CS 4756/5756: Robot Learning

Sanjiban Choudhury

PORTAL

Cornell Bowers CIS Computer Science


WHAT A TIME TO BE ALIVE!

AI tools that are OPEN!

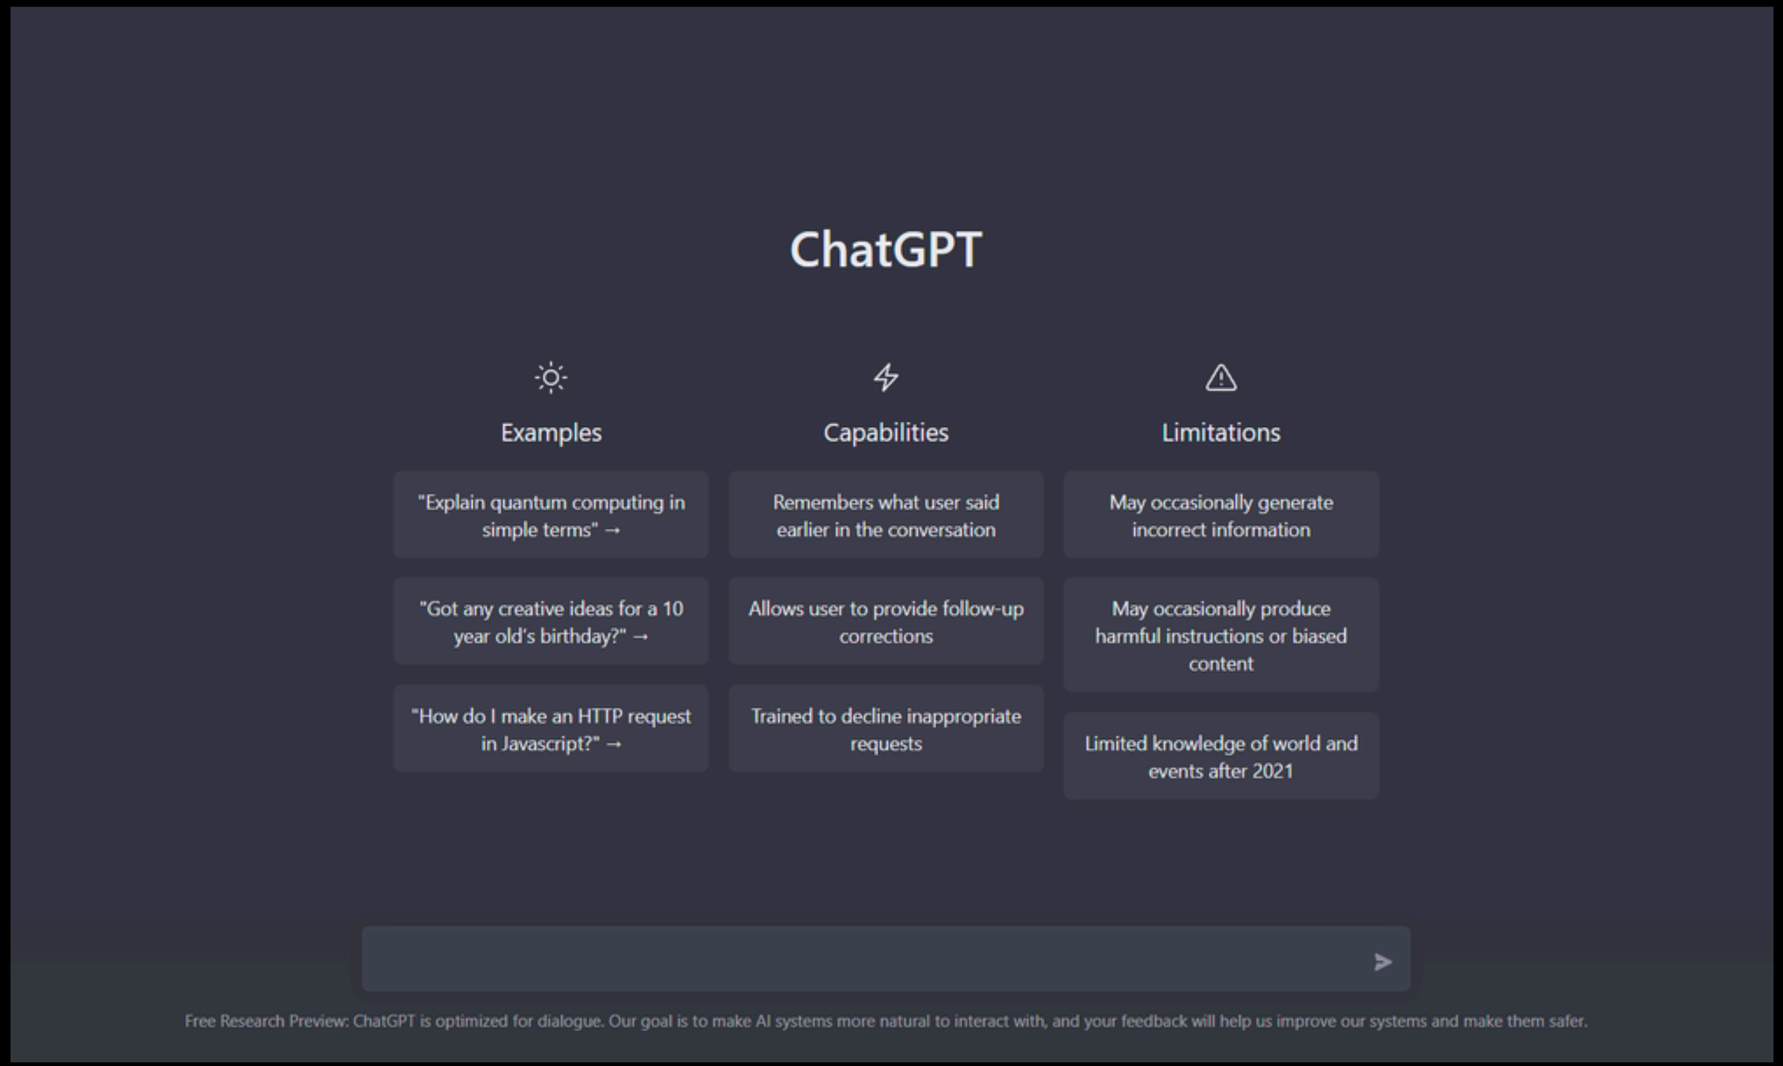

AI tools that are OPEN!

animate_double_pendulum.py

```python
import numpy as np
import matplotlib.pyplot as plt
from matplotlib.animation import FuncAnimation

# Define the equations of motion for the double inverted pendulum
def equations_of_motion(y, t, L1, L2, m1, m2, g):
    th1, th2, dth1, dth2 = y
    dydt = [dth1,
            dth2,
            (-g*(2*m1 + m2)*np.sin(th1) - m2*g*np.sin(th1 - 2*th2) - 2*np.sin(th1 - th2
                )*m2*(dth2**2*L2 + dth1**2*L1*np.cos(th1 - th2)))/(L1*(2*m1 + m2 - m2*
                np.cos(2*th1 - 2*th2))),
            (2*np.sin(th1 - th2)*(dth1**2*L1*(m1 + m2) + g*(m1 + m2)*np.cos(th1) + dth2
                **2*L2*m2*np.cos(th1 - th2)))/(L2*(2*m1 + m2 - m2*np.cos(2*th1 - 2*th2)
                ))]
    return dydt

# Set up the initial conditions and other parameters
y0 = [np.pi/2, np.pi/2, 0, 0]
L1 = 1
L2 = 1
m1 = 1
m2 = 1
g = 9.81

# Solve the equations of motion using odeint
from scipy.integrate import odeint
t = np.linspace(0, 10, 1000)
sol = odeint(equations_of_motion, y0, t, args=(L1, L2, m1, m2, g))

# Animate this pendulum ...
```

AI tools that are OPEN!

AI tools that are OPEN!

AI tools that are OPEN!

```
seed_and_prompt_sequence = [
        3764, 'in the beginning there was nothing, just darkness',
        1537, 'special effects render of the big bang',
        6573, 'HD photo of a large amount of spiral galaxies',
        1791, 'early planet formation in the solar system',
        9973, 'the Hadean earth was bombarded with asteroids and massive volcanic eruptions',
        736,  'panoramic view of earth with ocean surrounding newly formed land and volcanos',
        3639, 'hydrothermal vents at the bottom of the ocean',
        3559, 'bacteria under a microscope',
        4724, 'bacteria under a microscope',
        3359, 'ammonites floating in the ocean',
        6344, 'the first reptile to leave the ocean and crawl onto the land',
        6344, 'the first reptile to leave the ocean and crawl onto the land',
        6813, 'massive brachiosaurus walking amidst a green mountain range',
        6678, 'the exctinction of the dinosaurs be a huge meteorite',
        7450, 'small mammals thriving in a cave',
        9766, 'small, prehistoric mammals living in the jungle',
        5009, 'group of monkeys in the forest',
        7287, 'HD photograph of neanderthal, the first man',
        6008, 'cave painting',
        208,  'cavemen tribe gathered around a fire at night looking at the stars',
        2222, 'maasai tribe hunting on the savanna with spears',
        571,  'homo sapiens using stone tools',
        632,  'a small, tribal village with huts',
        1332, 'at the dawn of civilization, small villages emerged',
        2496, 'ancient egypt, the first massive civiliation',
        1869, 'the height of the roman empire, incredible architecture, by Greg Rutkowski',
        7559, 'medieval town square',
        1265, 'medieval city',
        6628, 'the skyline of New York city',
```

AI tools that are OPEN!

Where are the robots?

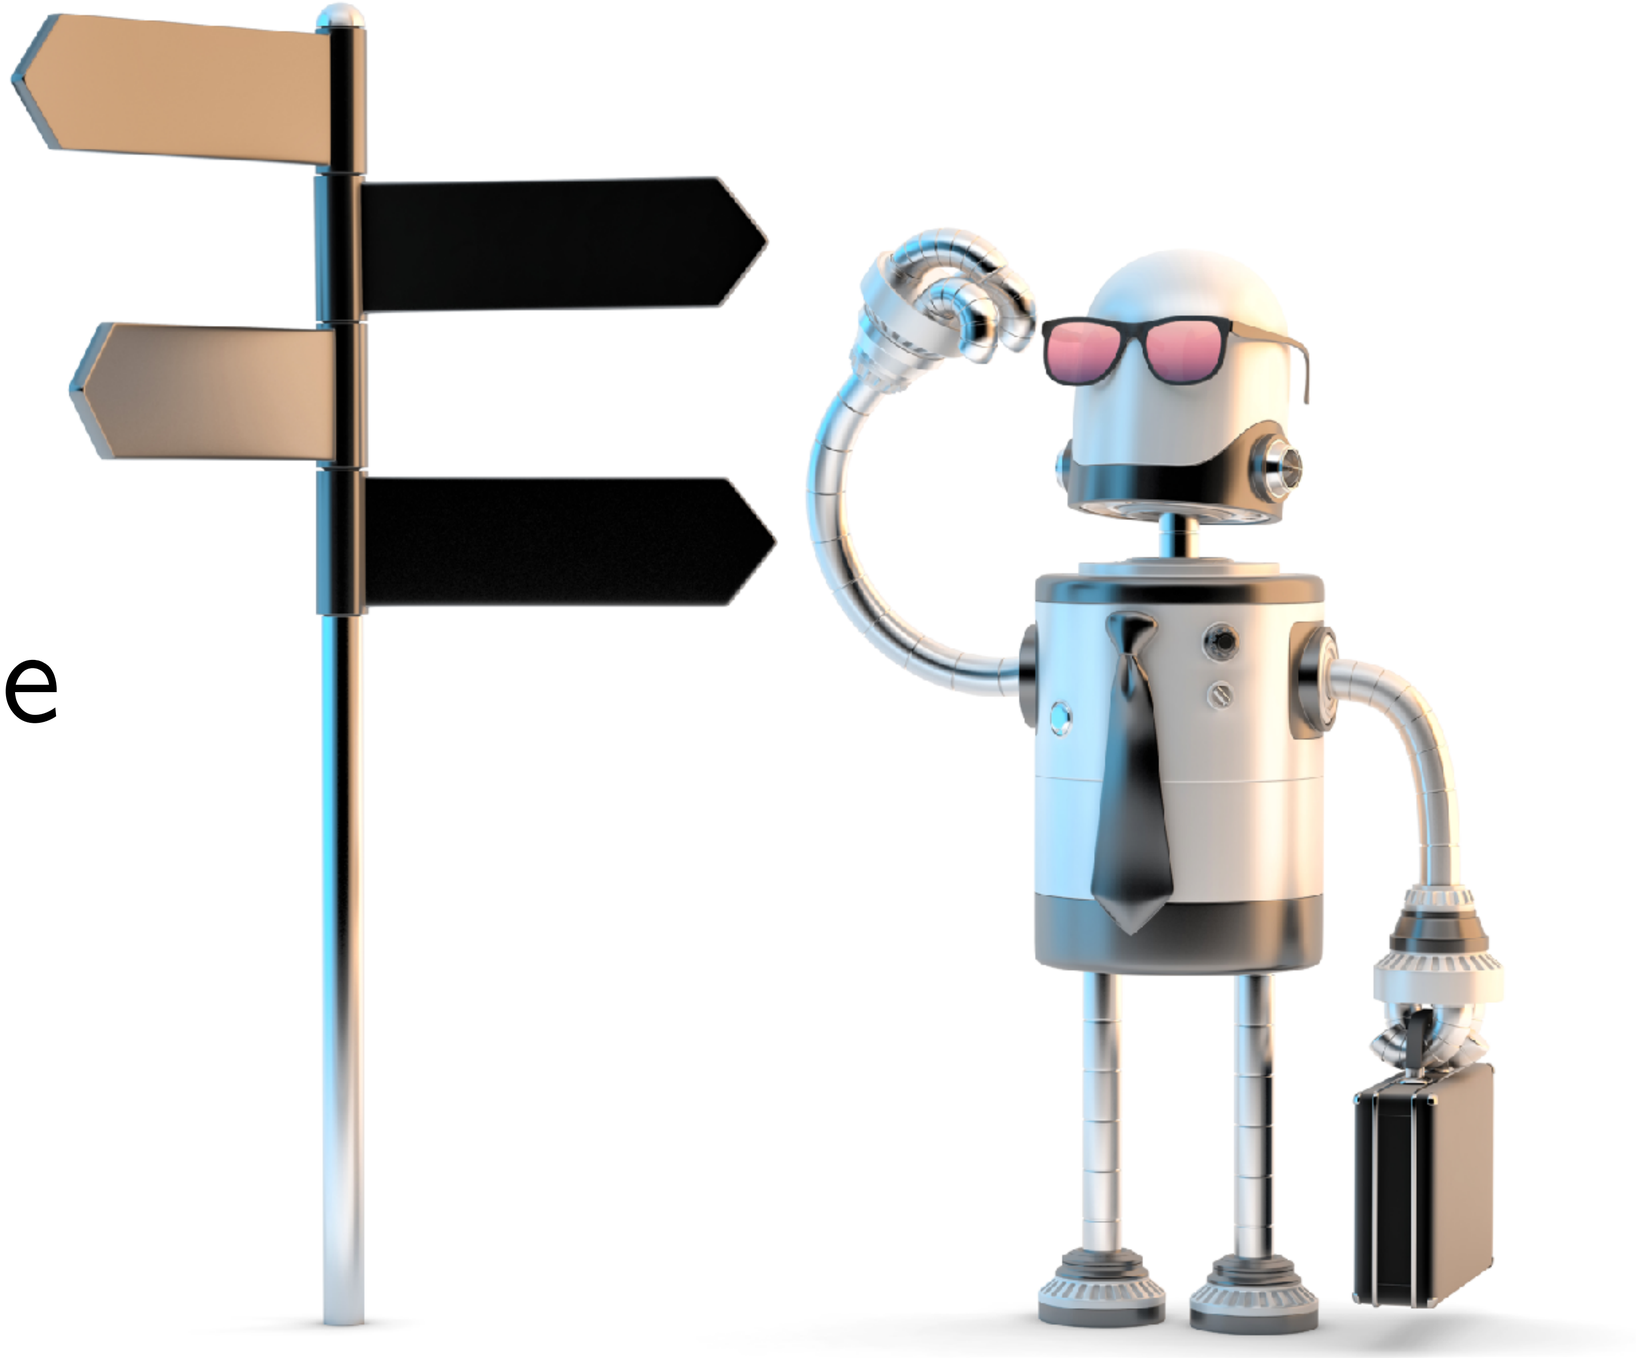

Robots are not far behind!

Robots are not far behind!

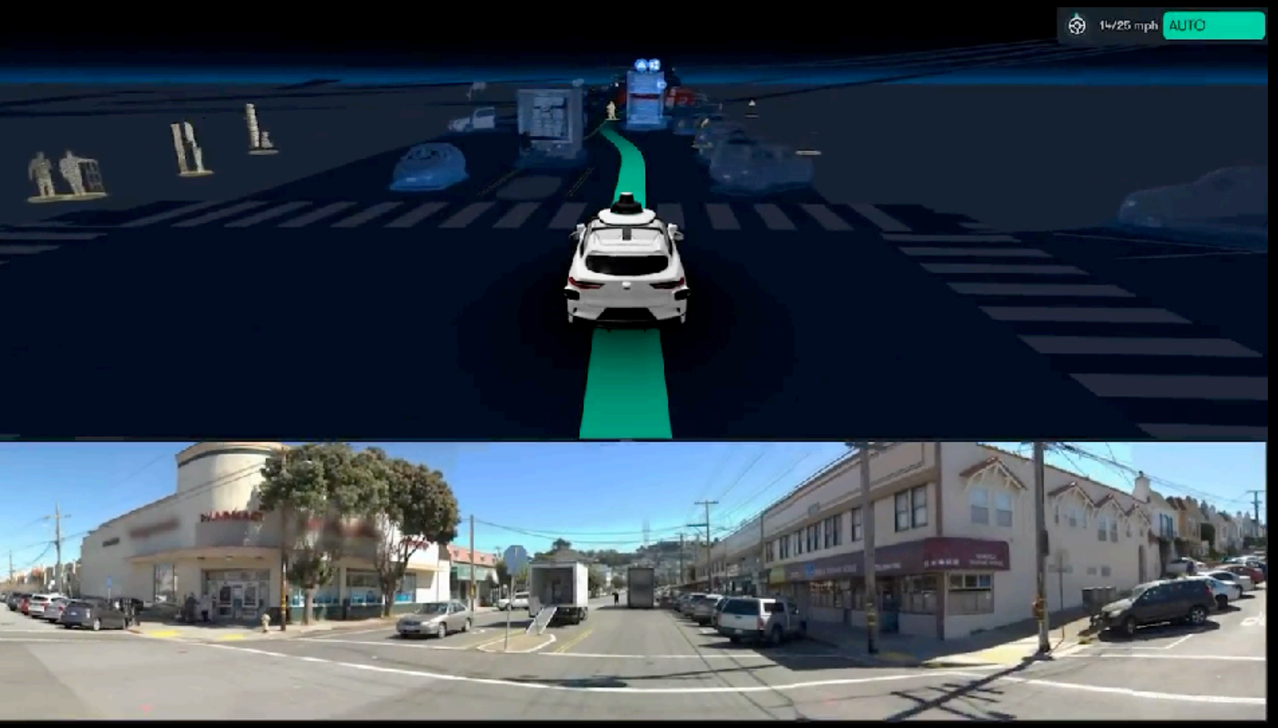

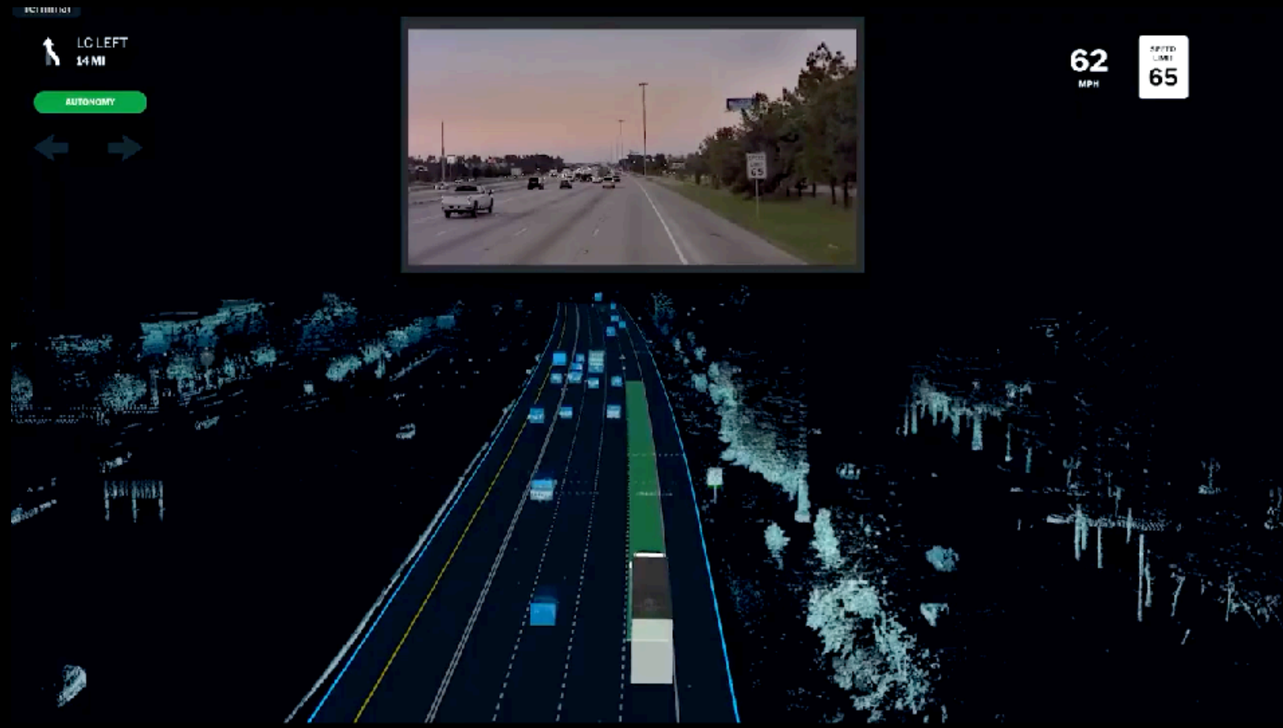

WAYMO

Aurora

Self-driving companies going driverless ...

