

GDC

March 21-25, 2022
San Francisco, CA

Walk Lizzie, walk!

Emergent physics-based animation through Reinforcement Learning

Jorge del Val Santos

Senior Research Engineer – Embark Studios

#GDC22







*“Allow anyone to create living creatures
... and play with life”*

I WANT A LIVING



LIZARD

How do we do that?

Come on... we know this stuff

- We create the lizard asset.
- We rig it.
- We tell a couple of our animators to spend some time animating movements.
- Voila!



But I told you... physics based!

- Oh right...
- Ok we do the same and just use inverse kinematics or something.
- We fine-tune it until it looks right, maybe add a few forces.

And if it stumbles or falls?

- We can just add a few imaginary forces to make it stay upright.

But I told you... physics based!

- Oh right...
- Ok we do the same and just use inverse kinematics or something.
- We fine-tune it until it looks right, maybe add a few forces.

And if it stumbles or falls?

- We can just add a few imaginary forces



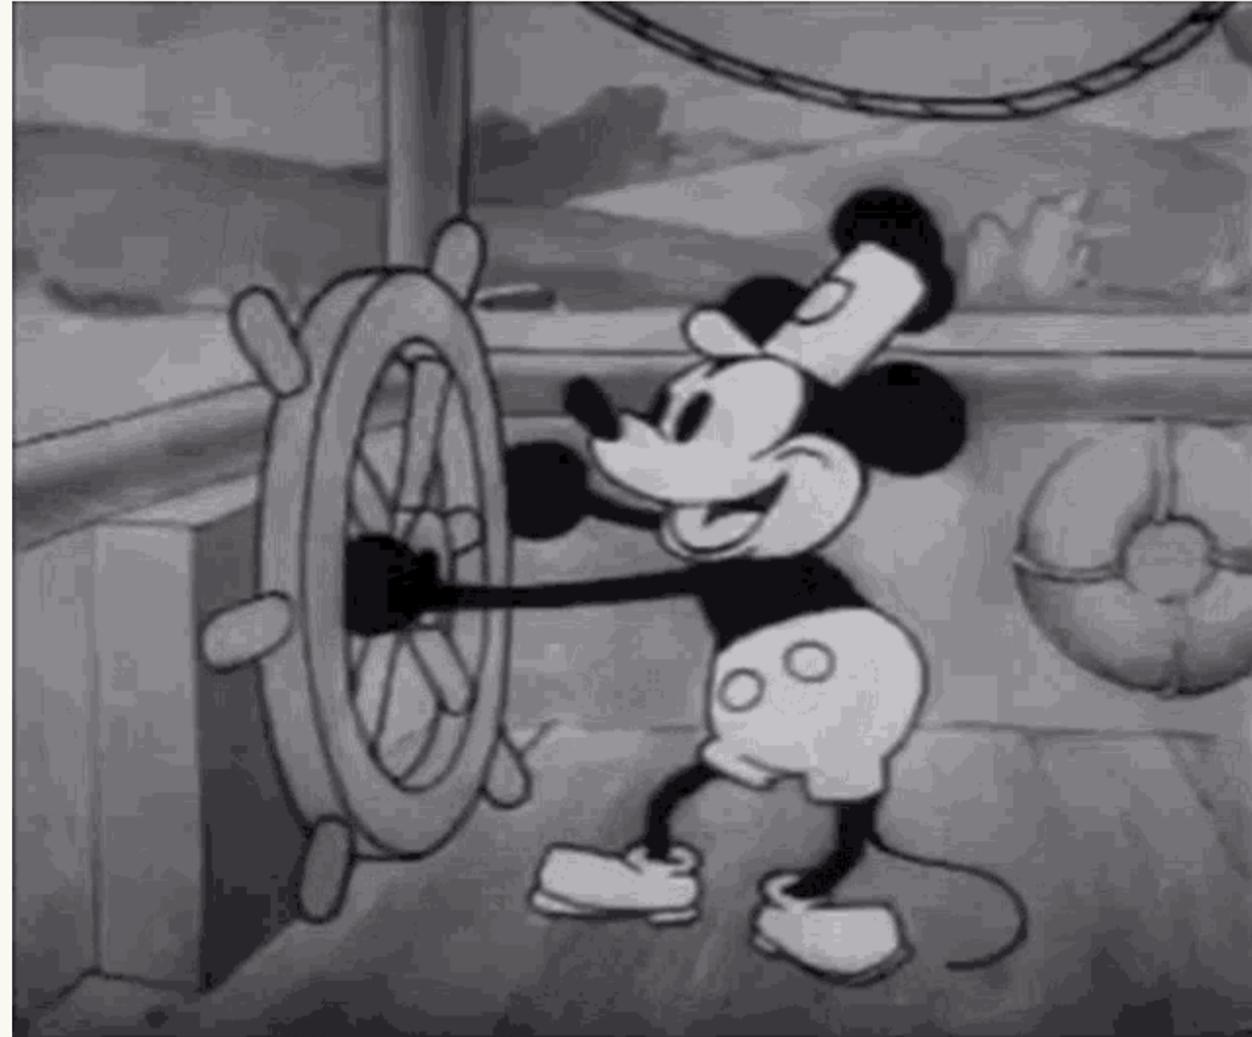
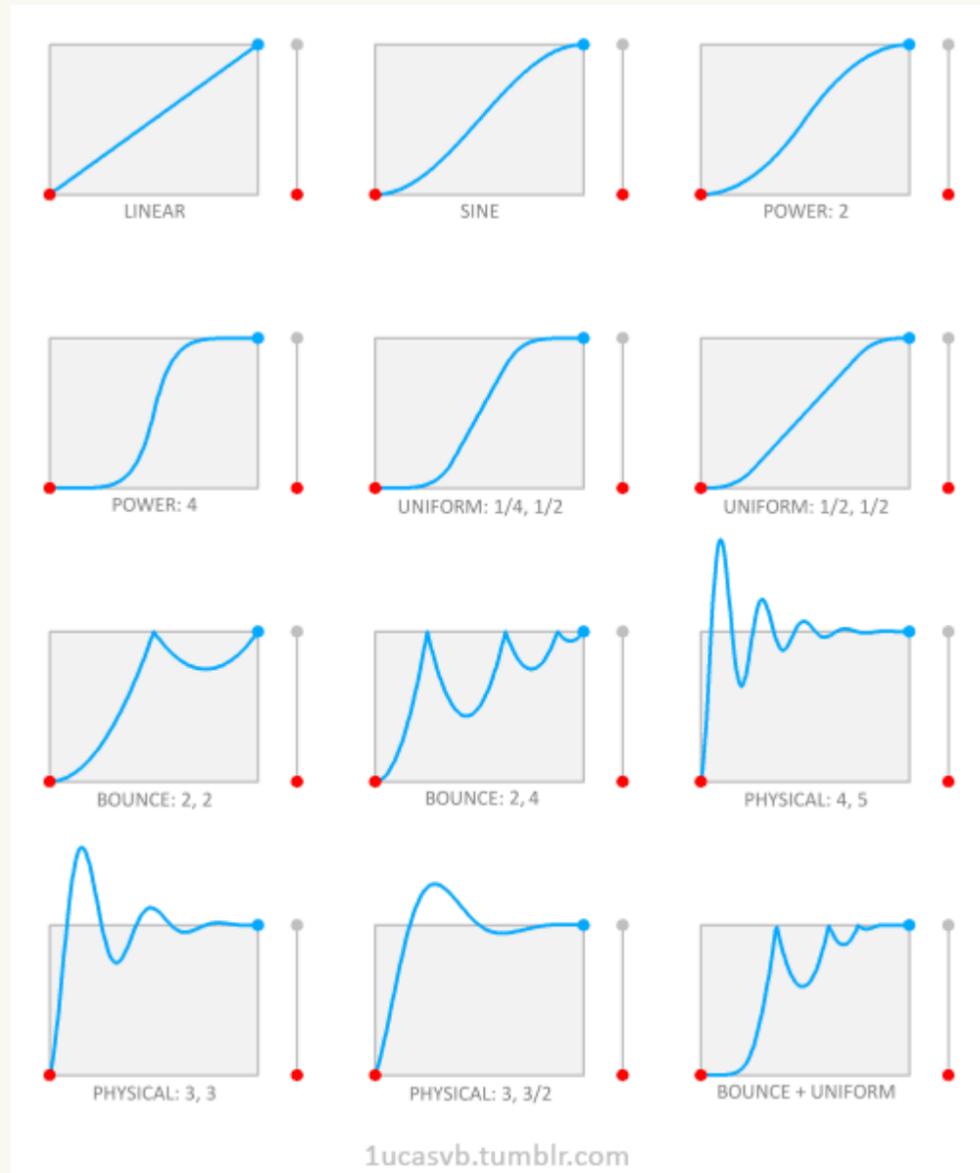
Cool...

Hey around 20 users more also want to create their living creatures. One wants a biped, another a dragon, another a dog, another a...

...



Animation takes a lot of work



Steamboat Wille, clip by Walt Disney

Animation takes a lot of work



Animation takes a lot of work

The embark animation team capturing Bence making cool moves.

Animation state machine by Fozzy28 from the Unreal forums.

Animation takes a lot of work

Animation state machine by Fozzy28 from the Unreal forums.

Traditional techniques are not scalable to multiple creatures!



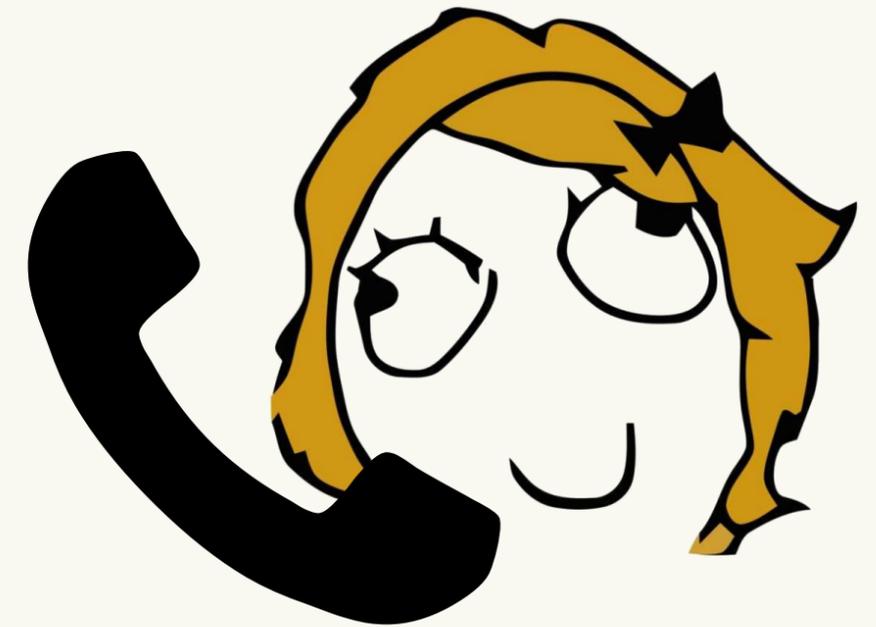
The embark animation team capturing Bence making cool moves.

Animation state machine by Fozzy28 from the Unreal forums.

And also:

- Hey I threw a box at a creature and it just kept doing exactly the same!
- In a hill it walks as if there was no hill, and when other things interact it doesn't care.

...



And also:

- Hey I threw a box at a creature and it just kept doing exactly the same!
- In a hill it walks as if there was no hill, and when other things interact it doesn't care.

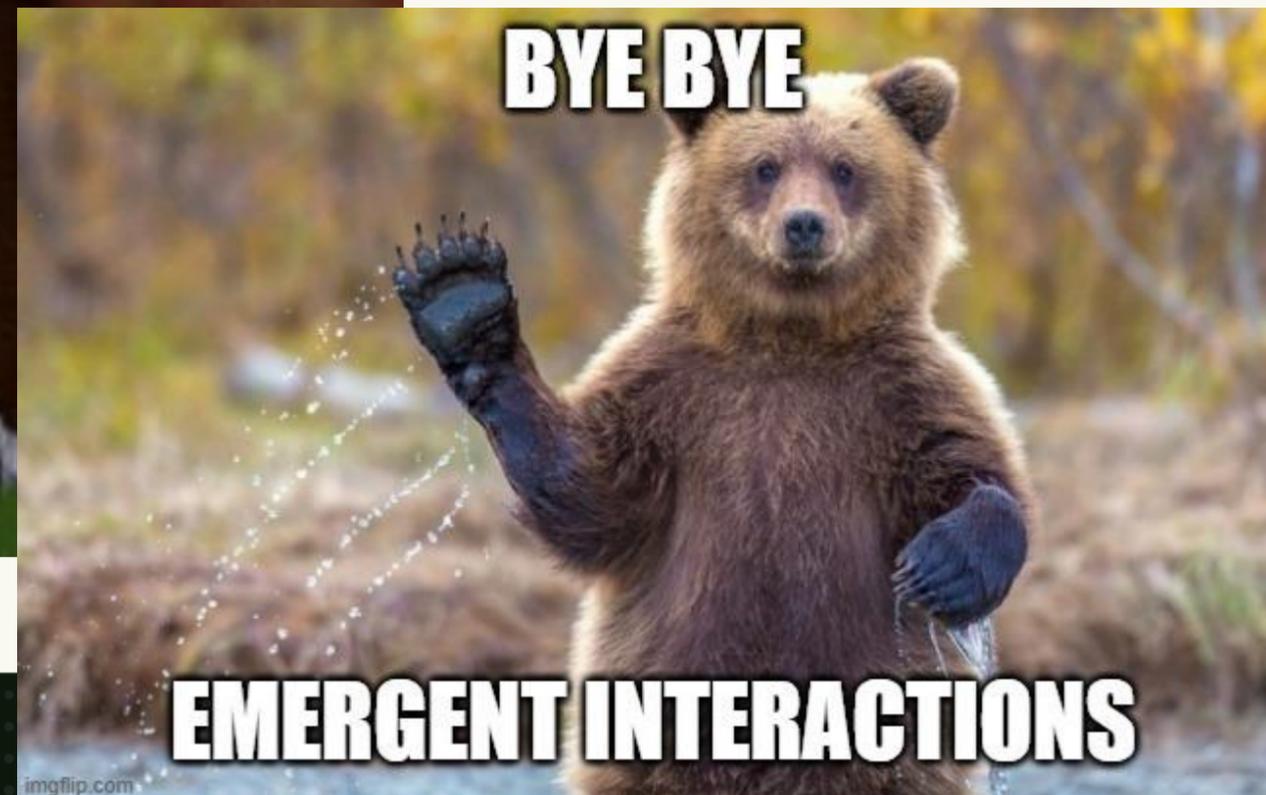
... **Results will not truly react or adapt to the environment unless specifically programmed to.**



And also:

- Hey I threw a box at a creature and it just kept doing exactly the same!
- In a hill it walks as if there was no hill, and when other things interact it doesn't care.

... **Results will not truly react or adapt to the environment unless specifically programmed to.**



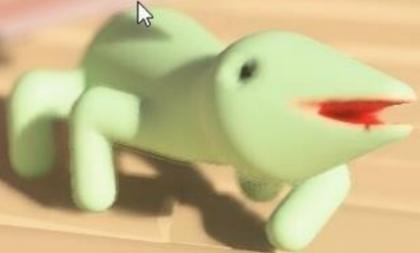
*“Allow **anyone** to create living creatures
... and play with life”*

let's try again: how do we do that?

Our solution:

*We allow them to **learn***

Hello!

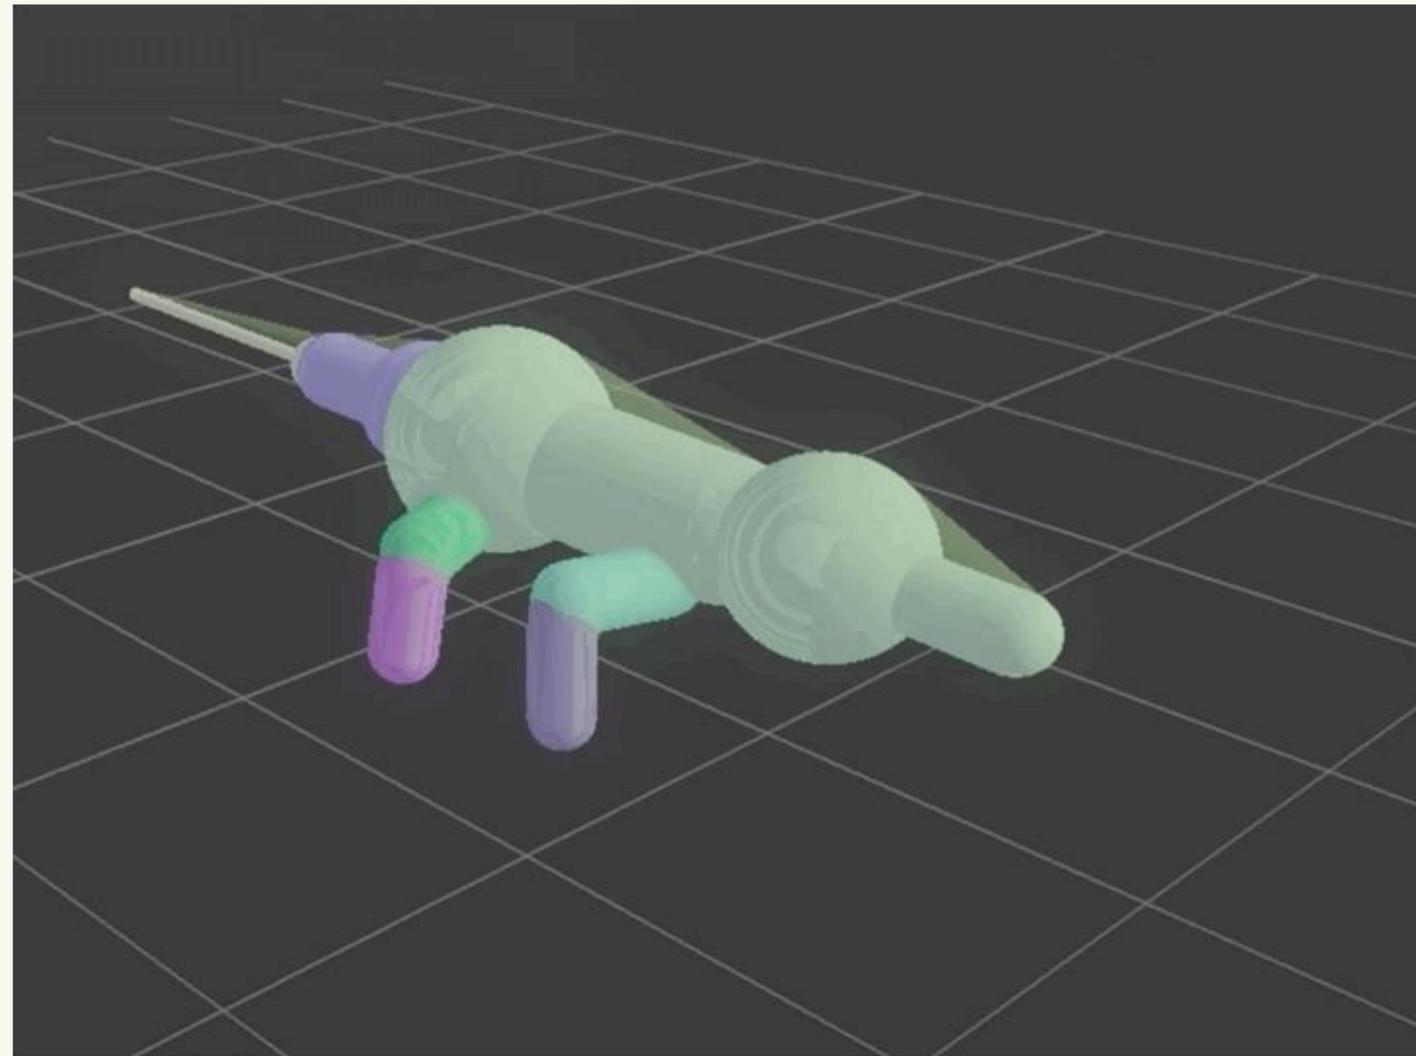


lizzie

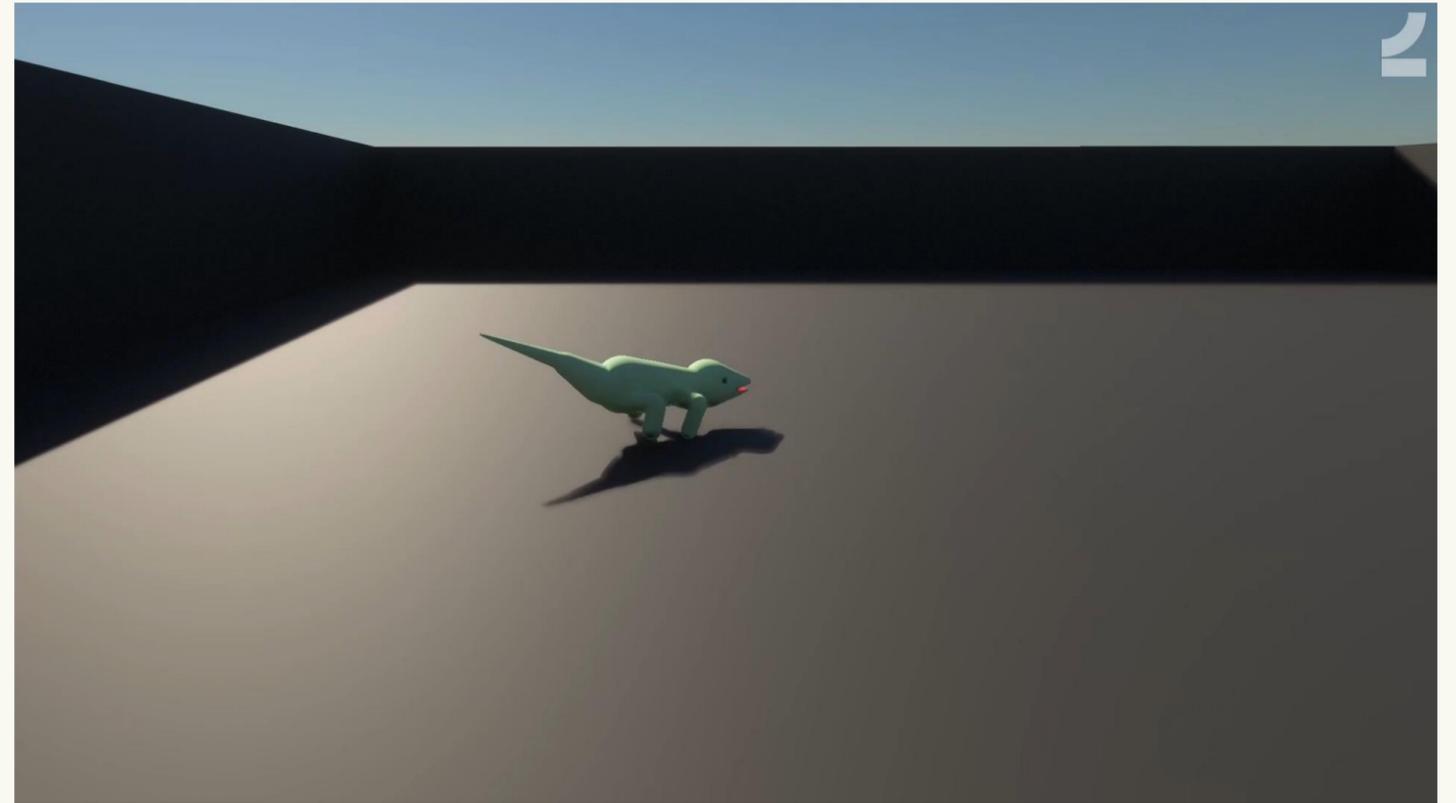
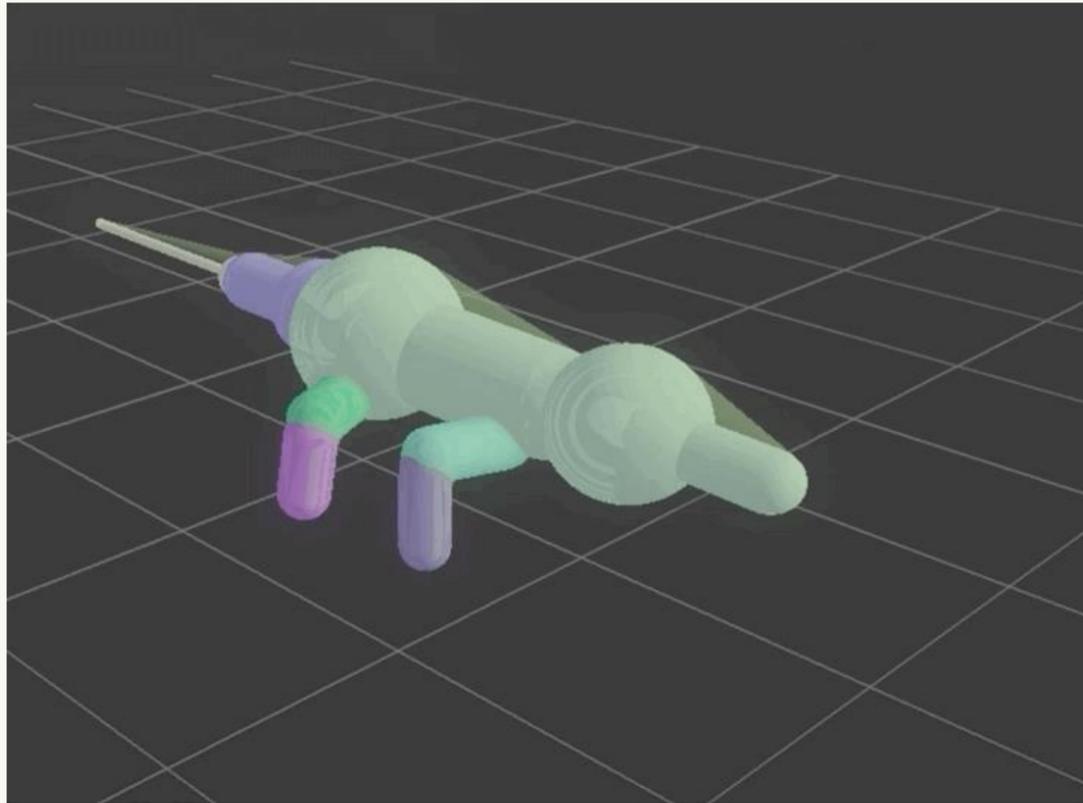


“Creature”

- Creature is a set of primitives and collision shapes stacked together by joints.



