CS 4756/5756: Robot Learning

Sanjiban Choudhury

PORTAL

Cornell Bowers CIS Computer Science


WHAT A TIME TO BE ALIVE!

2024 continues to be an exciting year for Machine Learning

Different things we can do with GPT-4/LLMs/Transformers?

Three advances in the last semester that I am excited about:

#1: Real-time, multi-modal interactions

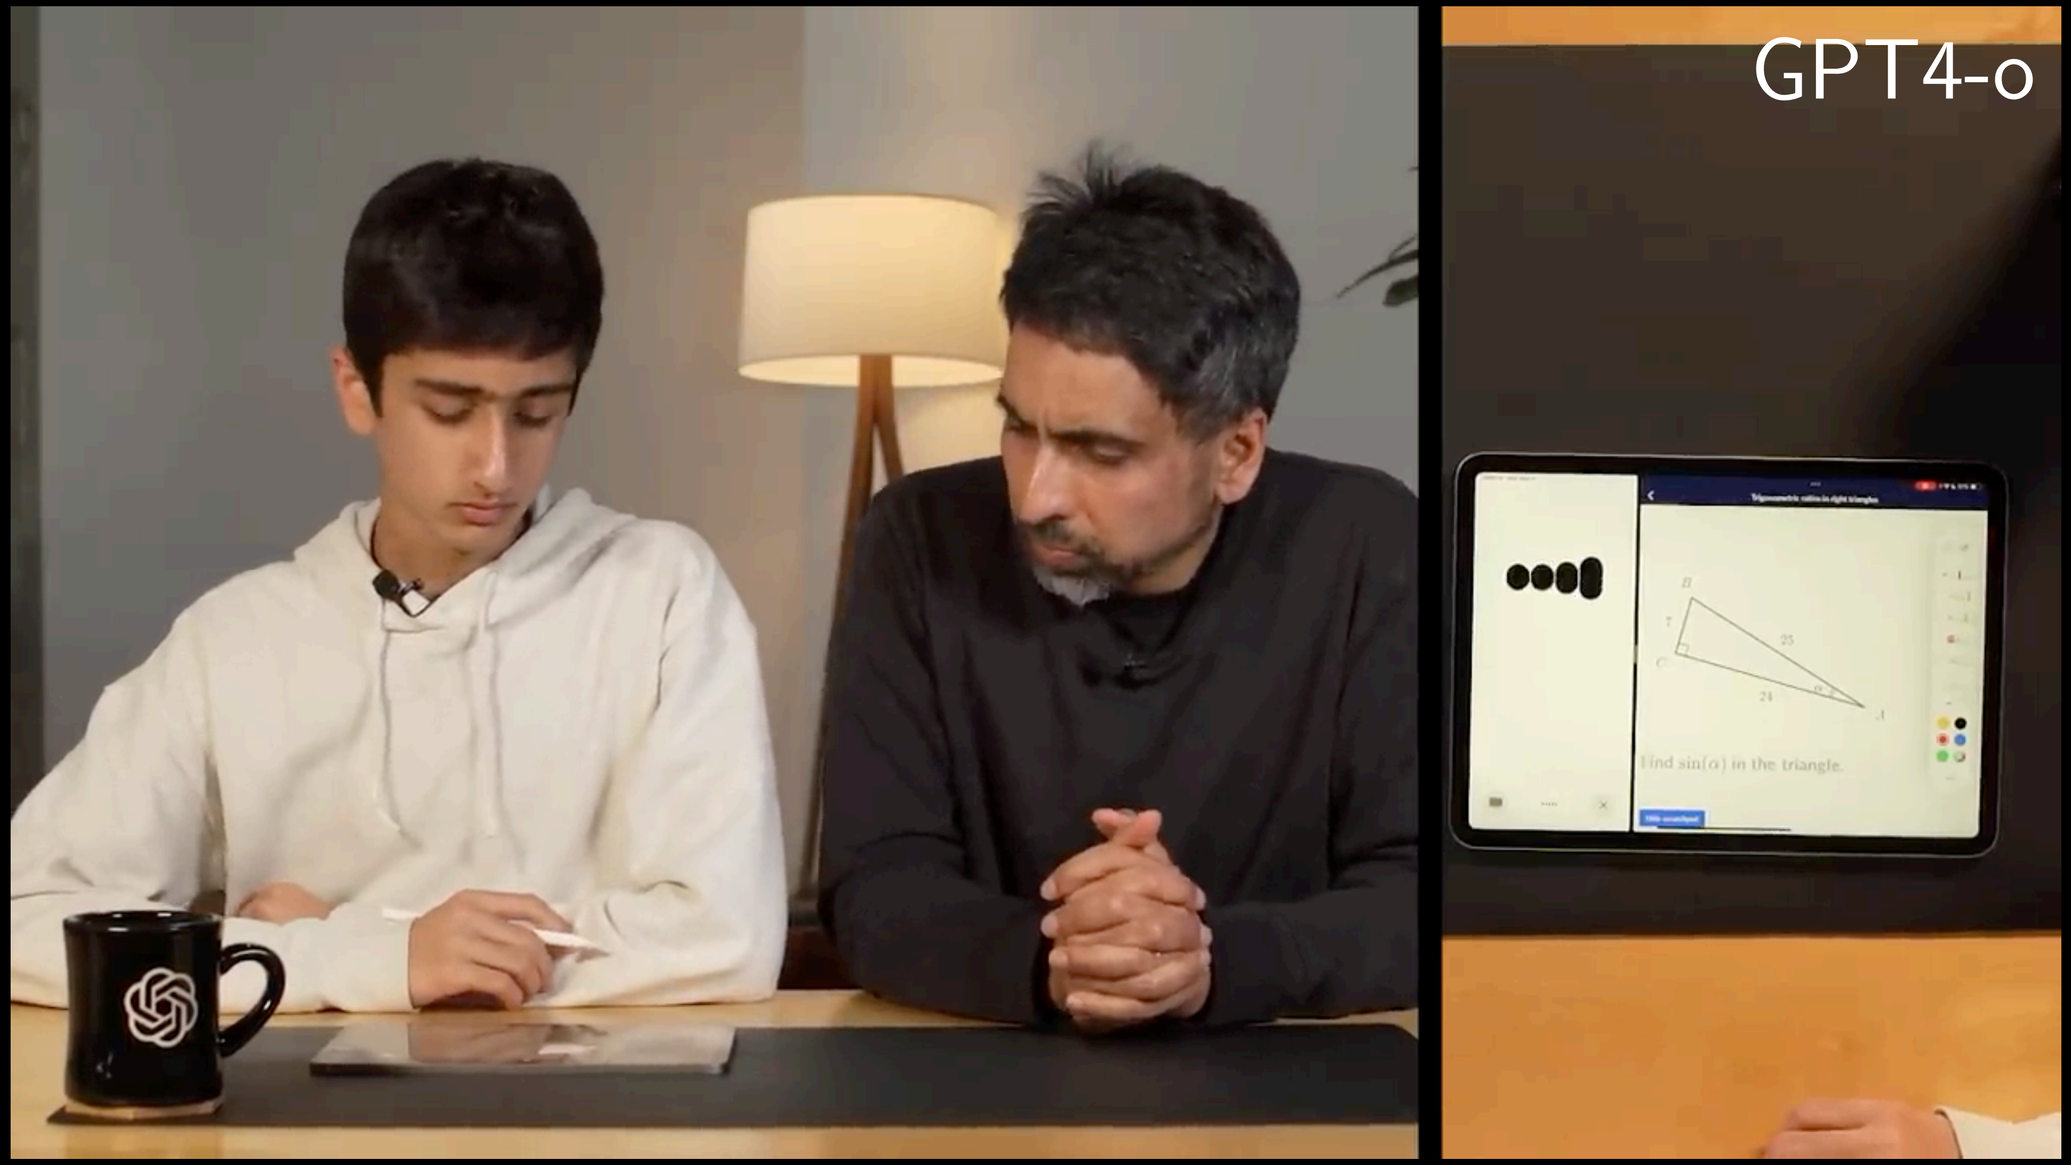