Robots are not far behind!

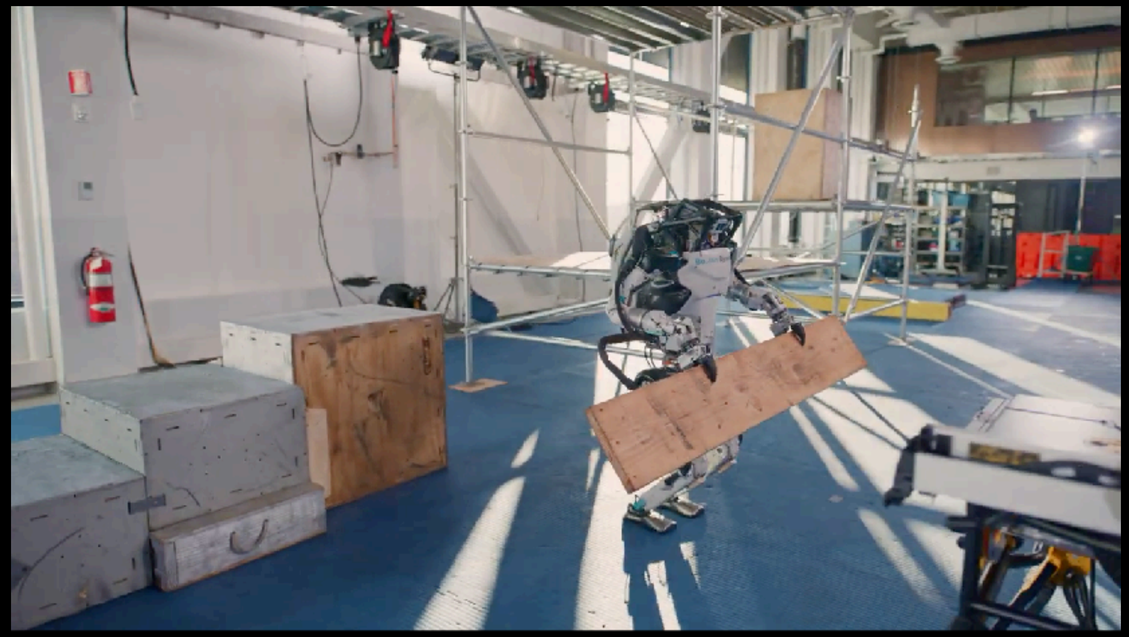

Boston Dynamics are starting to sell their robots ...

Robots are not far behind!

Drones are getting more reliable ...

But ...

... robots are not in millions of homes yet.

Why?

The way we program robots today is ... rigid!

Engineers hand-craft behaviors

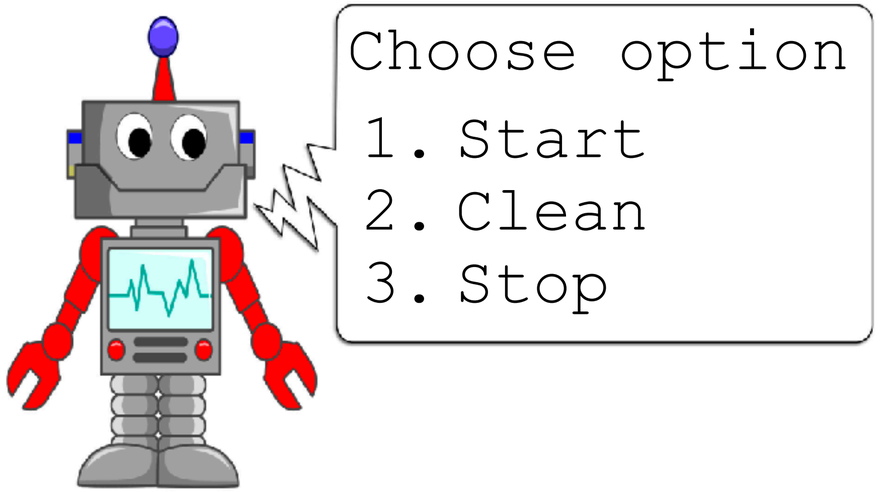

Ship robot

Frustrate users!

Not flexible enough to be used by everyday users for everyday tasks

This restricts robots to a CLOSED world

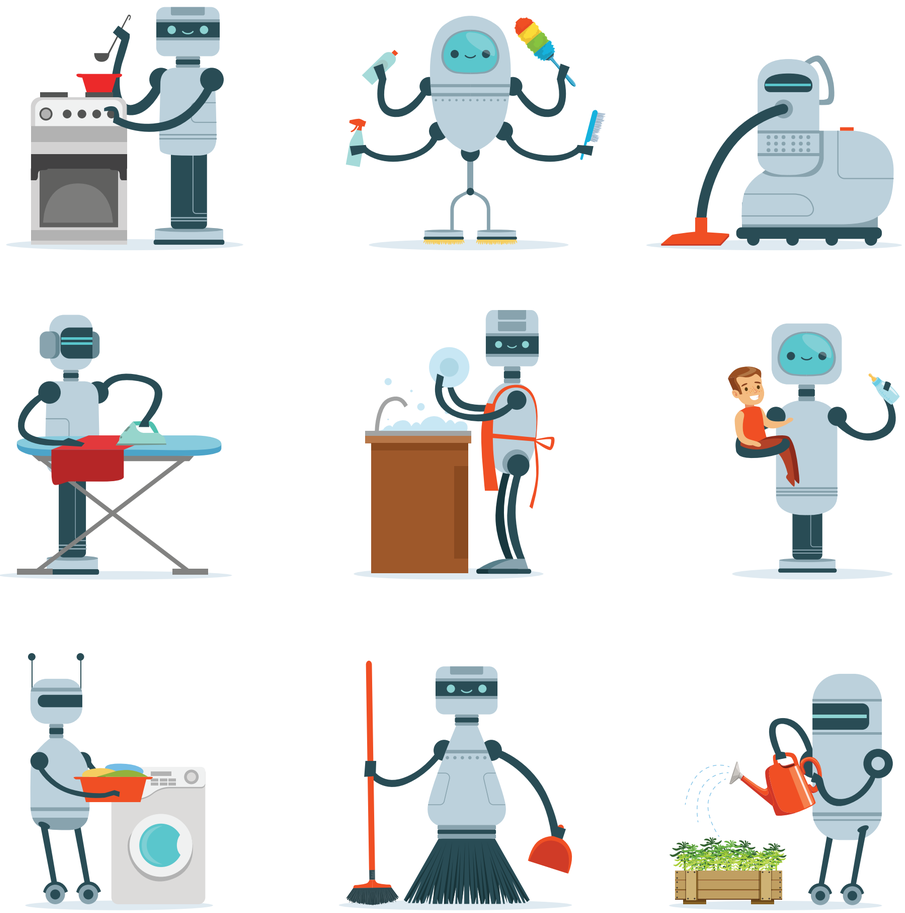

The Dream

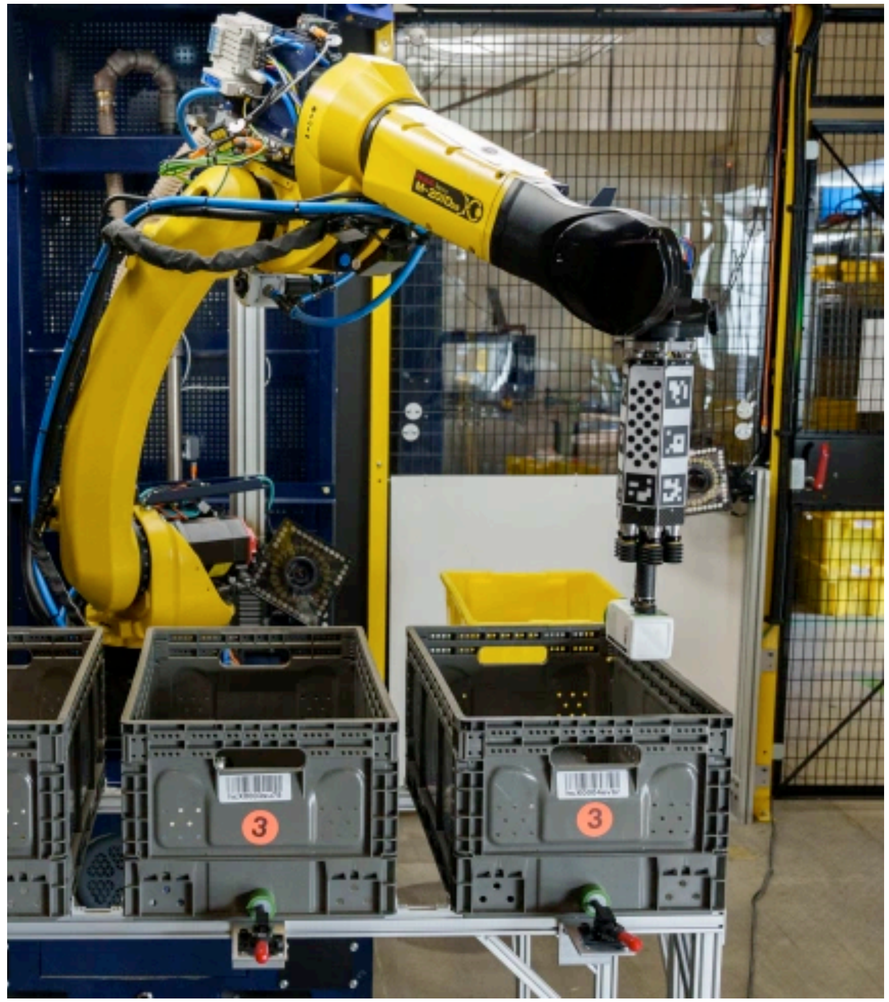

Reality

How can we get robots out of the factory into the OPEN WORLD?

Activity!

Build robot apprentice to help grandma!

Demonstrations,
Language

Interactive
feedback

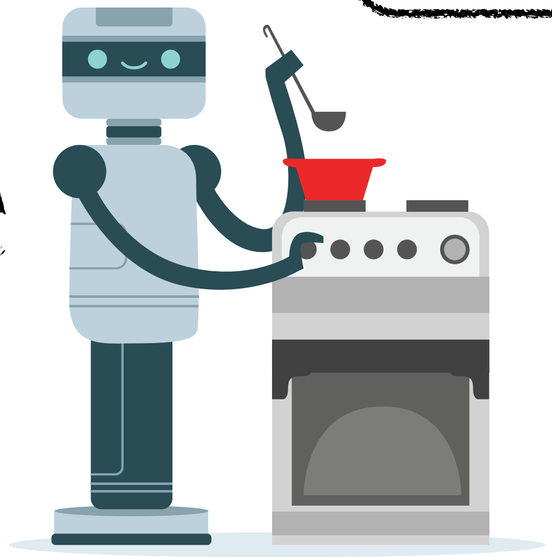

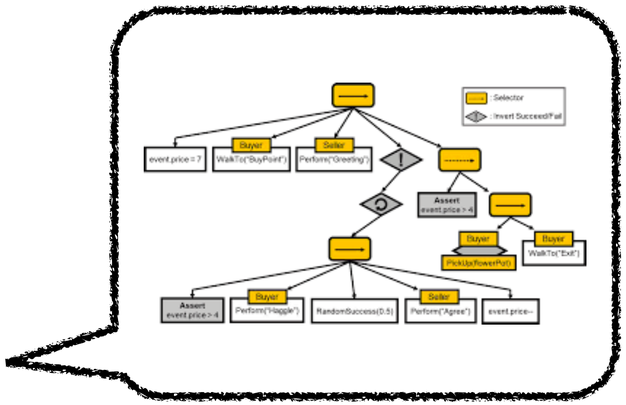

Think-Pair-Share

Think (30 sec): Think of *ALL* the challenges to building such a robot

Pair: Find partners

Share (45 sec): Partners exchange ideas

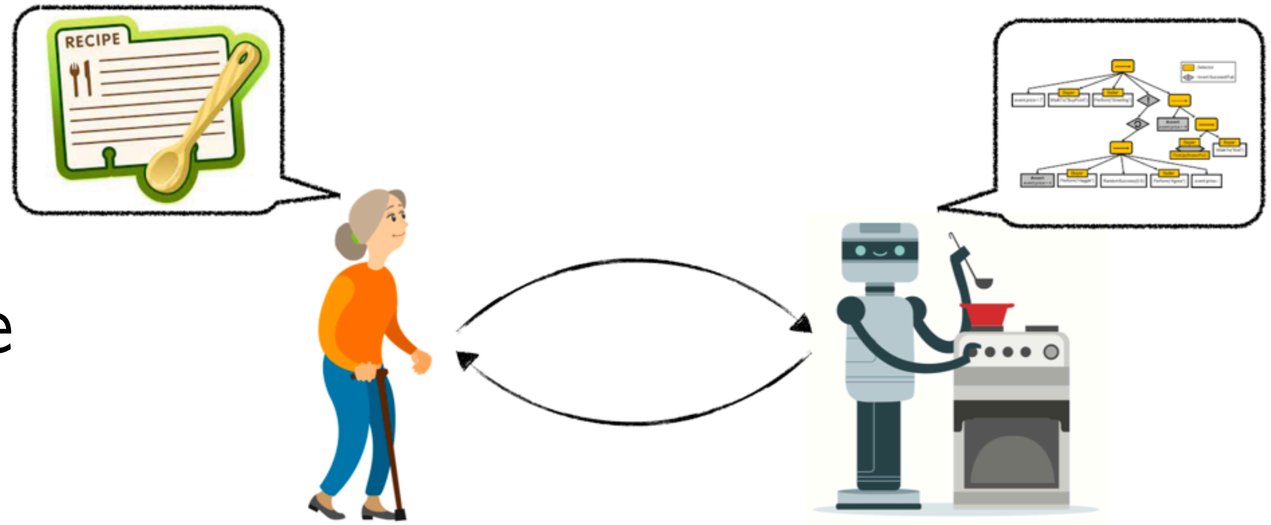

Learn from

interactions from both

humans and the world

How should robots learn from interactions?

WHY this course?






Take *any*
robot application

Formulate as a Markov
Decision Problem (MDP)

Solve MDPs using an
all-purpose toolkit
(Imitation/Reinforcement learning, Model based/free)

Deploy learners in real-world
(Safety, distribution shift, value alignment)

"Sanjiban"

He / Him

Why I work on robots

To better understand humans ...

"Sanjiban"

He / Him

Undergrad: First robot soccer team in India!

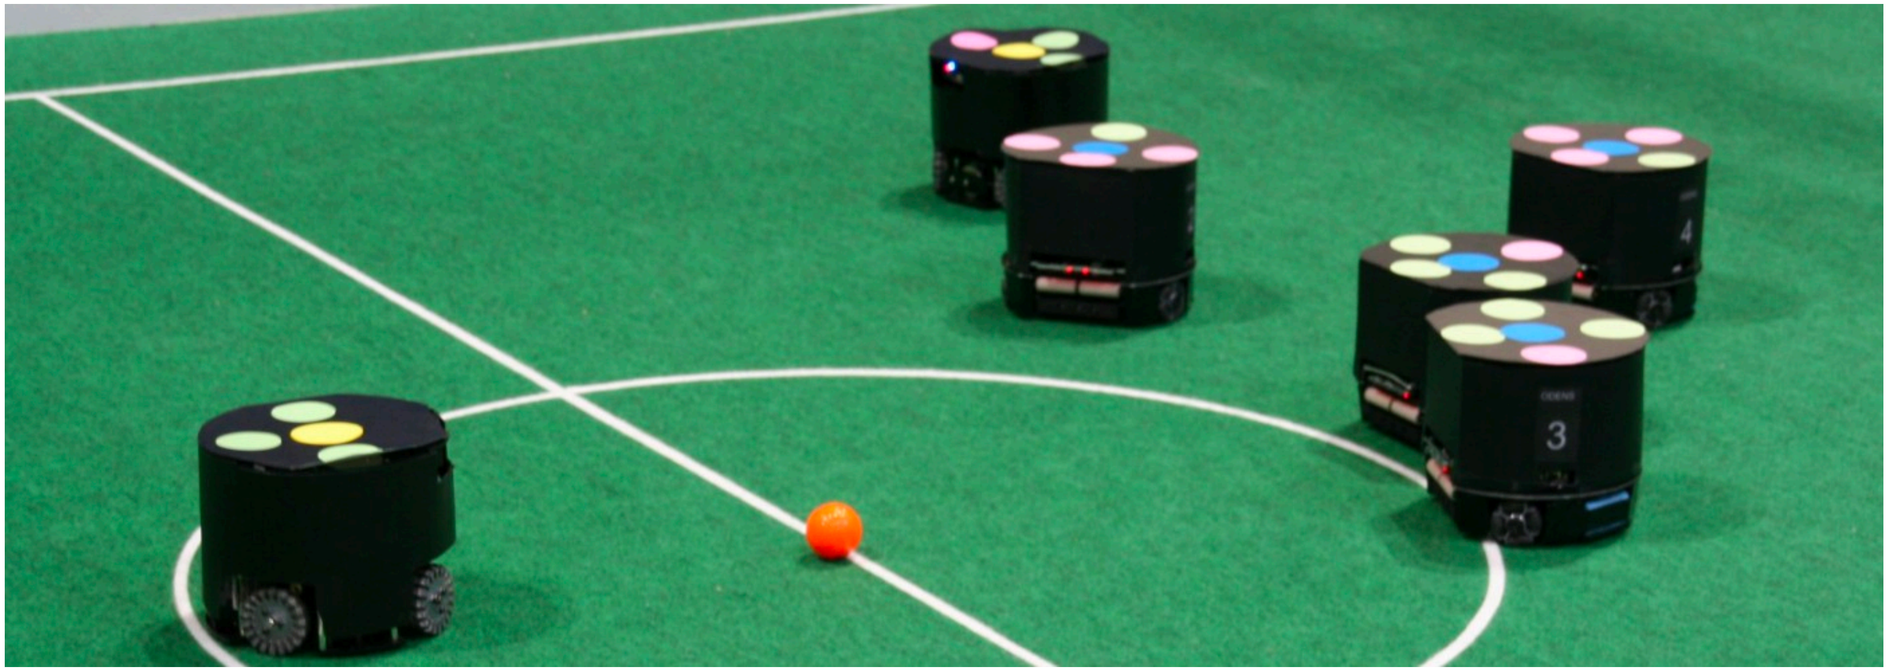

Learn strategies like humans?