The goal



The problem

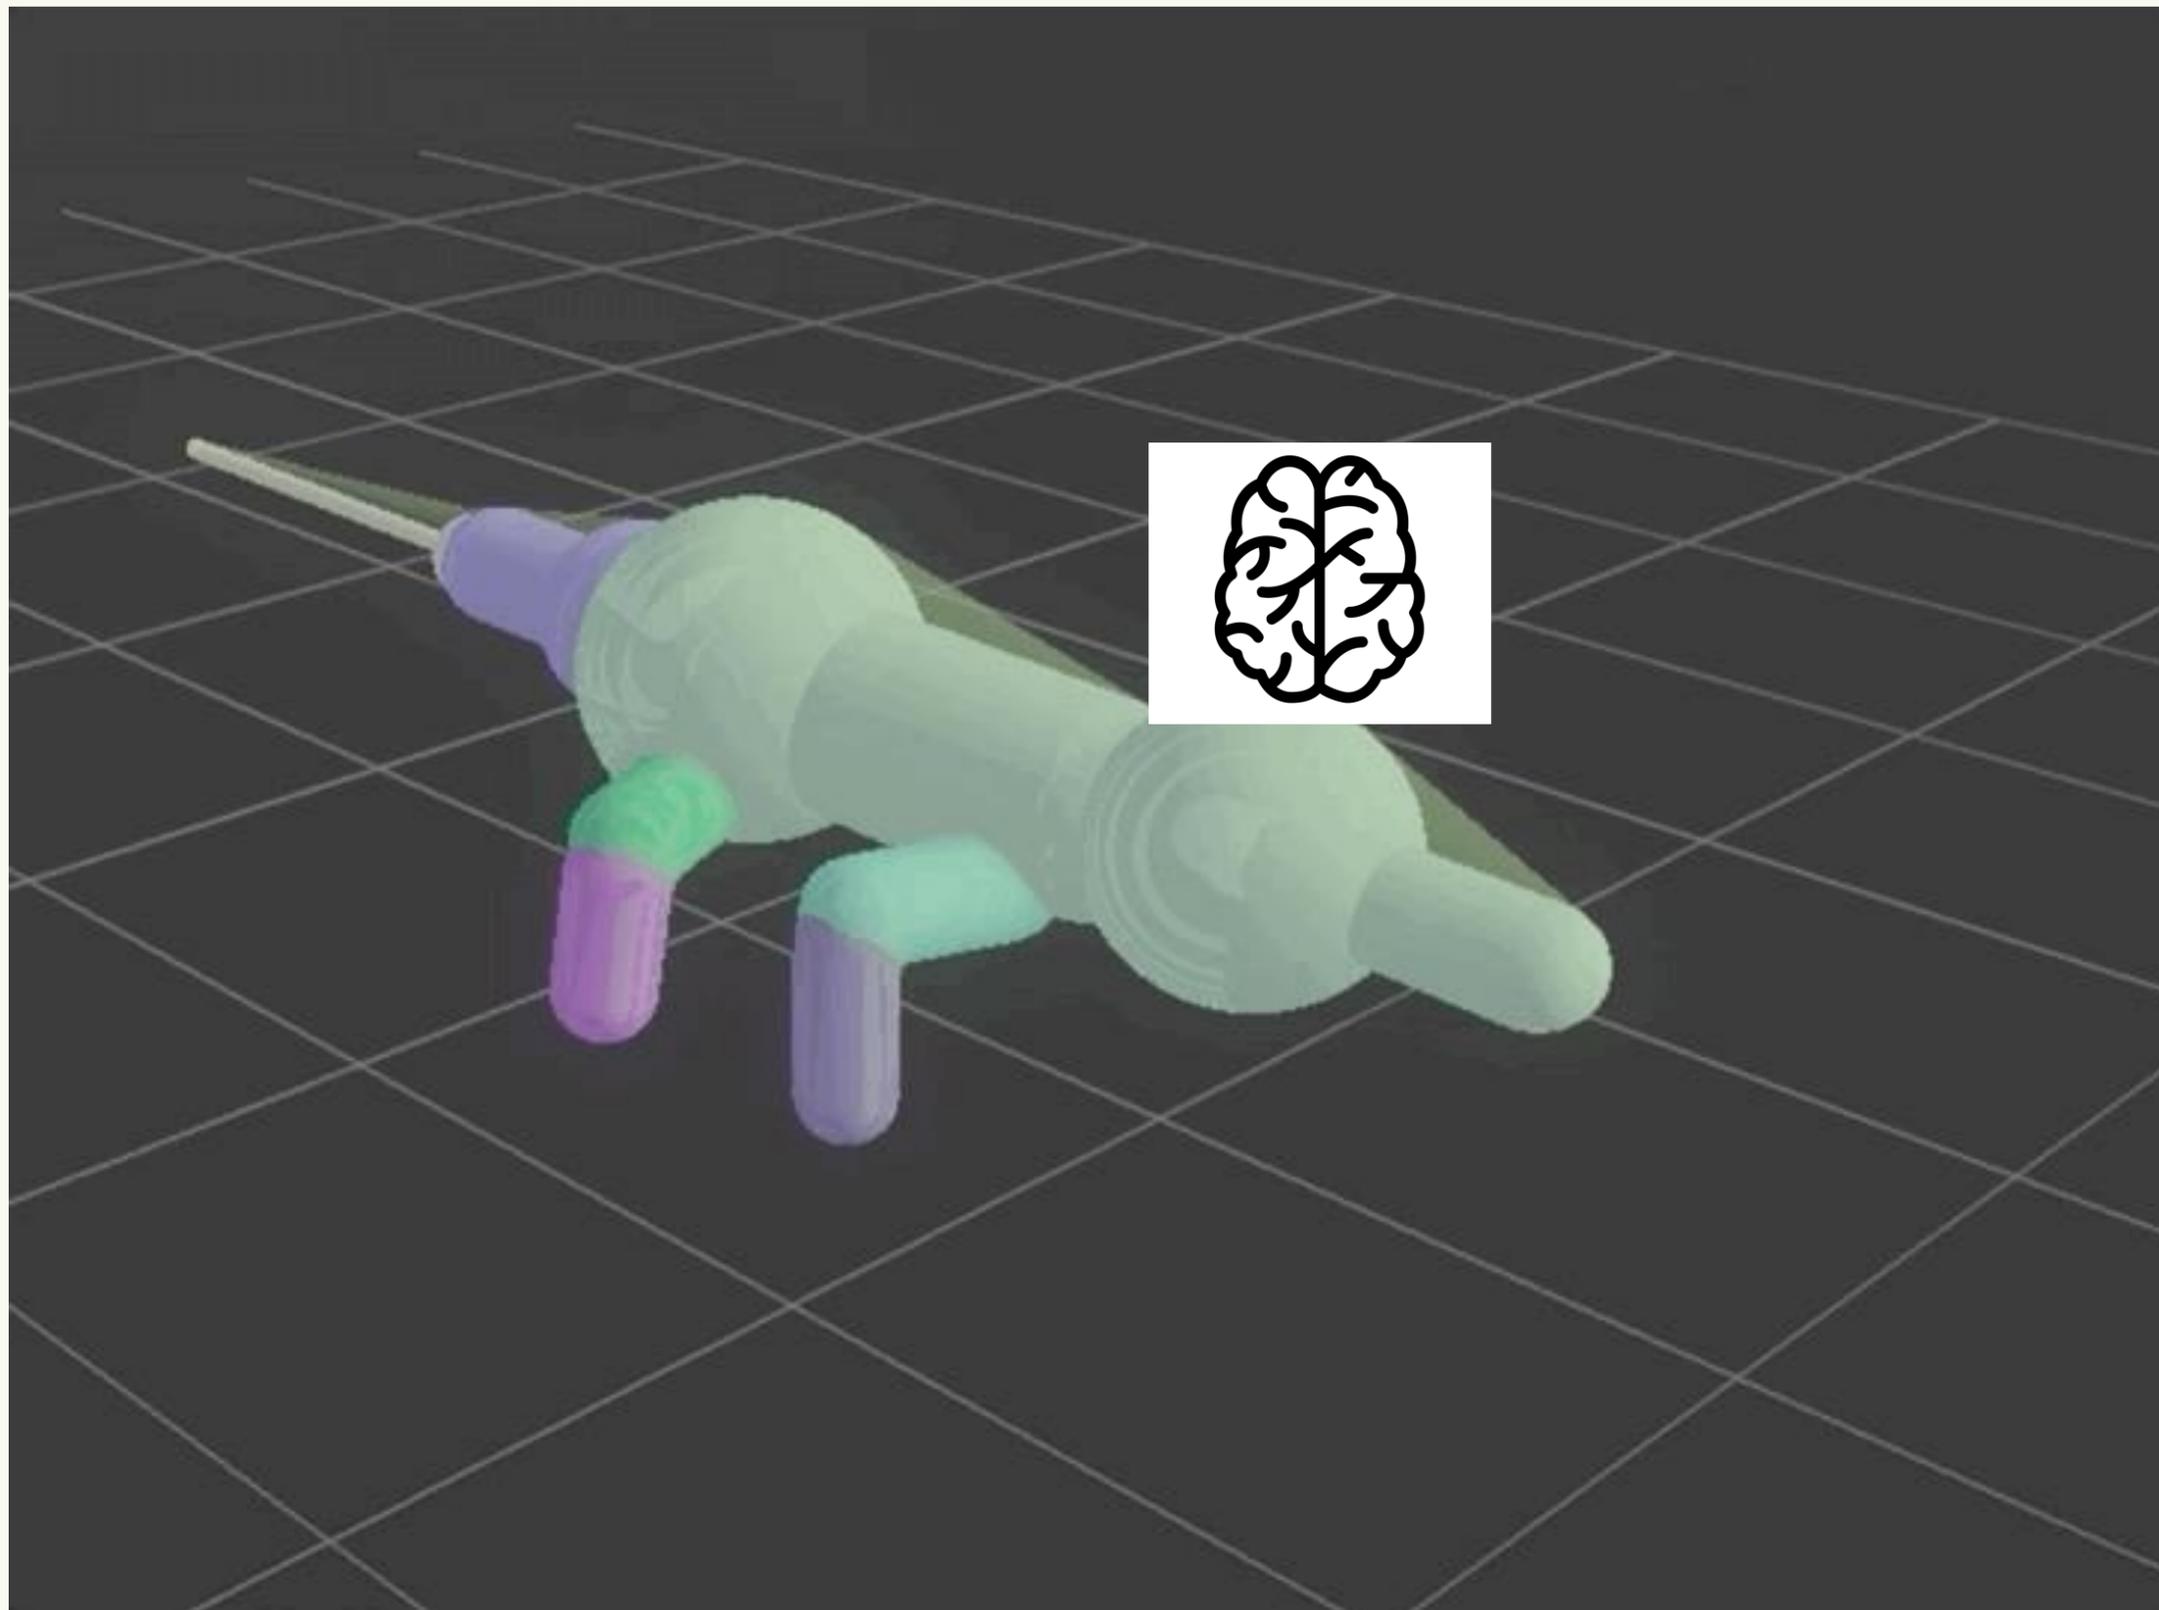
Animation

Robotics



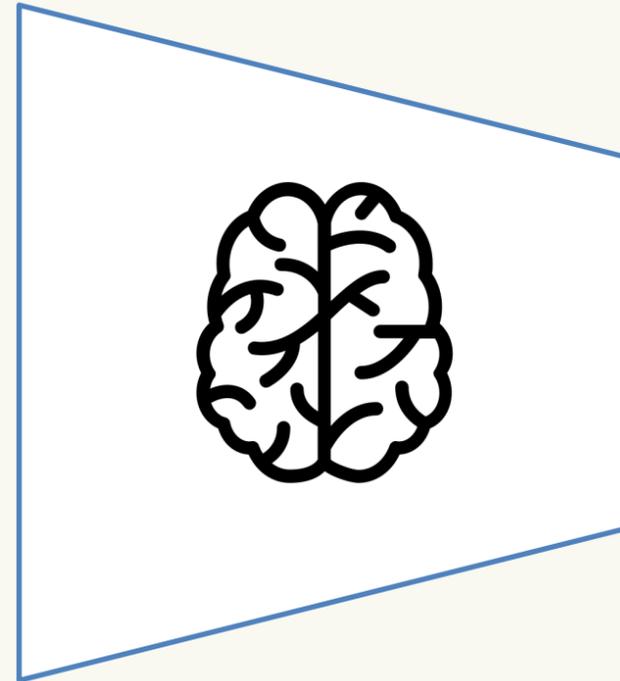
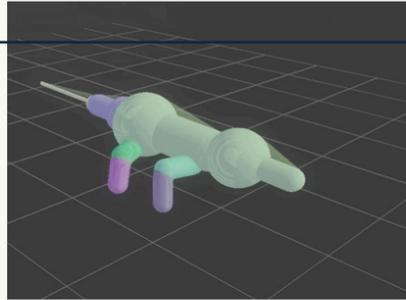
Animation state machine by Fozzy28 from the Unreal forums.

Boston dynamics



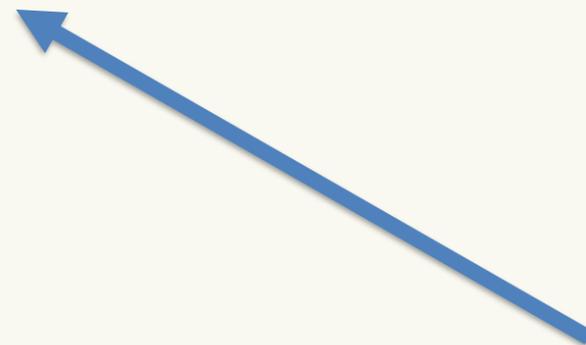
Observations

[
 joints orientation,
 body parts location,
 collisions,
 goal vector,
 terrain map,
 ...,
]

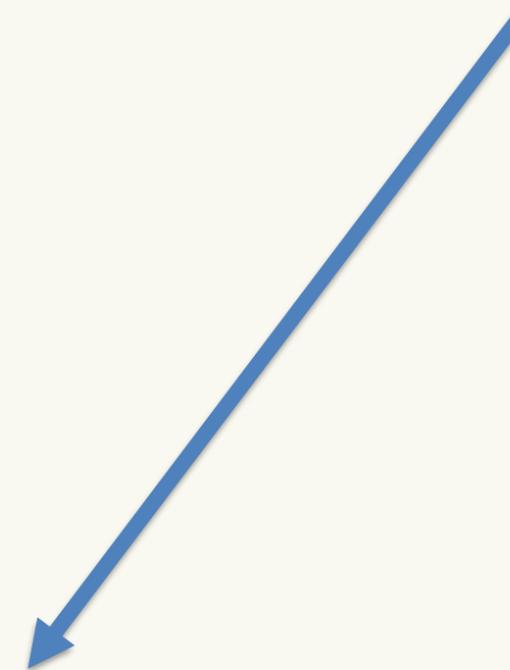


Actions

[
 controllable joints
 desired orientations,
 ...,
]

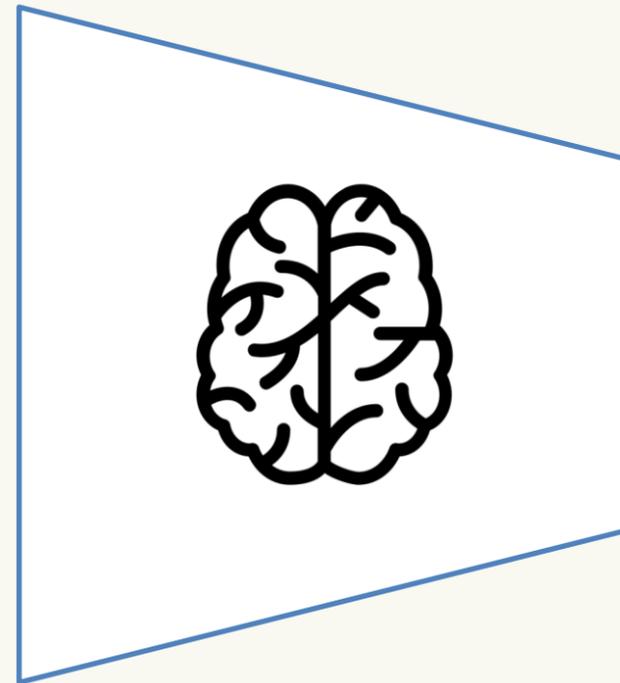
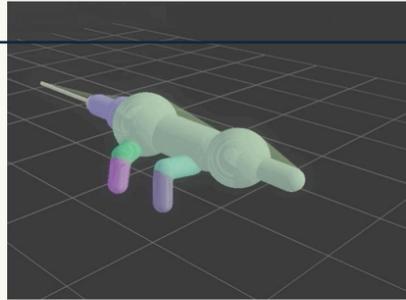


World
physics
step



Observations

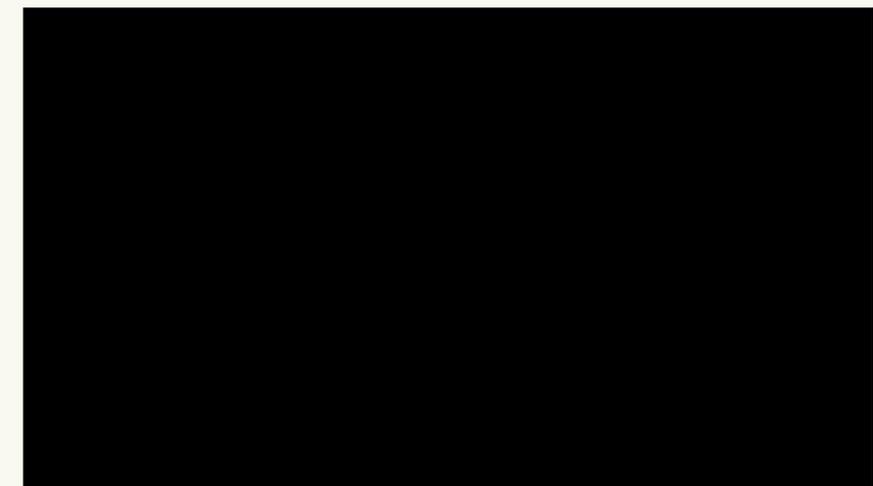
[
 joints orientation,
 body parts location,
 collisions,
 goal vector,
 terrain map,
 ...,
]



Actions

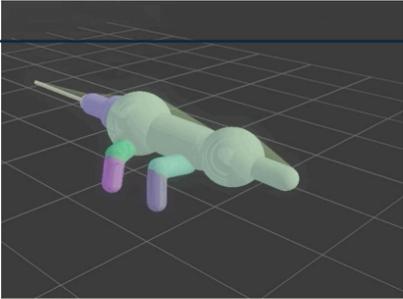
[
 controllable joints
 desired orientations,
 ...,
]

World
physics
step



Observations

[
 joints orientation,
 body parts location,
 collisions,
 goal vector,
 terrain map,
 ...,
]



Actions

[
 controllable joints
 desired orientations,
 ...,
]

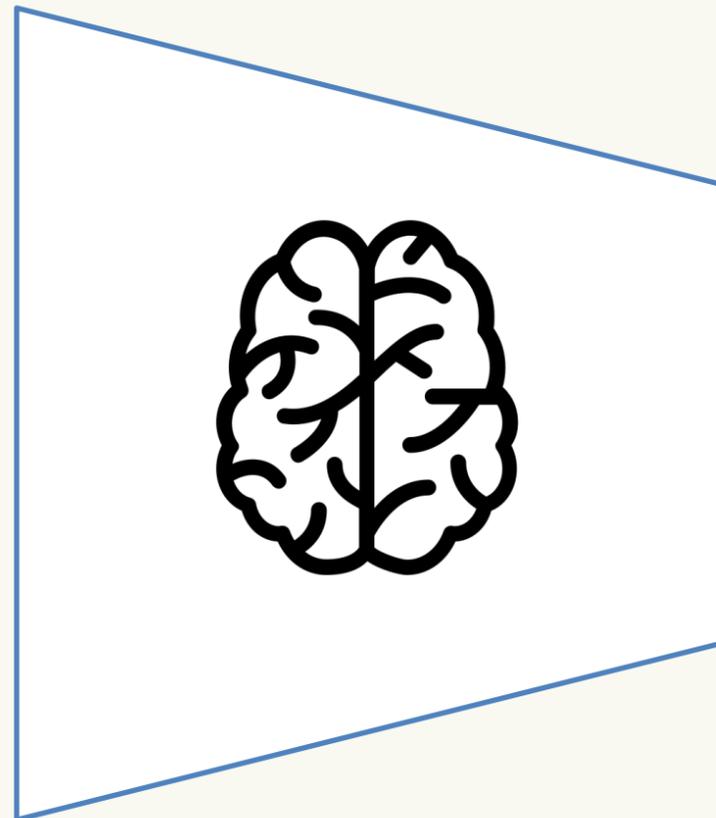
World
physics
step



The BRAIN / CONTROL POLICY

Observations

o

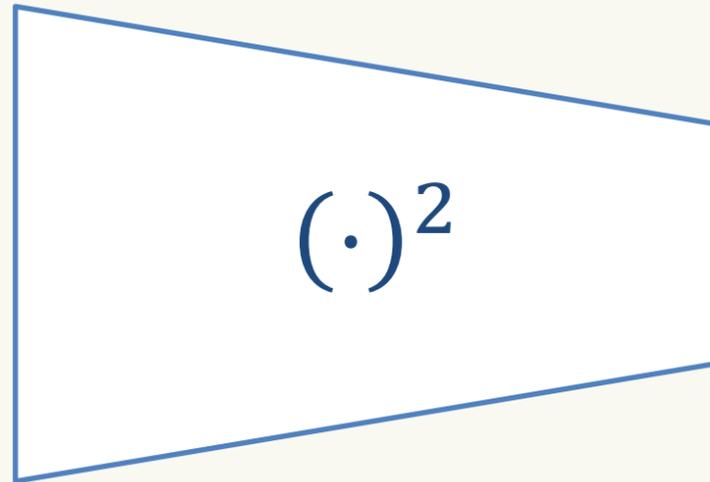


Actions

a

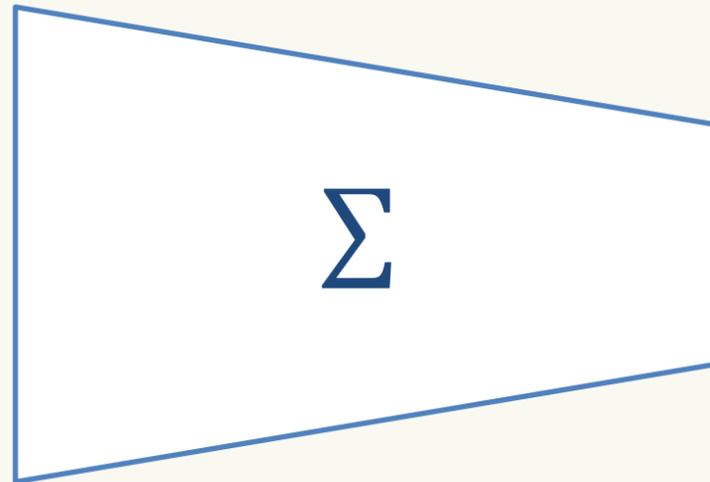
Wait, input-output relationship?