#2: Video Generation

SORA

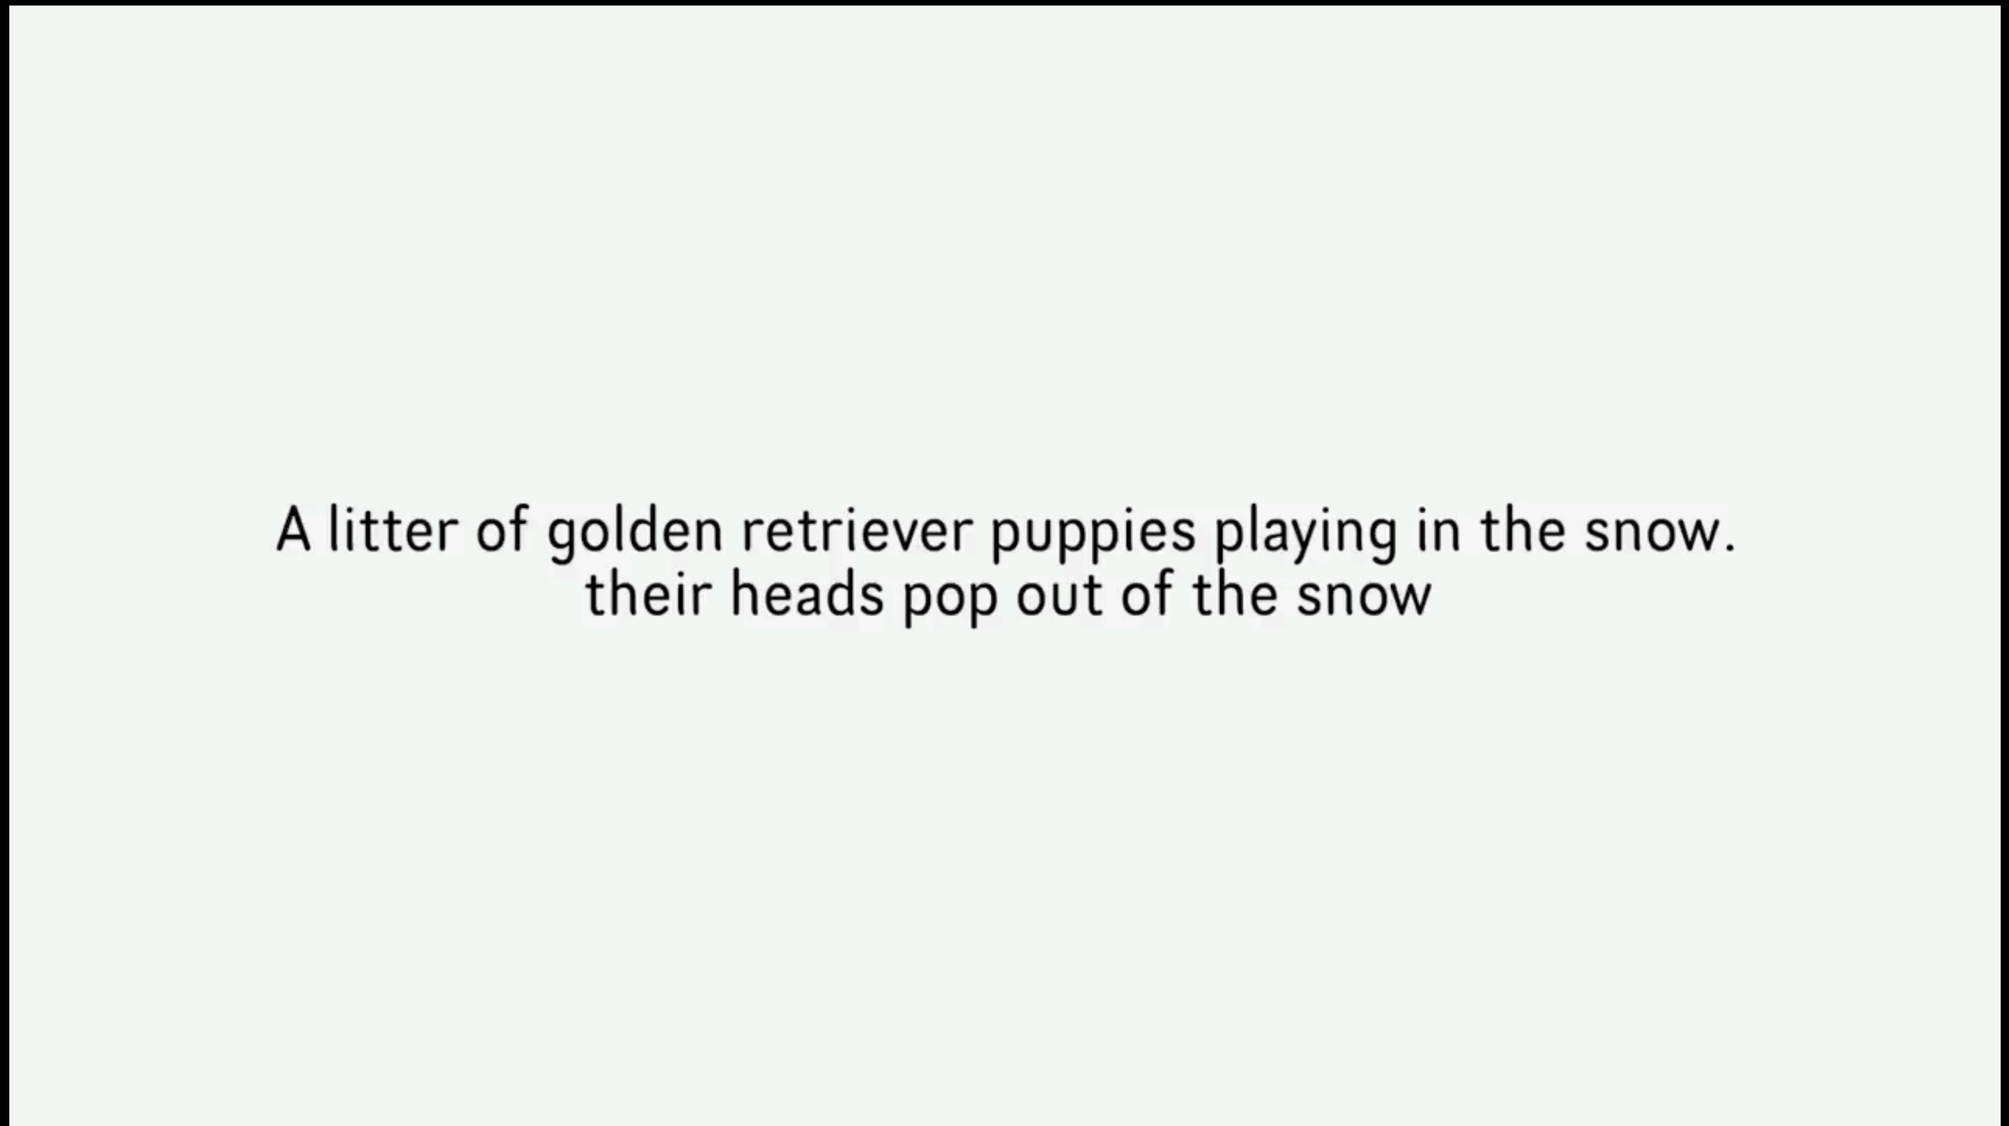

#3: LLMs that act

SWE-Bench

What algorithms are powering these advancements?

Reinforcement Learning from Human Feedback (RLHF)