"Sanjiban"

He / Him

Carnegie Mellon
THE ROBOTICS INSTITUTE

PhD: Full-scale autonomous helicopter flight

Fly like a human?

"Sanjiban"

He / Him

PAUL G. ALLEN SCHOOL
OF COMPUTER SCIENCE & ENGINEERING

PostDoc: Mobile manipulator in the wild

Manipulate unknown objects like humans?

“Sanjiban”

He / Him

Aurora

Research Engineer: Self-driving

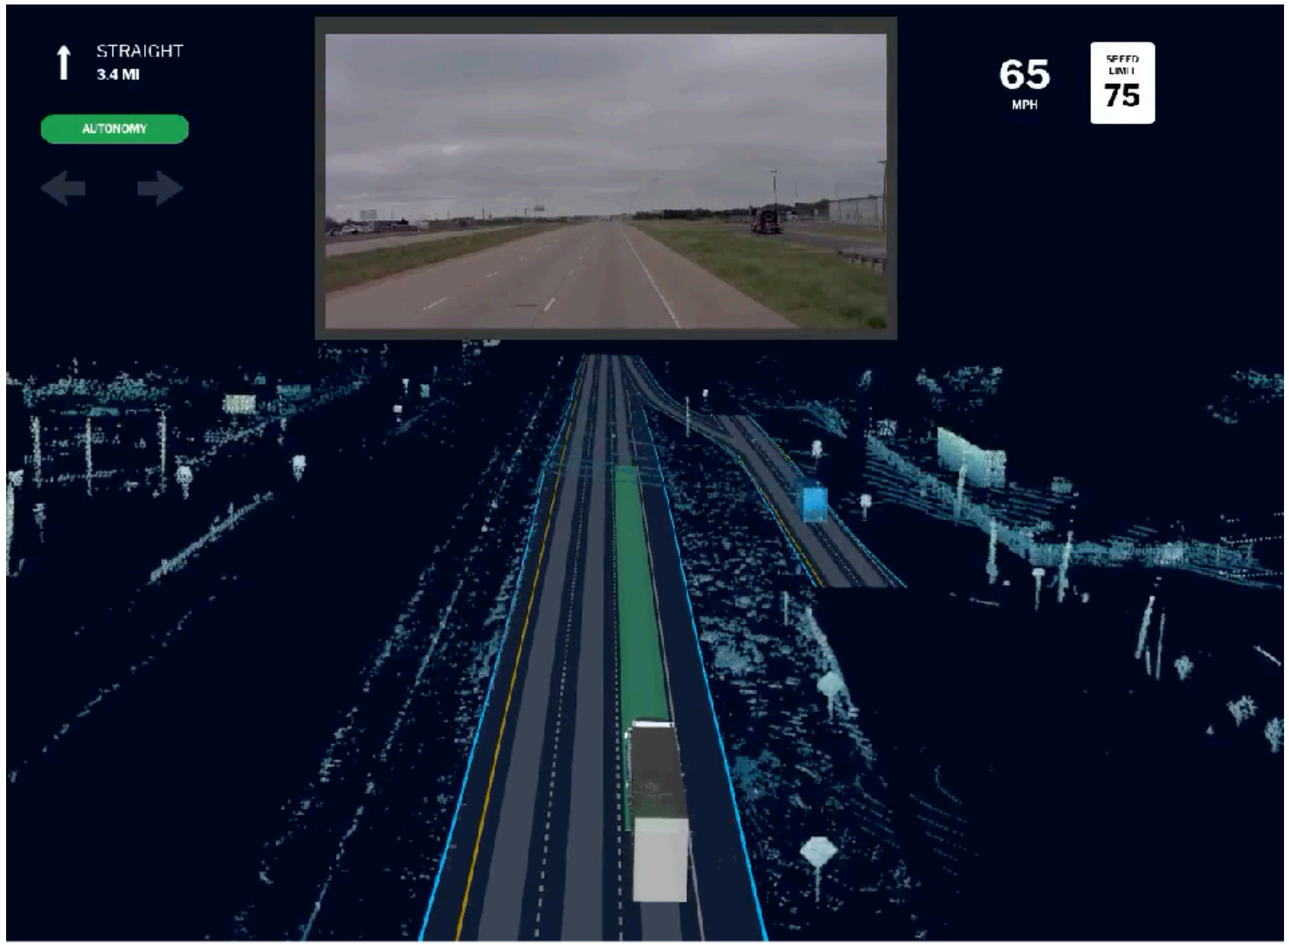

Drive alongside humans?

PoRTaL: People and Robots, Teaching and Learning

Belonging

How should robots learn from interactions?

Self-driving

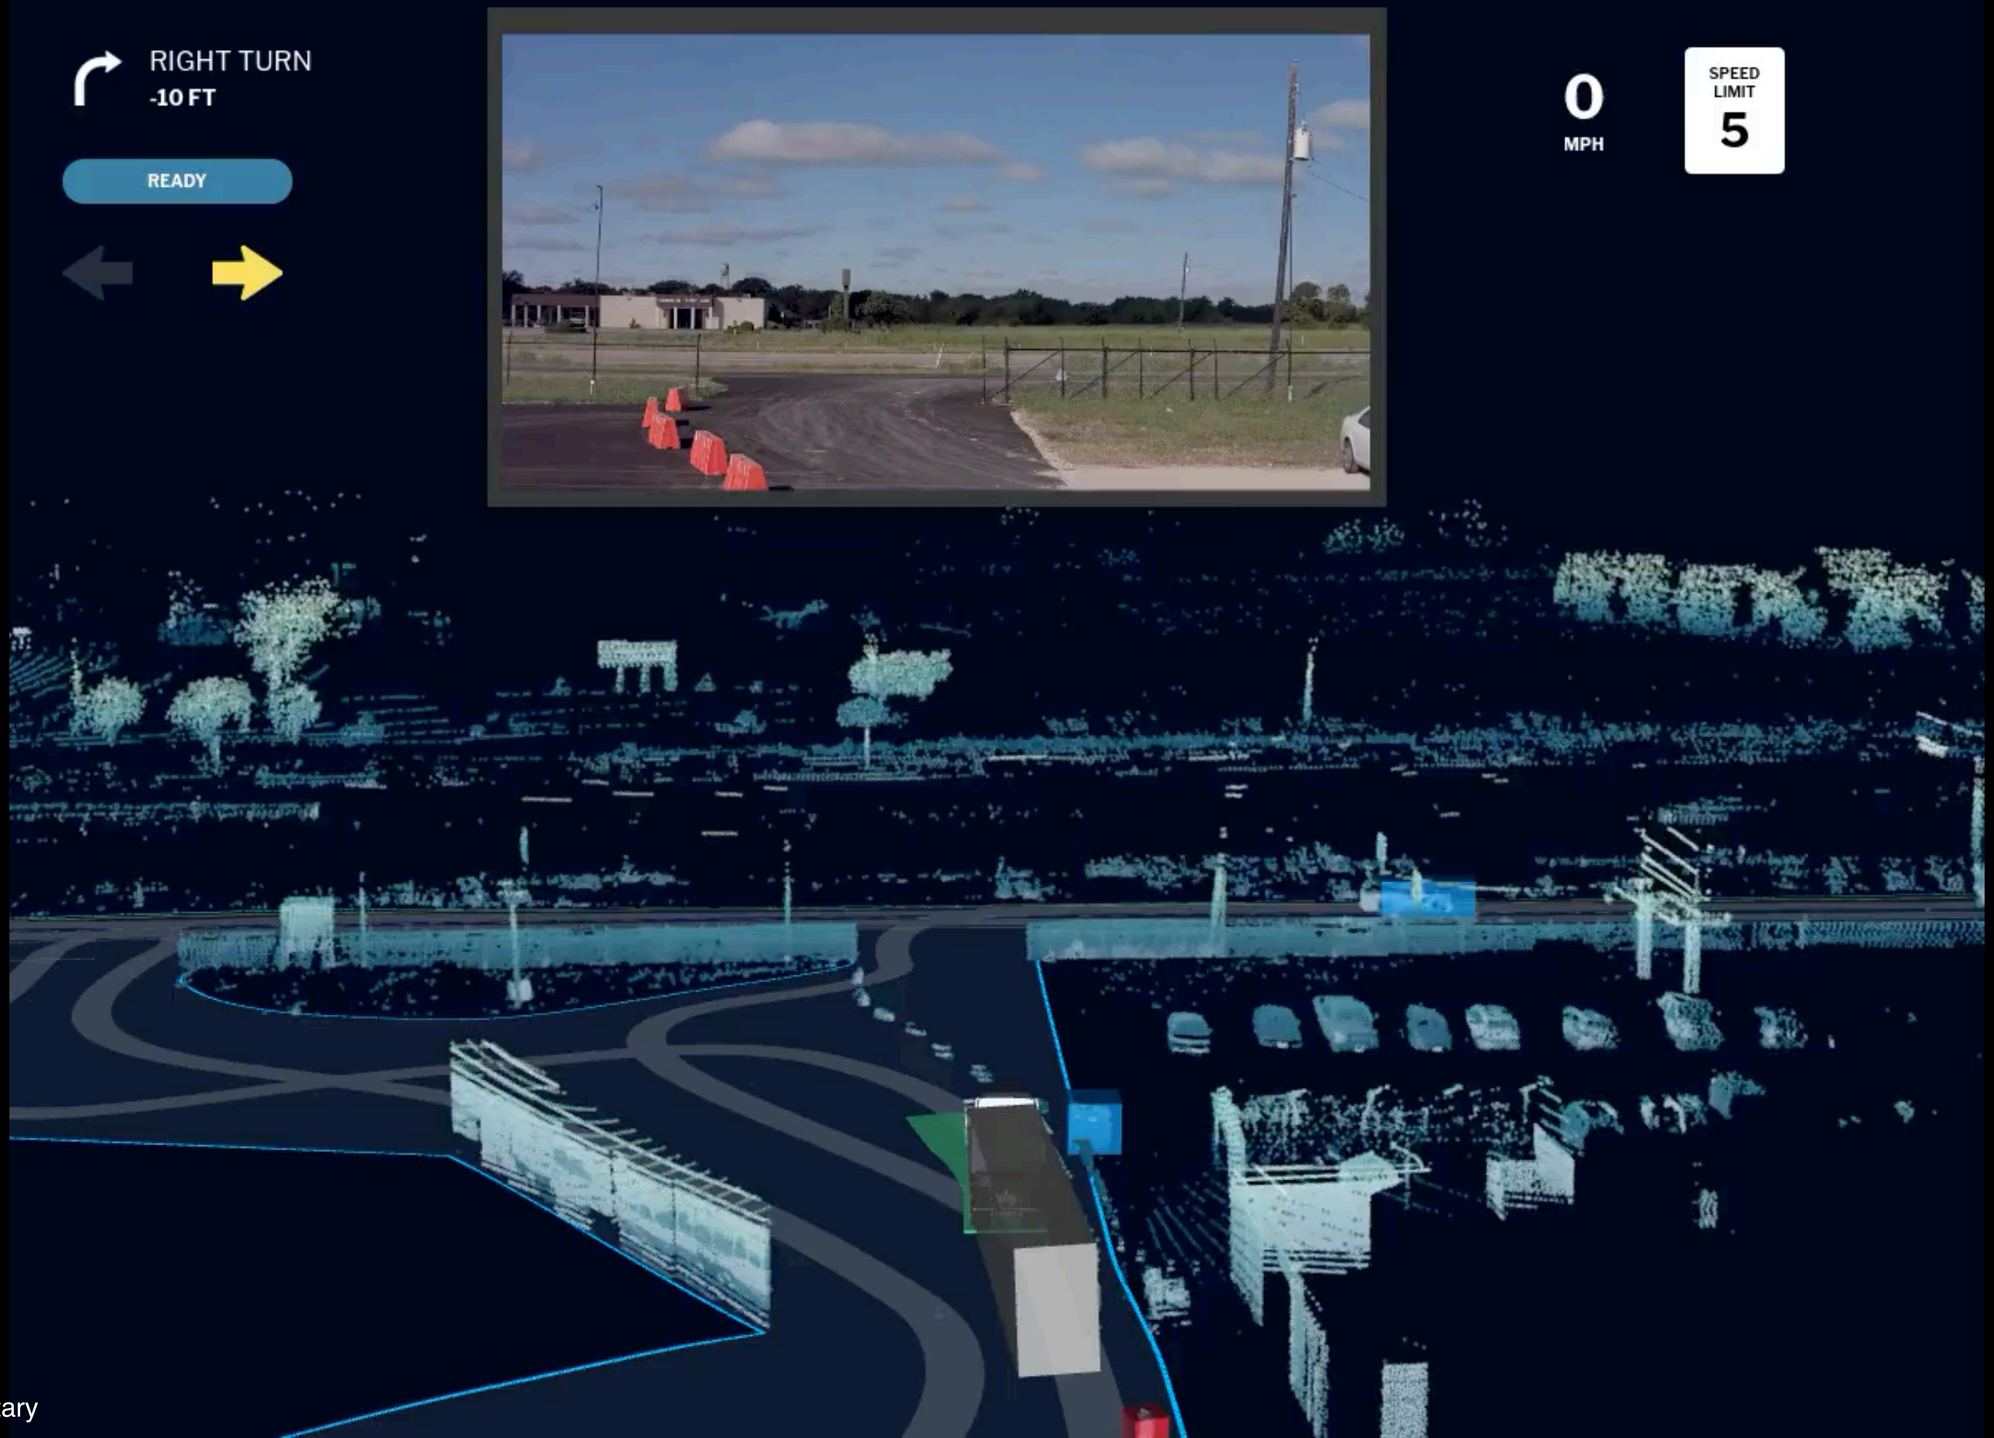

Lidar

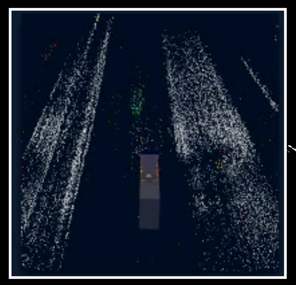

Radar

Camera

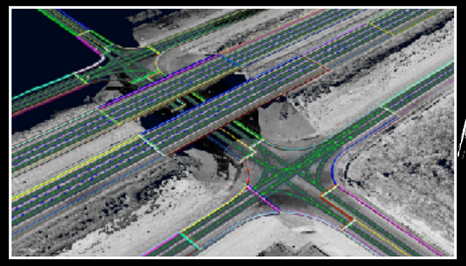

Maps

State

Perception

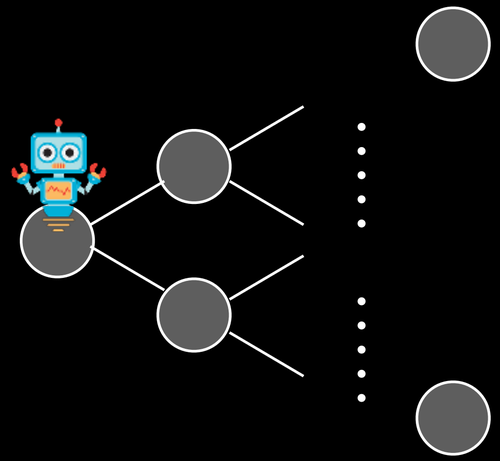

Prediction

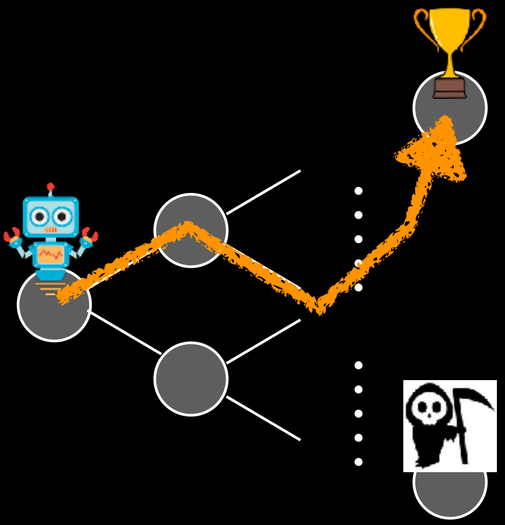

Decision Making

Lidar

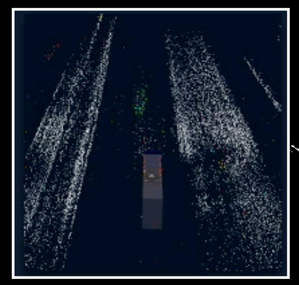

Radar

Camera

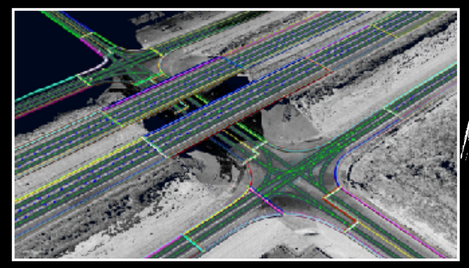

Maps

Learning!!