x



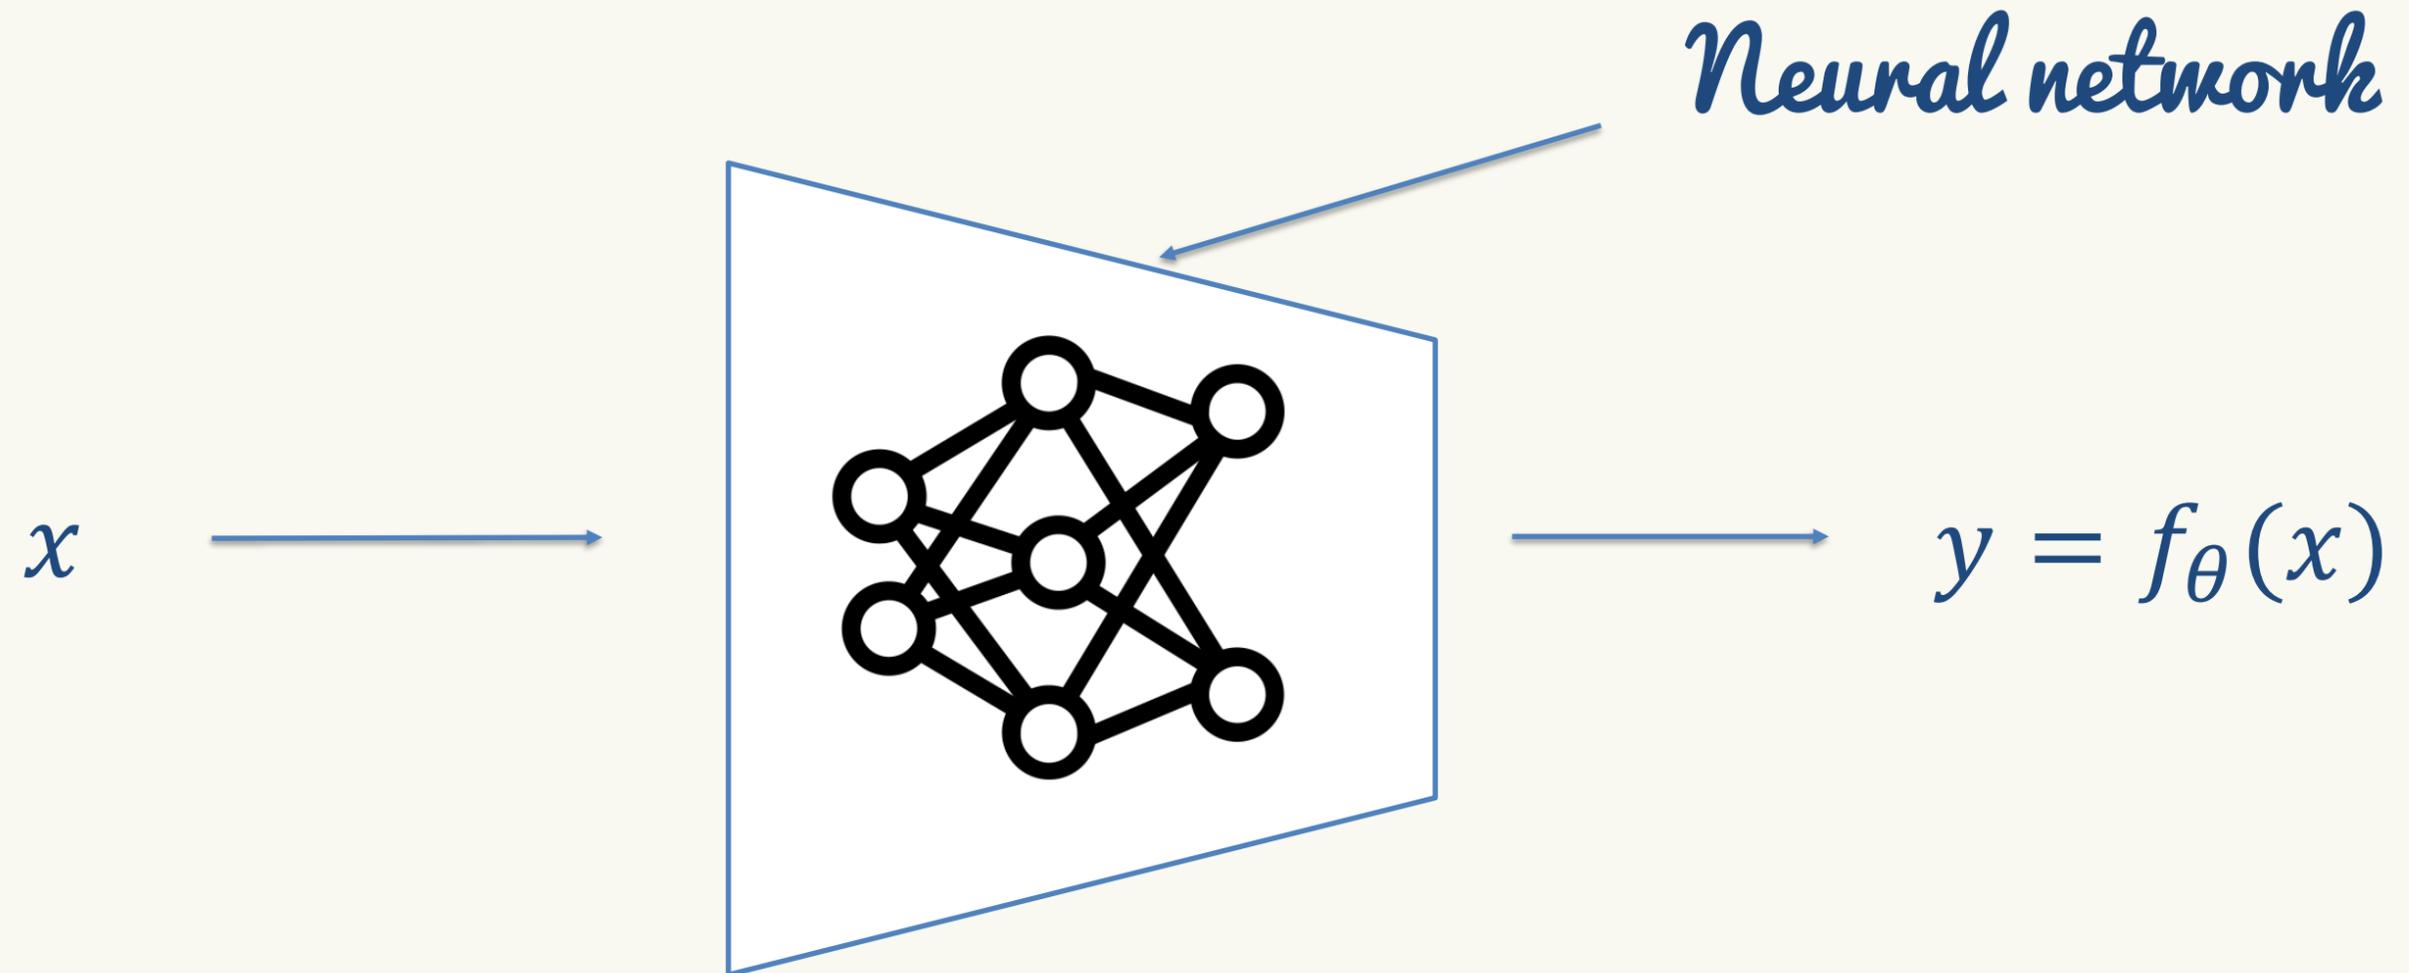
$$y = x^2$$

(x_1, x_2, x_3)



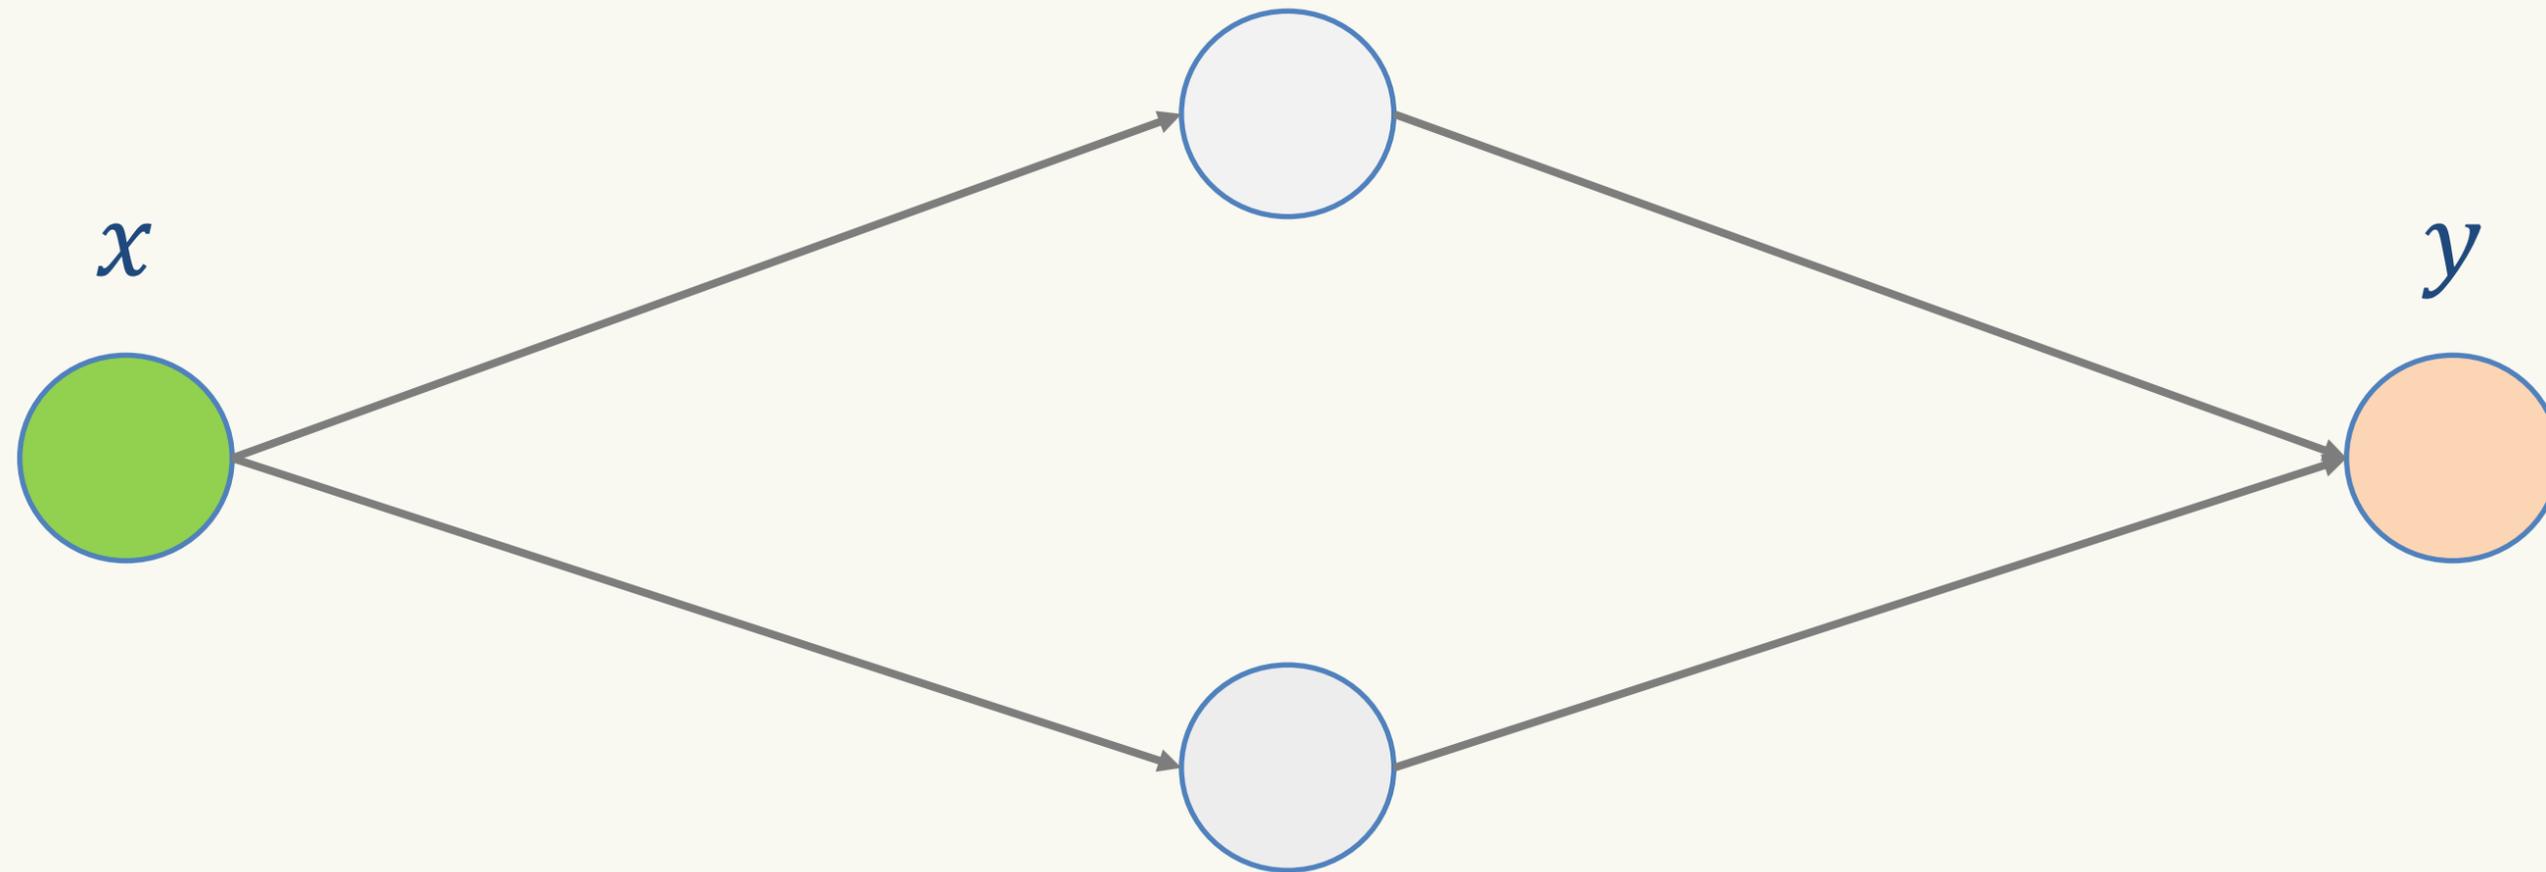
$$y = x_1 + x_2 + x_3$$

Wait, input-output relationship?

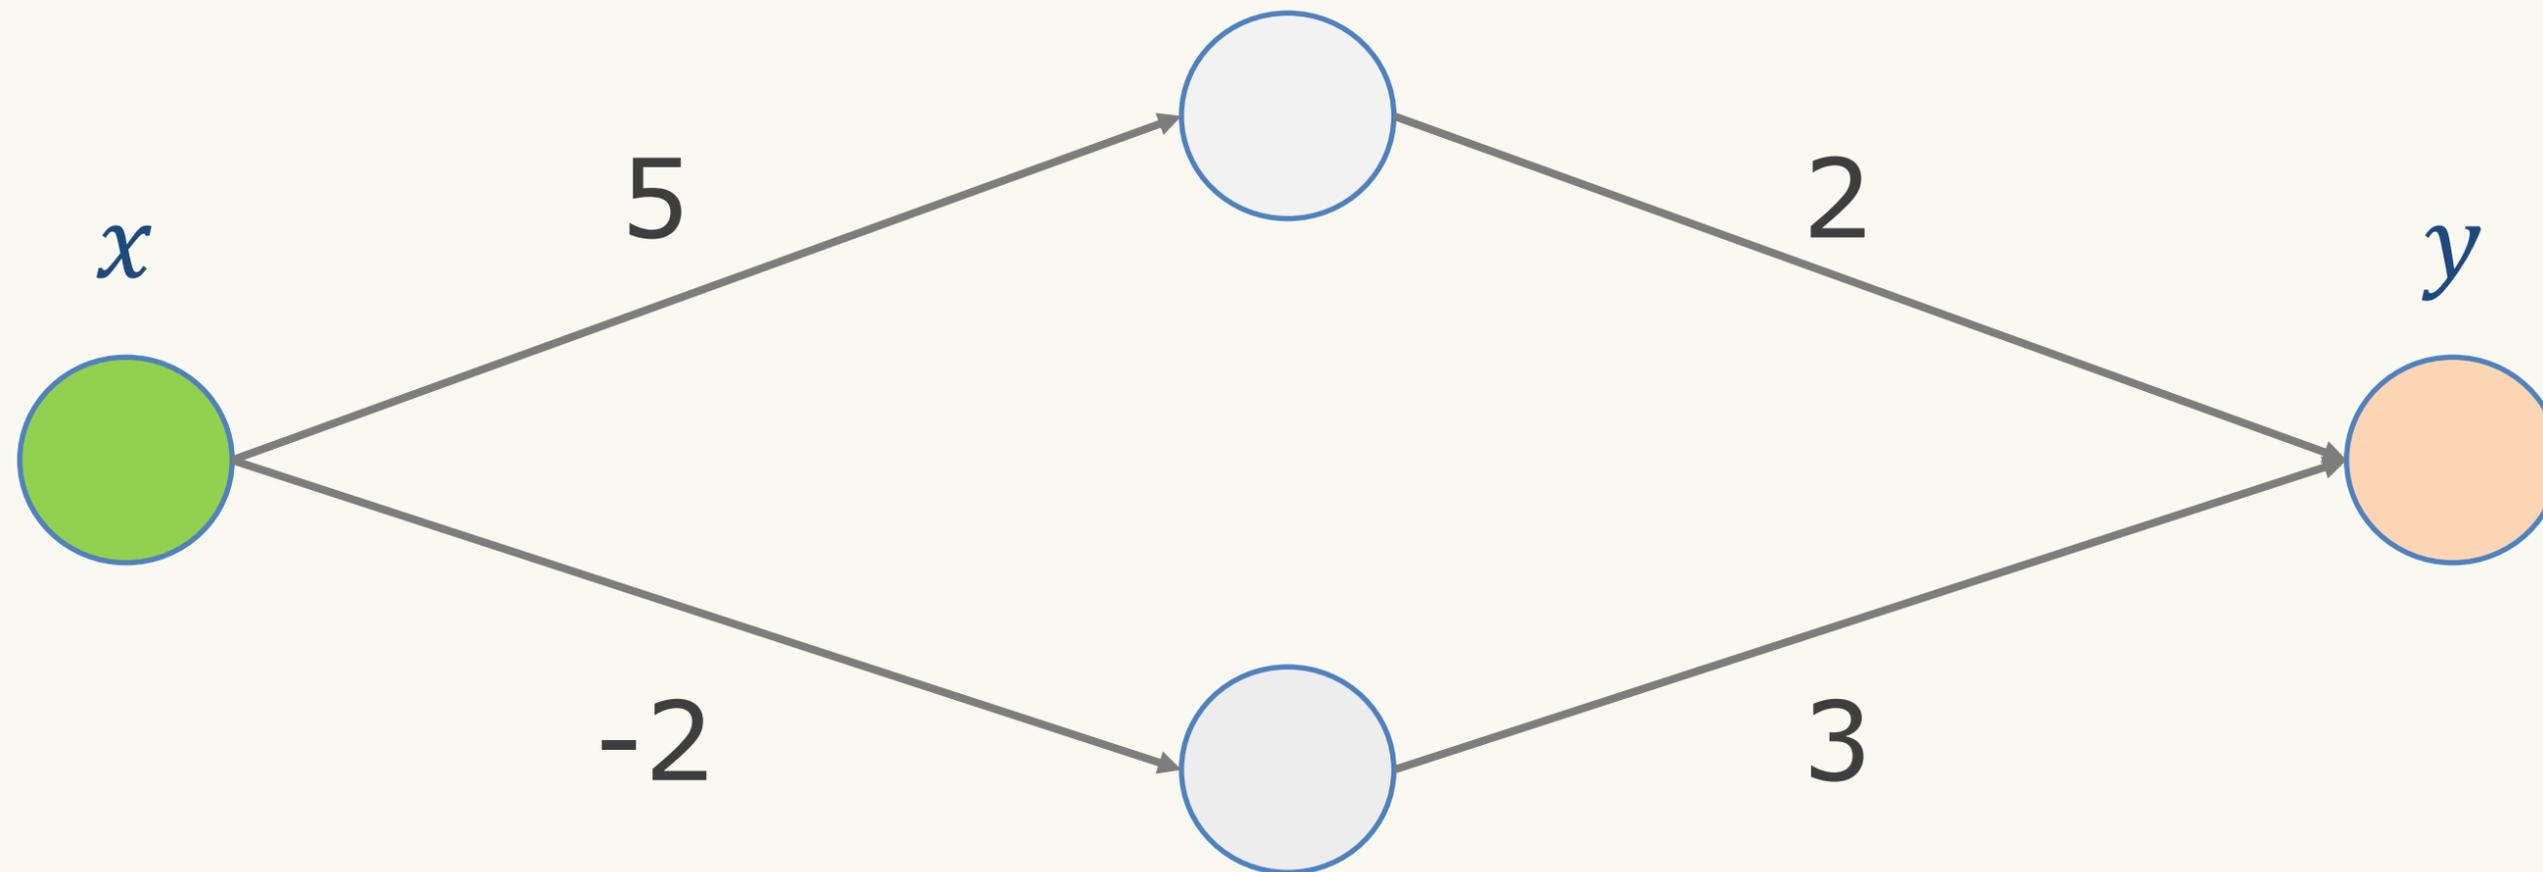


f_{θ} = "anything", given the right θ

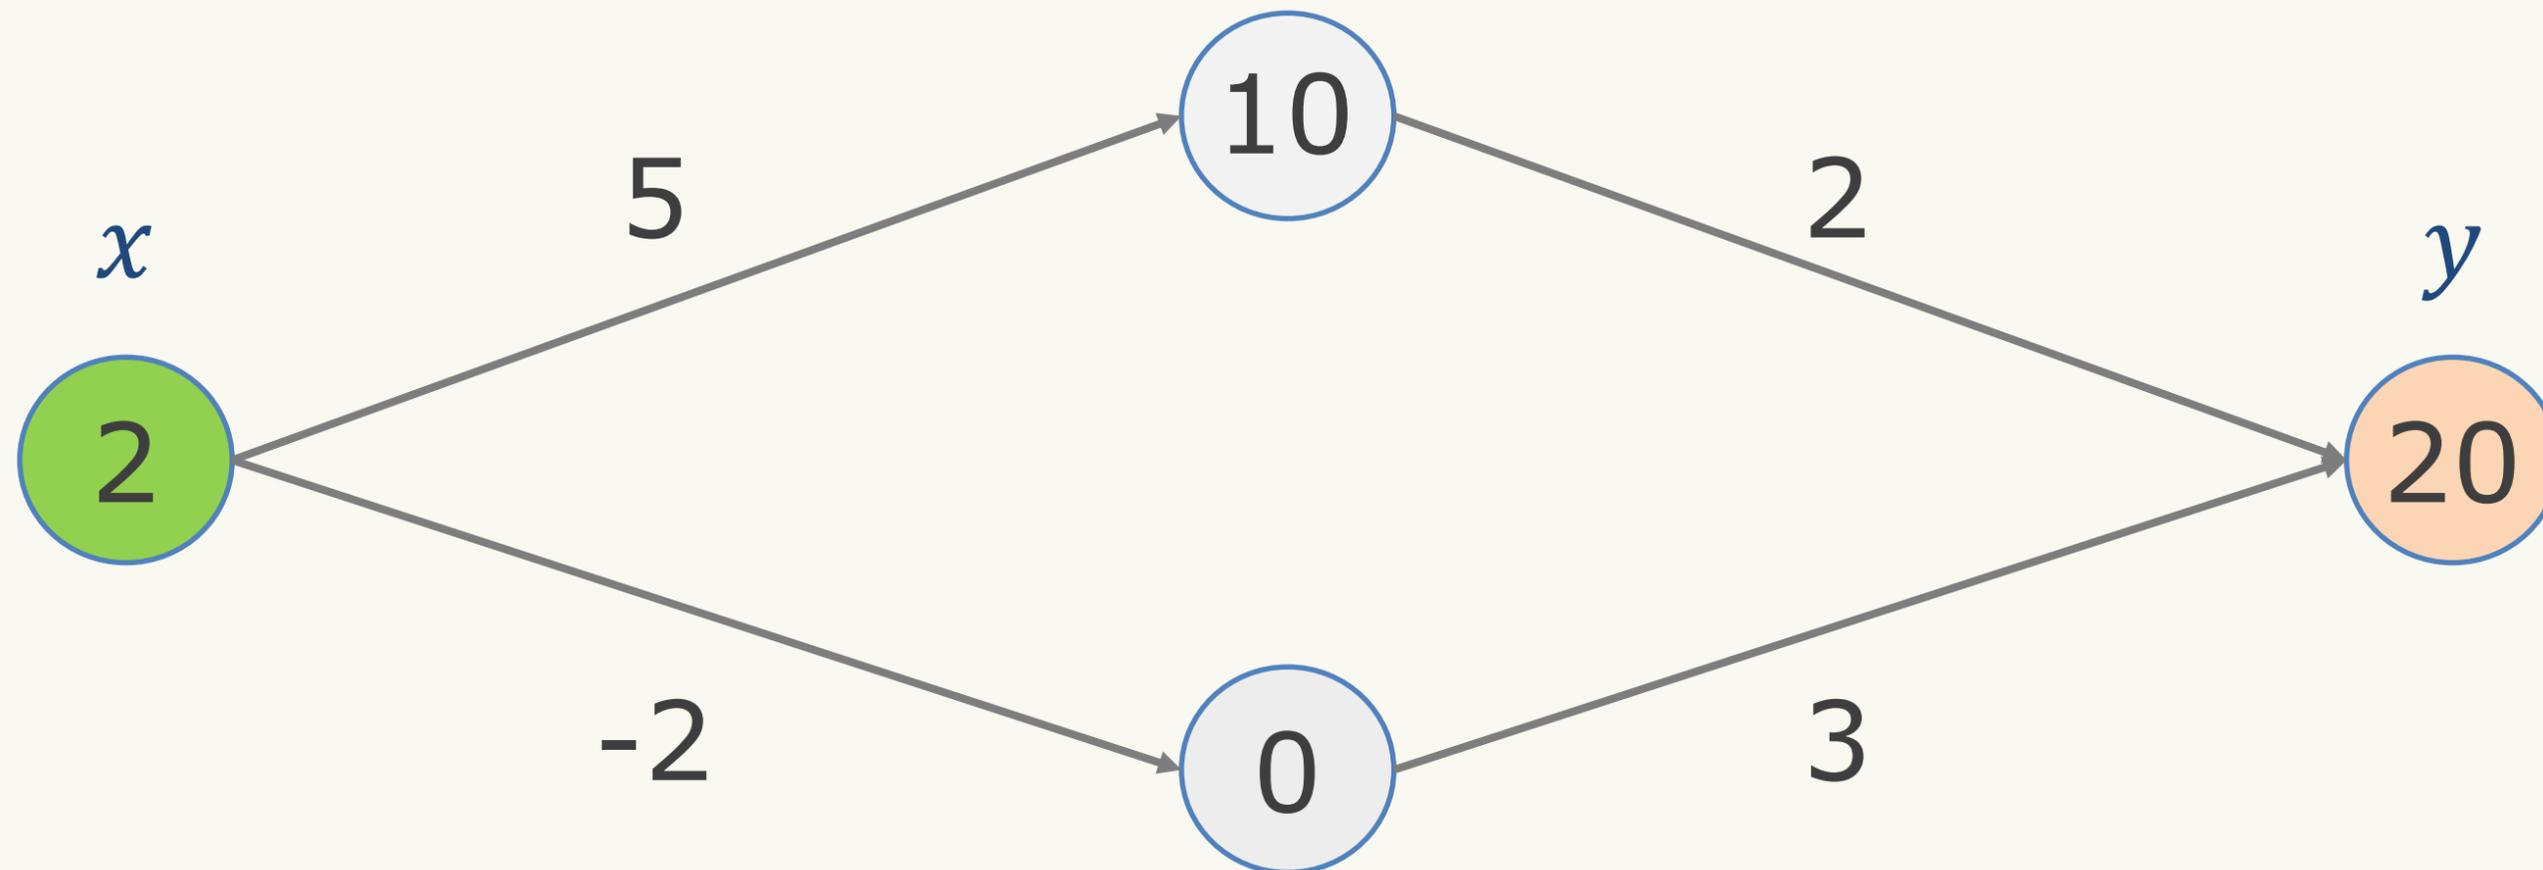
How neural networks work?