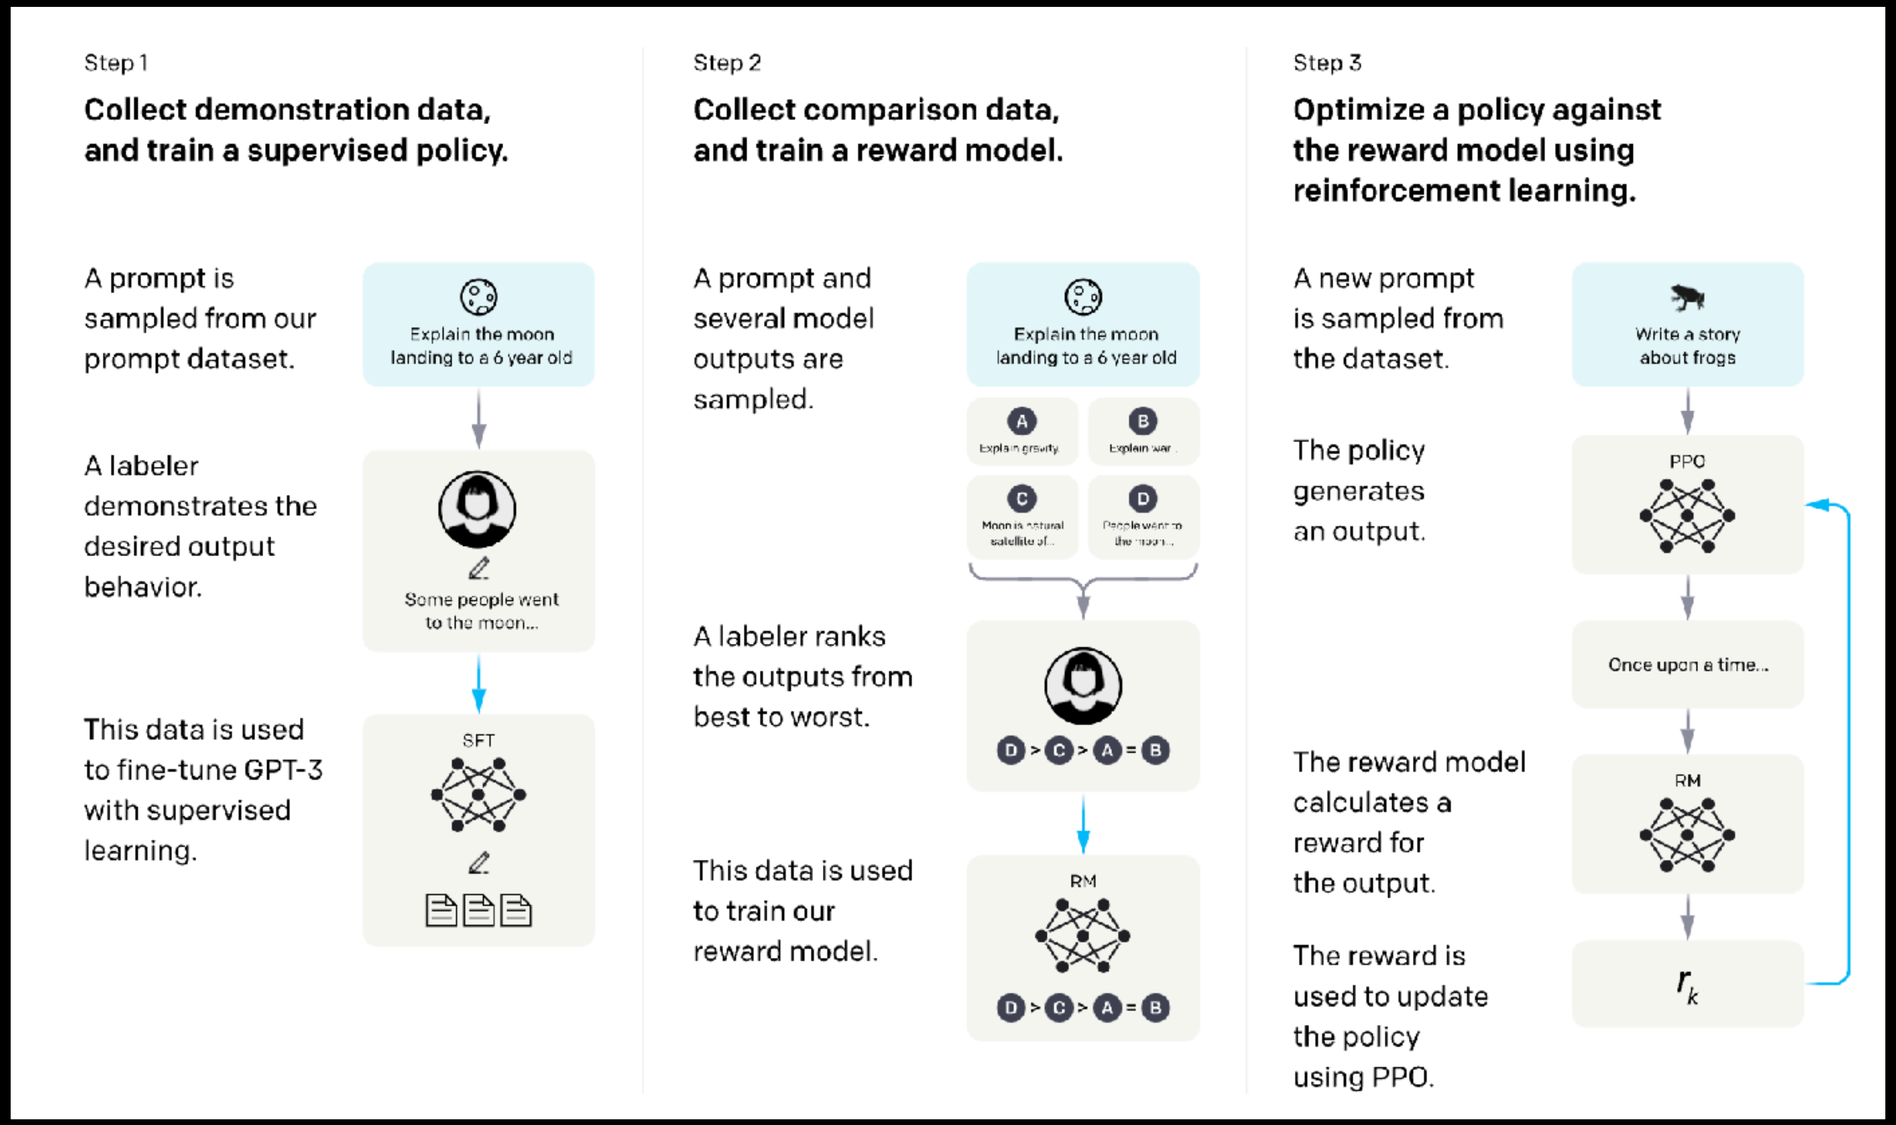

Open-source fine-tunable models

LLAMA

Alpaca

Vicuna

Mistral

Parameter Efficient Fine Tuning (PEFT)