State

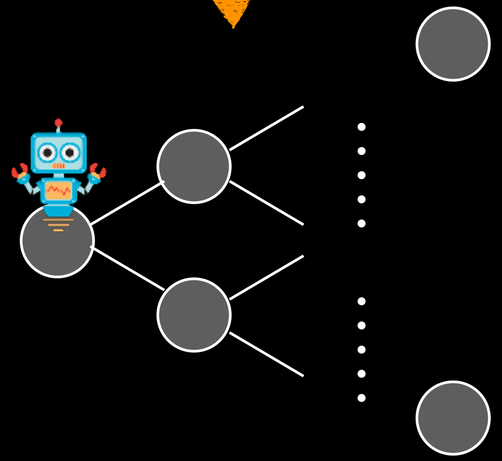

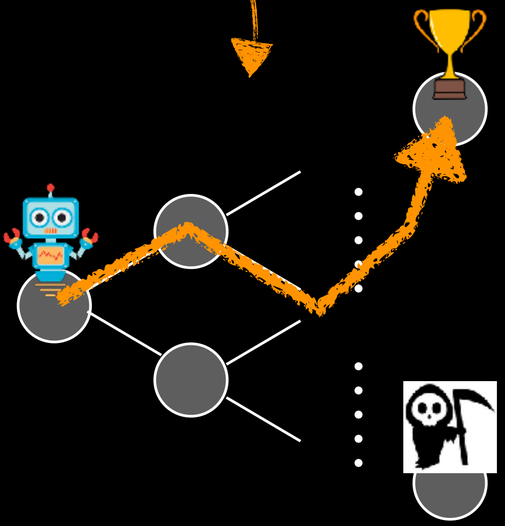

Perception

Prediction

Decision Making

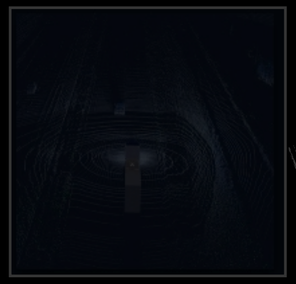

Lidar

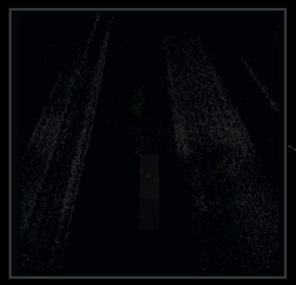

Radar

Camera

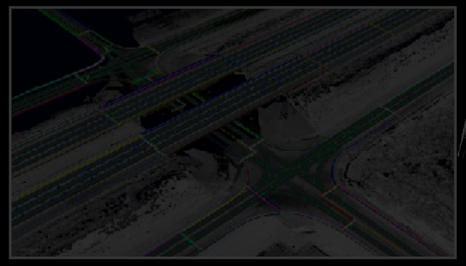

Maps

State

Perception

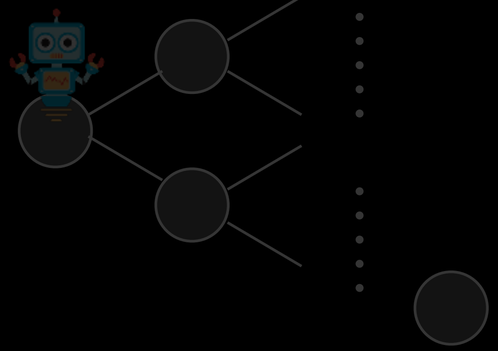

Prediction

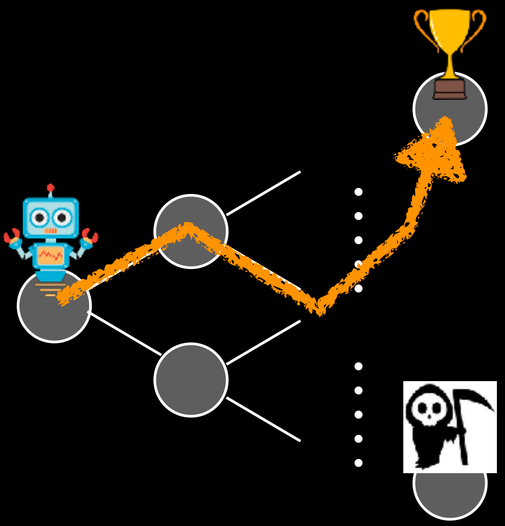

Decision Making

Activity: What is “good” behavior in a left turn?

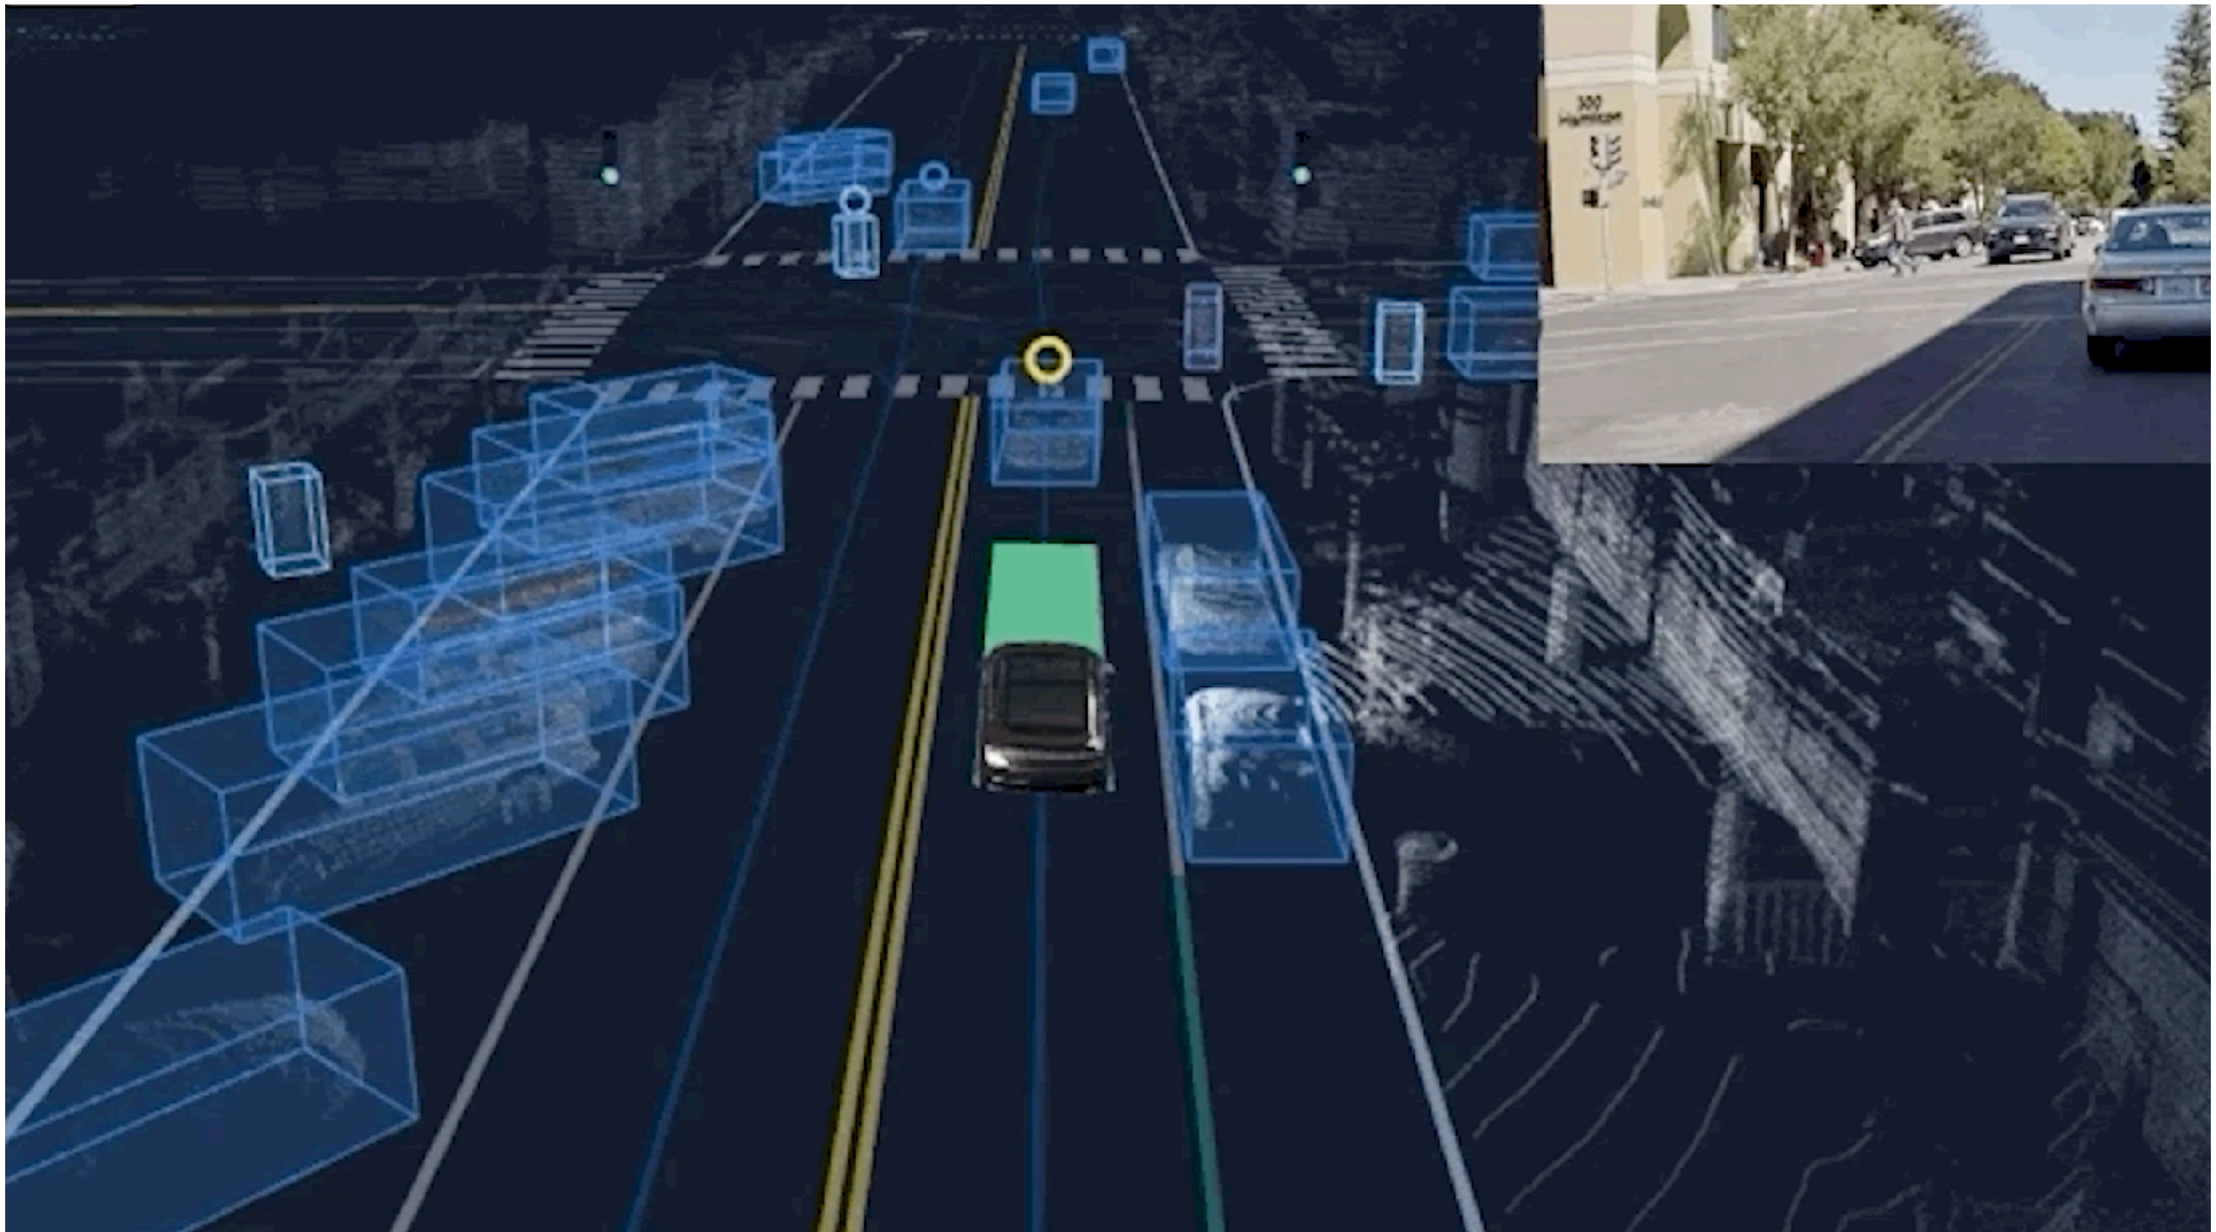

Activity!

Activity: What is “good” behavior in a left turn?

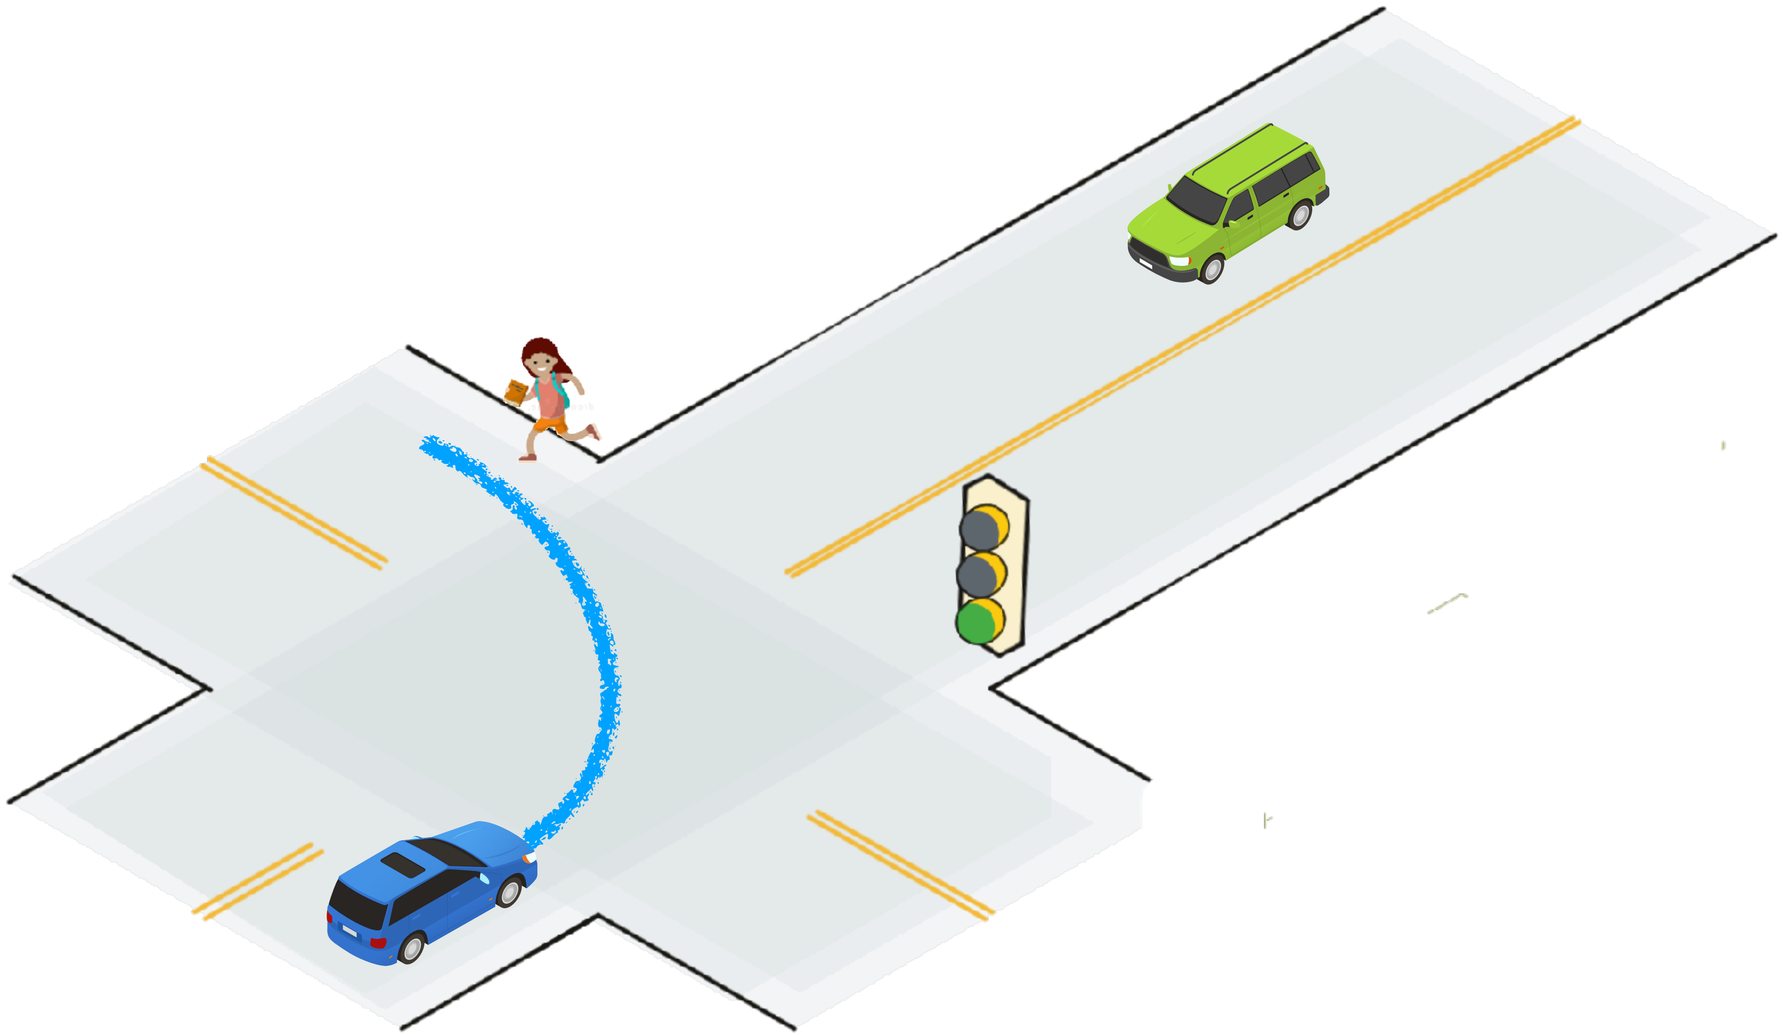

Lesson #1

Values are implicit in human driving!