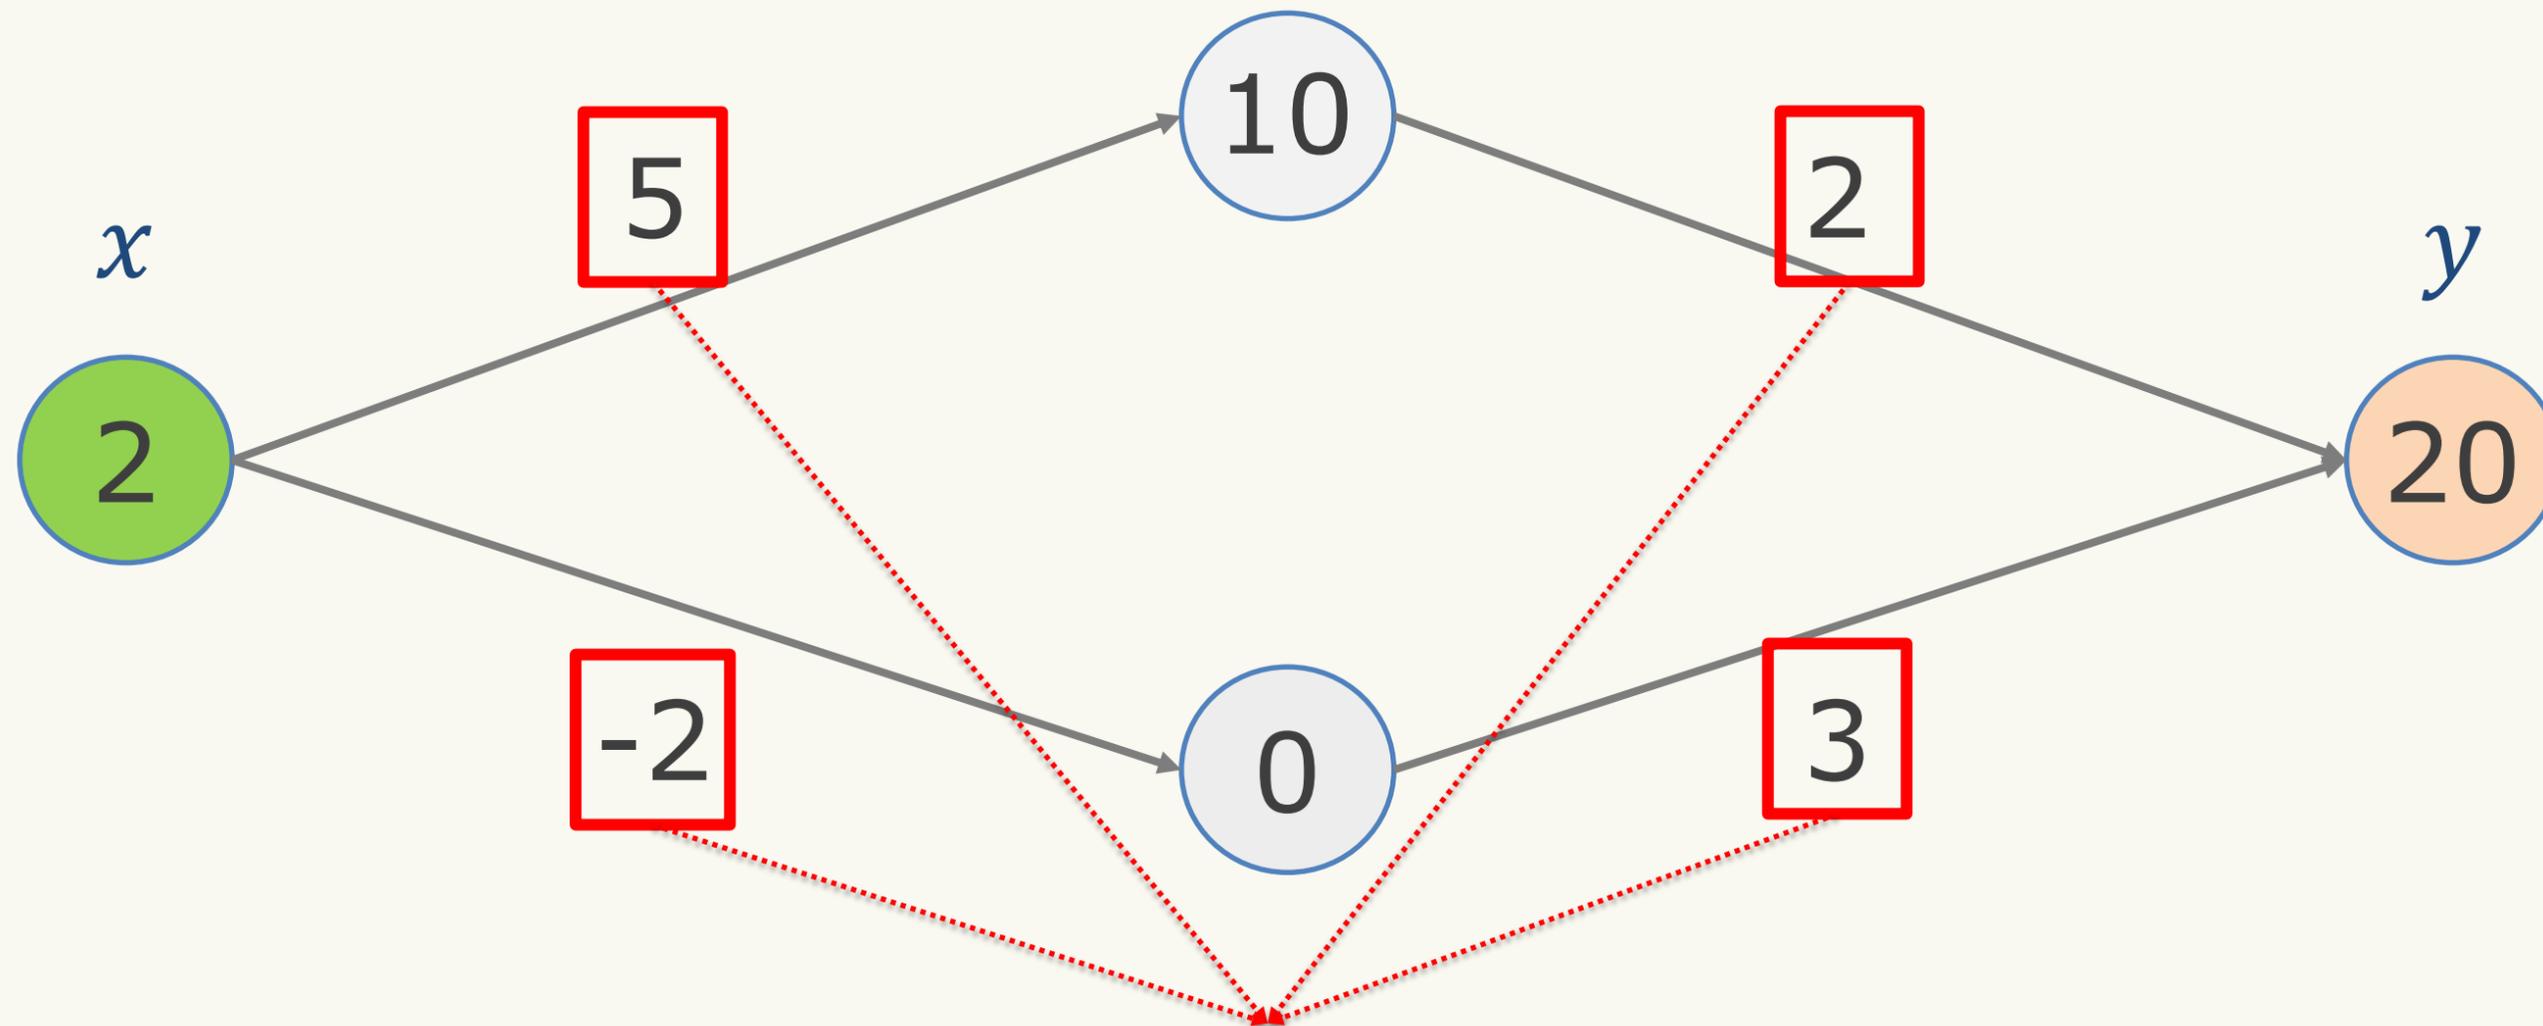
How neural networks work?



How neural networks work?

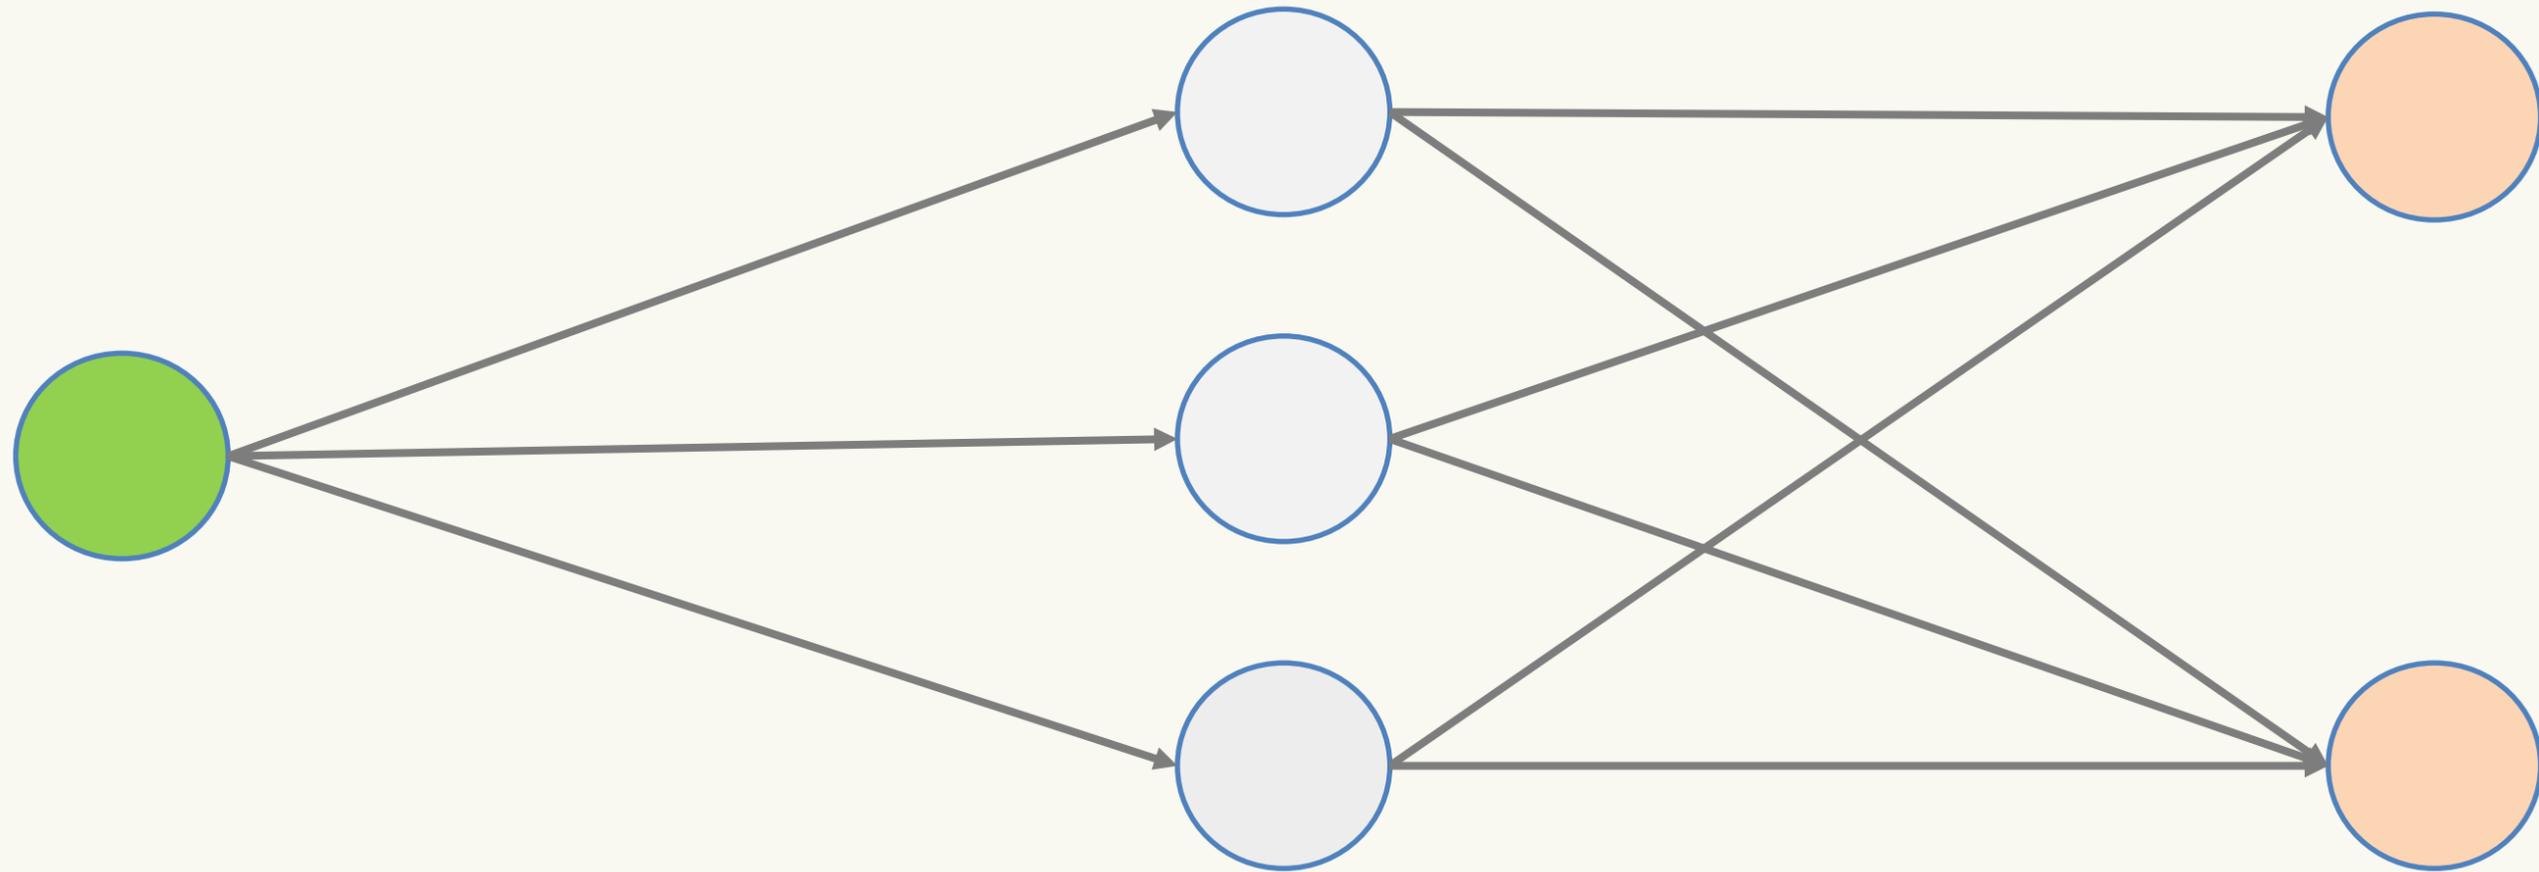


How neural networks work?

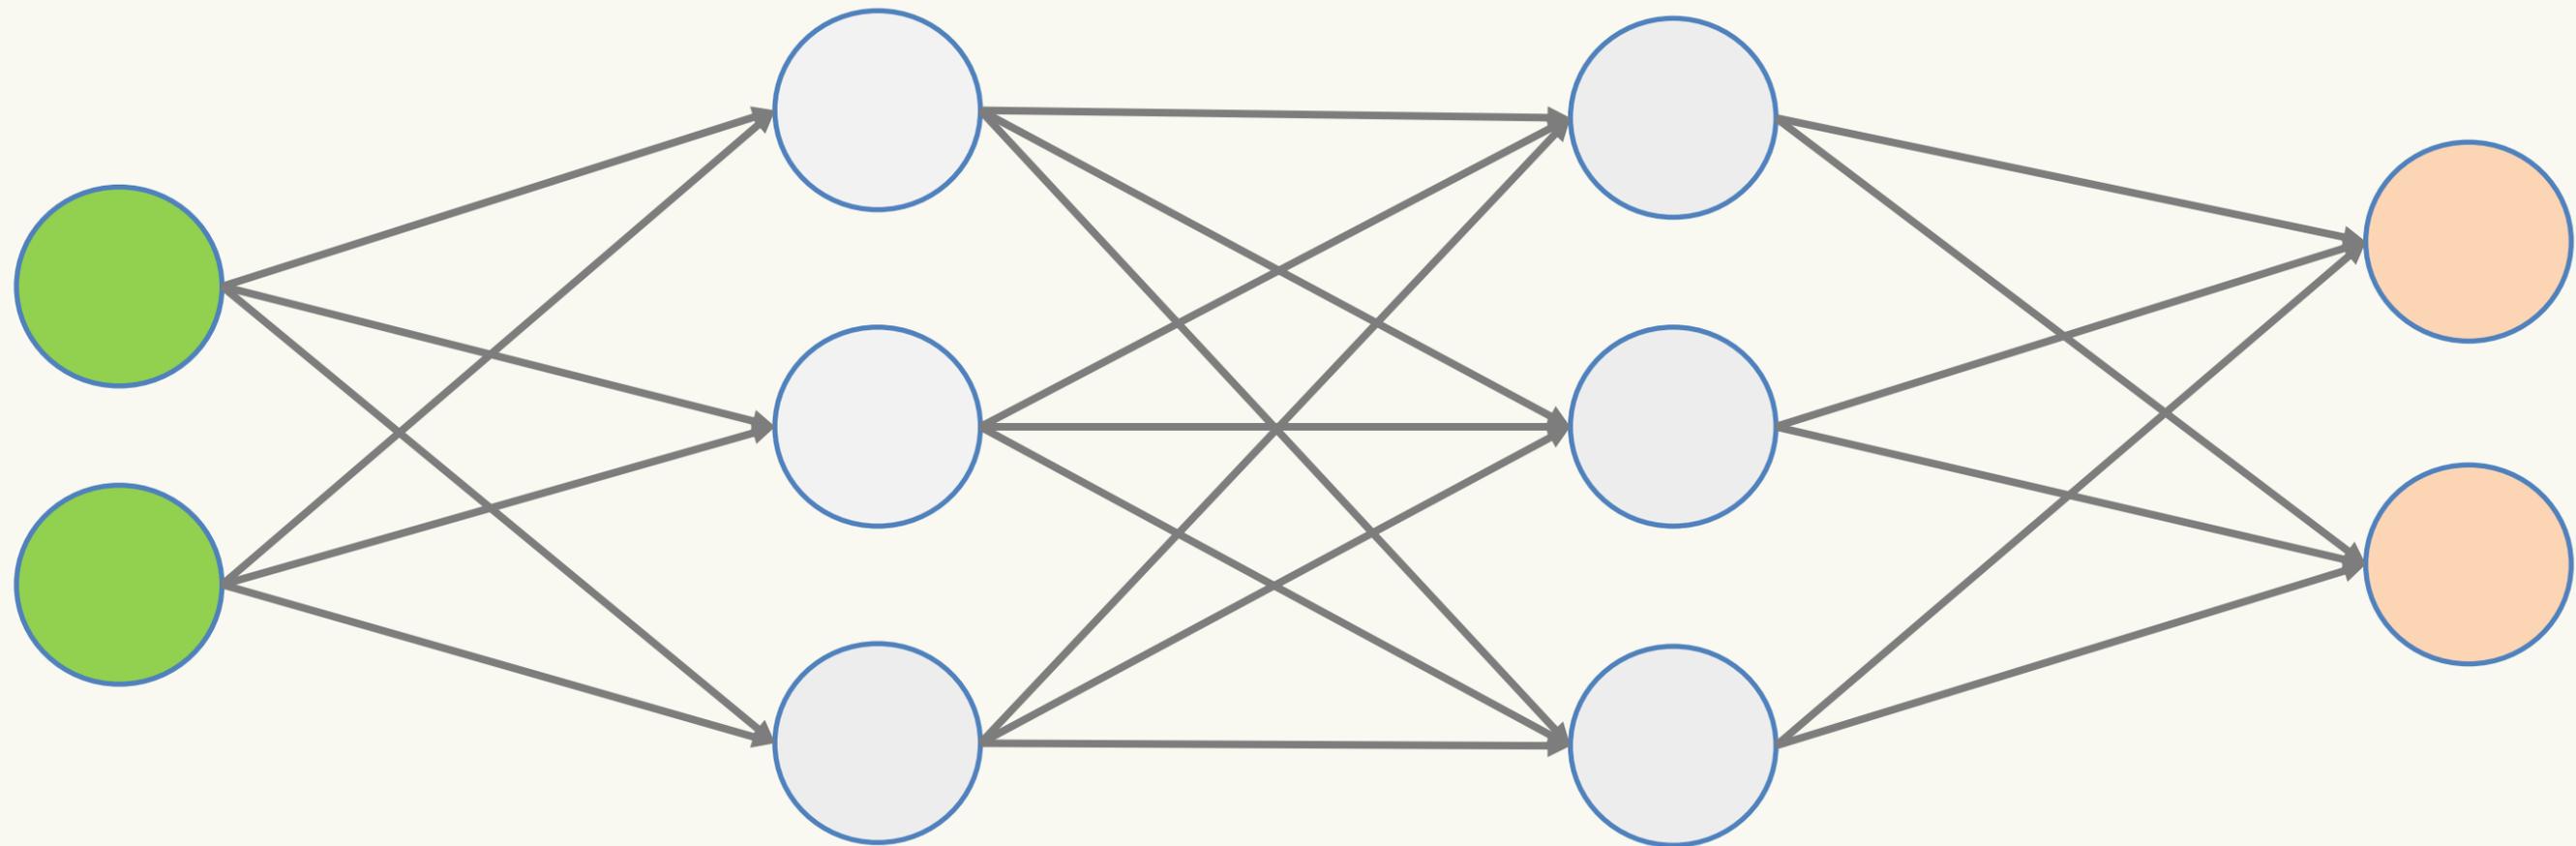


Parameters / weights! θ

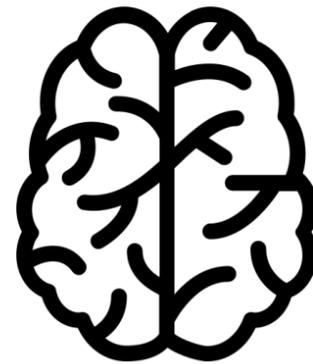
How neural networks work?



How neural networks work?



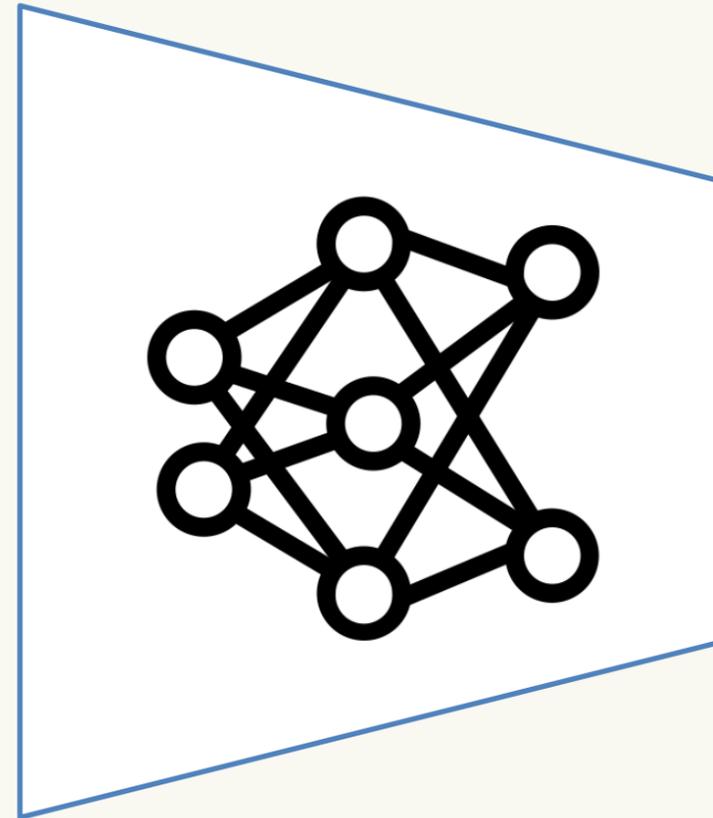
O (observations)



y = the "right" action!

f = "optimal controller"

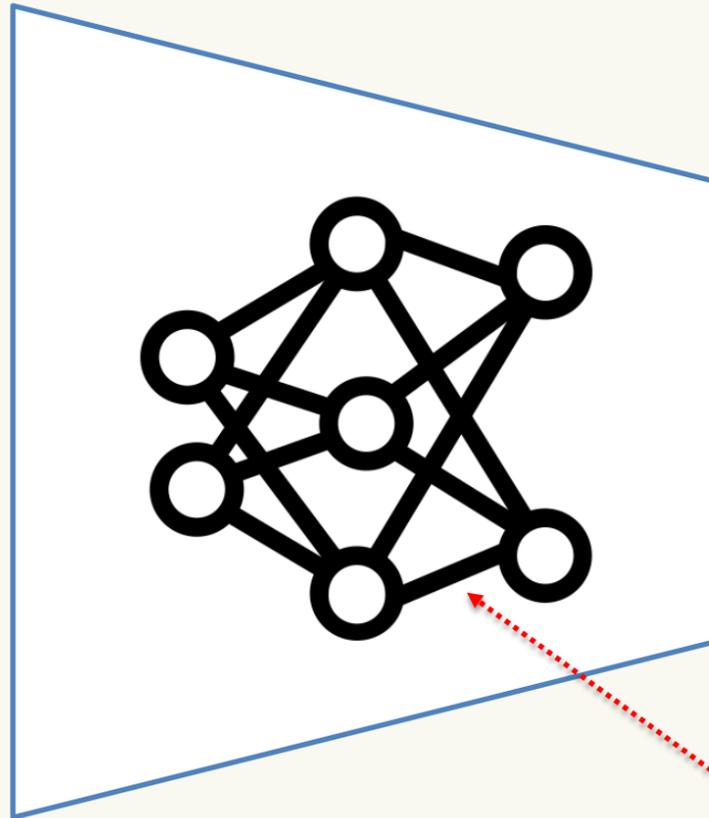
O (observations)



y = the "right" action!

f_{θ} = "optimal controller"

O (observations)

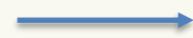


y = the "right" action!