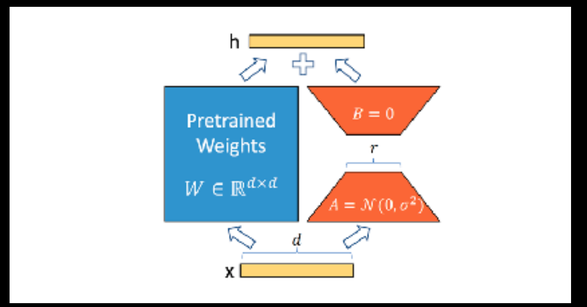

LORA

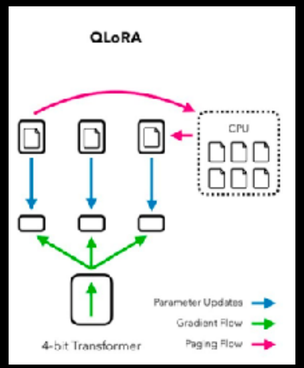

QLORA

Where are the robots?

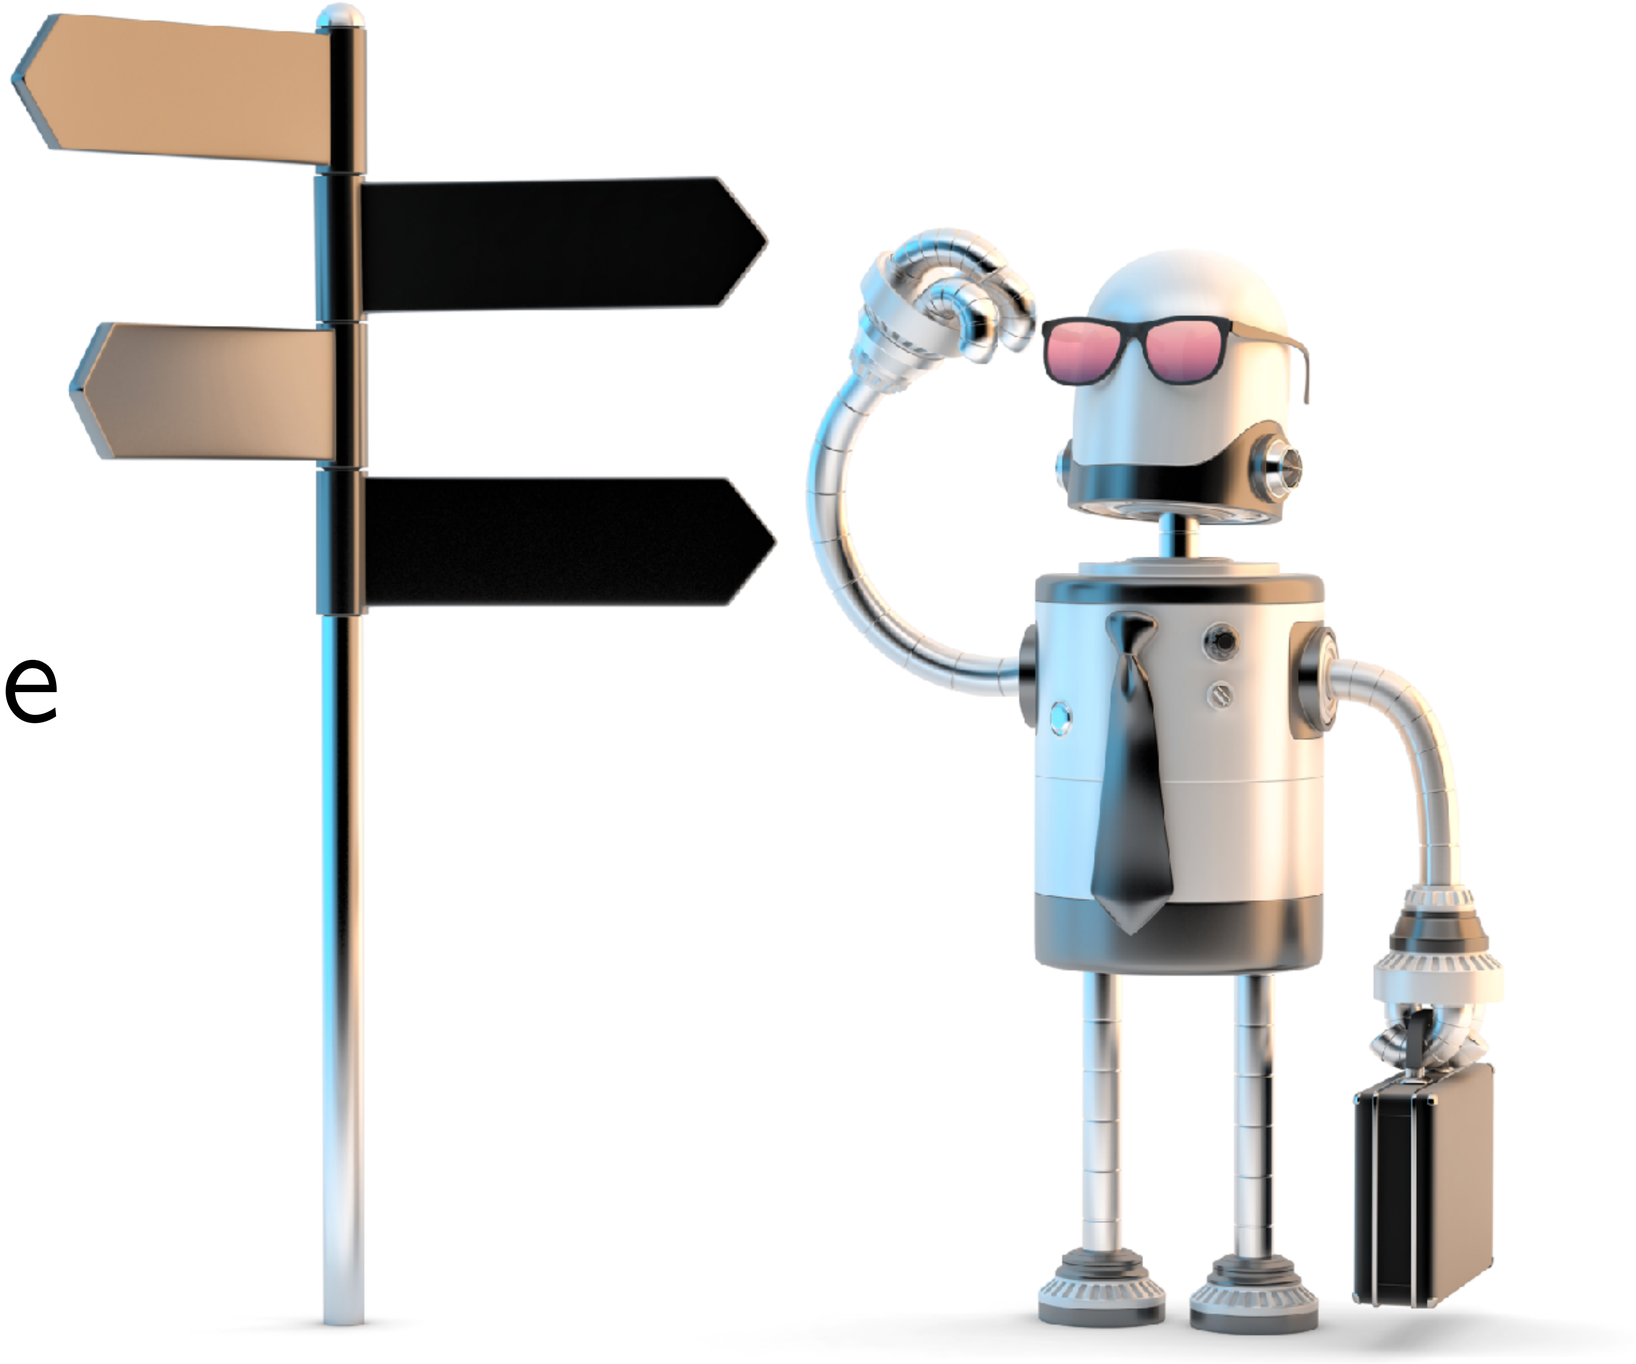

Rise of the Humanoids

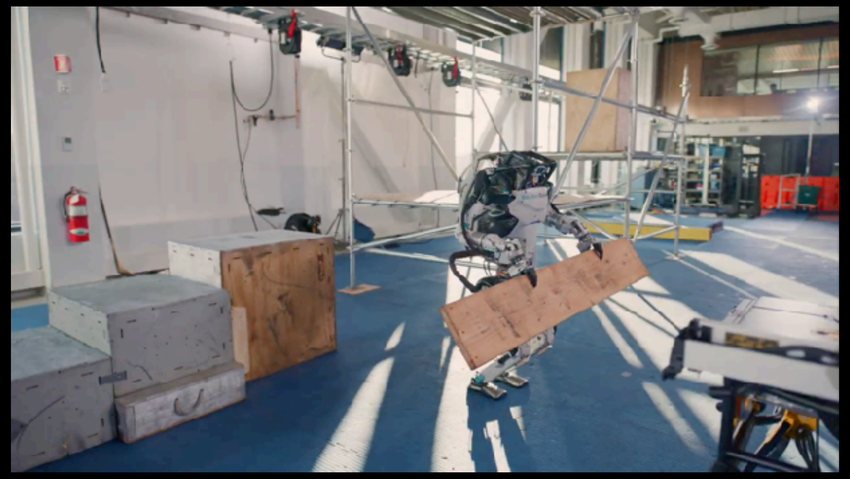