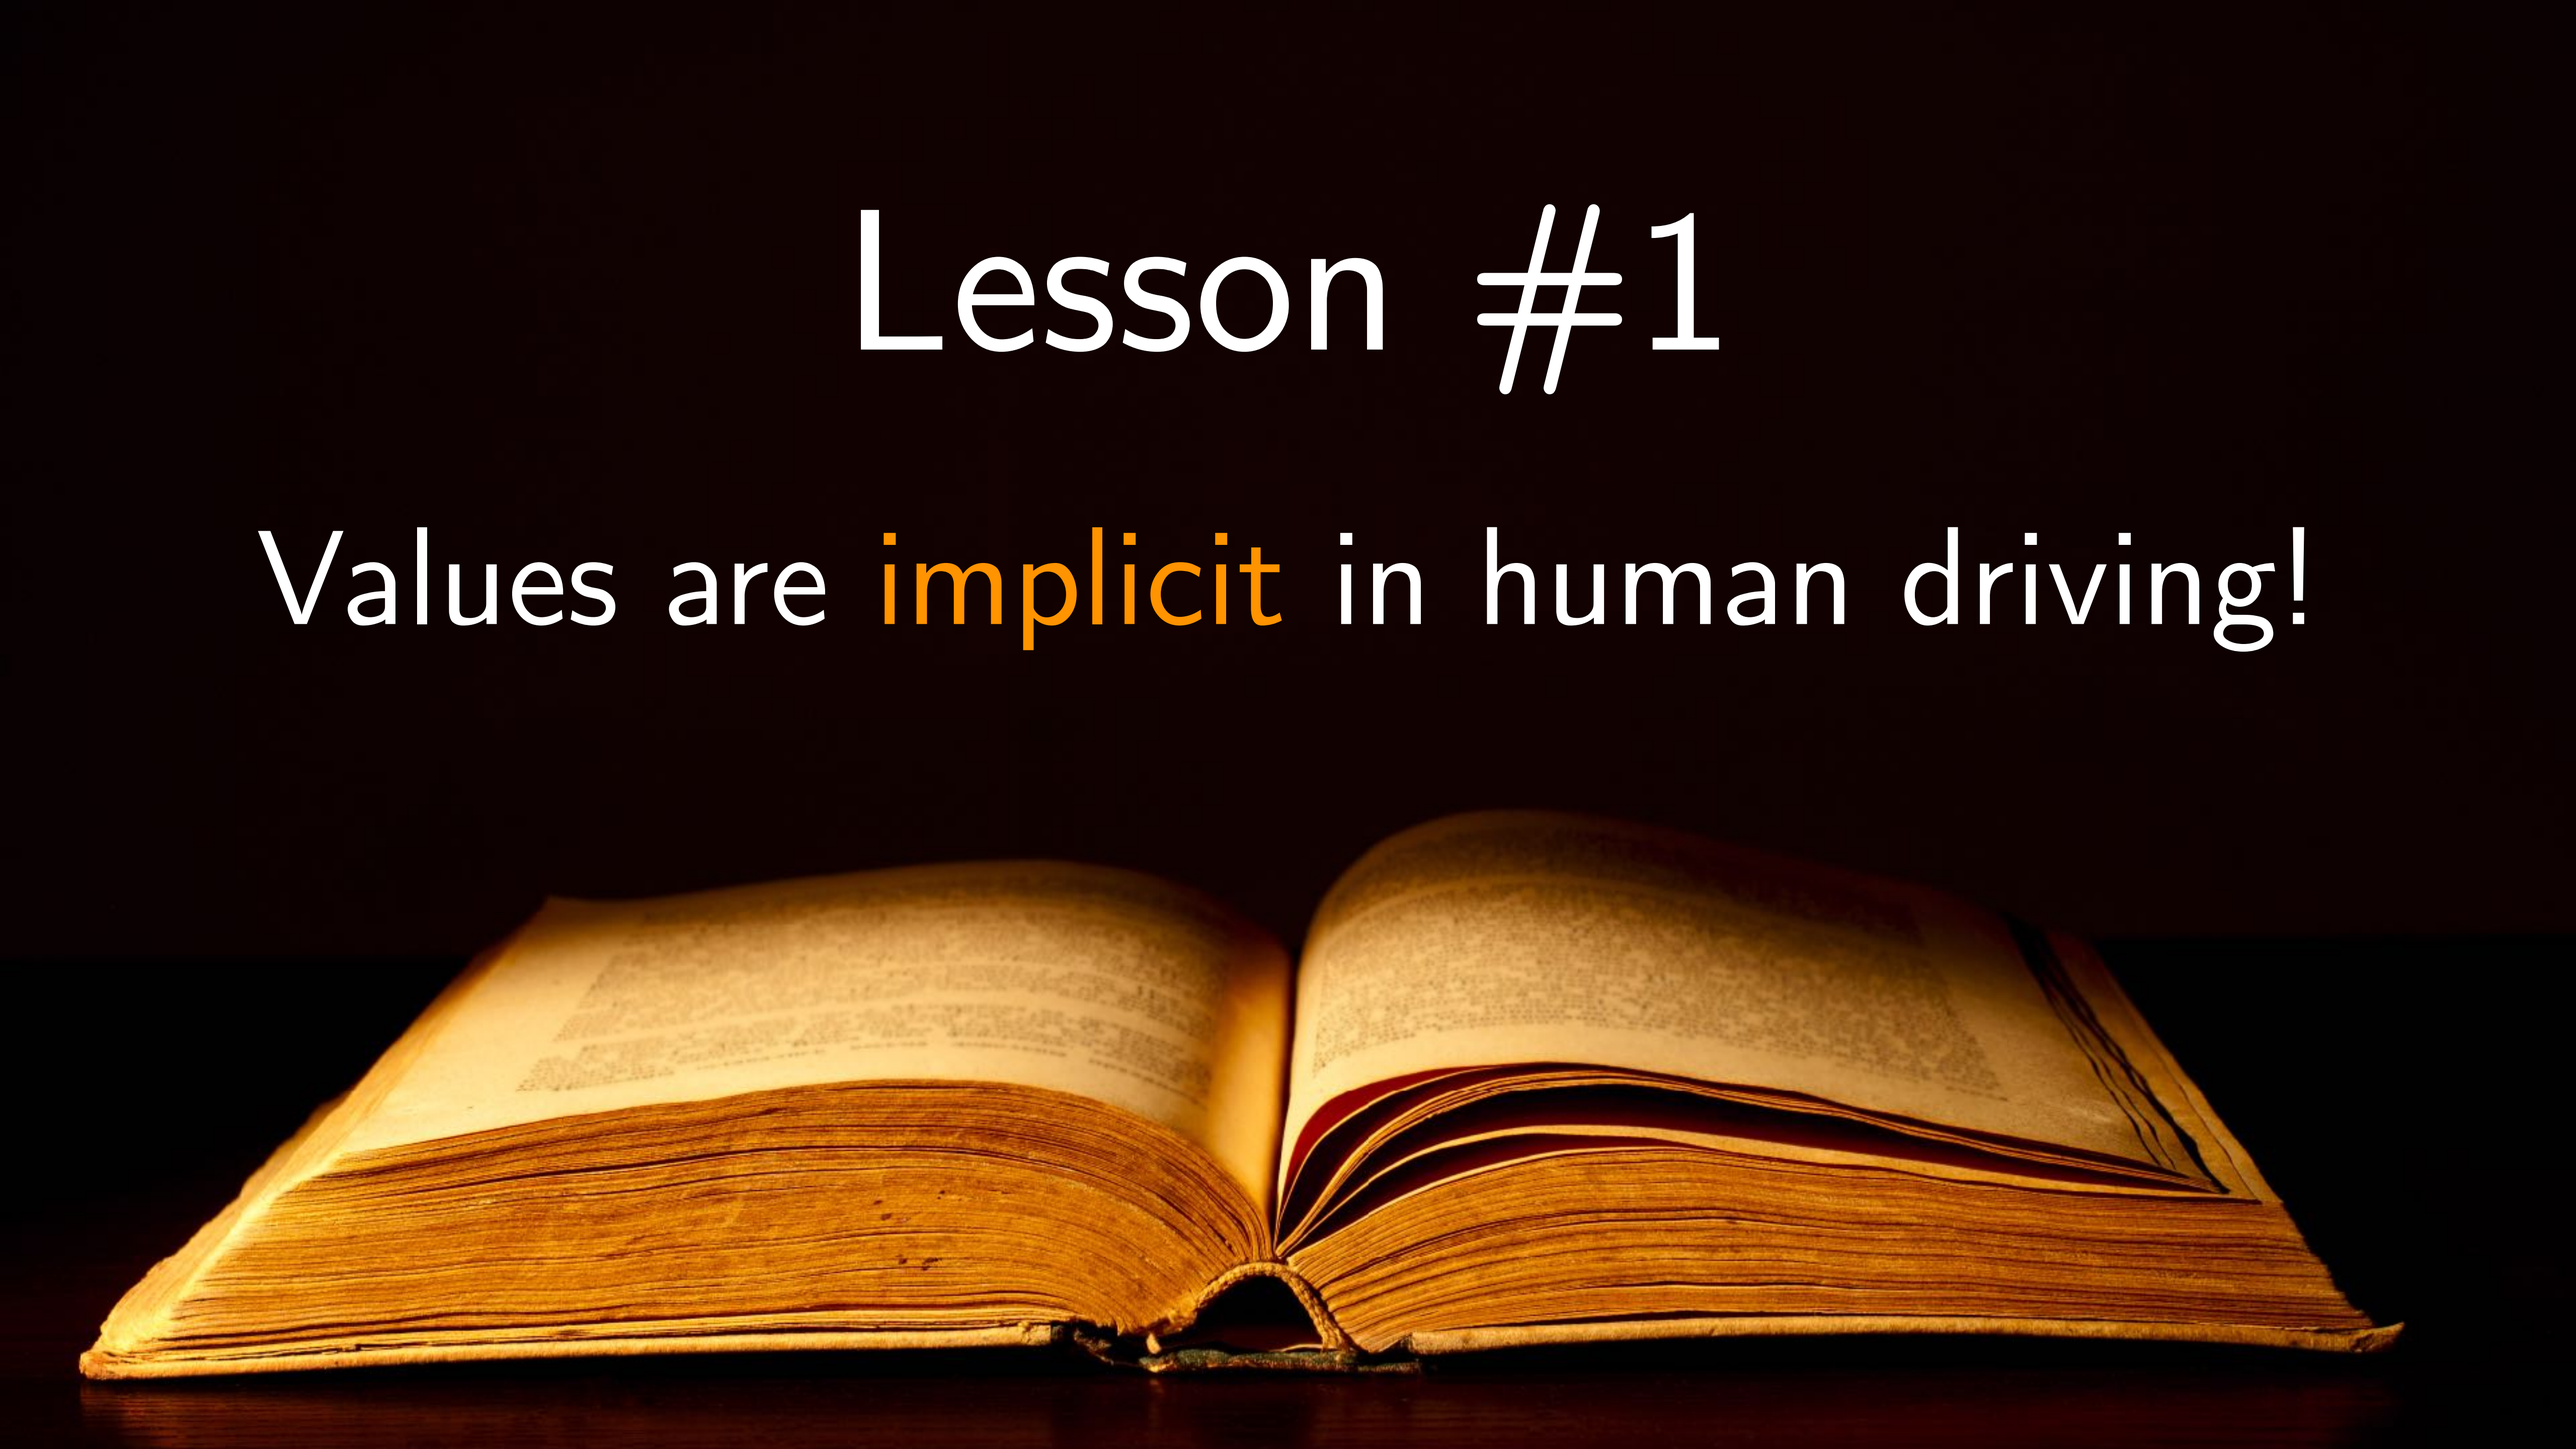

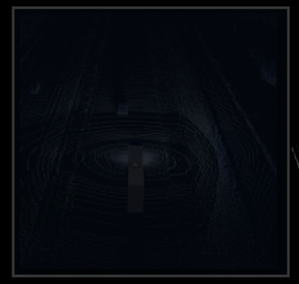

Lidar

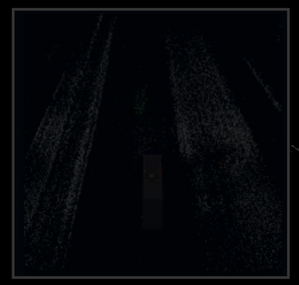

Radar

Camera

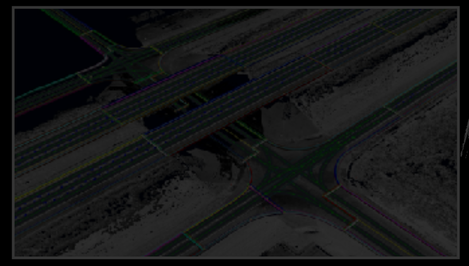

Maps

State

Perception

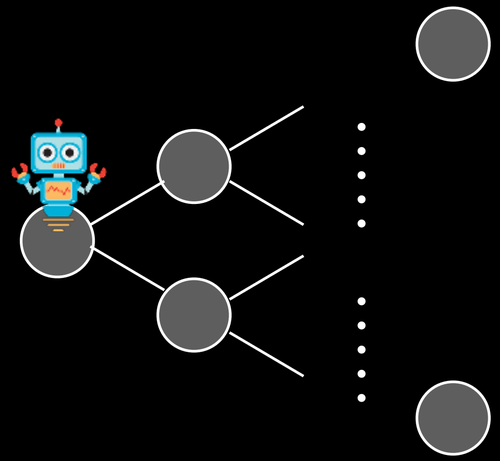

Prediction

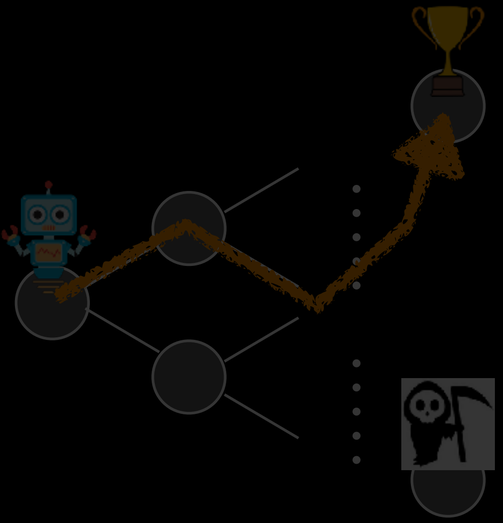

Decision Making

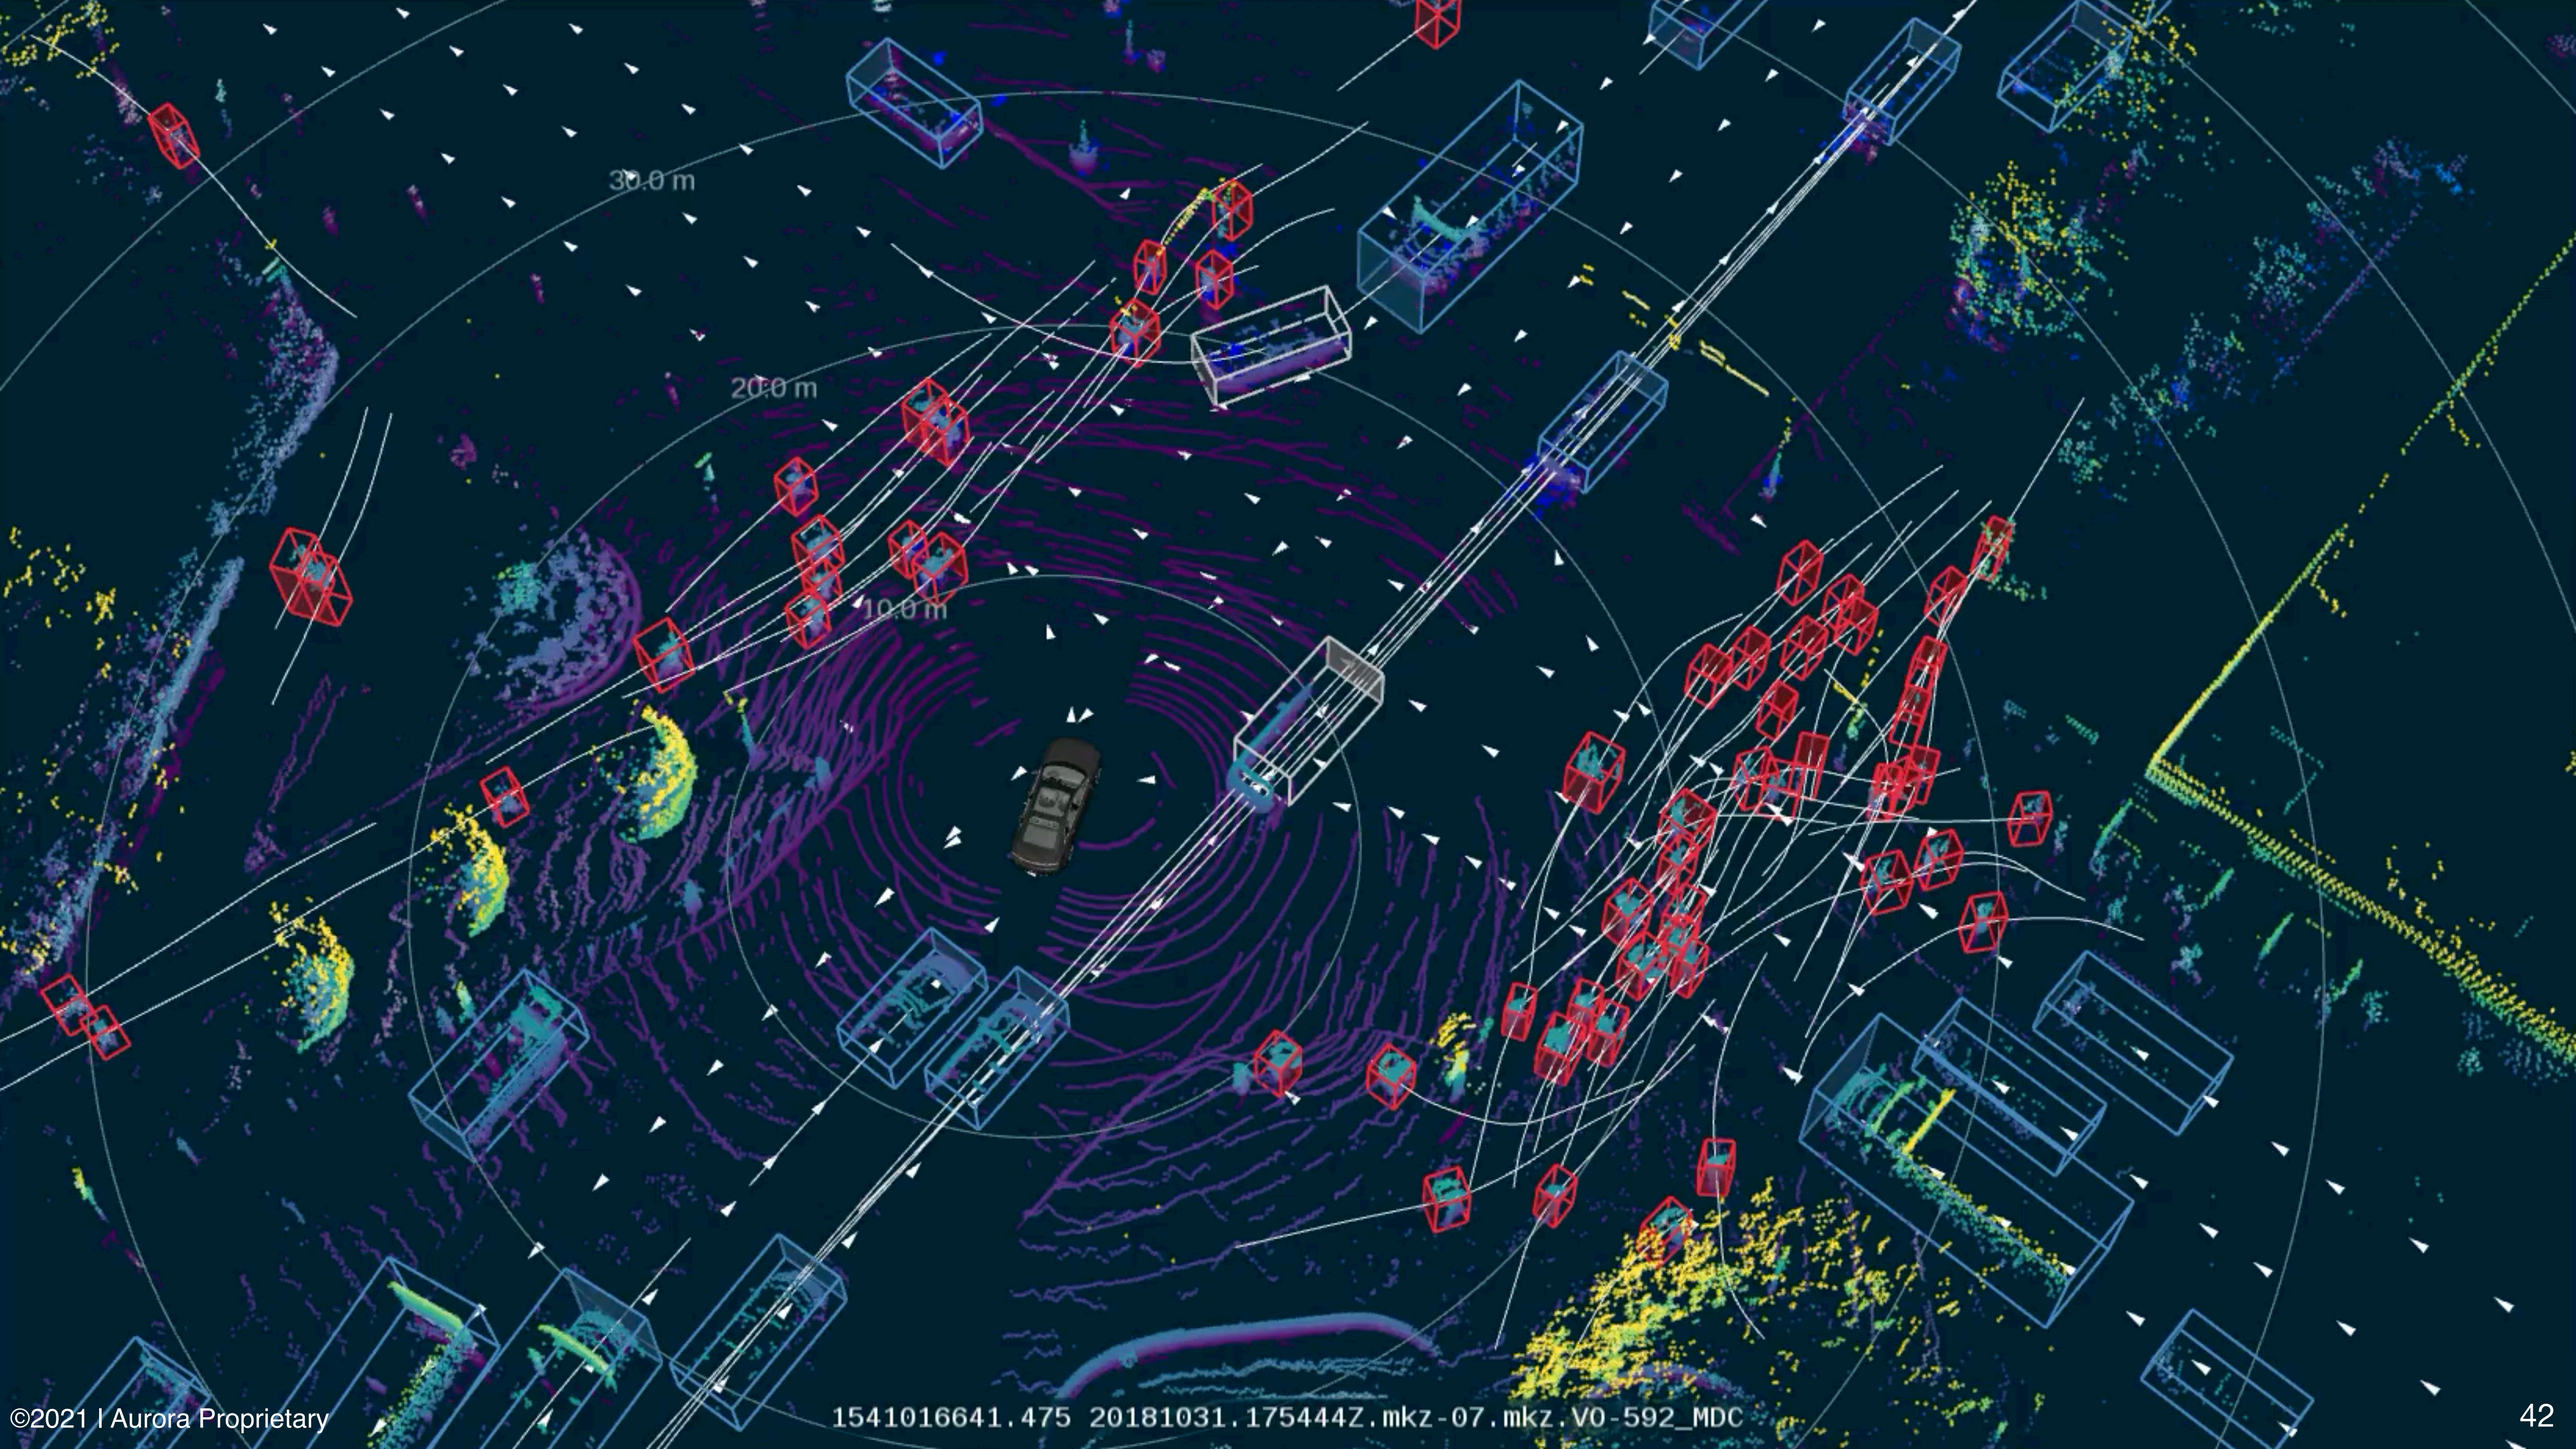

Brainstorm: How can we predict the pedestrian motion?