Parameters??

f_{θ} = "optimal controller"

O (observations)



y = the "right" action!

f_{θ} = "optimal controller"

Reinforcement learning

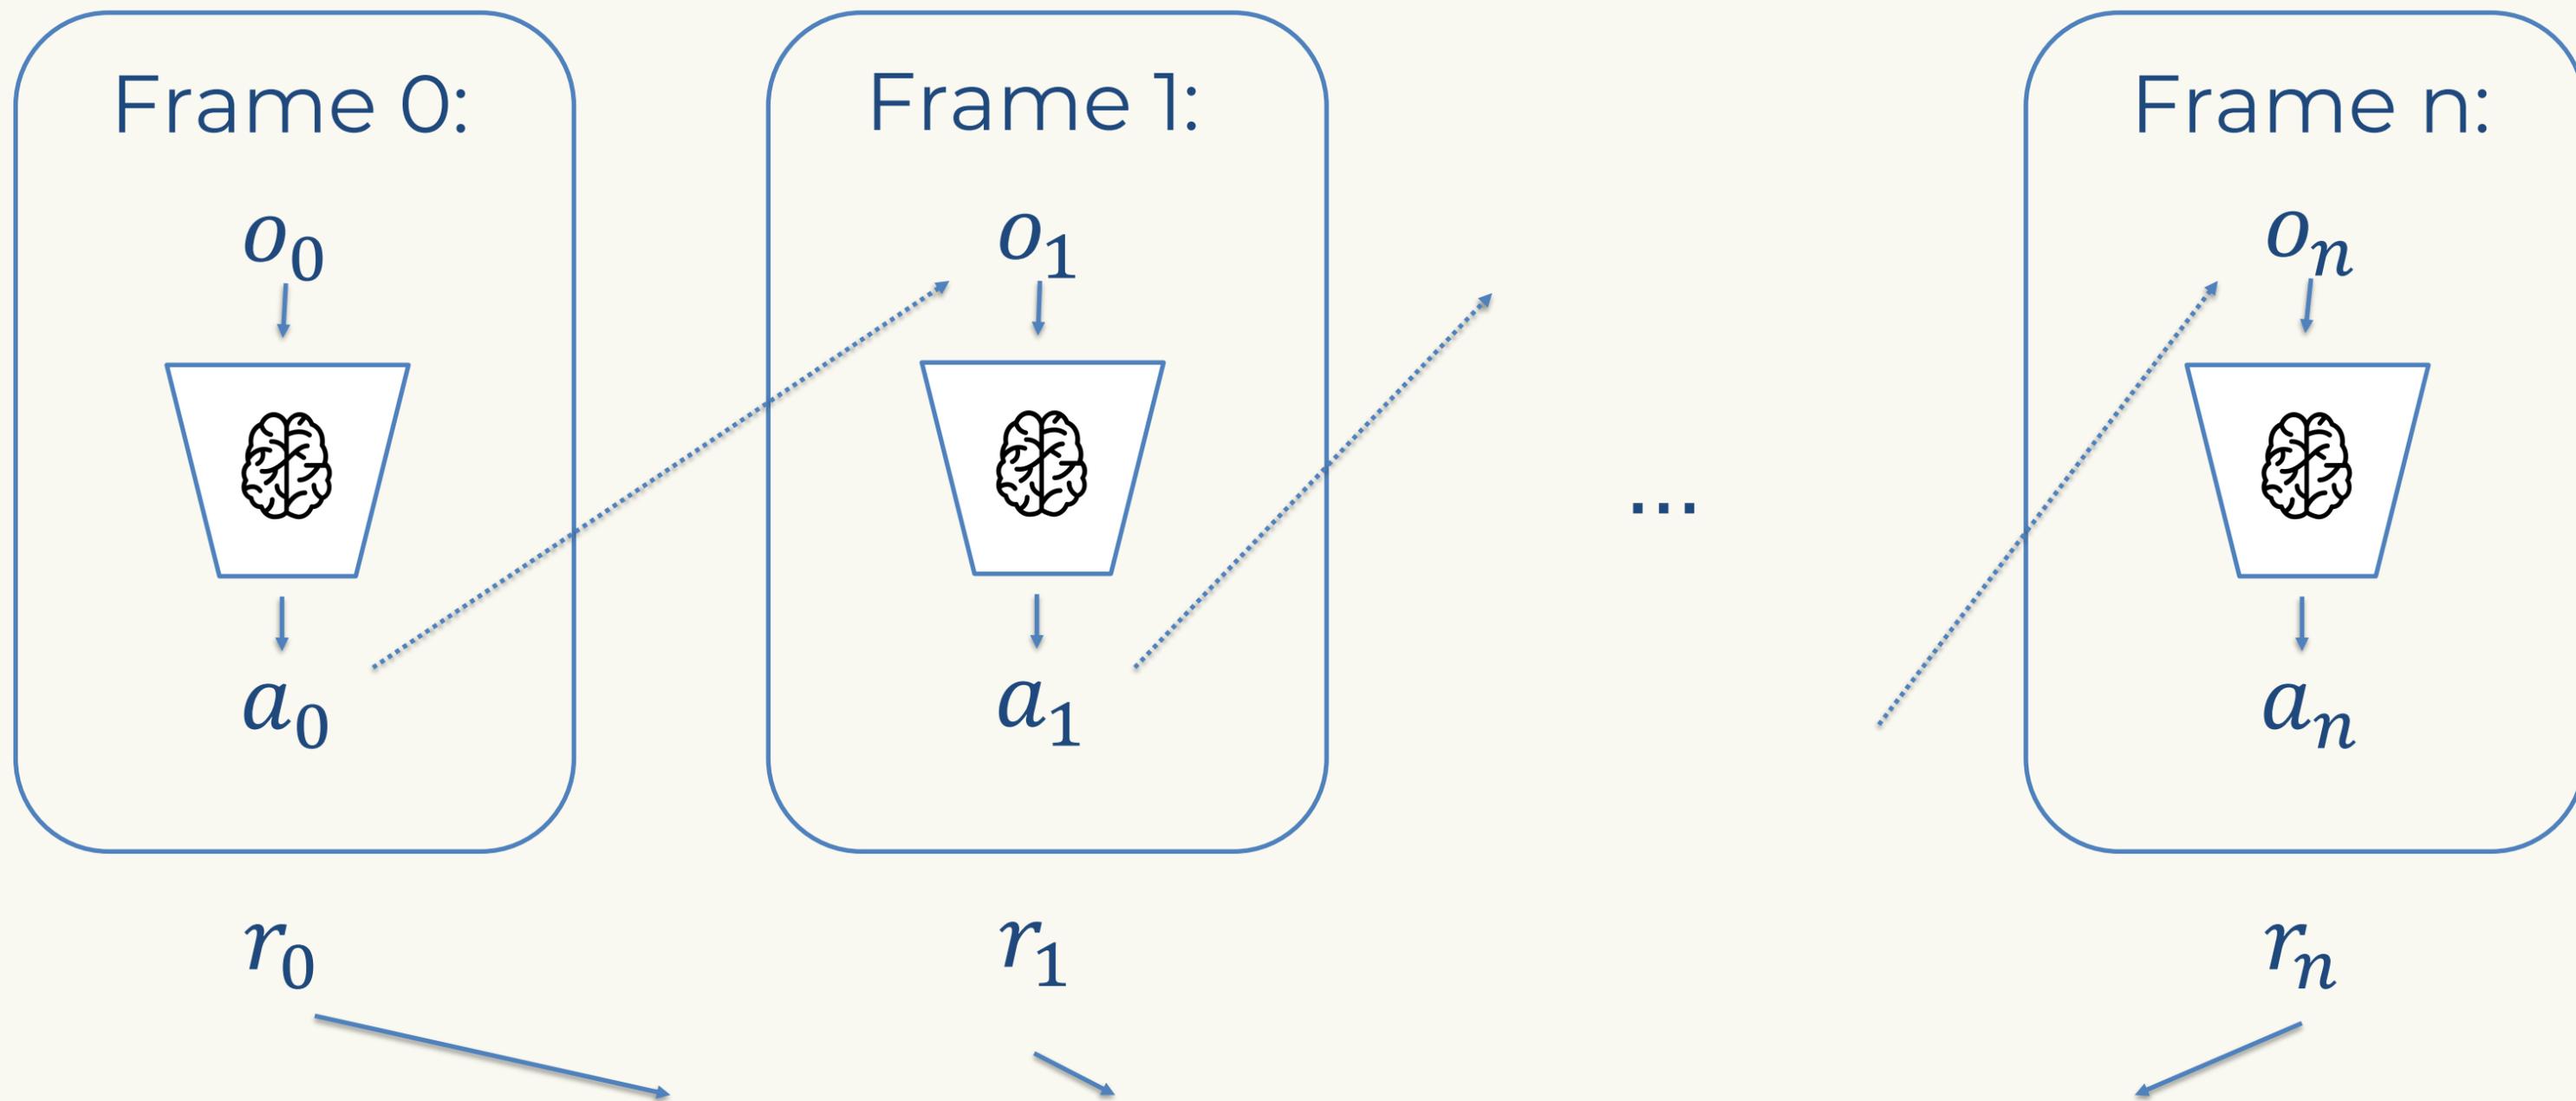
Defining “right” - it’s all about rewards

- A **reward function** evaluates a state and action frame by frame.
- E.g.
 - +2 if you get closer to the goal
 - +1 if you face the goal
 - -100 if you fall
 - +100 if you get to the goal

It's all about MAXIMIZING future rewards

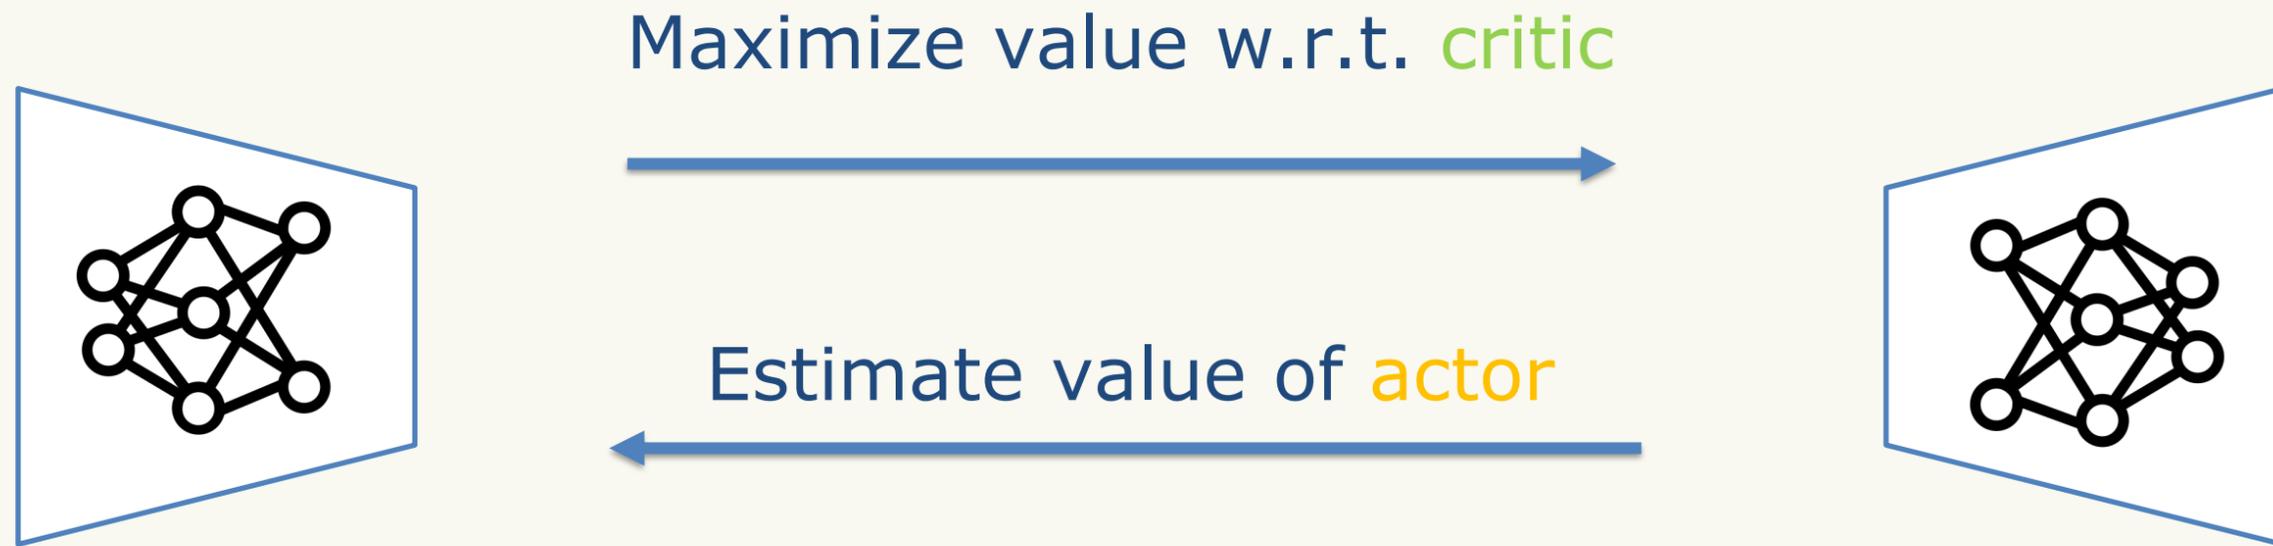
"Value" of π

"Find policy π which optimizes sum of future rewards"



“Value of  starting at o_0 ” = $r_0 + \gamma r_1 + \gamma^2 r_2 + \dots + \gamma^n r_n$

Two networks: actor-critic algorithms



Policy network (brain)

Actor

Value network

Critic

Outline

- We parametrize decision policies (brains) with a neural network.
- We use another network to evaluate the value of future rewards.
- We fit them together using **soft actor-critic** algorithm (SAC).

Soft Actor-Critic Algorithms and Applications. Haarnoja et al. 2018

Some results

Lizzie learning (t=0)

[Lizzie learning \(link\)](#)

Lizzie walking (t~4h)

[Lizzie walking \(link\)](#)

More user-created creatures

[More creatures \(link\)](#)

Even... bipeds

[Biped in flat terrain \(link\)](#)

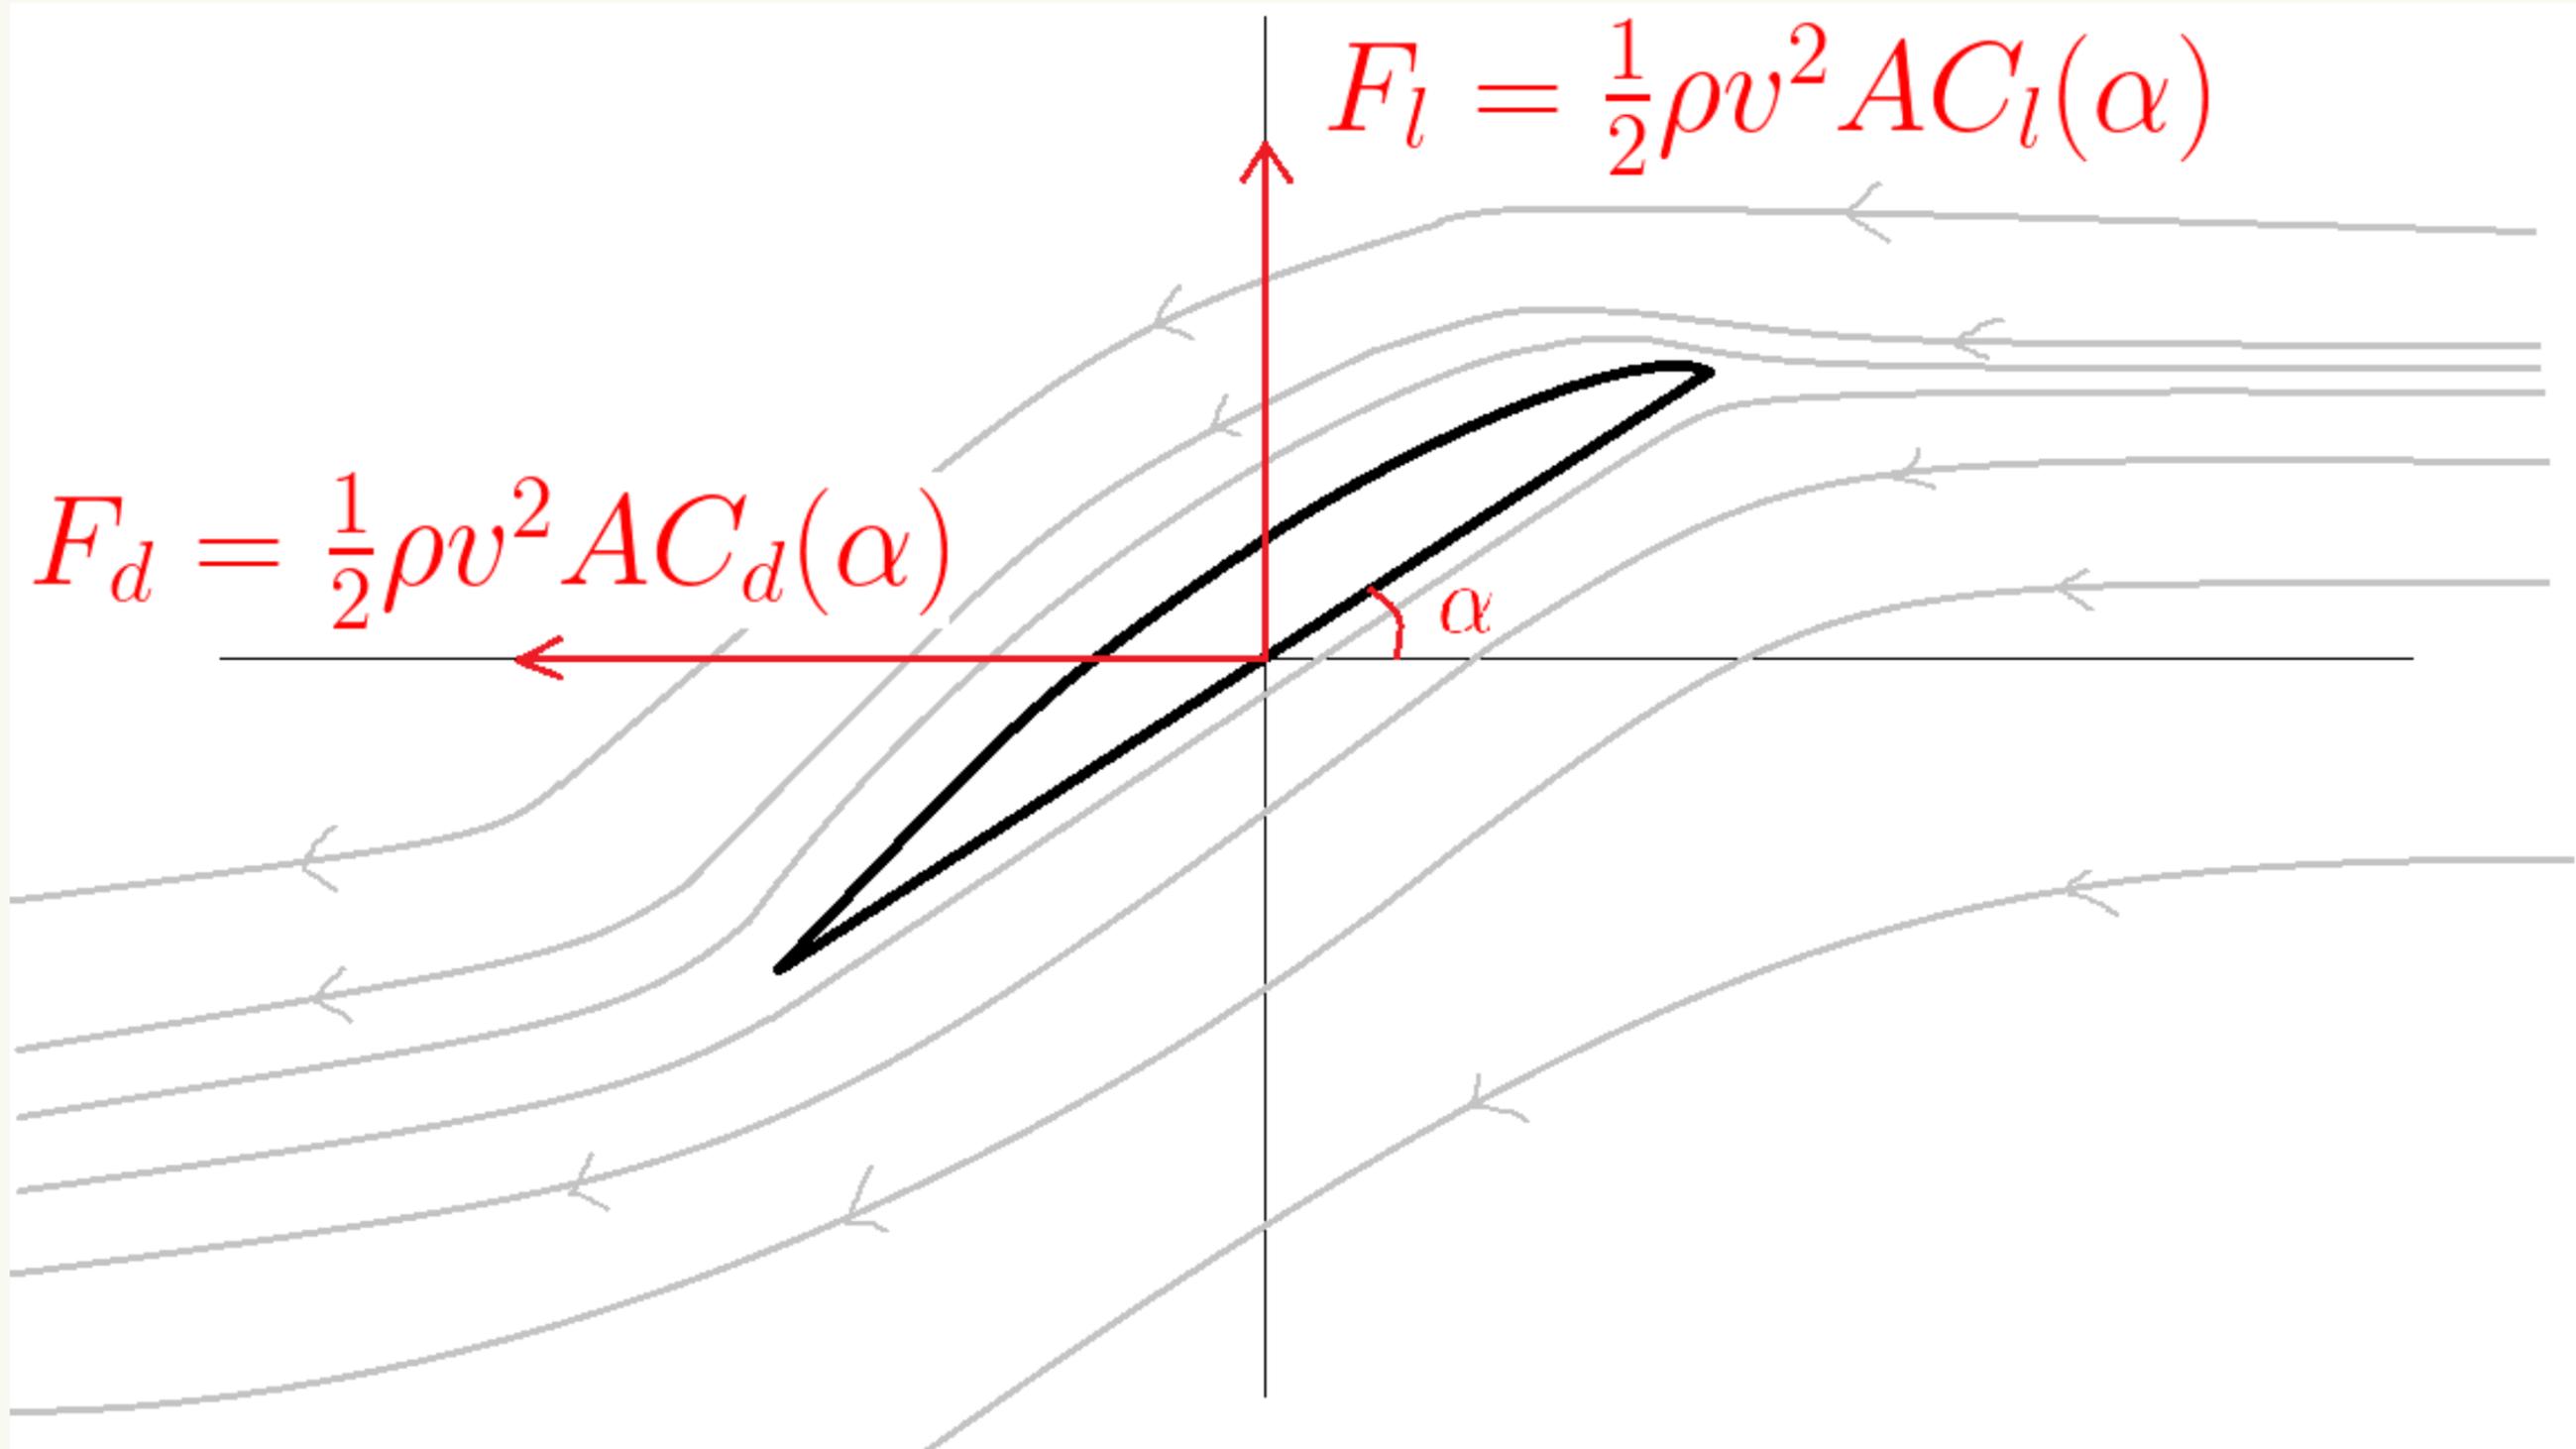
This is a basic setup. I will show improved ones later on.