Boston Dynamics

Tesla

Agility Robotics

Figure AI

Quadrupeds going strong

Boston Dynamics SPOT

Unitree GO 2

ANYmal

Self-driving continues driverless runs

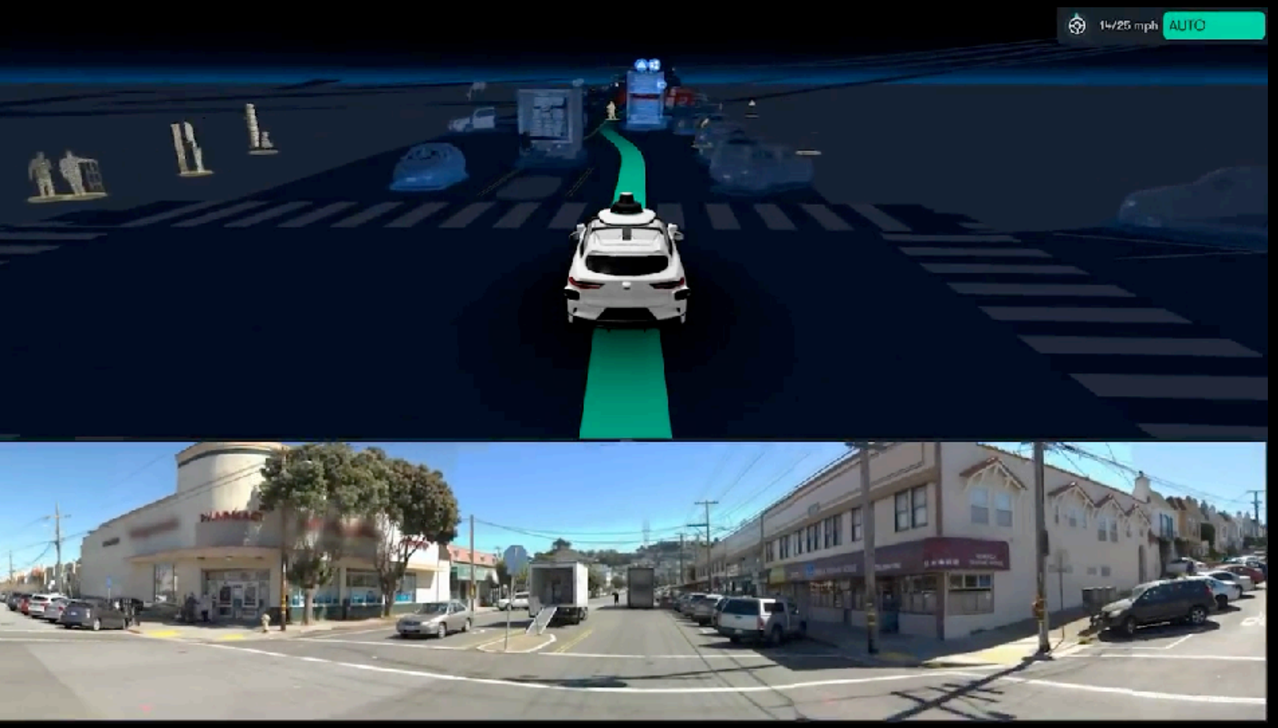