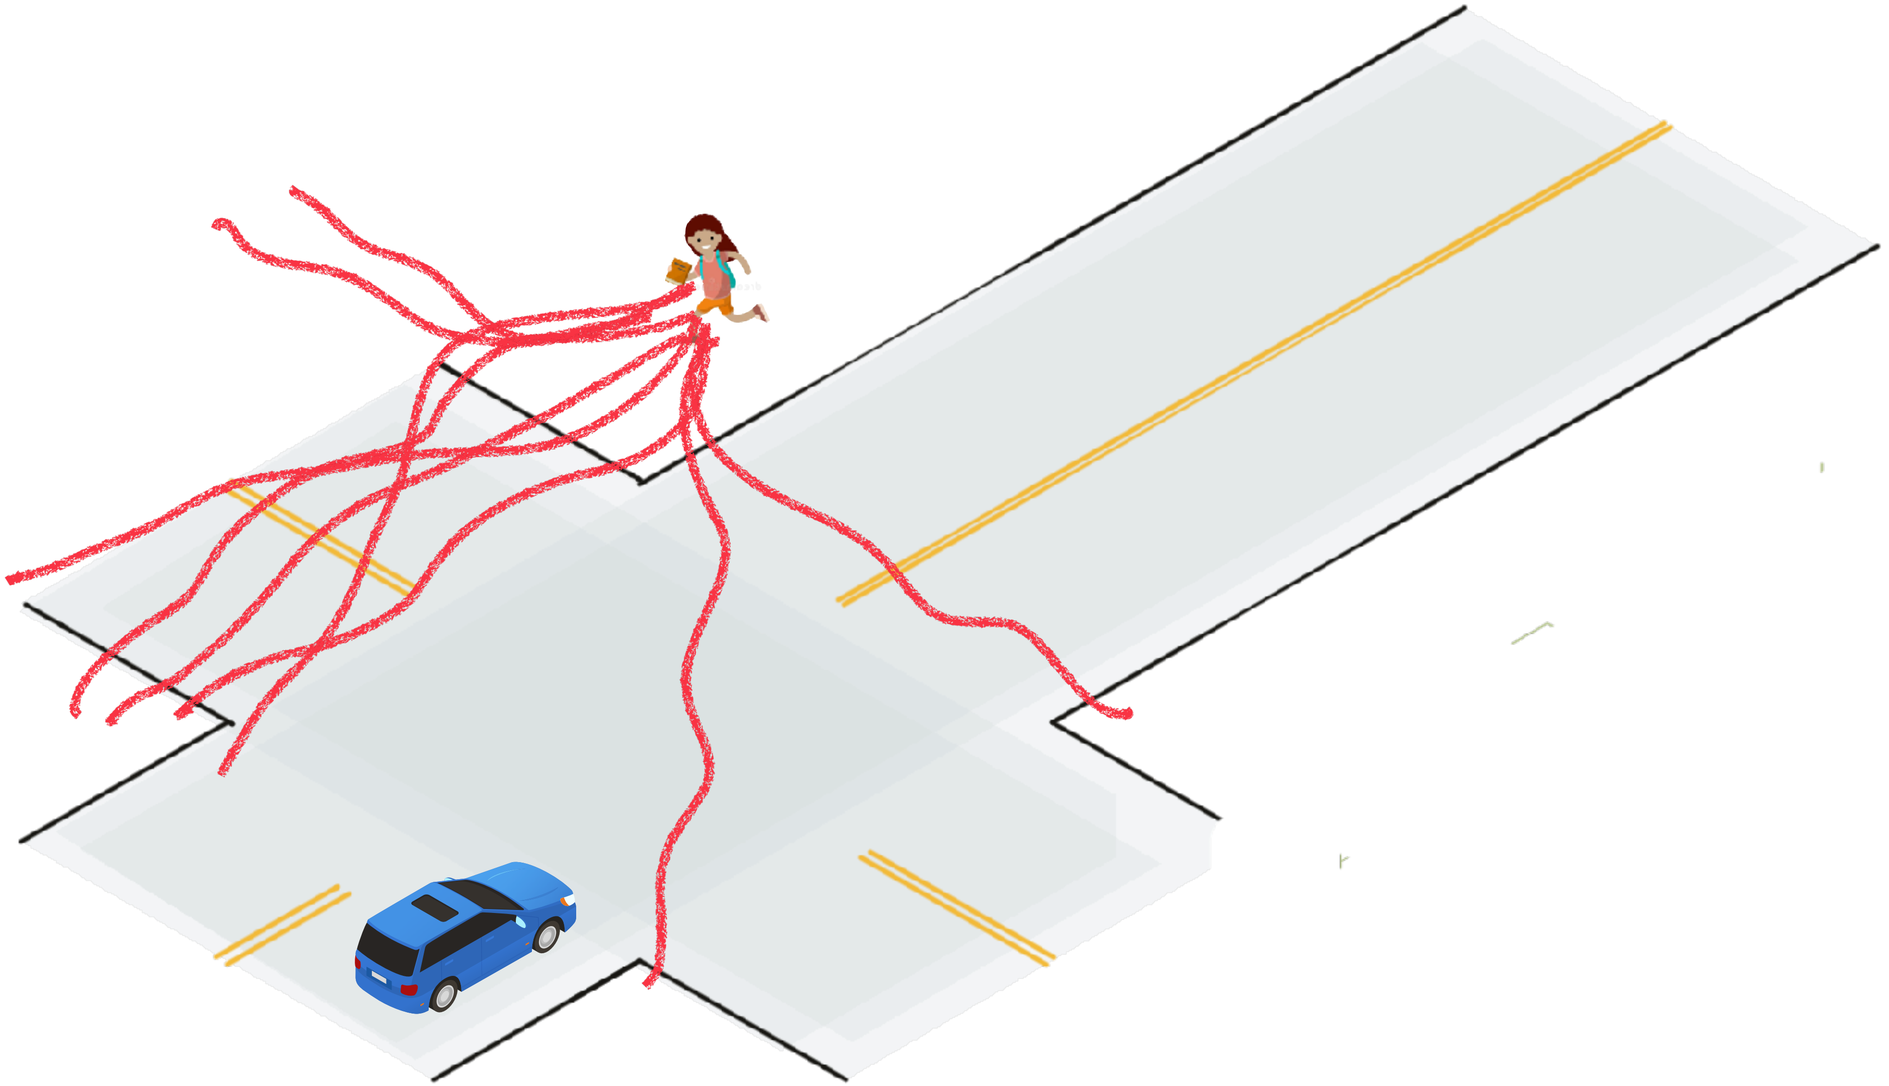

Lesson #2

Models are useful fictions

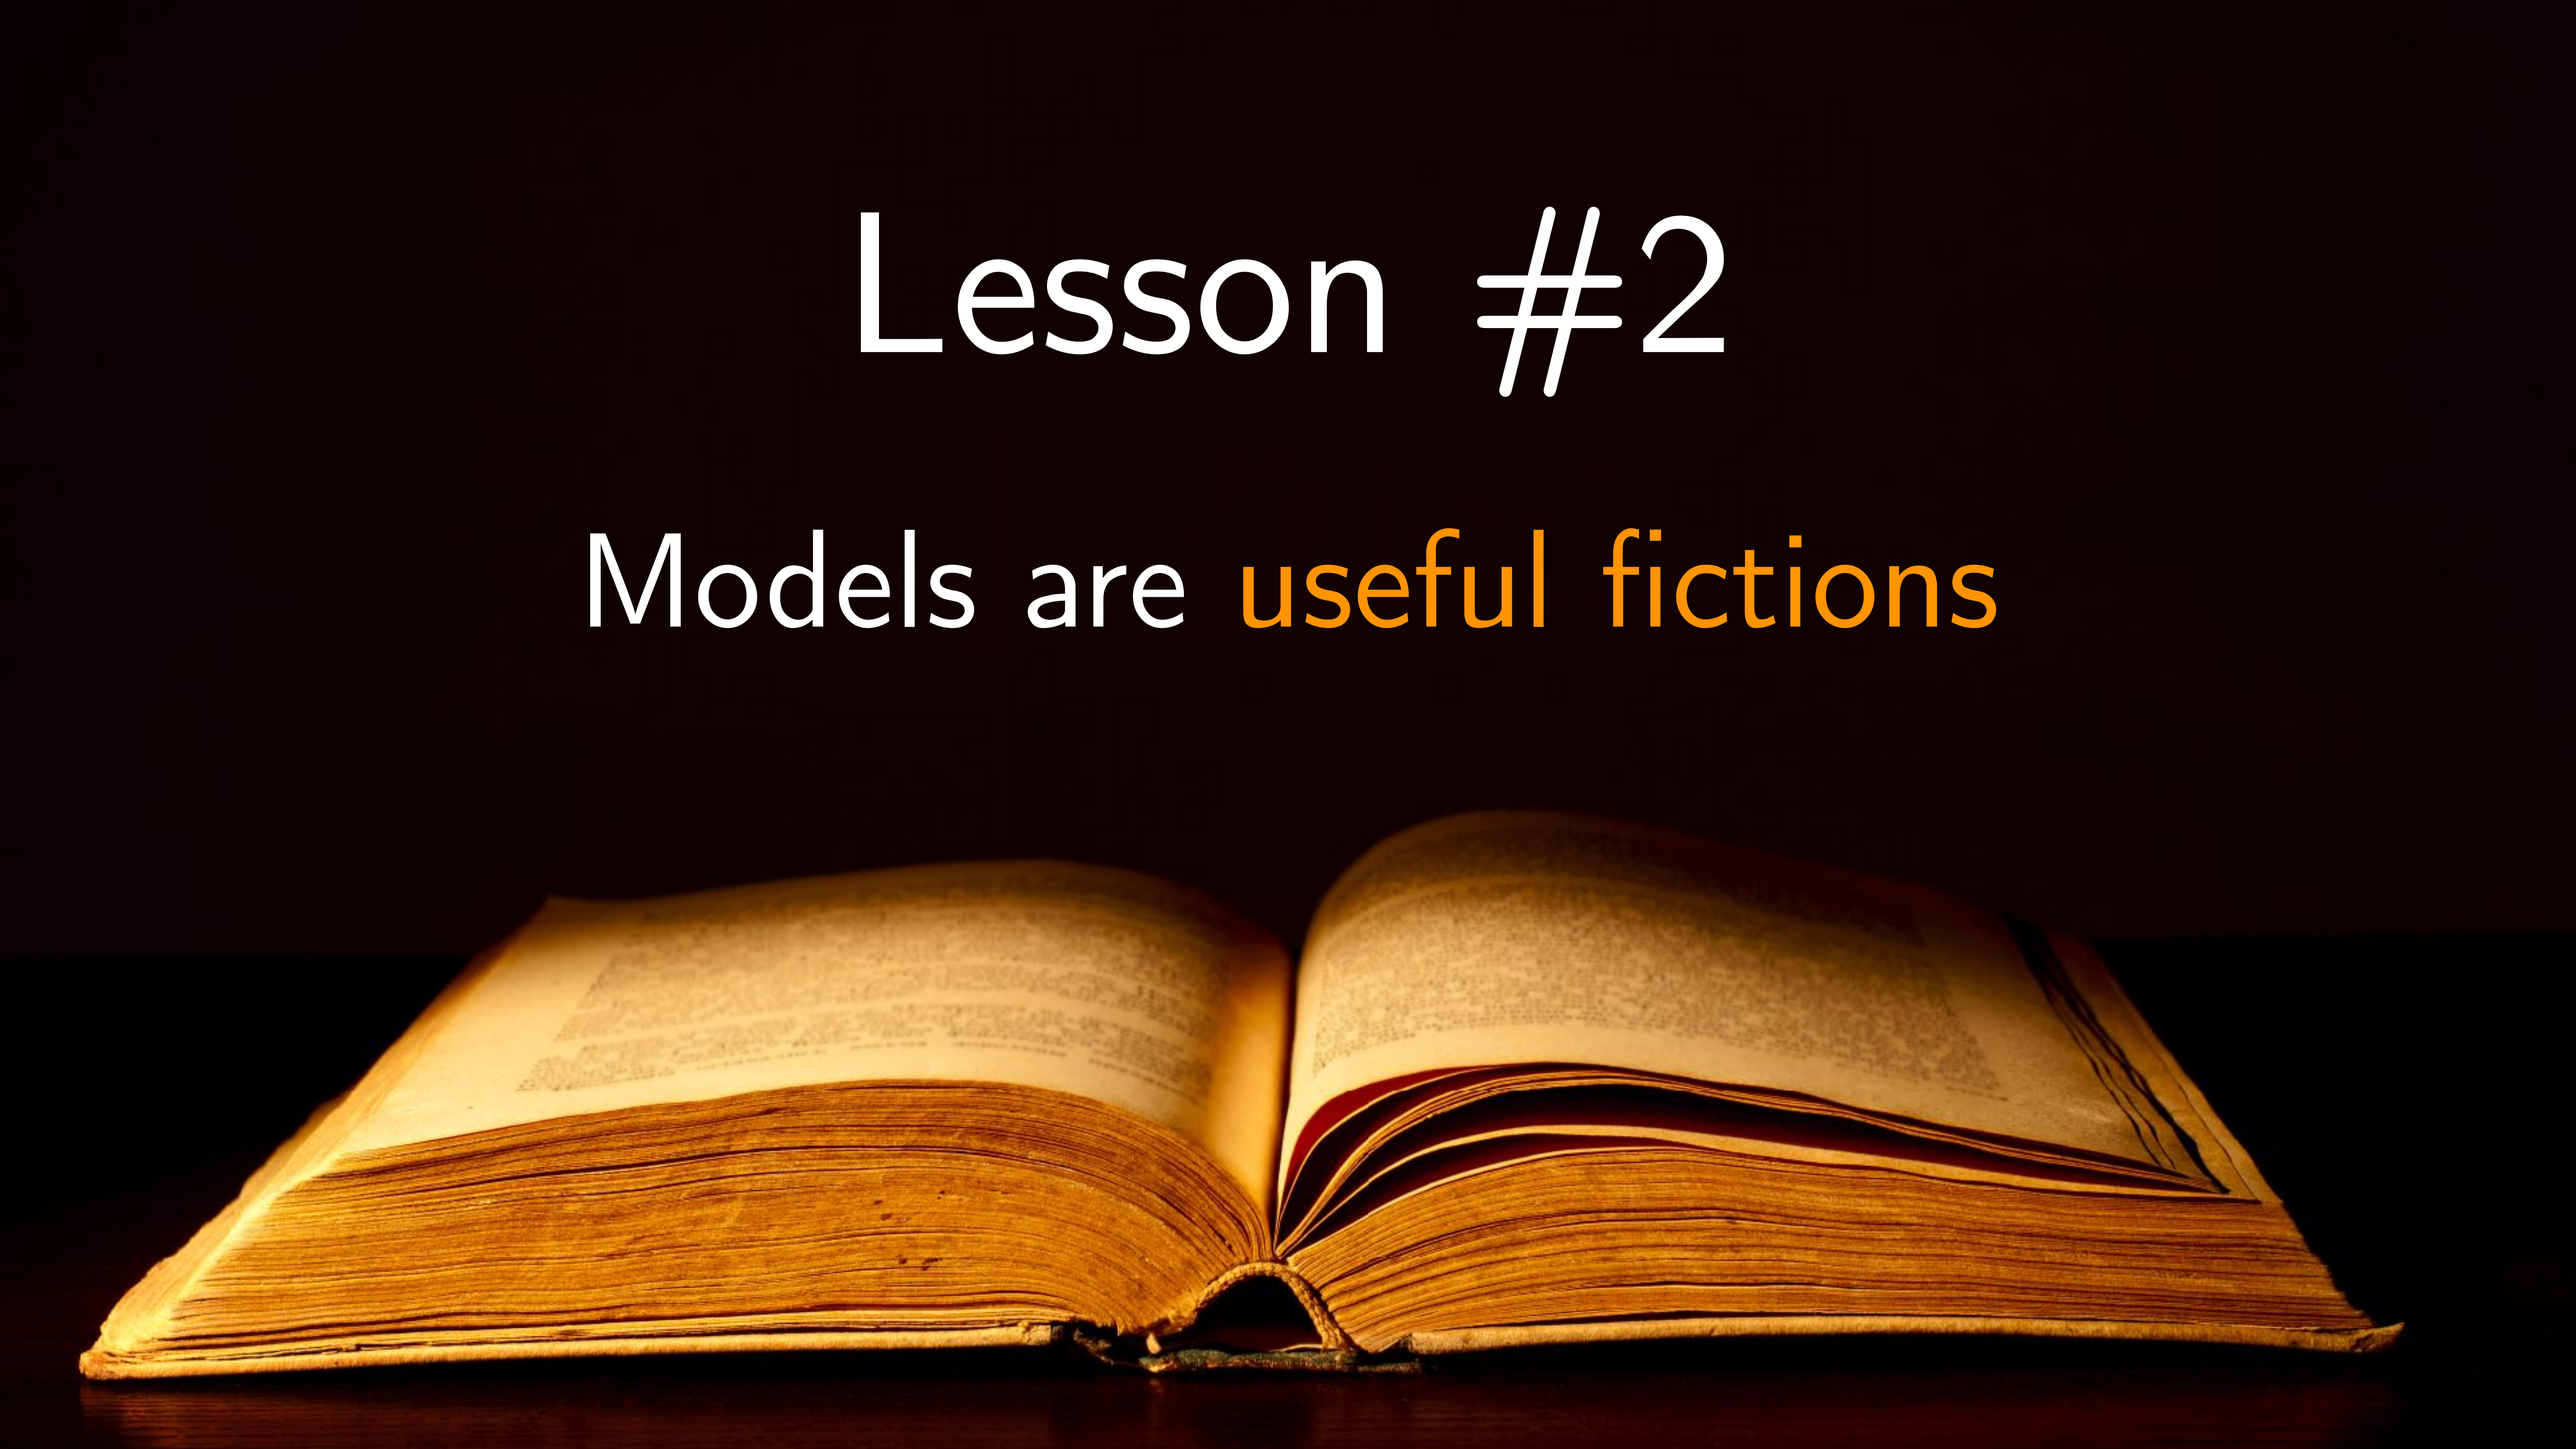

Lidar

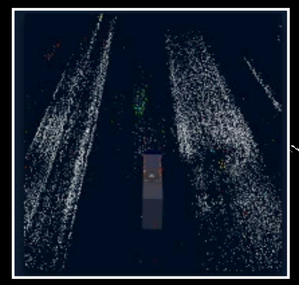

Radar

Camera

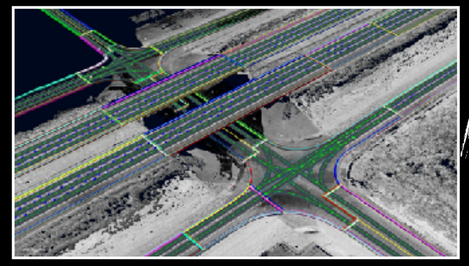

Maps

State

Perception

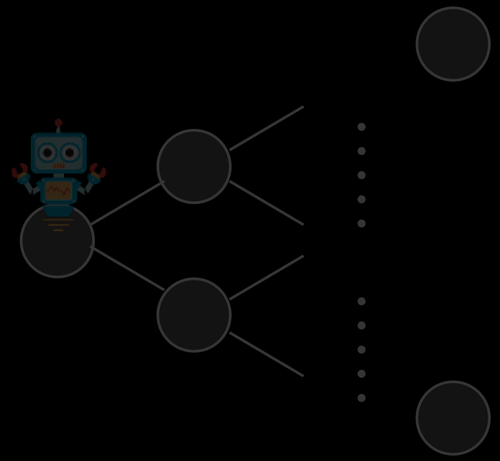

Prediction

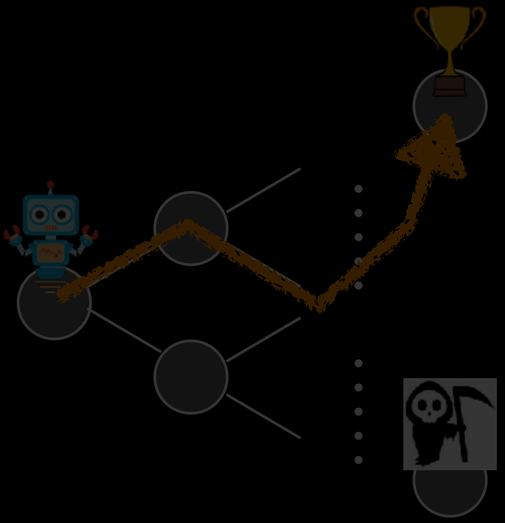

Decision Making

Fuse complementary sensors

LiDAR

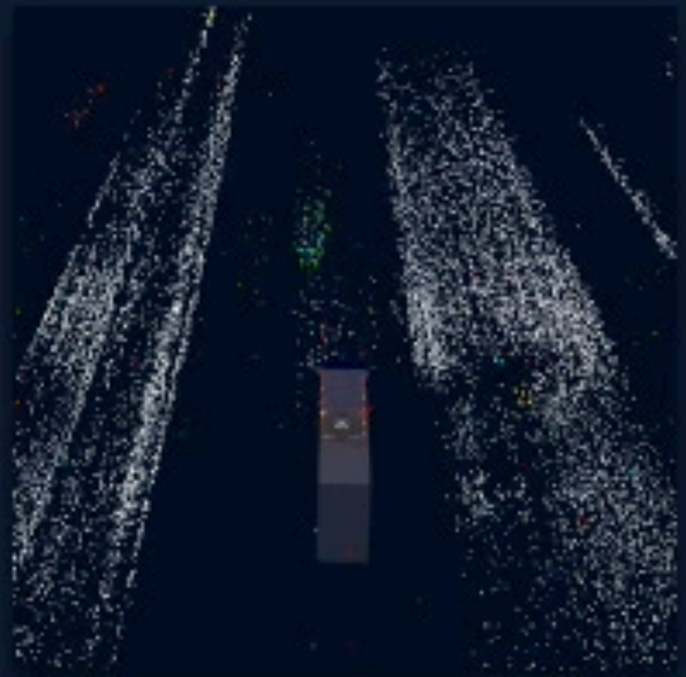