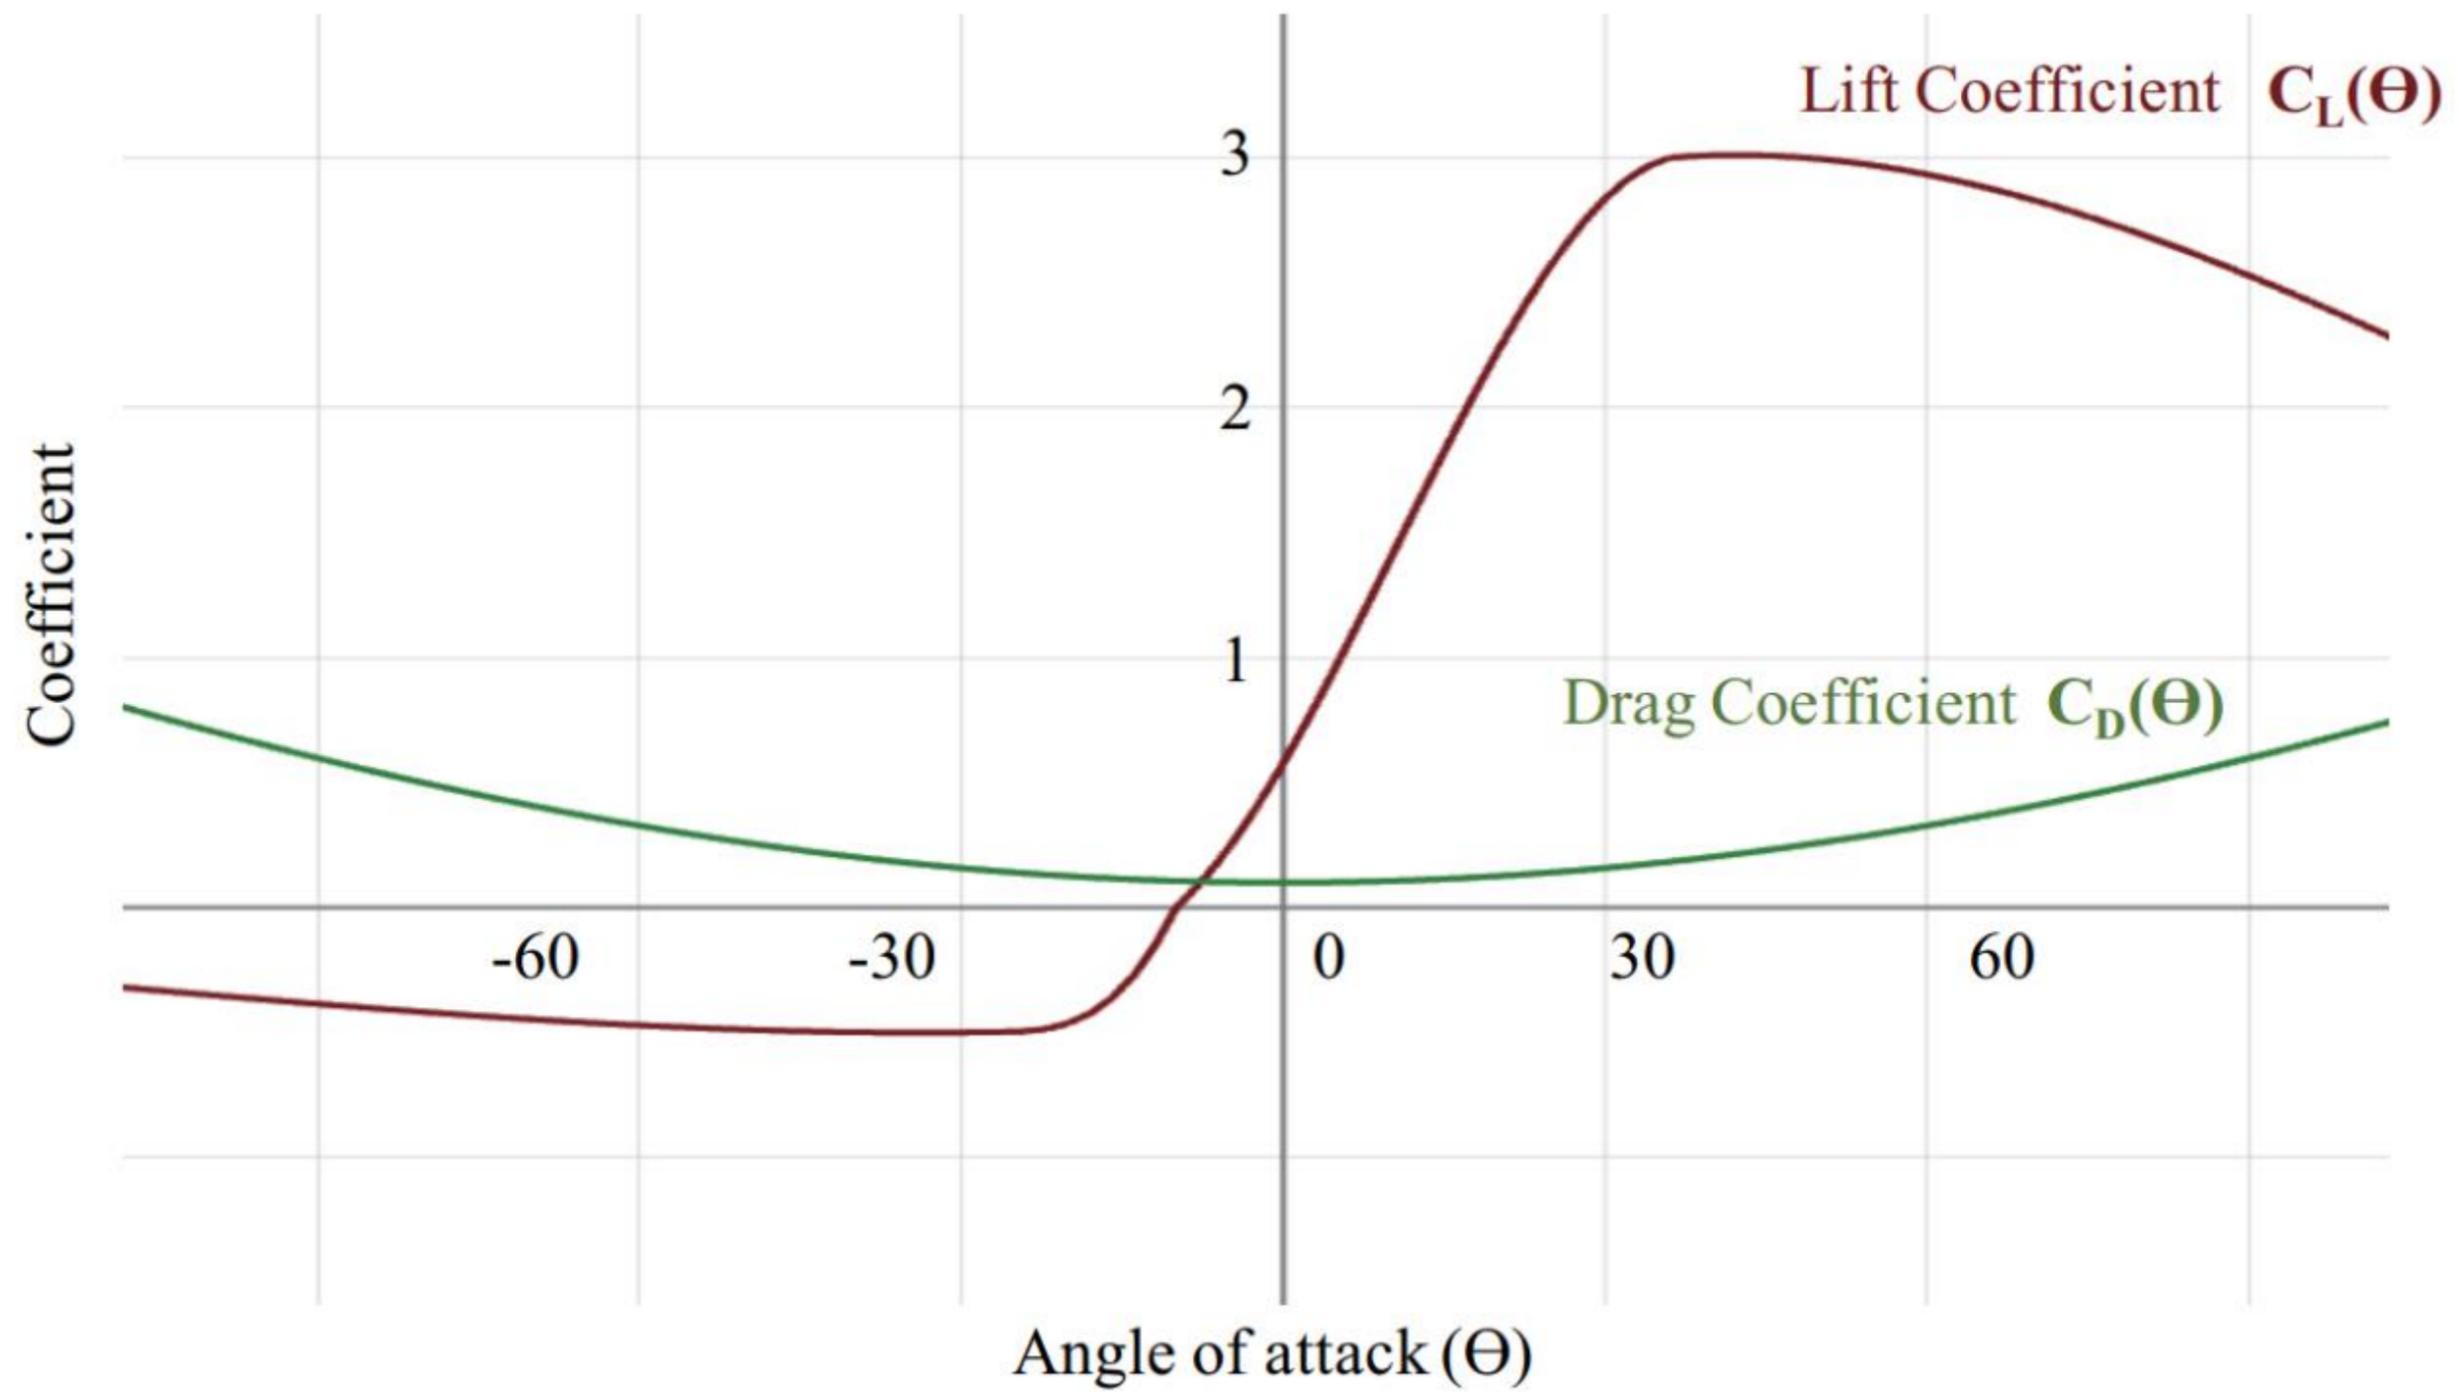
DOES THIS WORK



IF I CHANGE PHYSICS?

imgflip.com

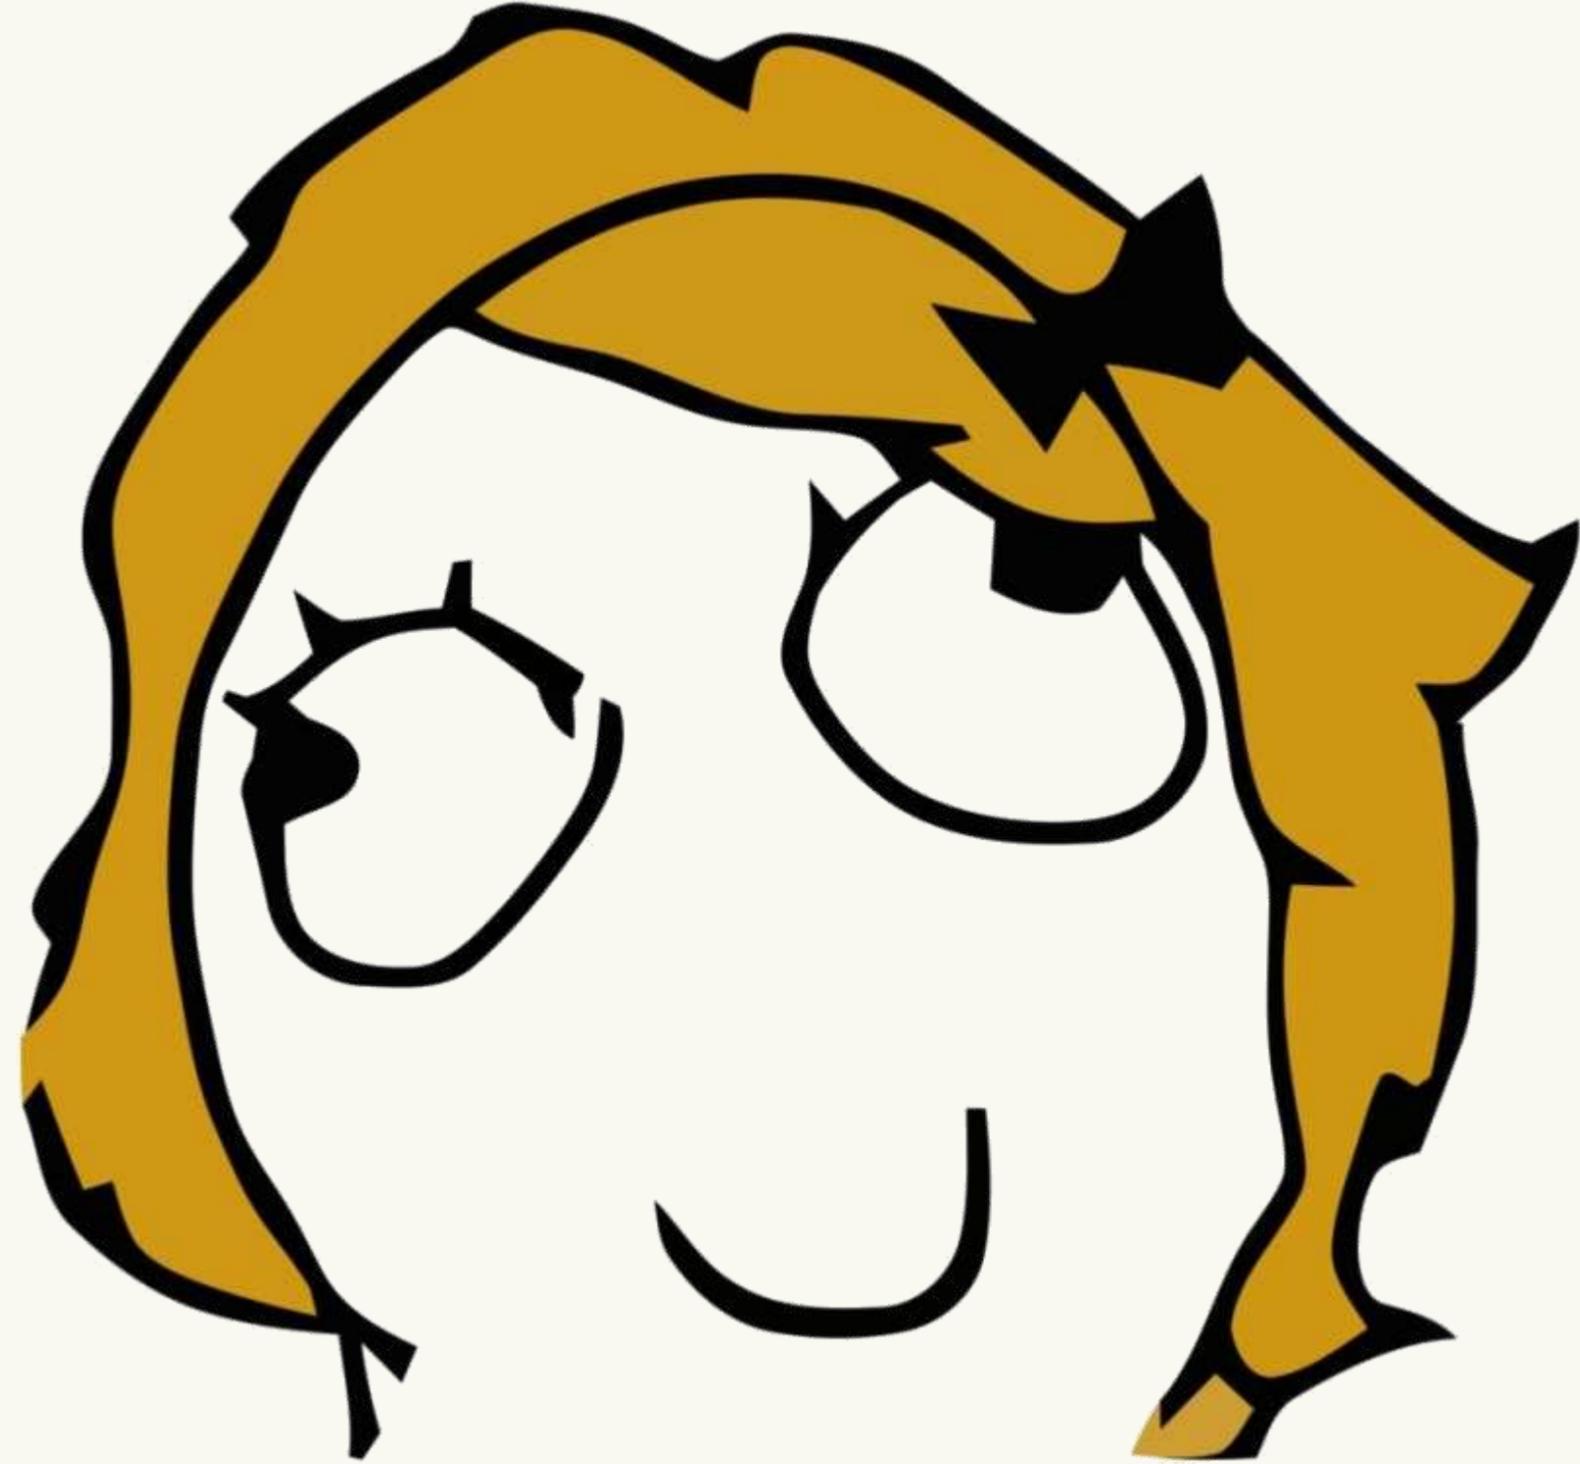


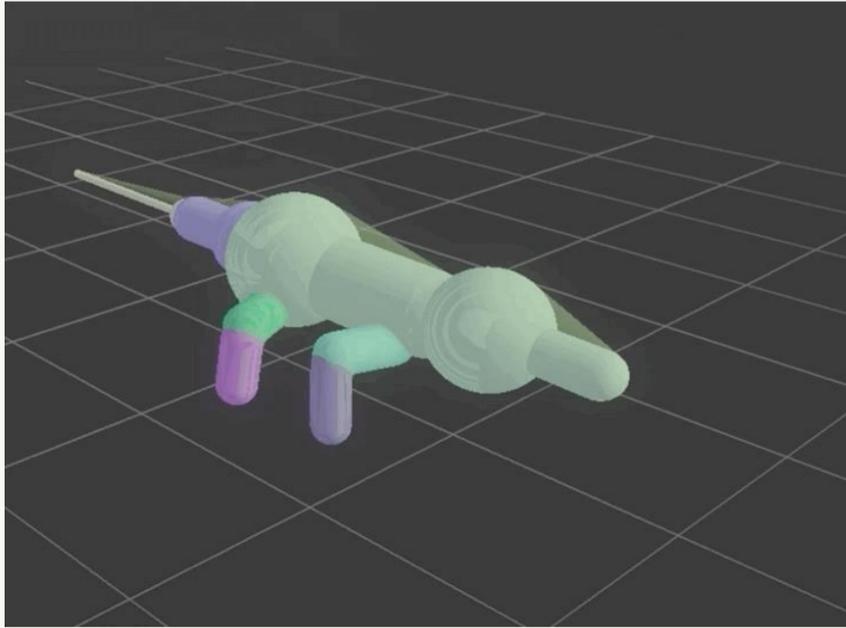


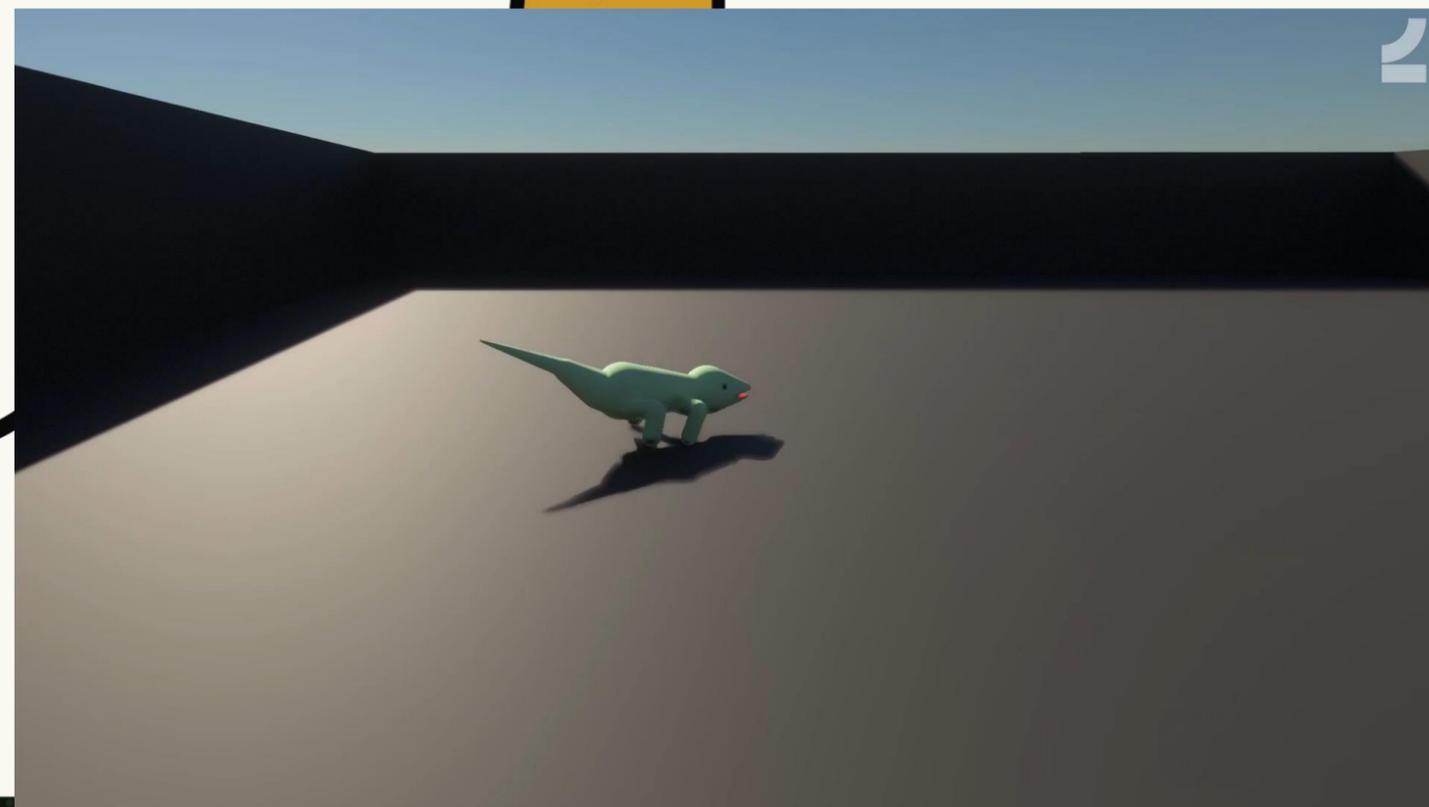
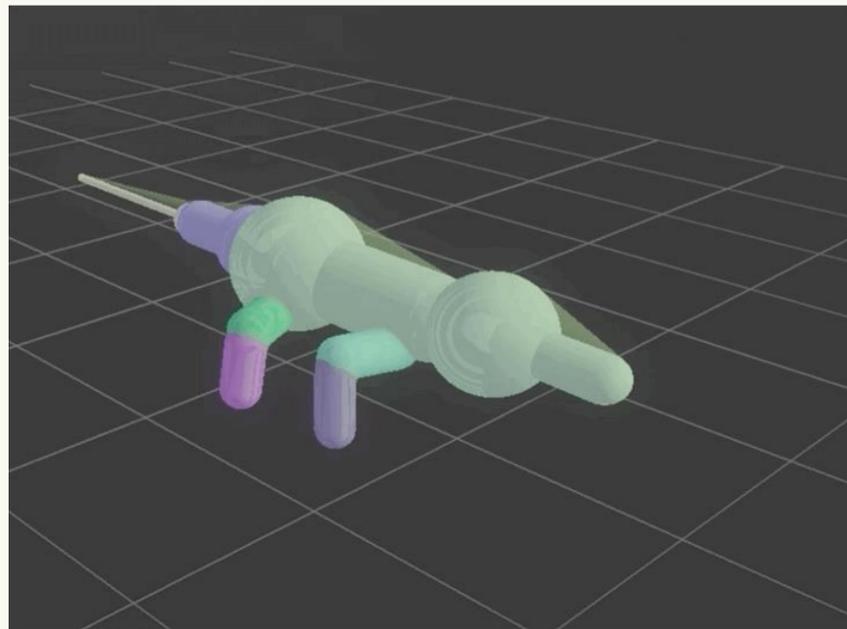
[Flying creature \(link\)](#)



Real time, physics-based flying creatures!