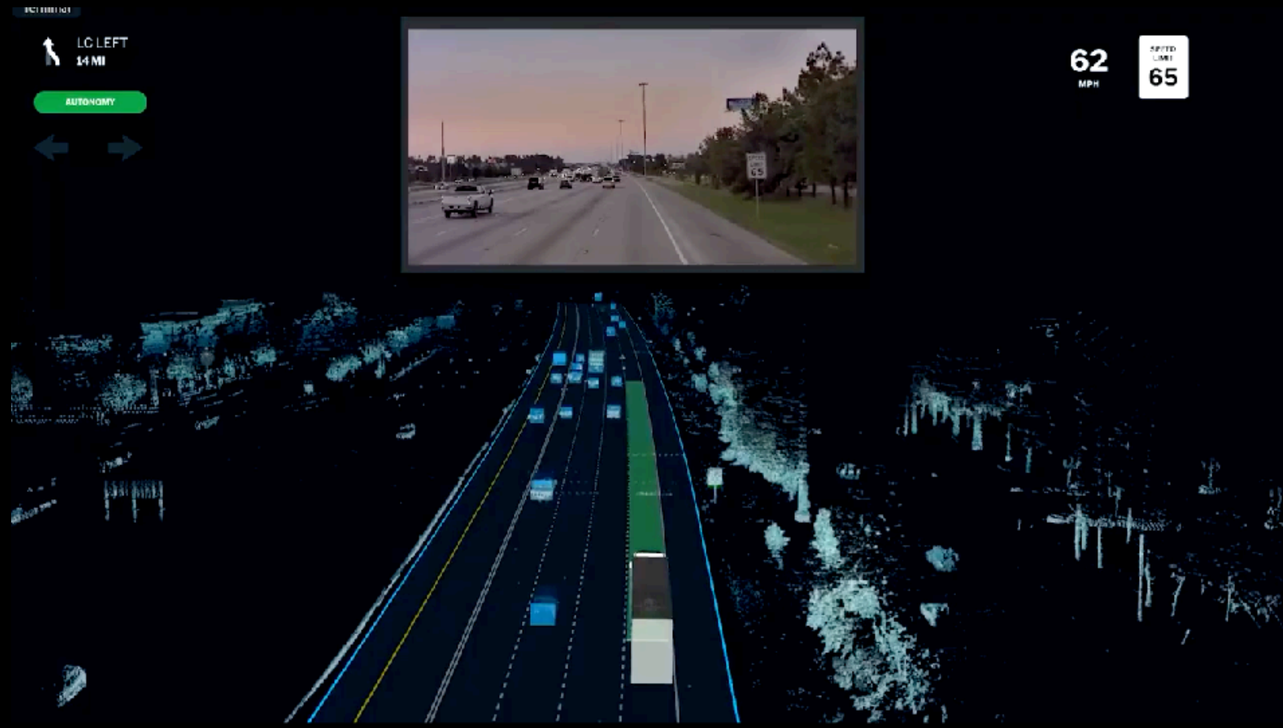

WAYMO

Aurora

But ...

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

Why?

What is so challenging about our homes?

Why can't we throw ML at this problem?

