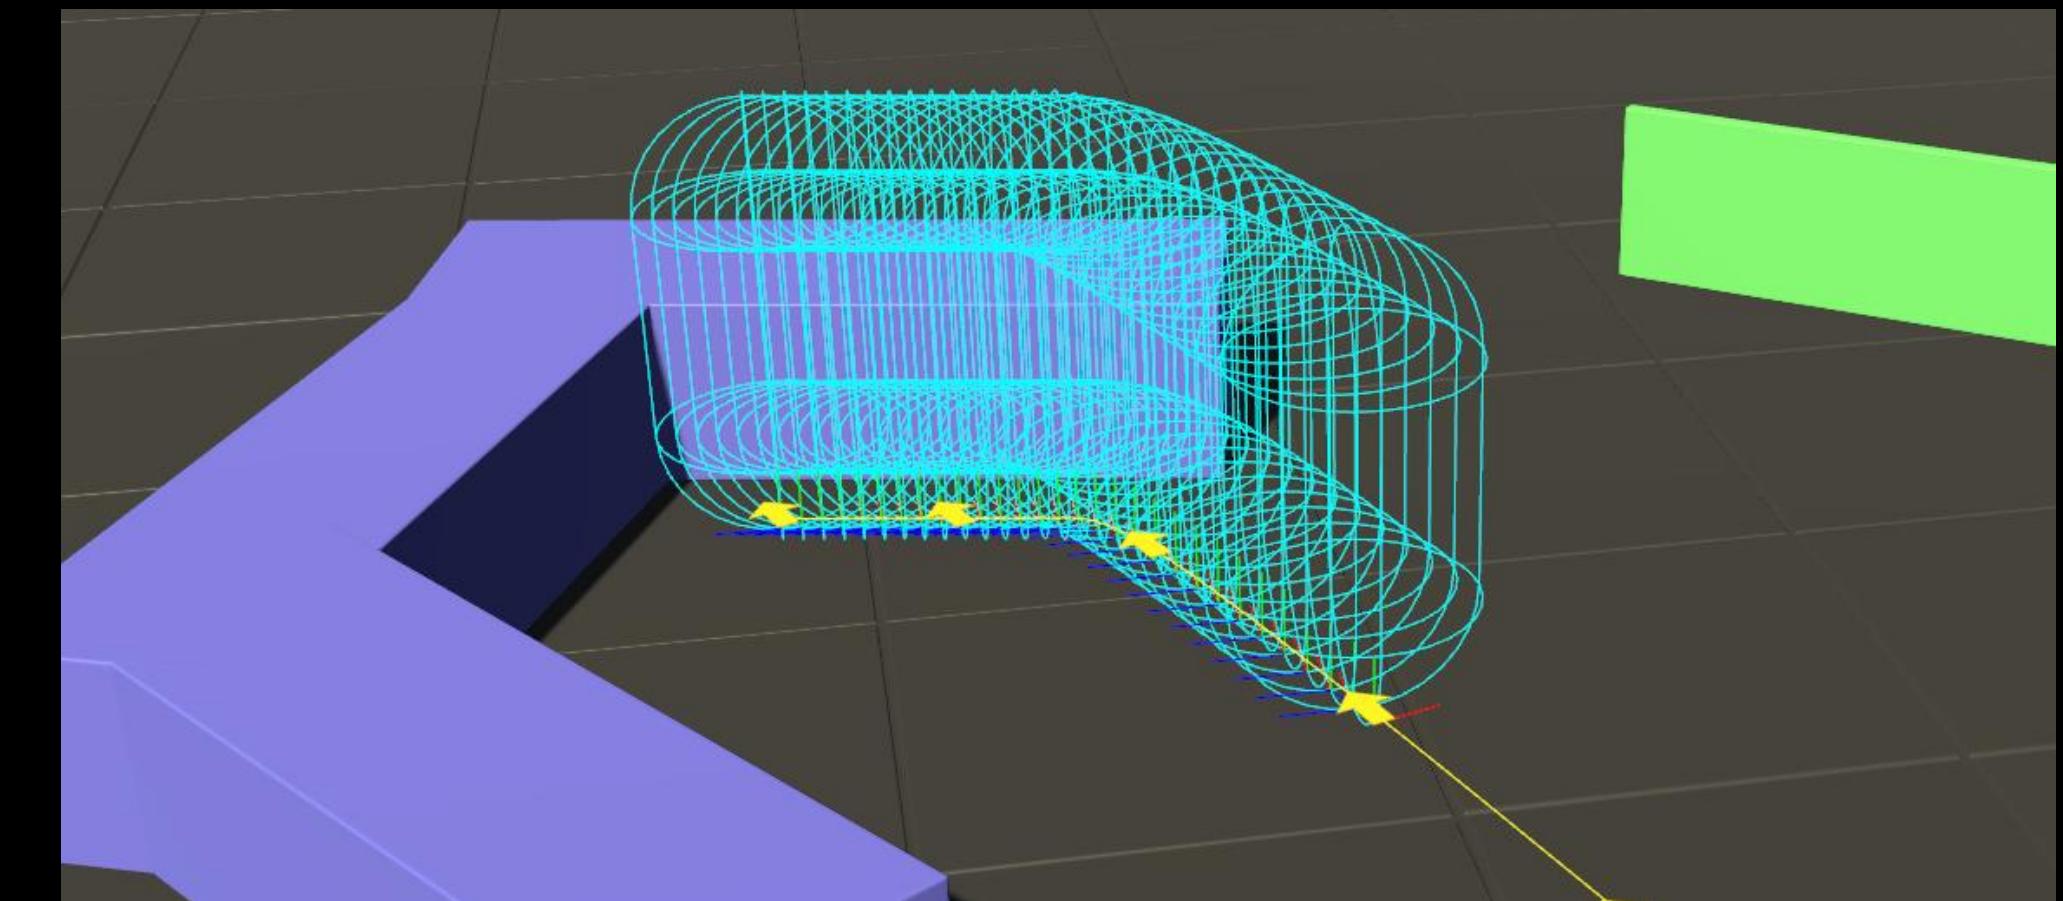
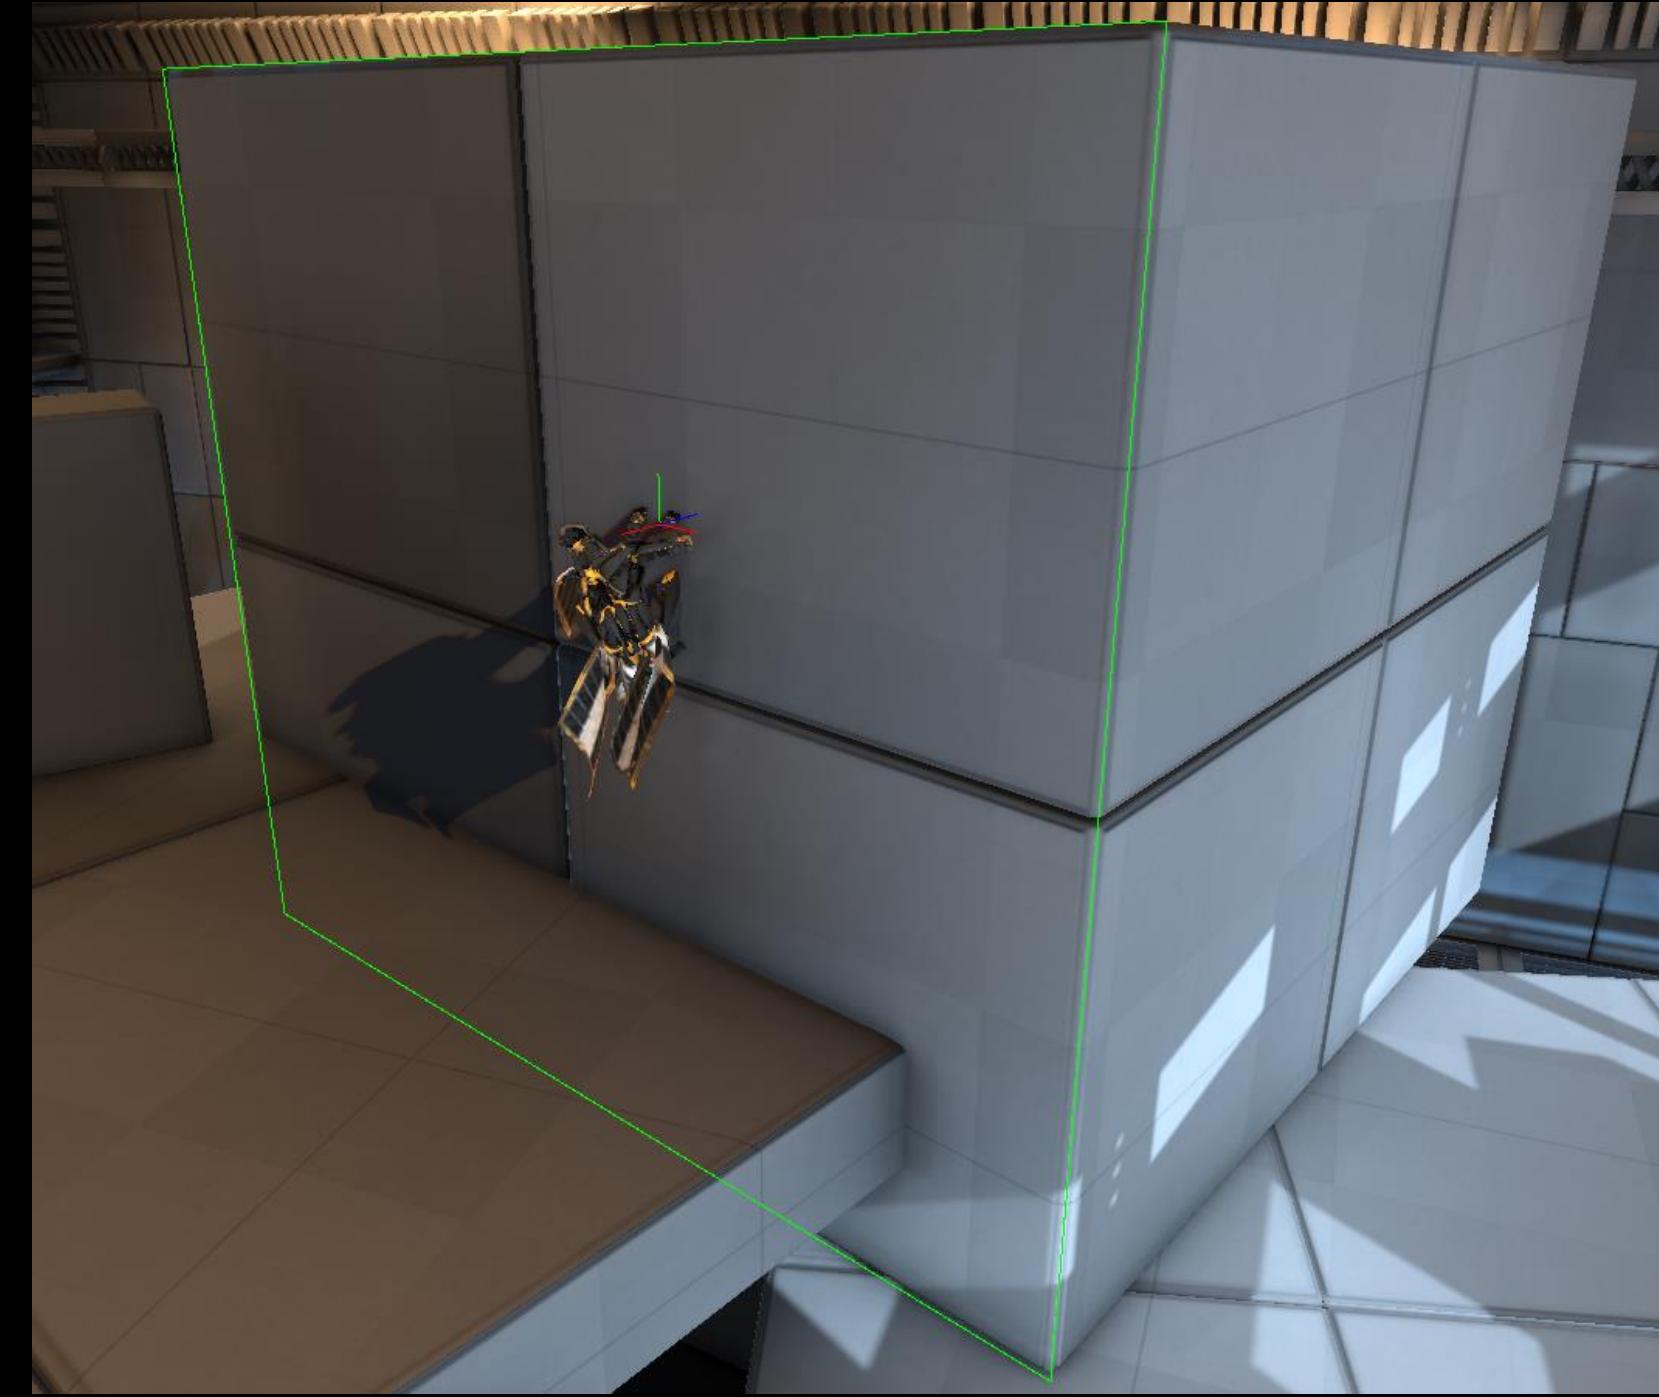