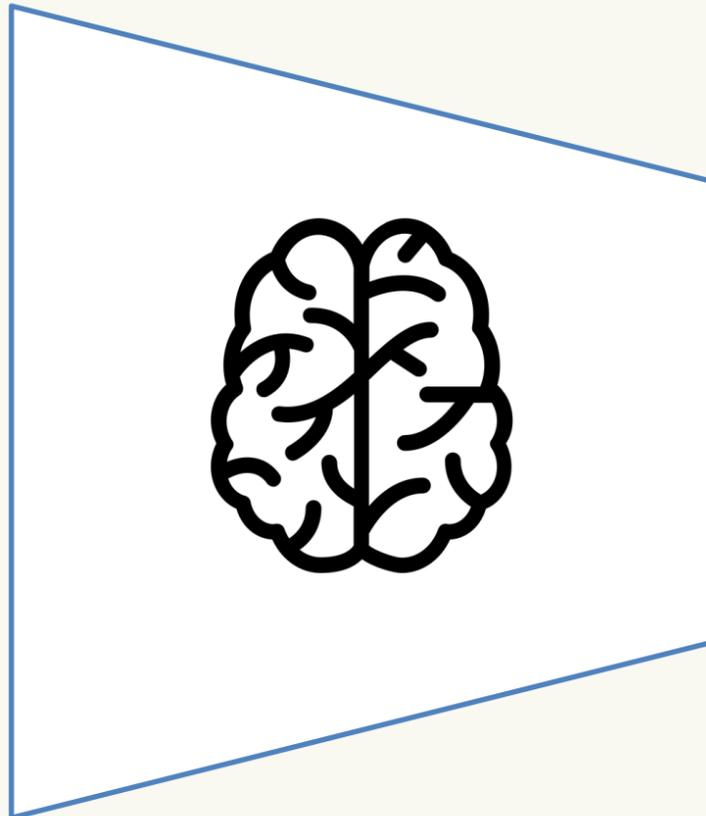




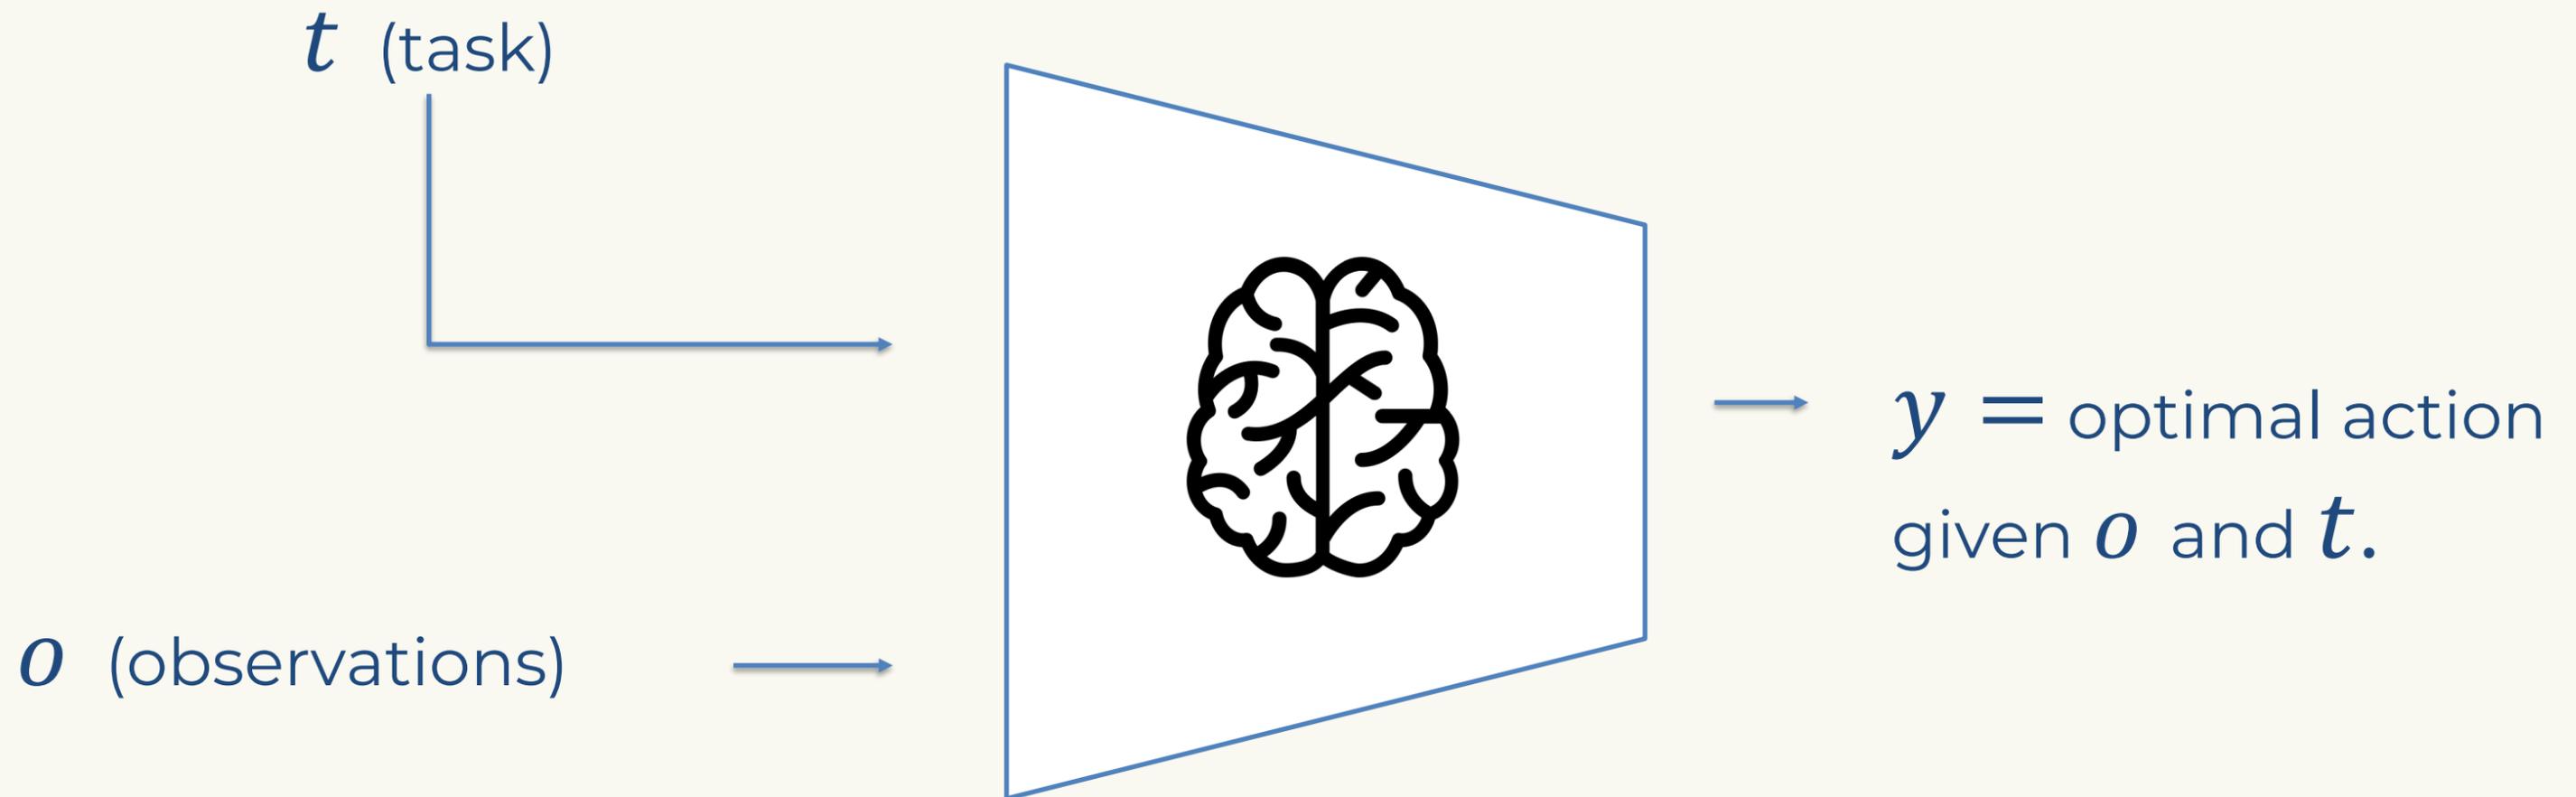
Can they do more than just walking?

High-level control

O (observations)



y = optimal action
given O .



Modular tasks and rewards

- Task #1
 - Reward module 1
 - Reward module 2
 - ...

- Task #2
 - Reward module 2
 - Reward module 3
 - ...

$$\text{Reward for task } \#t = \sum_{\text{module} \in \text{Task } t} r_{\text{module}}(s)$$

[Lizzie stopping \(link\)](#)

[Biped stopping \(link\)](#)

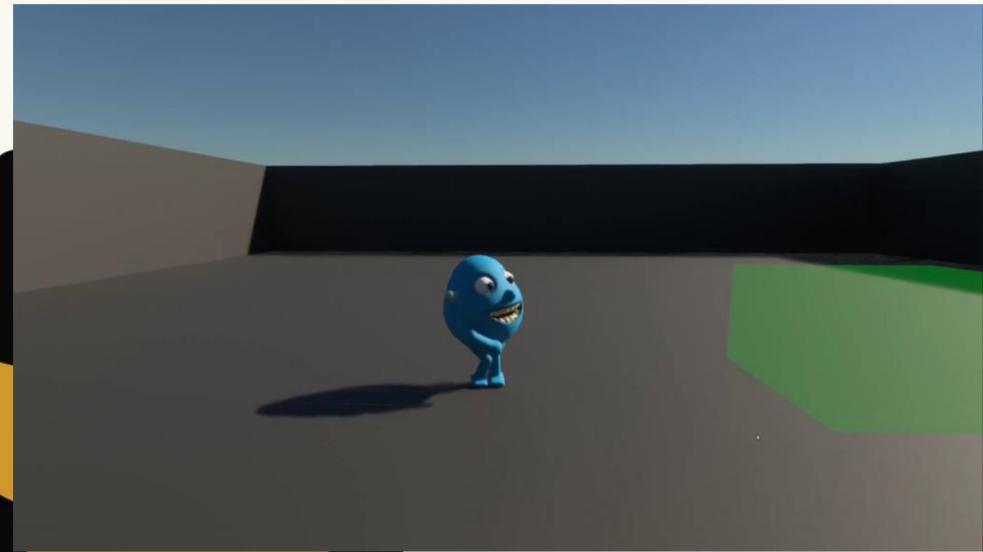
[Cowboy Lizzie jumping \(link\)](#)

[Cowboy Lizzie multiple tasks \(link\)](#)

Music: Grace Behind The Curtain by Silent Partner

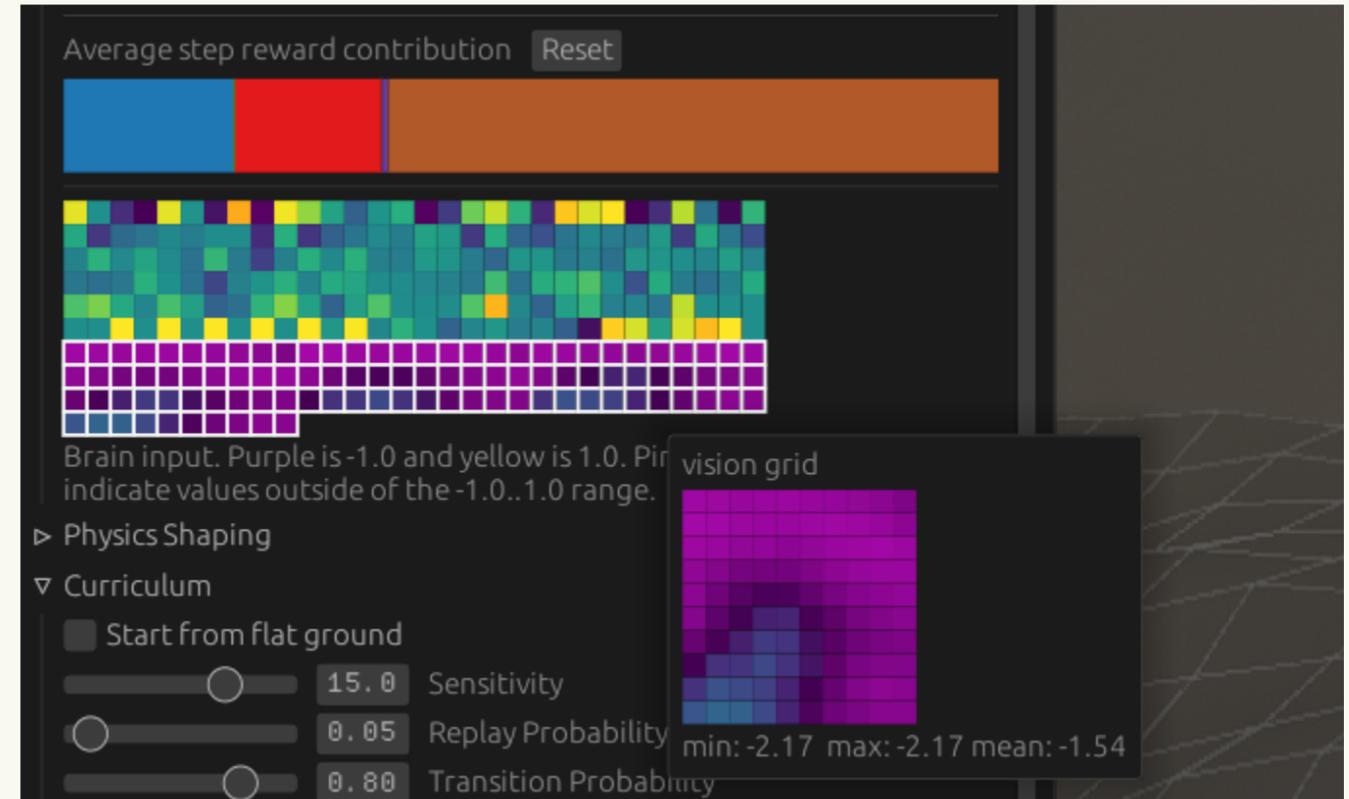
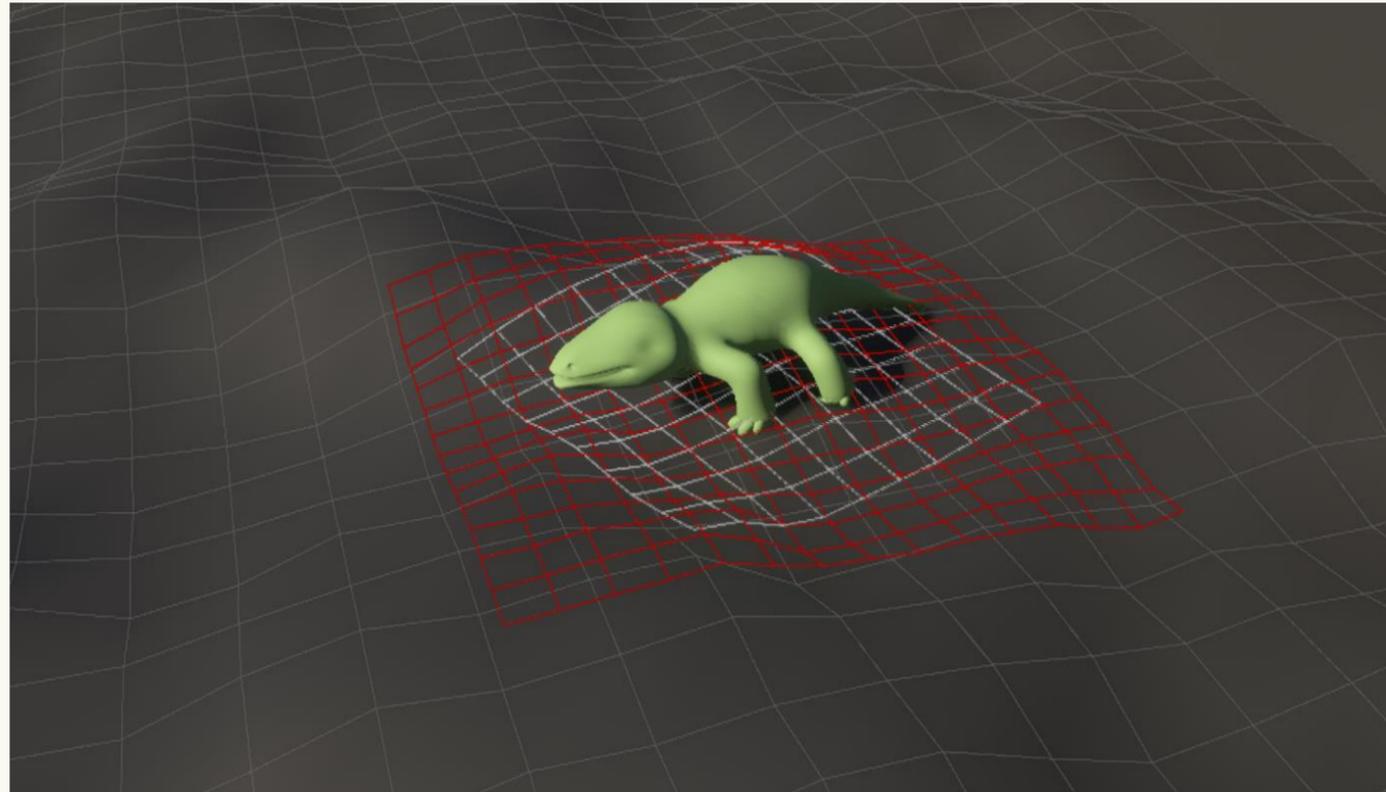
[Biped dancing \(link\)](#)