Radar

Camera

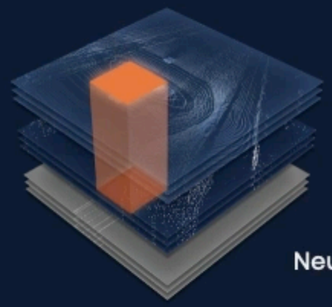

Neural Convolution Engine

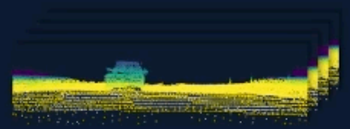

Range Conv. Engine

Image Fusion Engine

Euclidian Ray
Scatter Engine

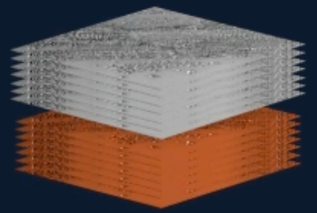

Neural Convolution
Engine

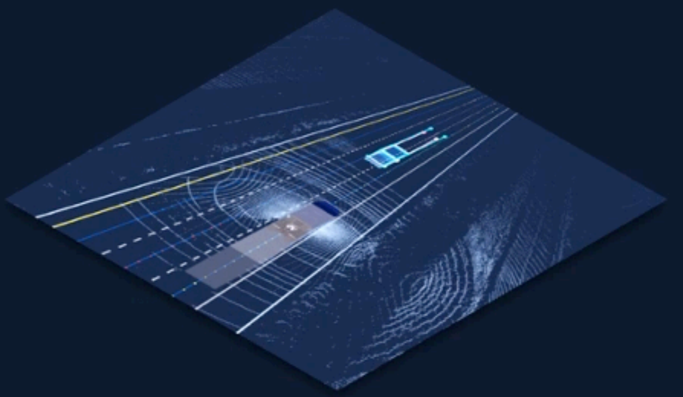

Lesson #3

Solve for the state

that explains all observations

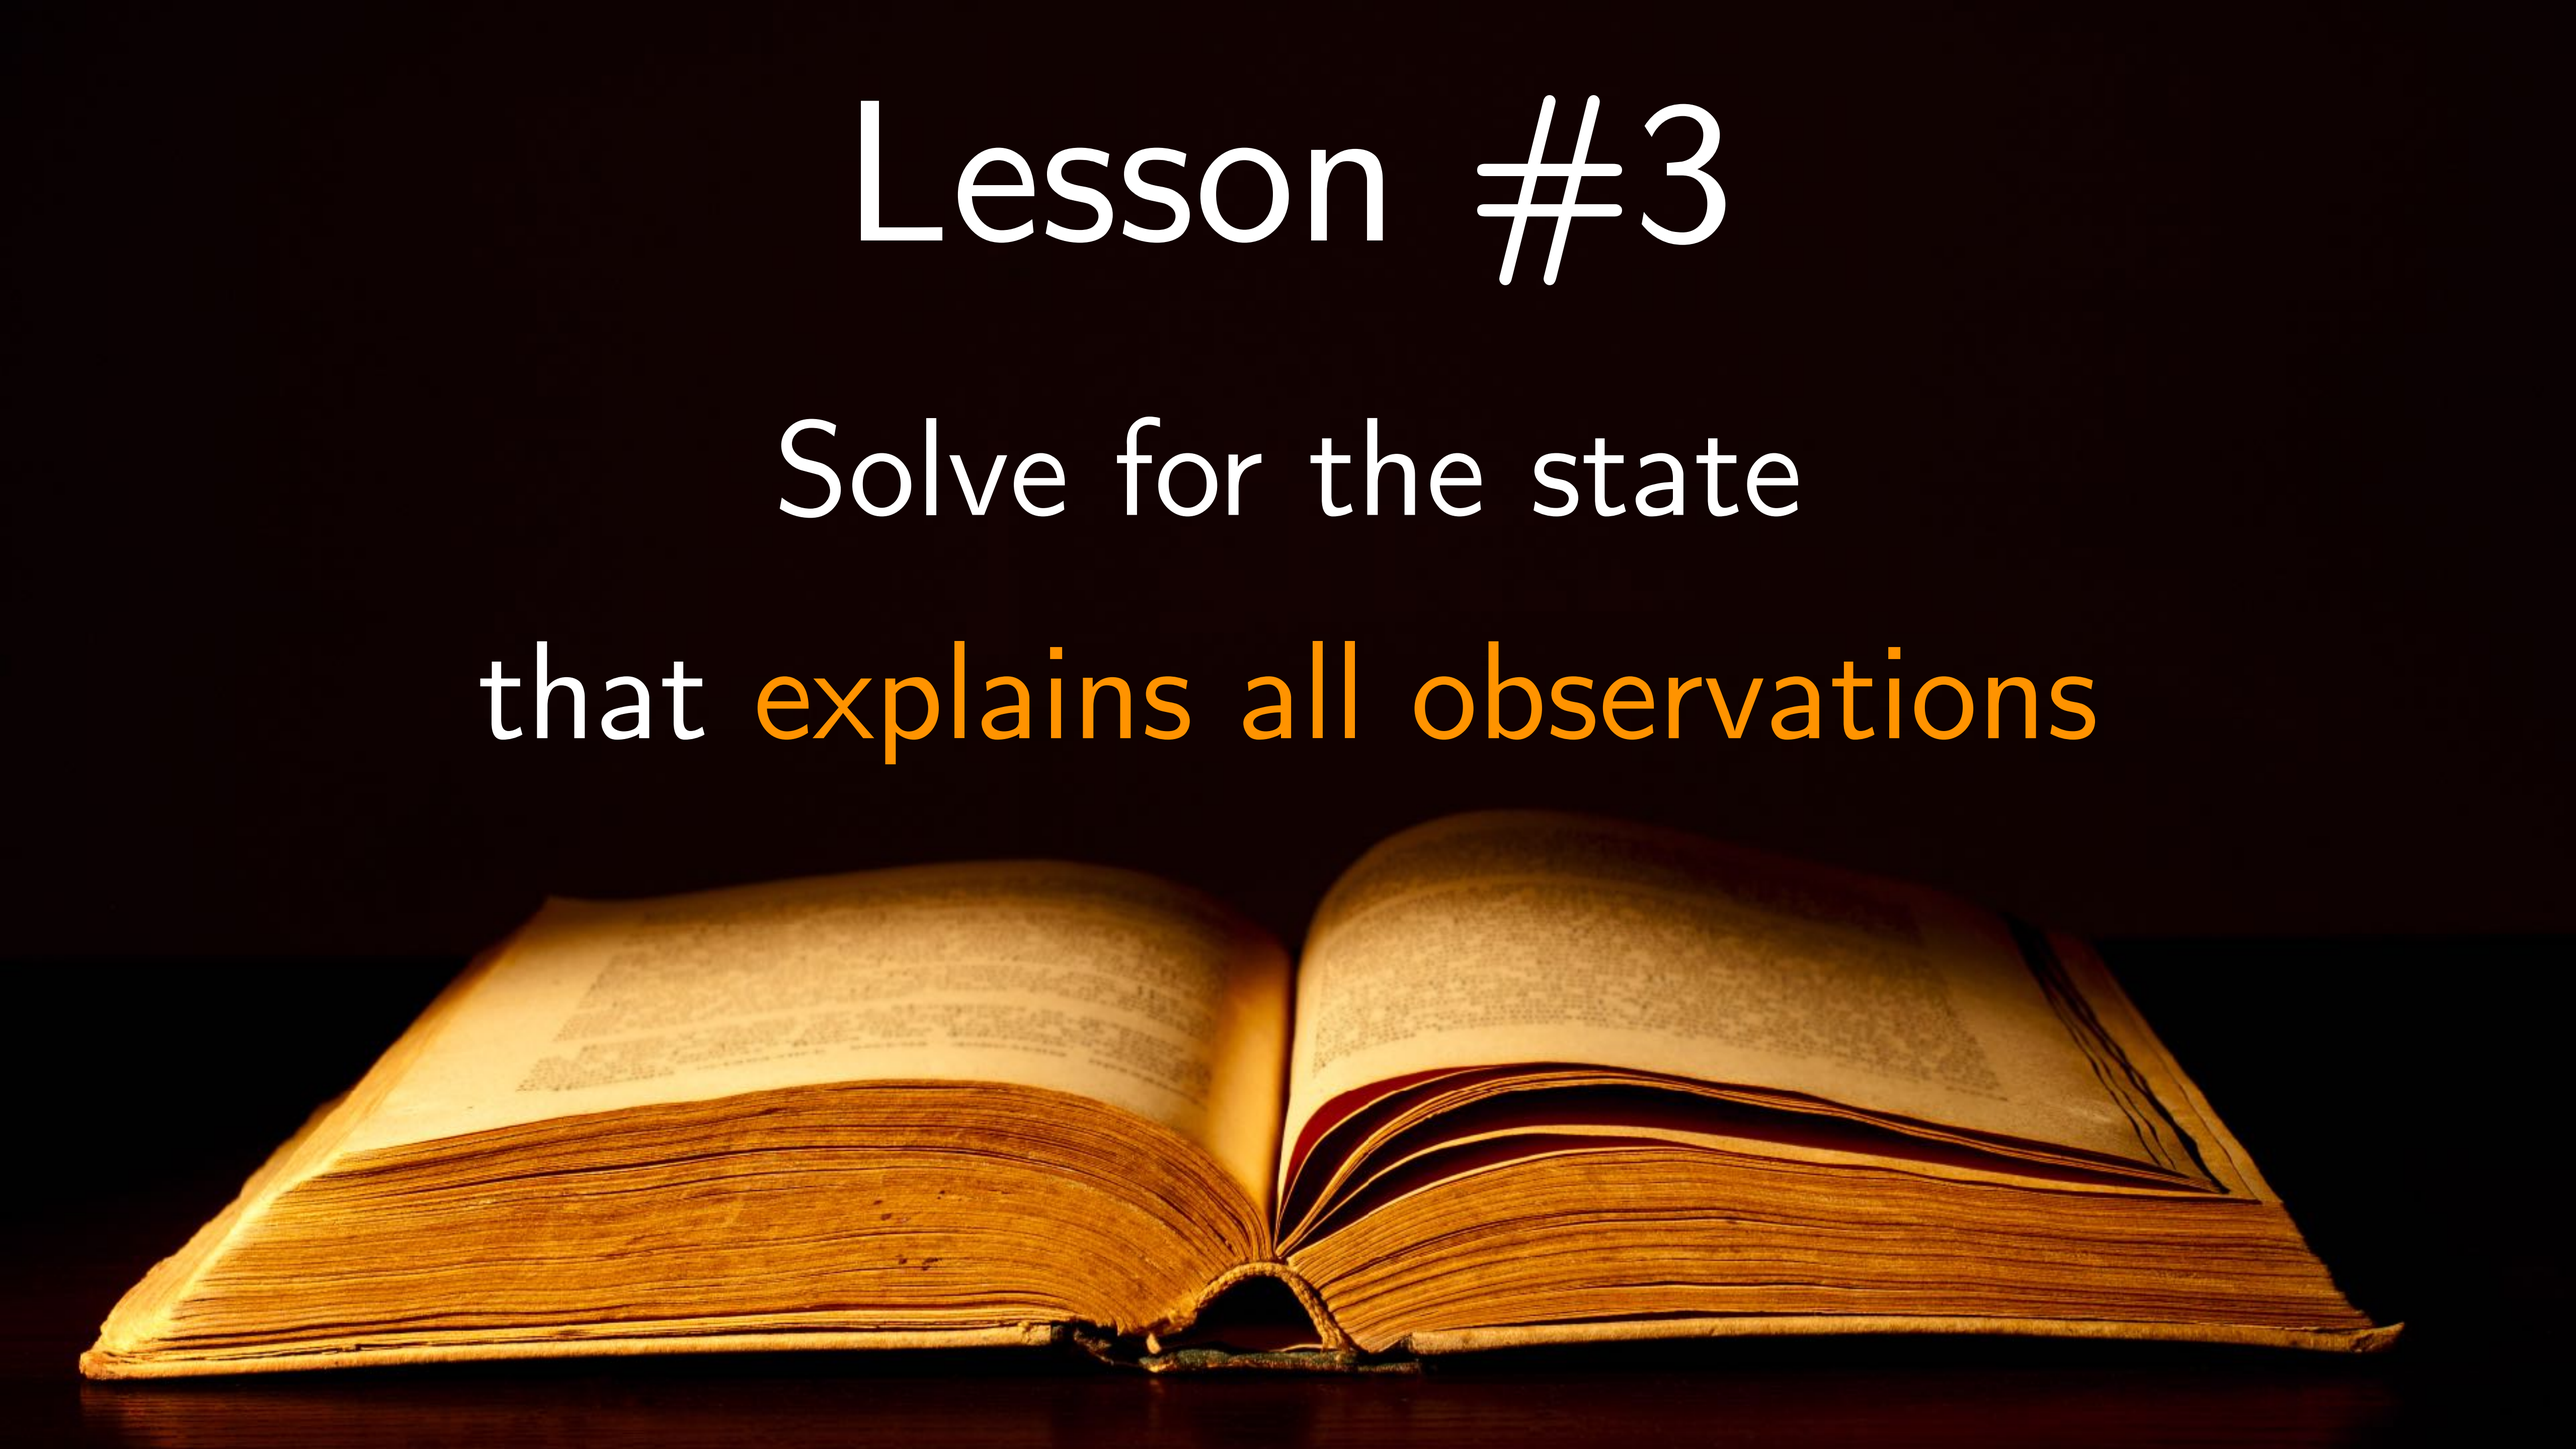

The journey ahead!