Snapshot() / Move() / Rewind()

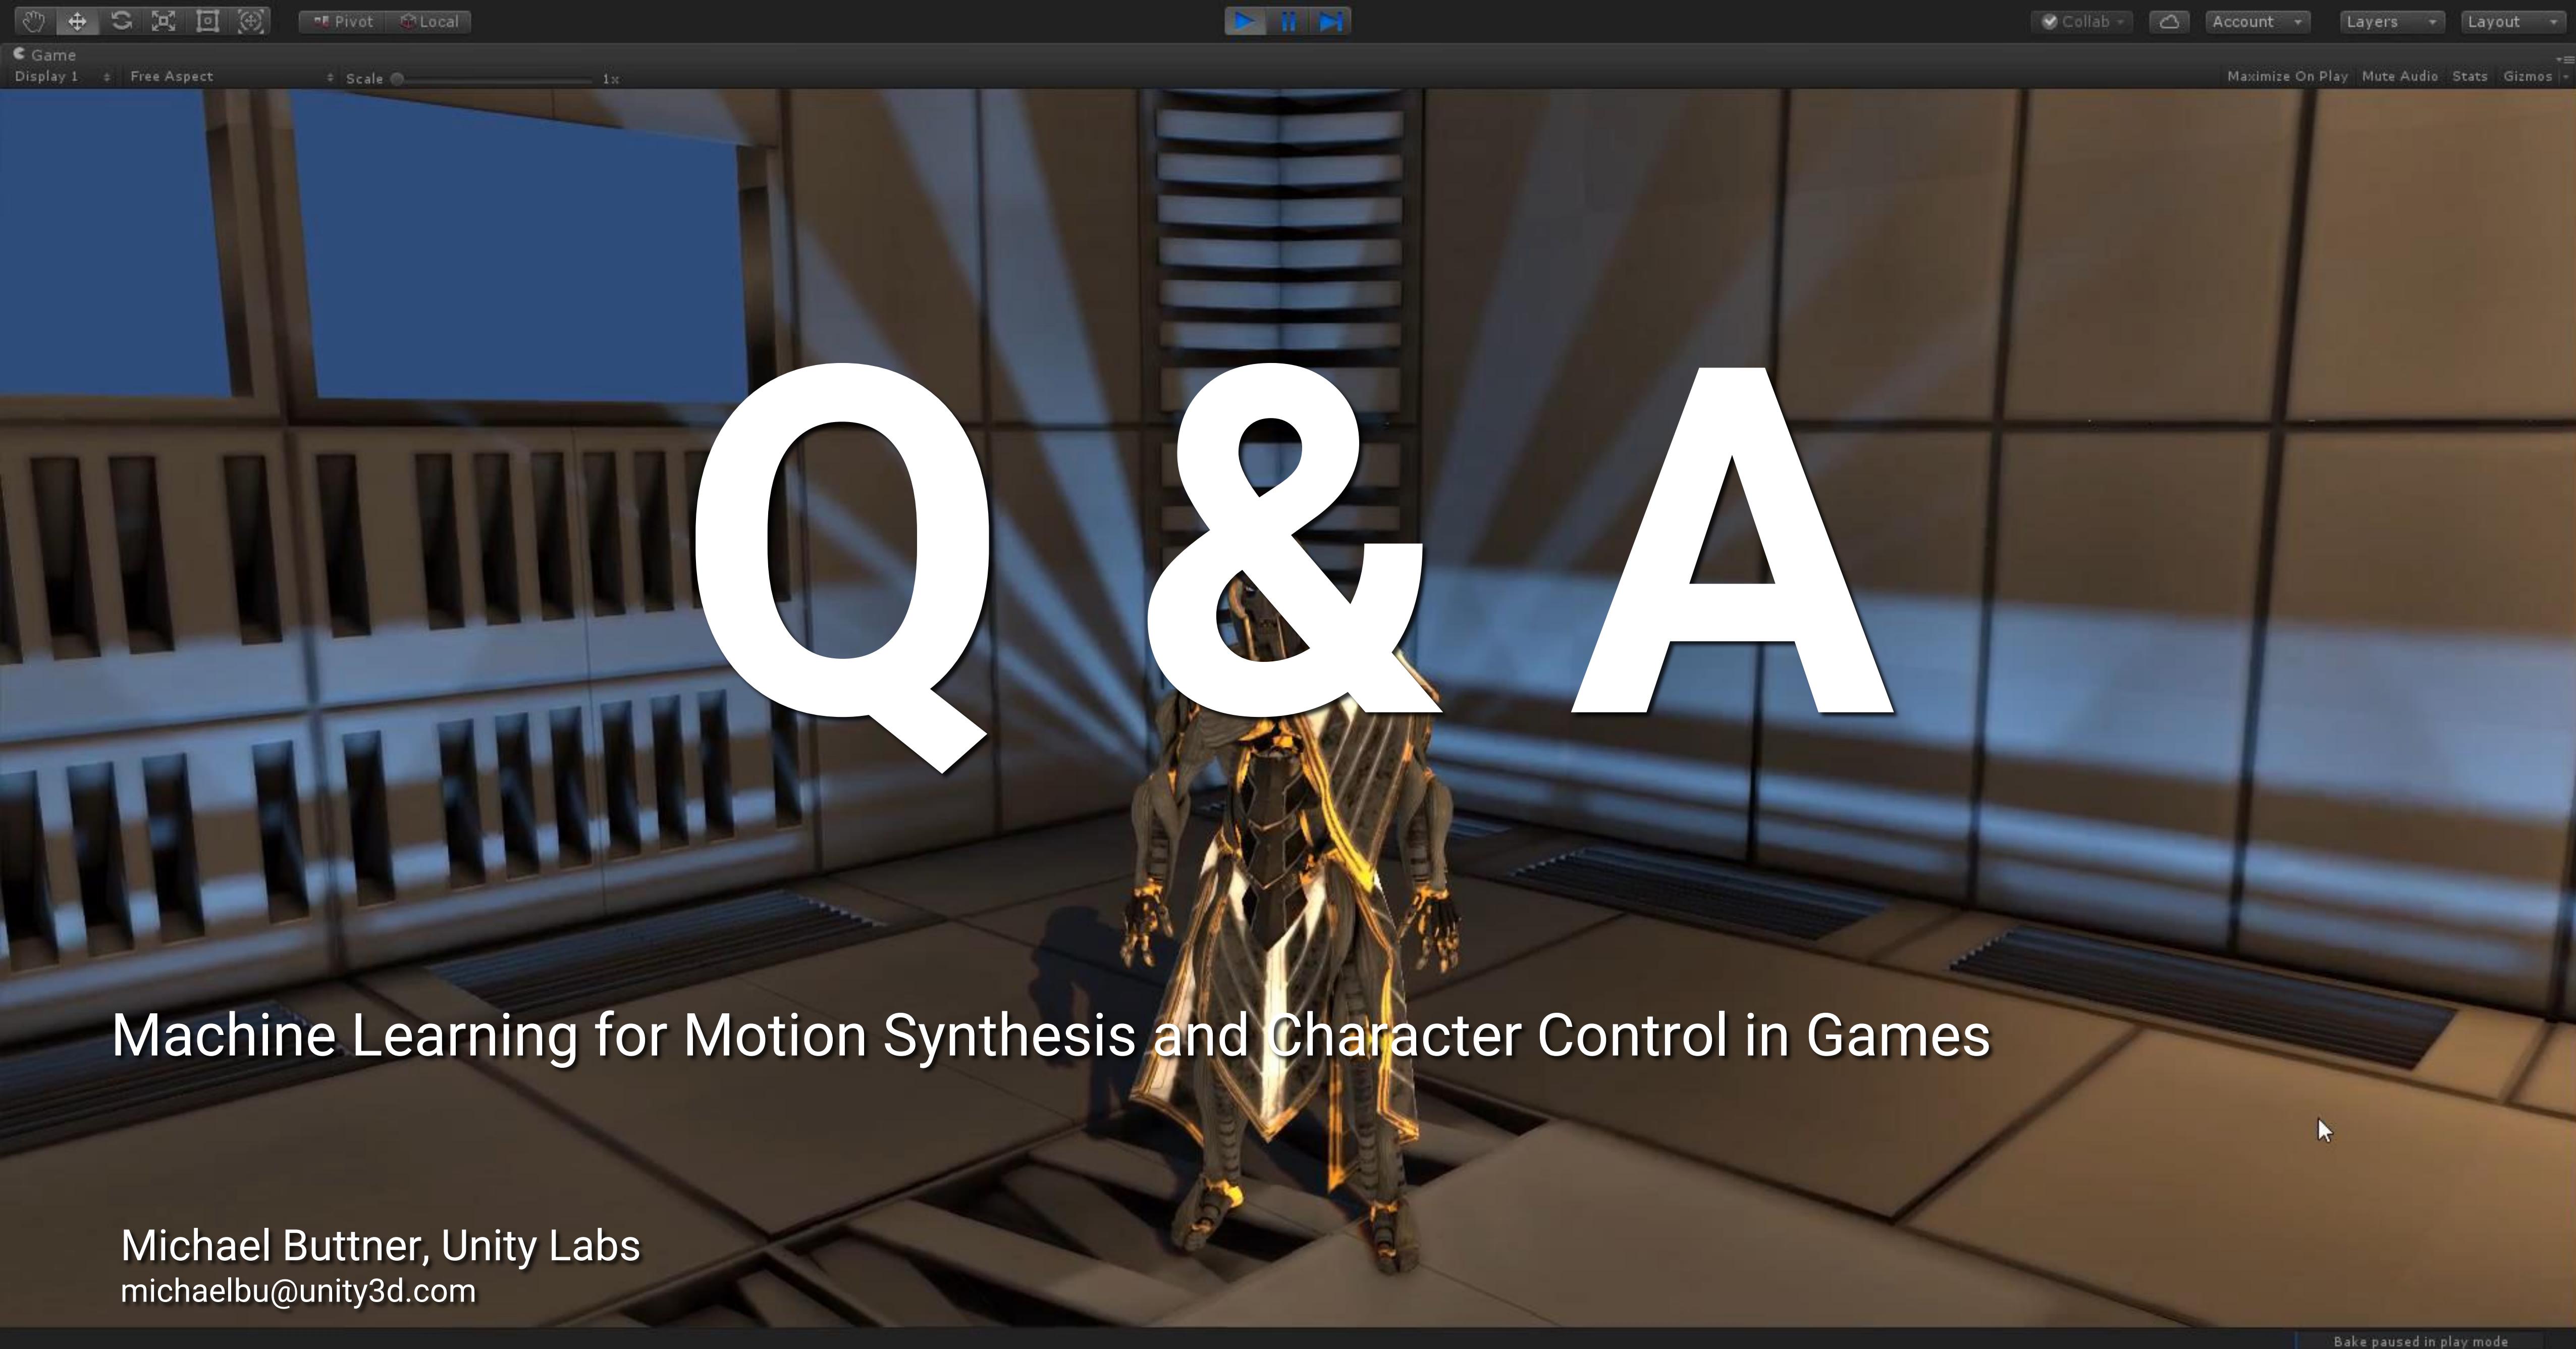


# World Model

It is important to note that this is an unavoidable complexity for any non-trivial character navigation

Any game will require this kind of structure in one form or another





Michael Buttner, Unity Labs  
michaelbu@unity3d.com

Bake paused in play mode