Music: Stylish Funky Beat [Funky Music] by MokkaMusic / Leno



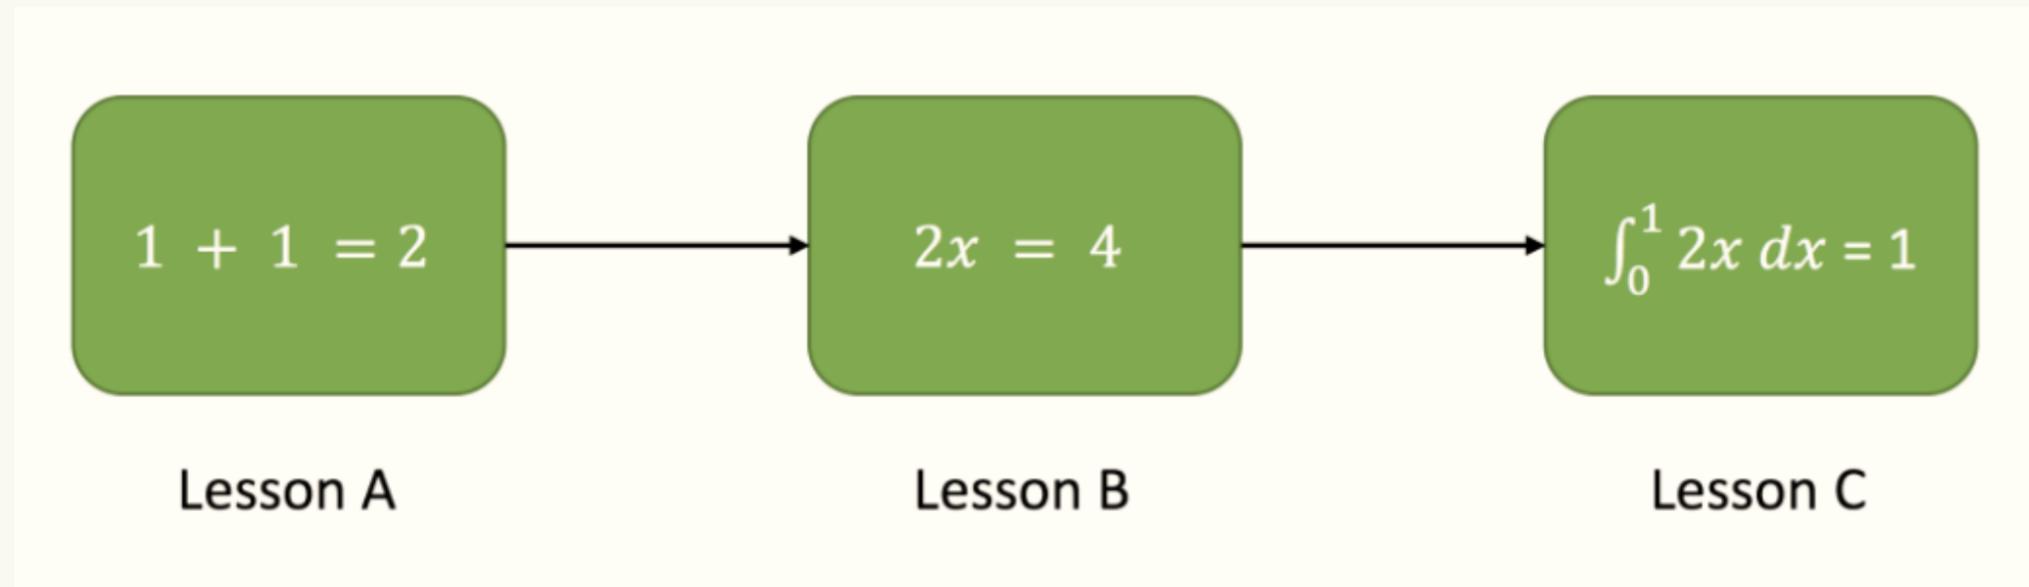
Handling hard environments

We observe terrain

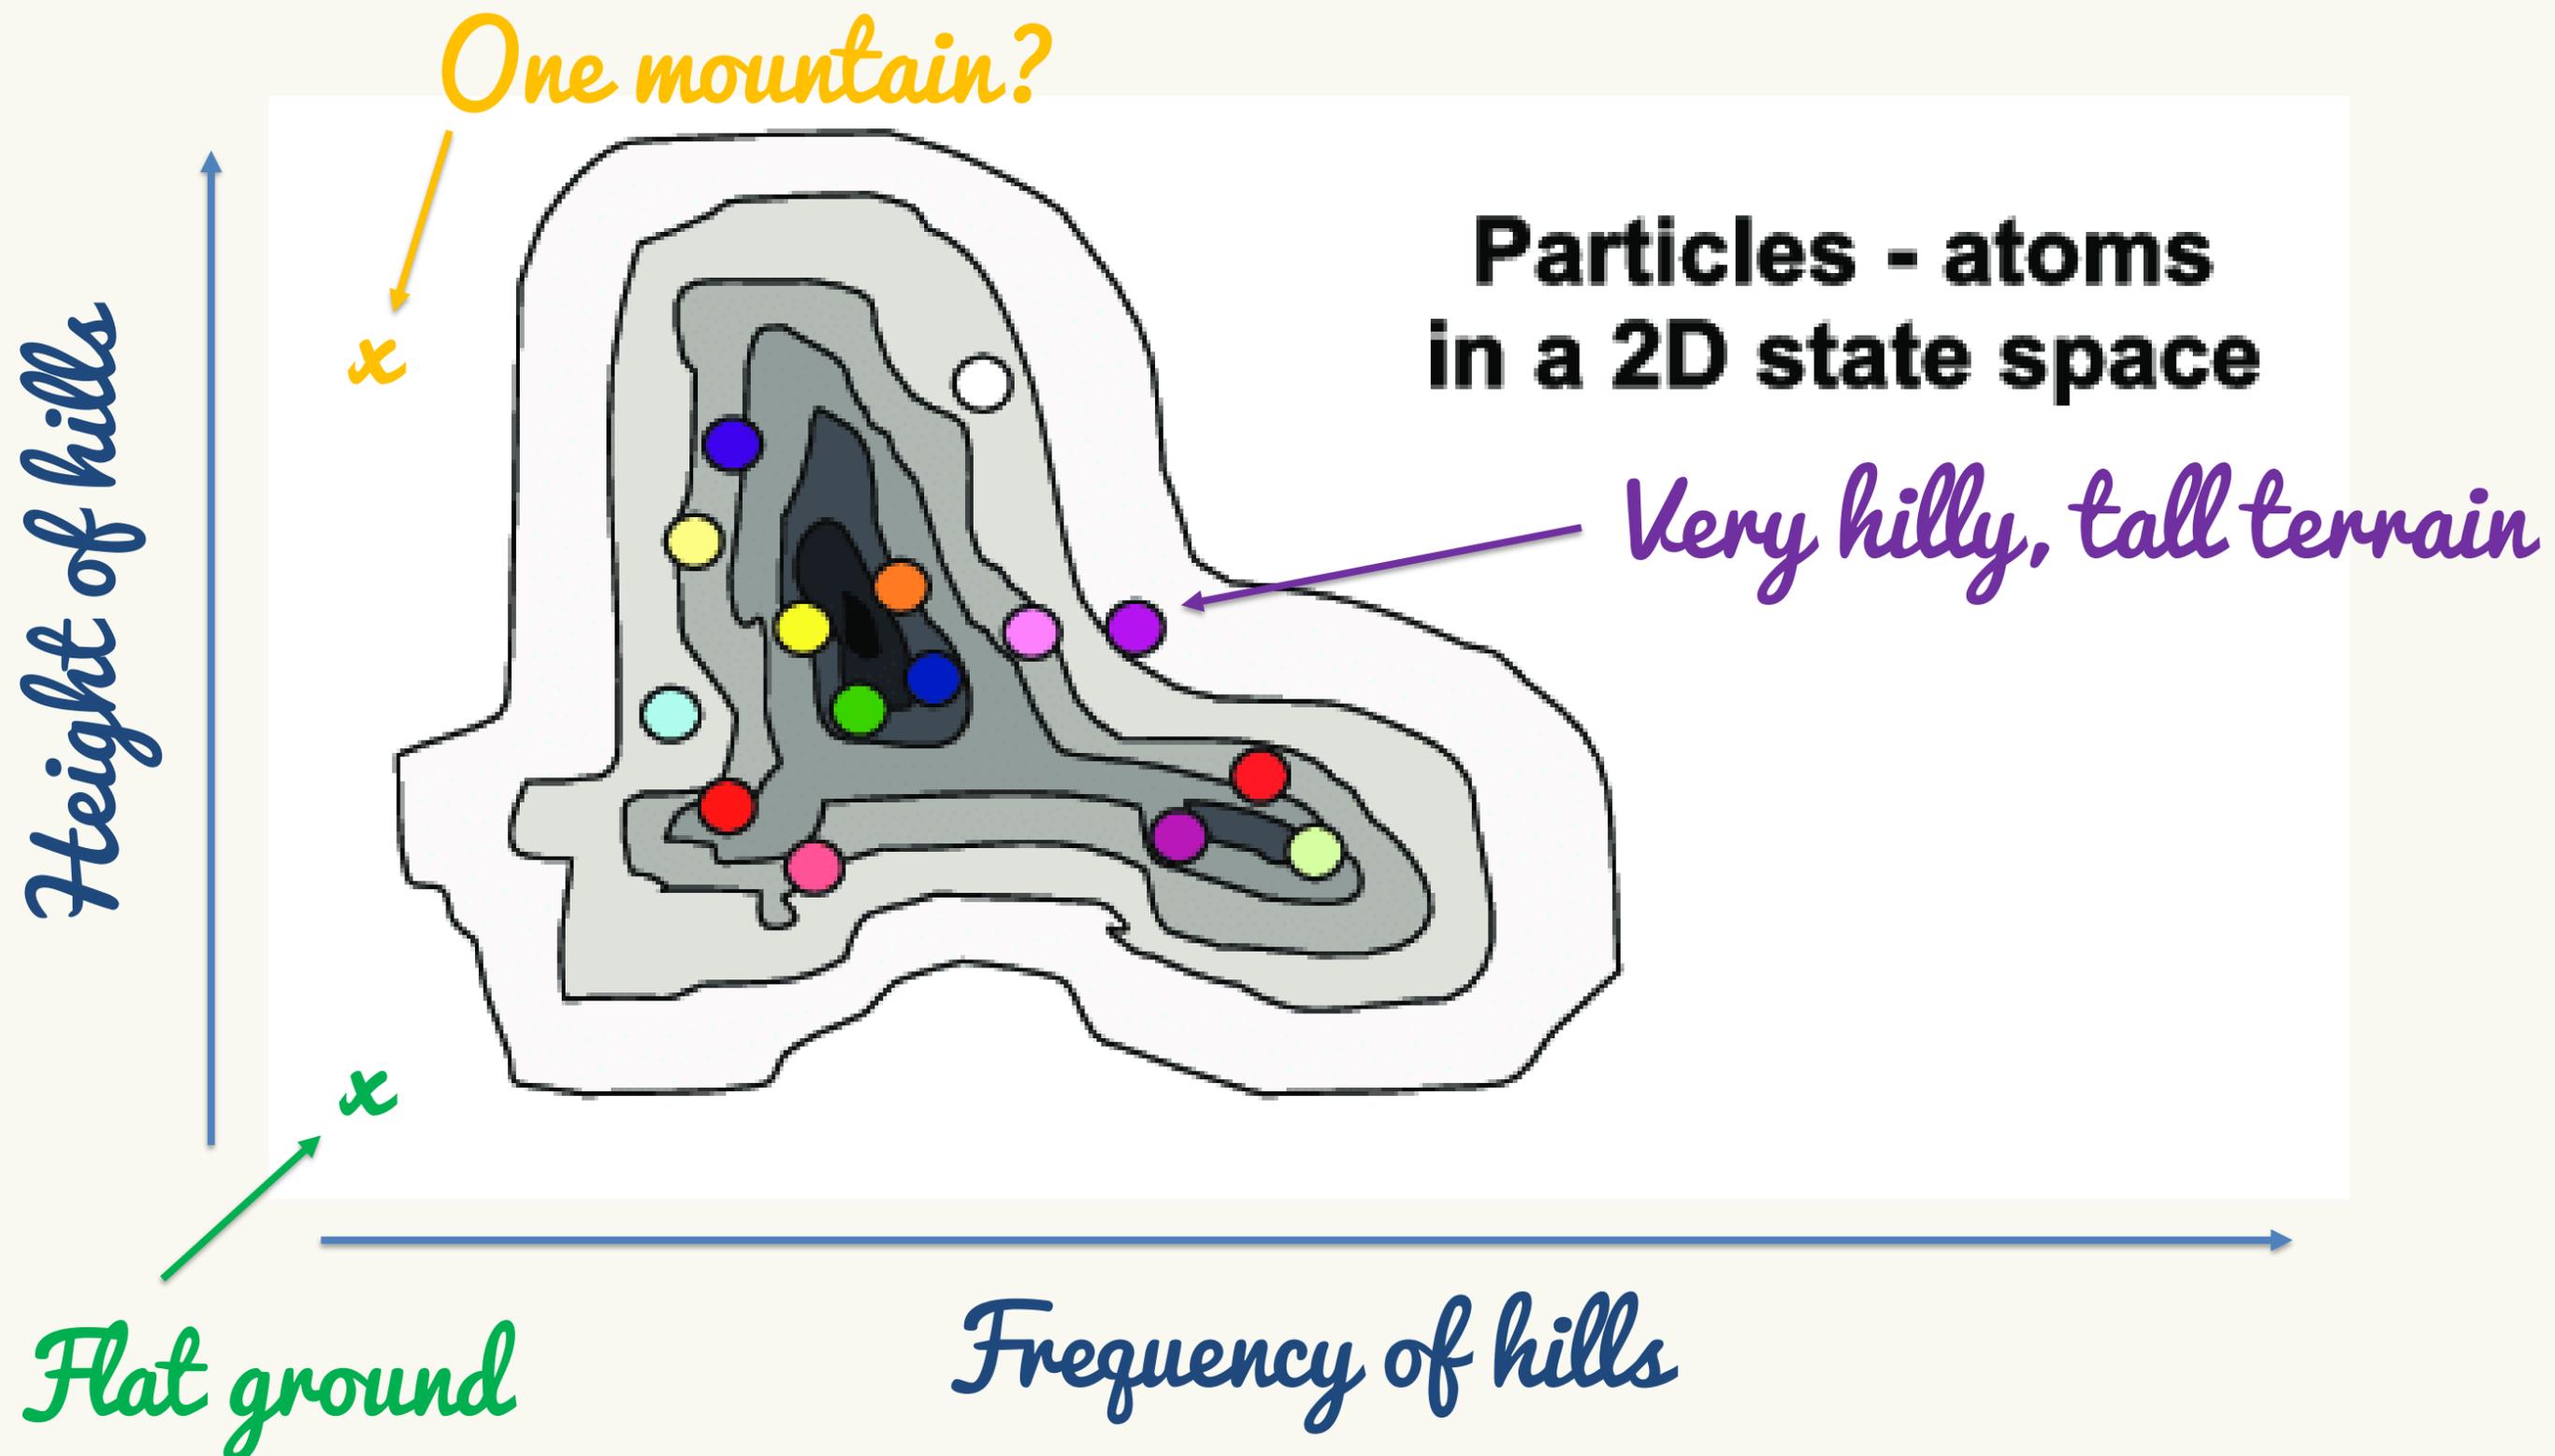


It works terribly bad if you just do this!

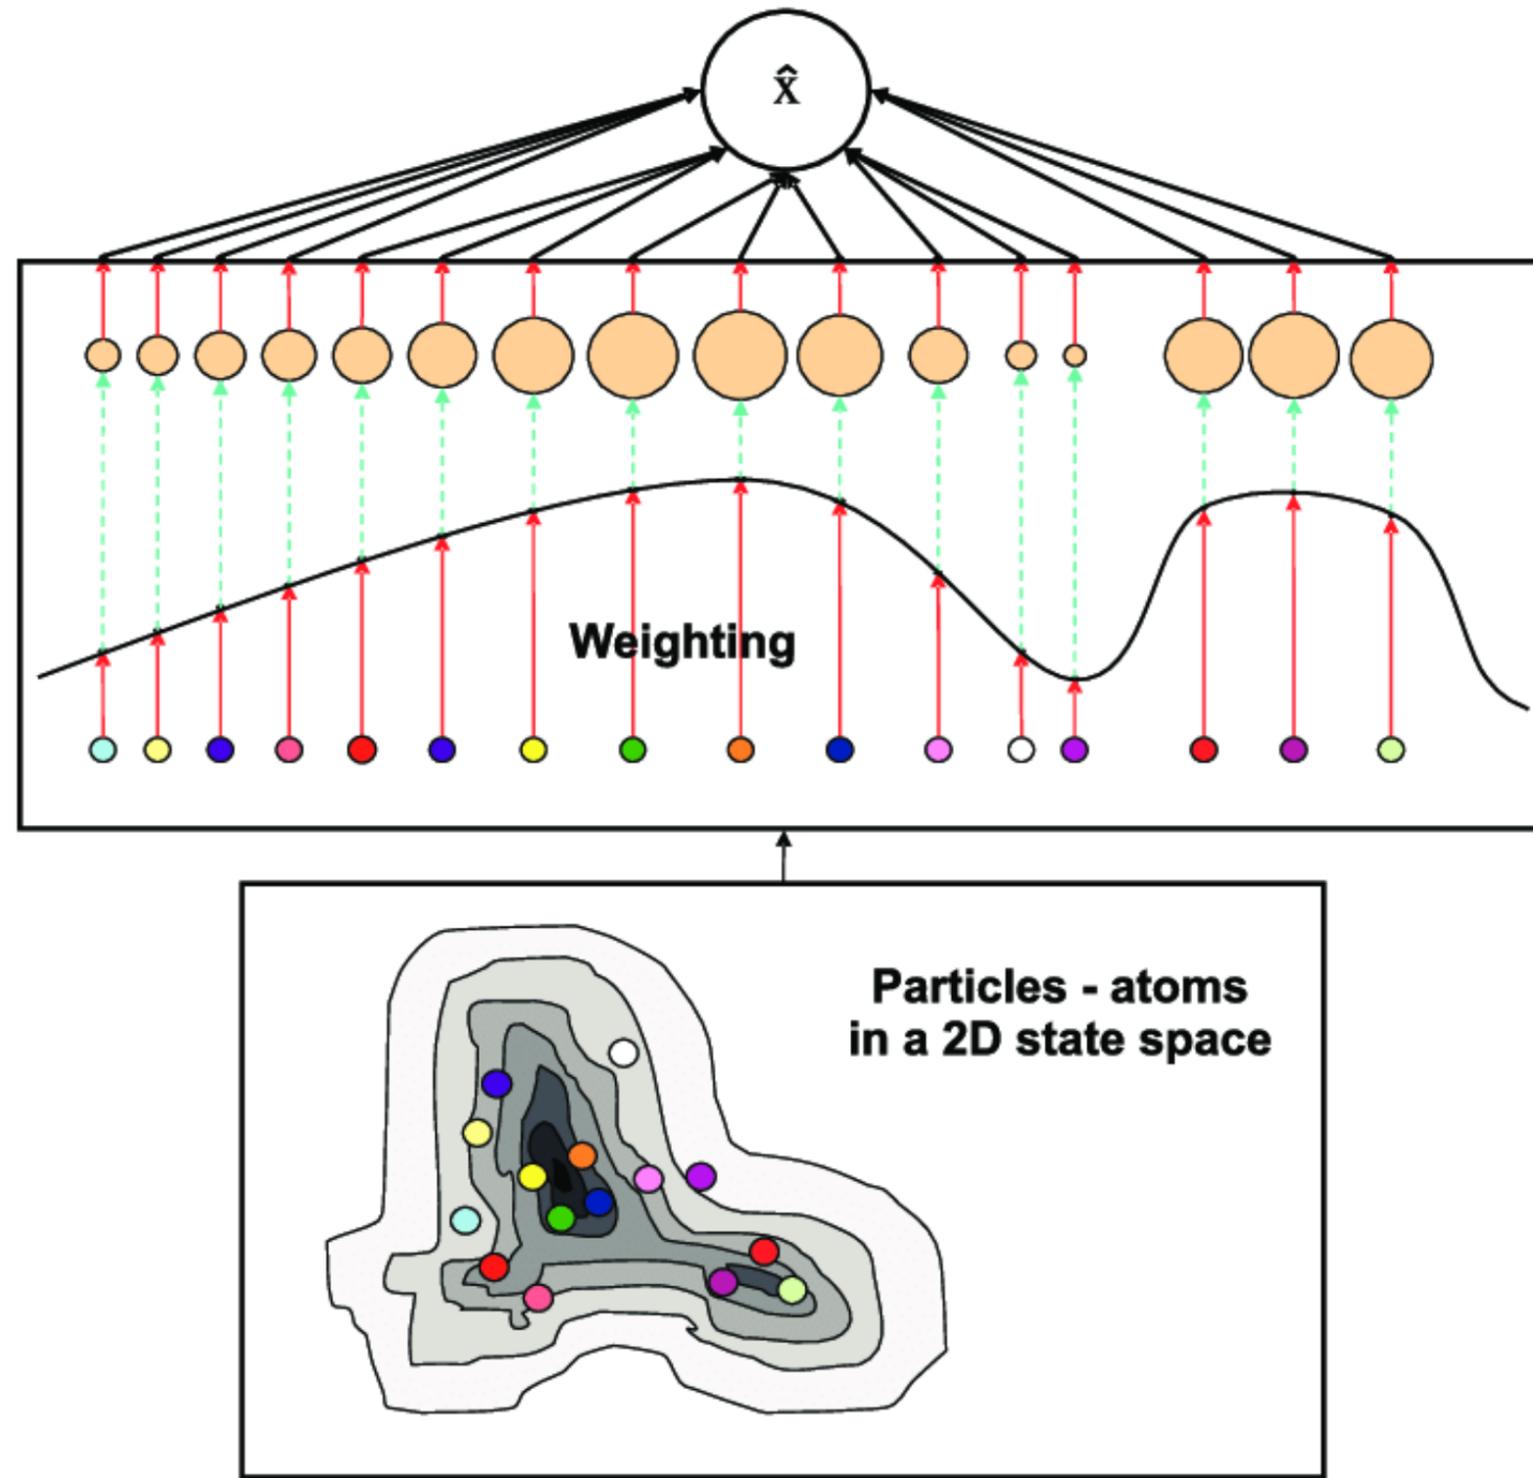
Curriculum



We use an *adaptive* curriculum



Learning Quadrupedal Locomotion over Challenging Terrain. Joonho Lee et al. 2020



Learning Quadrupedal Locomotion over Challenging Terrain. Joonho Lee et al. 2020

[Creatures in terrain \(link\)](#)

[Biped in hard terrain \(link\)](#)

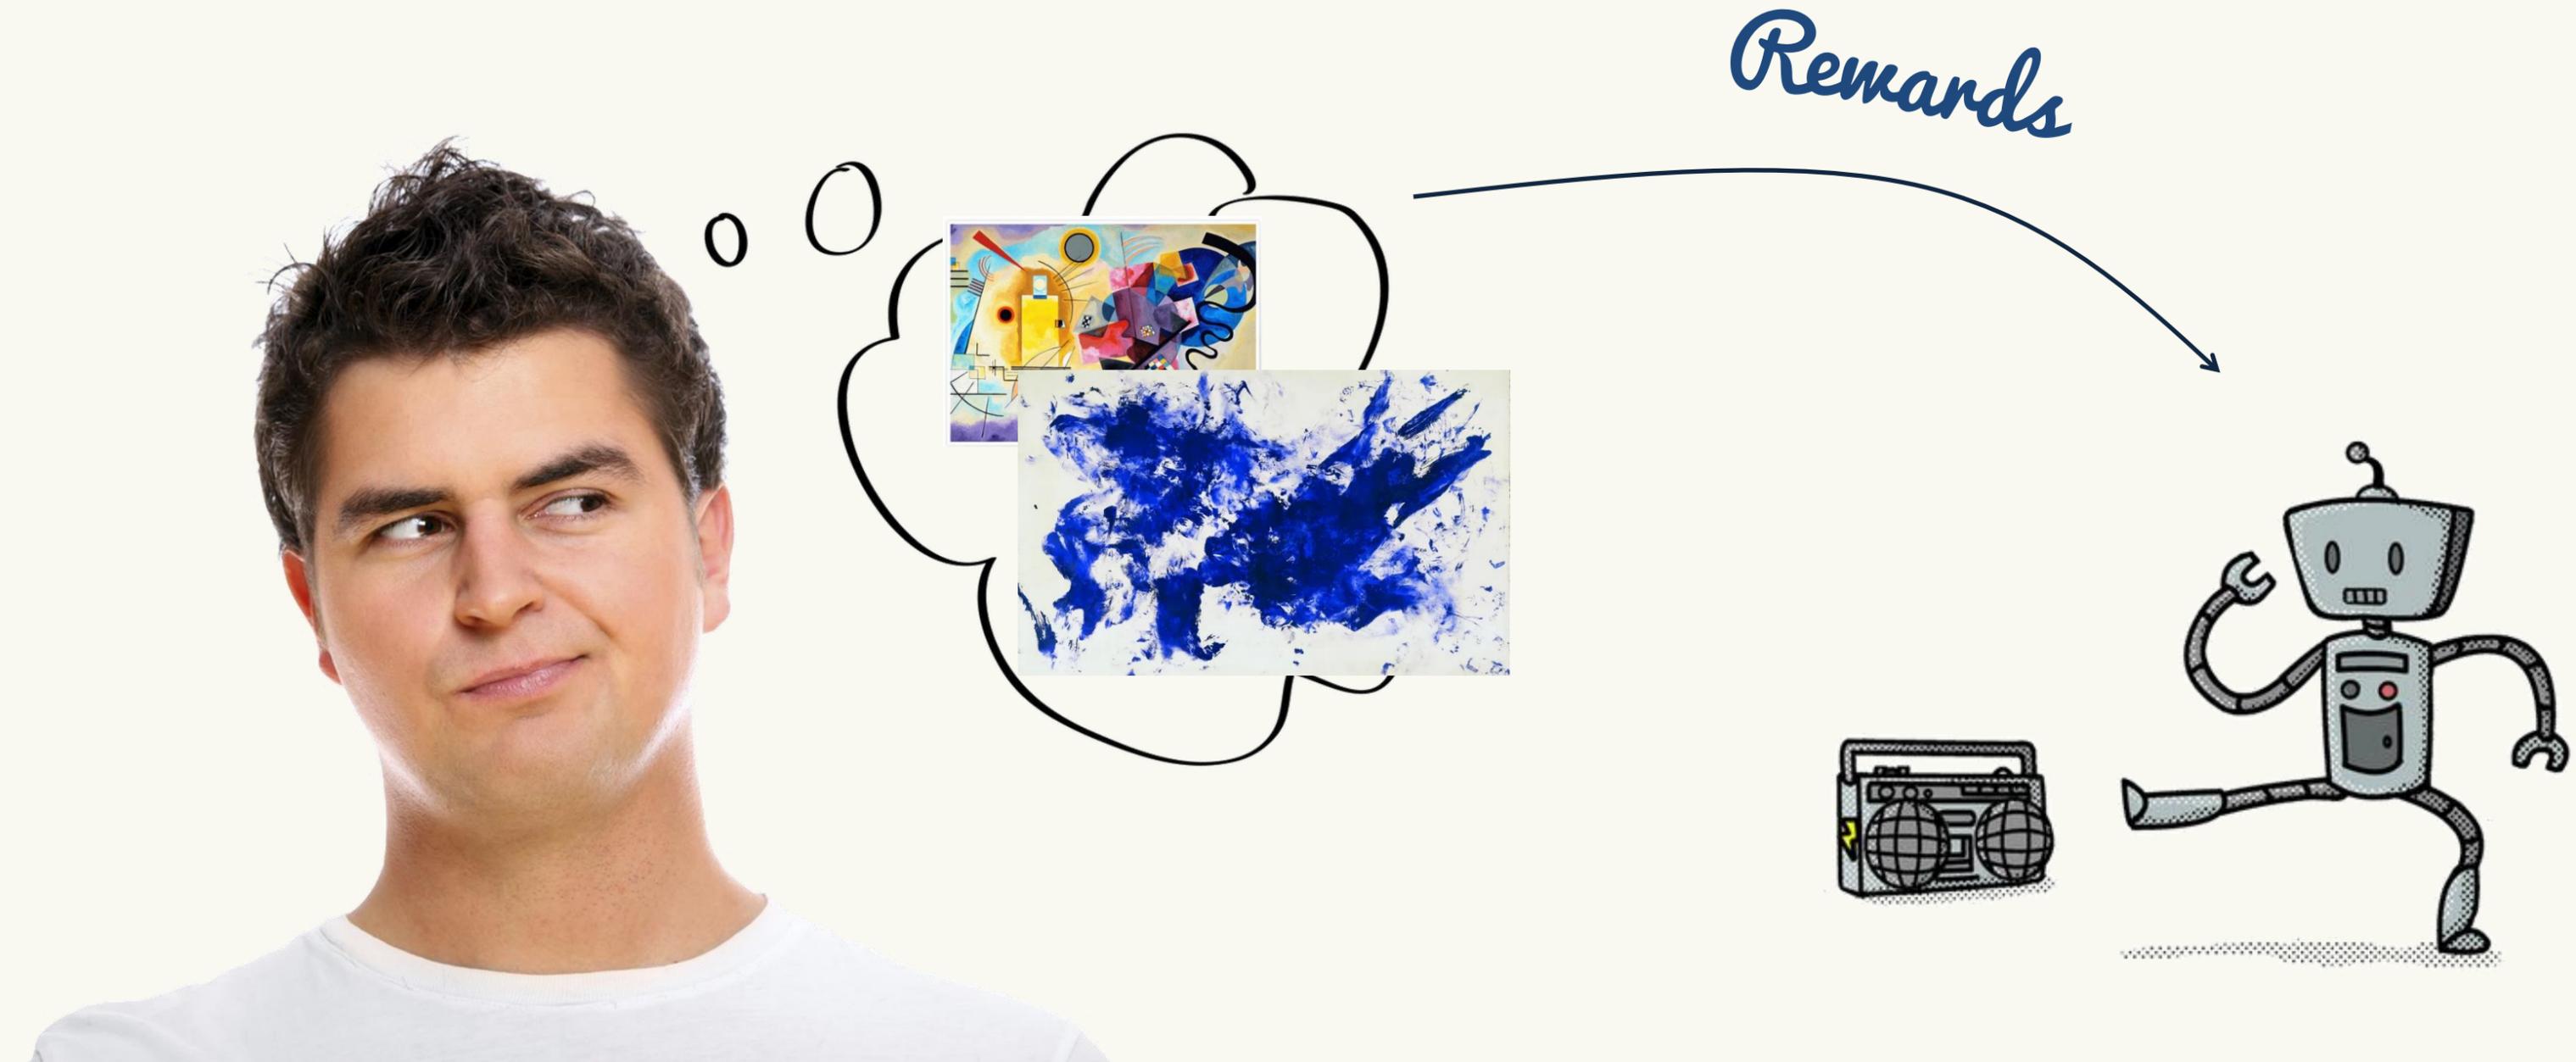
Application to giant robots – Arc Raiders

[Wasabi \(link\)](#)

[Arc Raiders Reveal & Gameplay Trailer \(link\)](#)

Challenges and beyond

Reward shaping



[Locomotion bloopers \(link\)](#)

Special mentions



Riley Miladi



Aaron Granberg



Emil Ernerfeldt

Special mentions



Jack Harmer



Martin Singh-Blom



Tom Solberg



Magnus Nordin

Some more people

Alex Camilleri

Anastasia Opara

Arvid Burström

Benjamin Bouvier

Celia Lewis

David Pedersen

Diego Fernández Goberna

Dirk de la Hunt

Doug Church

Ferruccio Cinquemani

Gray Olson

Henrik Rydgård

Jake Shadle

Johan Andersson

Johannes Hirche

John Pertoft

Jonas Hedman Engström

Jorge Fuentes

Junsik Hwang

Luc Baracat

Maik Klein

Marcus Grass

Marcus Rundström

Mathias Tervo

Max de Danschutter

Melody Ju

Michael Andersson

Michele Bianchi

Mikhail Novoseltsev

Niklas Nummelin

Nora Björklund

Oskar Nehlin

Paul Grevenson

Peter Cardwell-Gardner

Sara Lempiäinen

Tomasz Stachowiak

Viktor Zoutman

Thank you!!

Questions