Localization
& Mapping

World Models
& Forecasting

Open Challenges

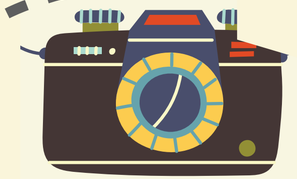

Robot
Perception

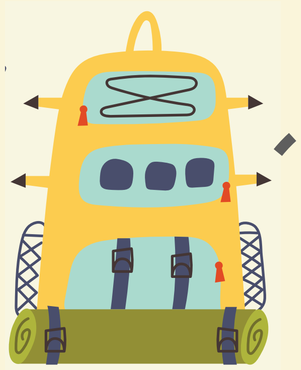

Fundamentals

Imitation
Learning

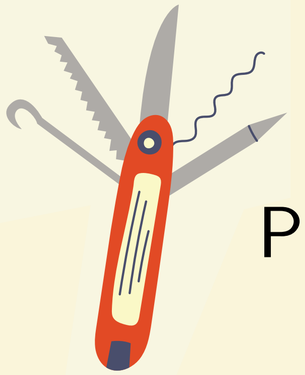

Planning &
Control

Reinforcement
Learning

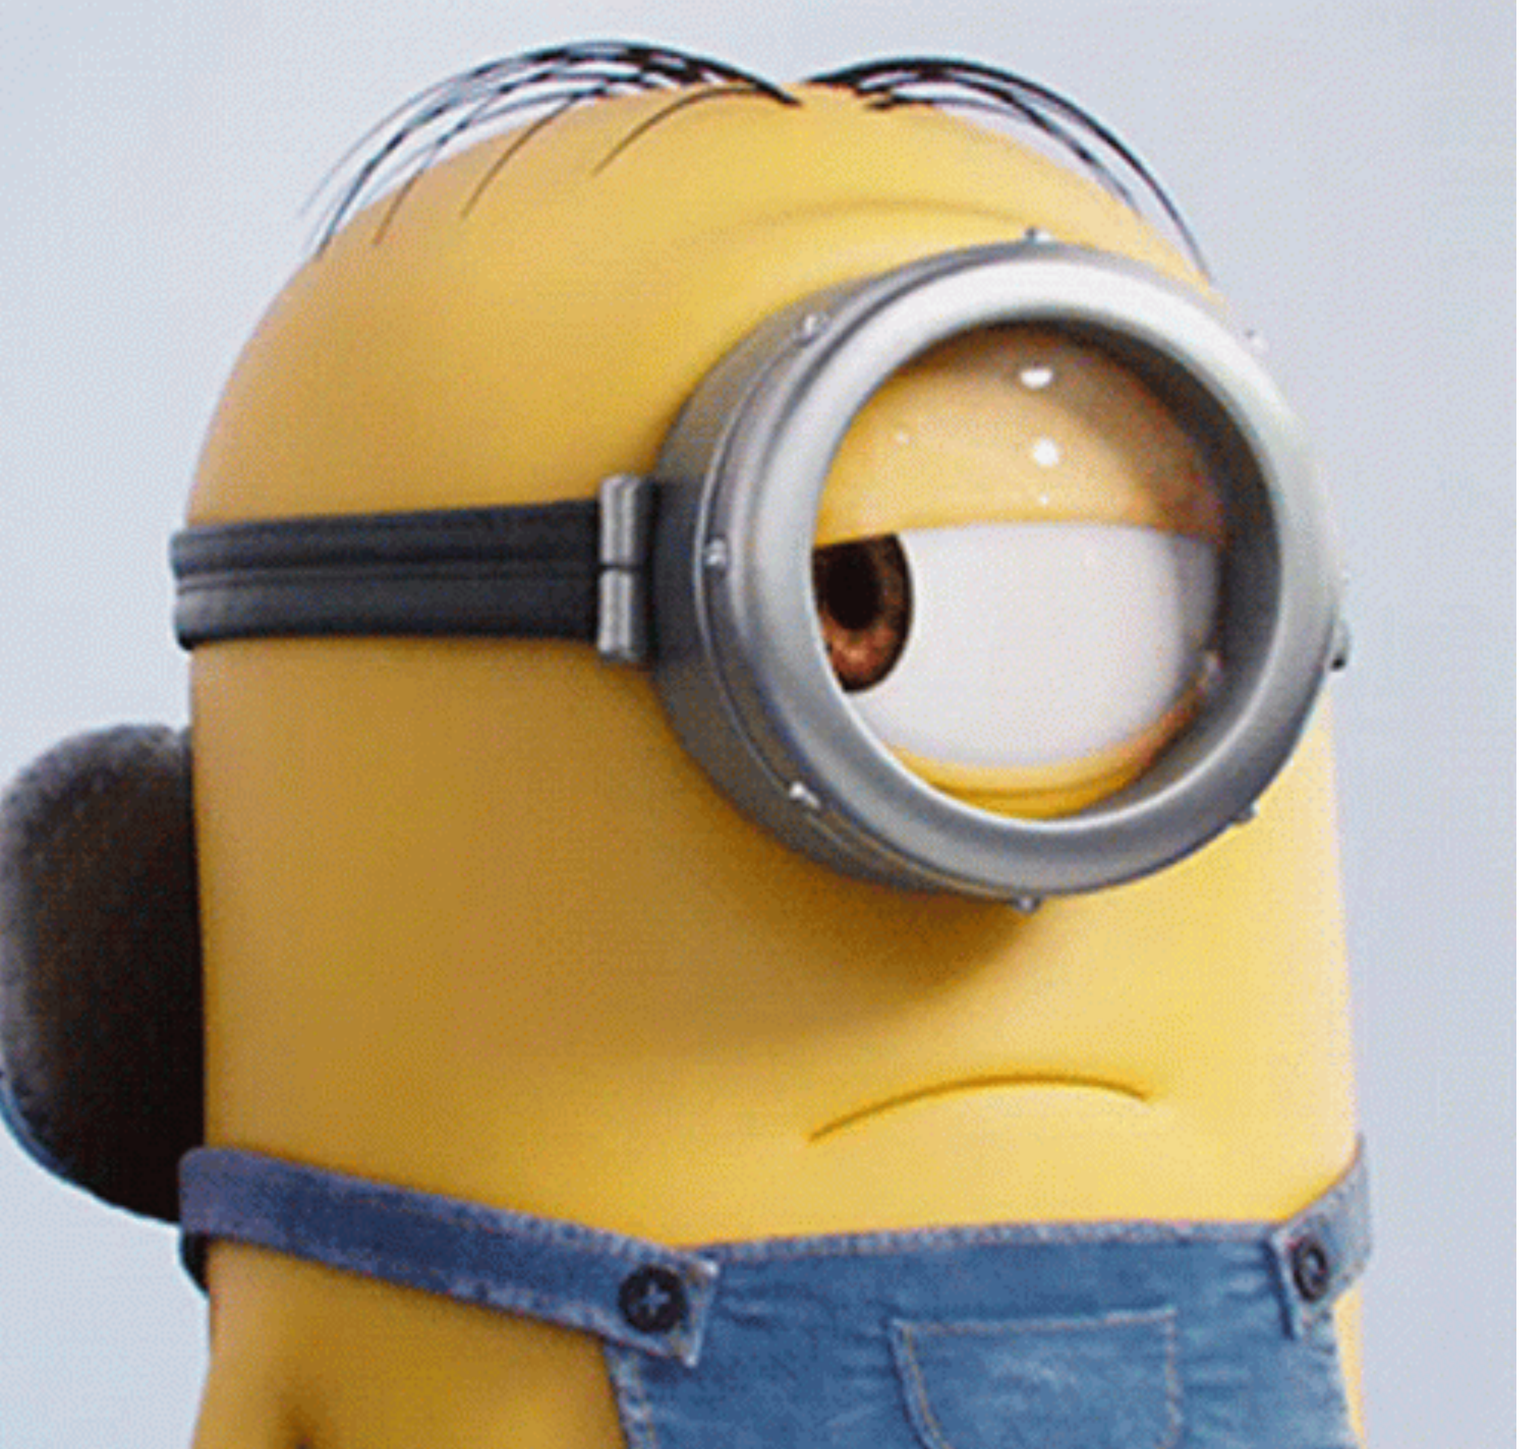

Logistics

Website is the ONE true hub

https://www.cs.cornell.edu/courses/cs4756/2023sp/

5 Assignments [65%]

Final Project [30%]

Participation [5%]

https://www.cs.cornell.edu/courses/cs4756/2023sp/

Announcement!

Assignment 0

https://github.com/portal-cornell/cs4756_robot_learning/tree/main/assignments/HW0

Link in website!

Checks familiarity with PyTorch!

Checks probability fundamentals!

Due Tuesday 1/31!

Graduate Version (CS5756)

If you are enrolled in CS 5756, every assignment has an **extra question** that you must solve.

This can be either an extra experiment or extra theory question.

Undergraduates (CS 4756) do not have to solve this question. But there is extra credit if you do!

Books and other resources

Work-in-progress book

Modern Adaptive Control and Reinforcement Learning,
James A. Bagnell, Byron Boots, and Sanjiban Choudhury

(Please feel free to send me feedback)

For other resources, keep checking website

Course Policies

All policies are posted on the Website!

Course Website: 3 TOTAL late days. Any assignment turned in late will incur a reduction in score by 33% for each late day

Academic Integrity: Any work presented as your own must be your own, with no exceptions tolerated. Submitting work created by ChatGPT, or copied from a bot or a website, as your own work violates academic integrity.

ChatGPT!

Use of text generation/editing systems such as ChatGPT: For each component of the workload, the vast majority of the intellectual work must be originated by you, not by text generation systems. It is OK to use aids for writing fluency --- but note that writing fluency is not part of the assessment rubrics below anyway.

i). Example of something that is allowed: You write the initial draft(s), review its contents. You then use some form of text generation system to proofread and improve the flow. You do not use the system's output to add extra content.

ii). Example of something that is definitely not allowed: You essentially use a text generation system to generate an early draft, even if you later post-edit and correct the output.

iii). Example of something that is OK but requires special treatment. You start with the procedure in i). But, the system output includes good points that you hadn’t thought of before, or makes you realize that a point you had made isn’t quite right. You may include the new material and/or make appropriate edits, but you should mention what specific system(s) you used and what changes you made based on it.

The Crew

Yuki Wang

- 1st Year PhD Student in CS
- A member of the PoRTaL lab
- Interested in high level task-planning for robots
- Fun fact: Because I learned Japanese in college, I invented a new language “JanKoreanglish”, which mixes Japanese, Korean, and English in one sentence to annoy my friends.

Juntao Ren

Interested in using IL/RL to efficiently teach robots how to make safe and optimal decisions in assisting humans.

Currently working on building a library of low-level skills for our robot Hal at PoRTaL :).

Fun fact: I used to play water polo in high school.

- Kushal Kedia, PhD Student
- Research Interests: Motion Planning, Human-Robot Interaction
- Fun Fact about me: I love collecting merchandise for my favourite sports team, Chelsea! Let’s chat about soccer :)

Abhishek Masand

MEng Computer Science’23

Interests: Large Language Models, Reinforcement Learning, Representation Learning, NLP

Current Research: Goal directed influence and reasoning for large language models. Also interested in exploring various strategies to control the influence of large language models on downstream tasks, and improving unsupervised language quantification reward functions.

Previous Work: Recently worked as an AI Engineer at Microsoft, where he collaborated with the Office of the CTO and Microsoft Research NLP group to develop libraries to facilitate the use of LLMs and applications demonstrating OpenAI models.

Fun Fact: I started out as a consultant and over time became a researcher

tl;dr

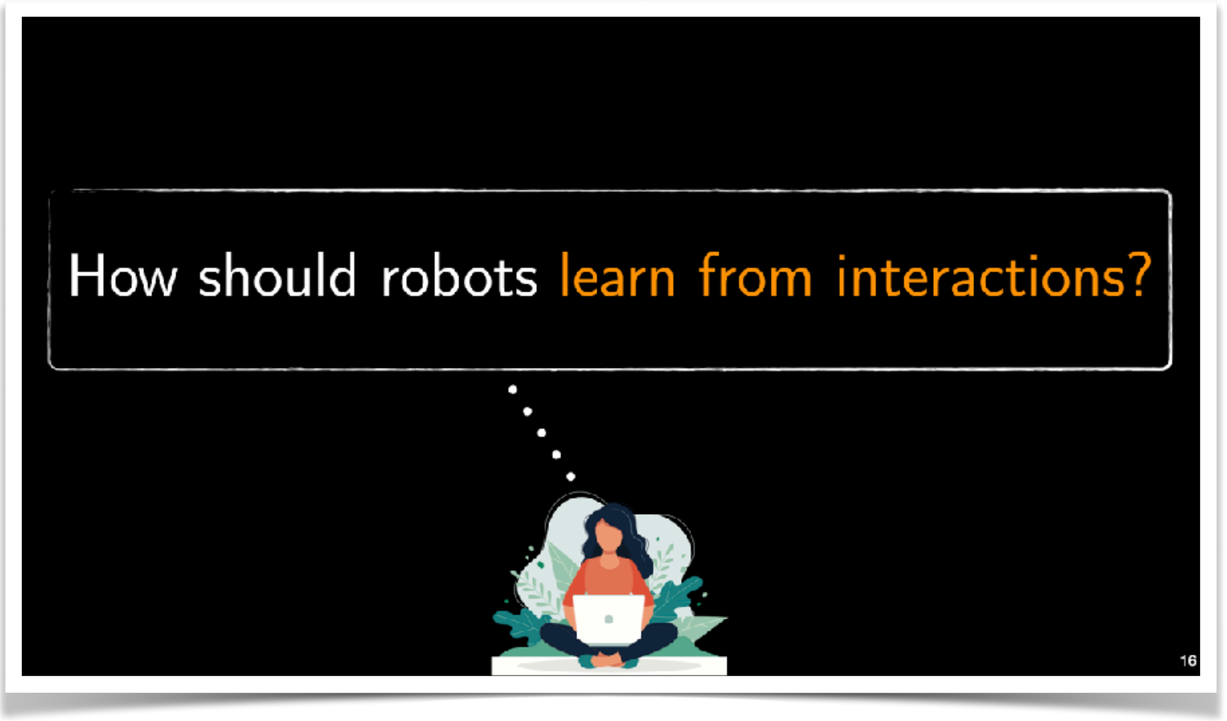
How should robots learn from interactions?
16

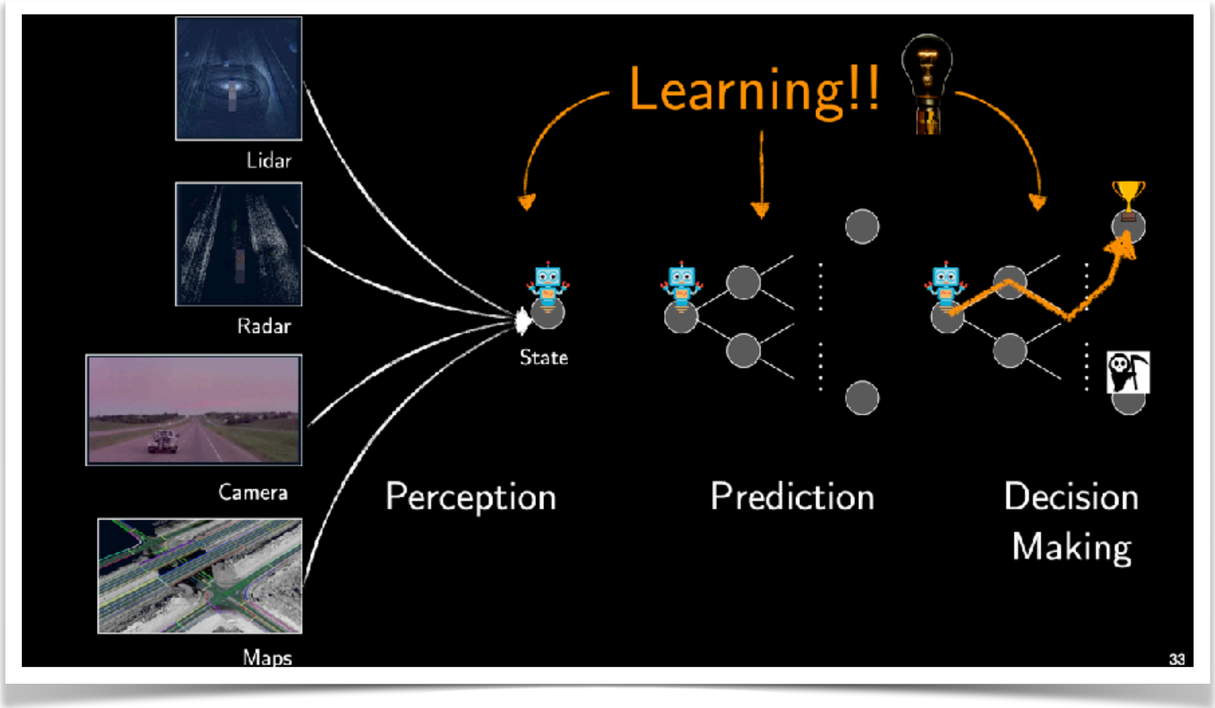
Learning!!
Lidar
Radar
Camera
Maps
State
Perception
Prediction
Decision Making
33

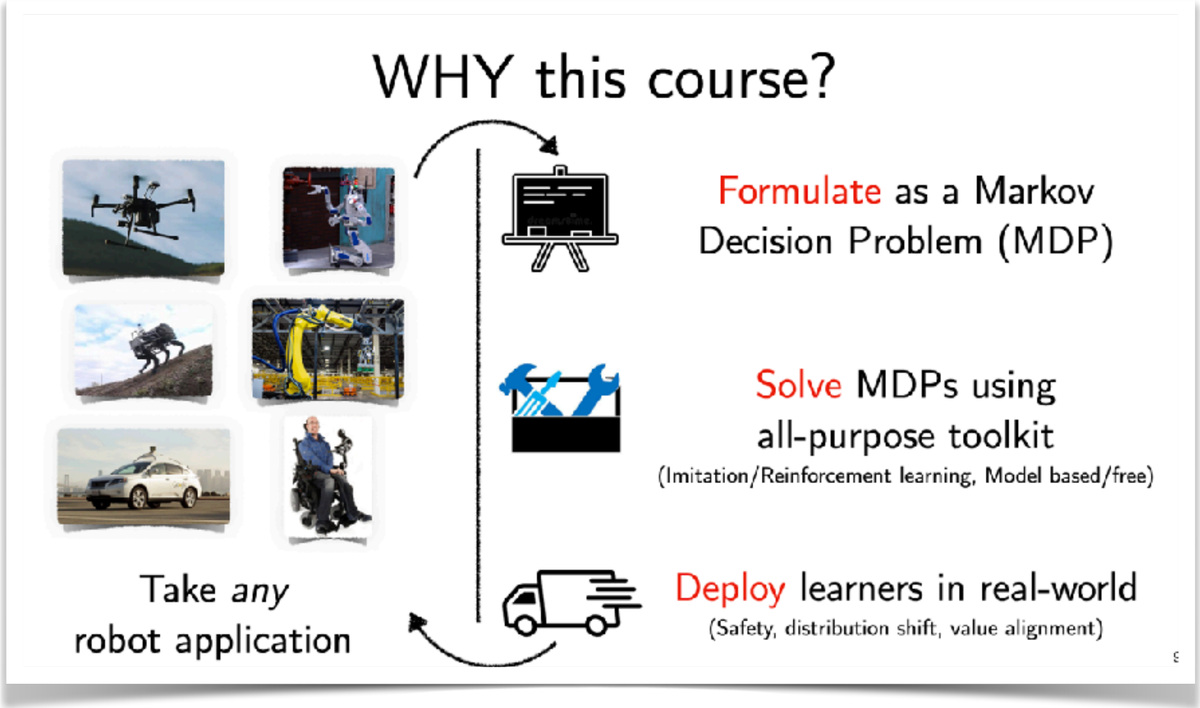
WHY this course?
Formulate as a Markov Decision Problem (MDP)
Solve MDPs using all-purpose toolkit
(Imitation/Reinforcement learning, Model based/free)
Deploy learners in real-world
(Safety, distribution shift, value alignment)
Take any robot application