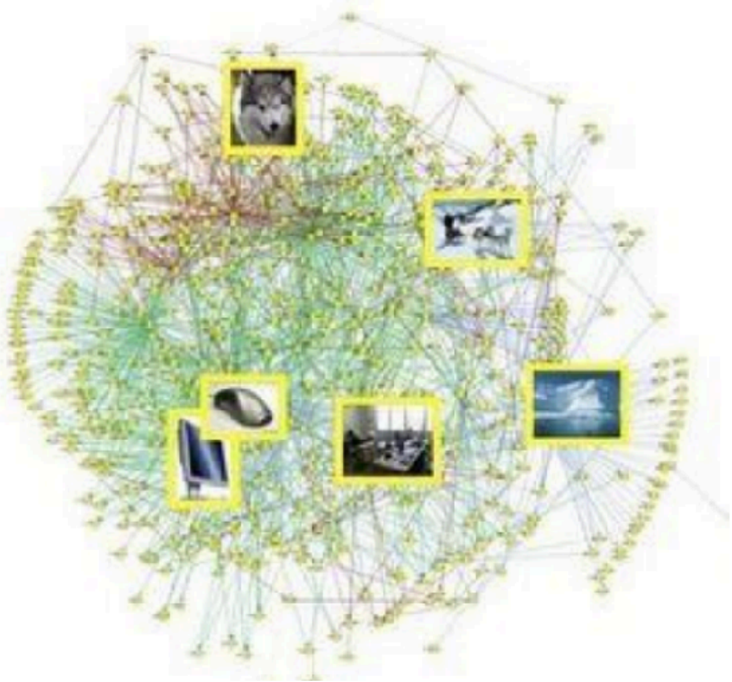

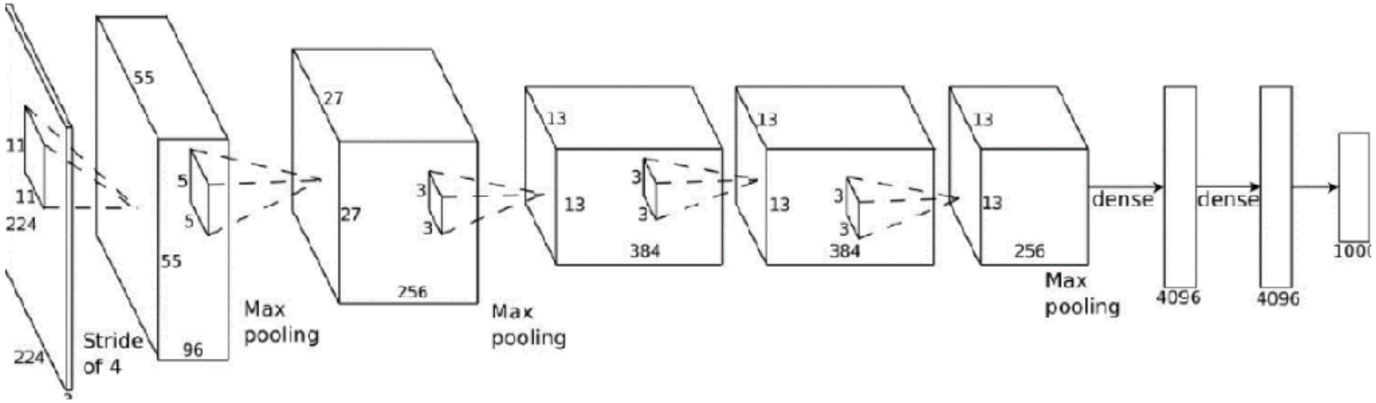

Big Data

Big Models

Credit: **Sergey Levine** "Offline RL lecture"

Big Data

1M trajectories, 22 robots, 21 different institutions

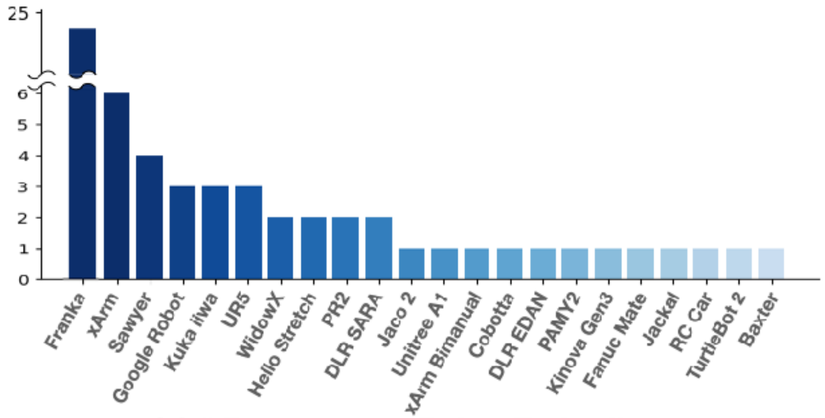

(a) # Datasets per Robot Embodiment

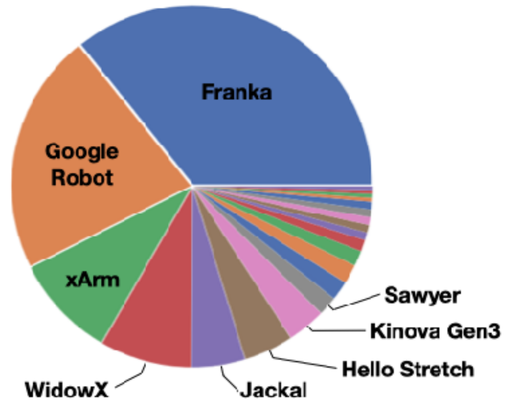

(b) # Scenes per Embodiment

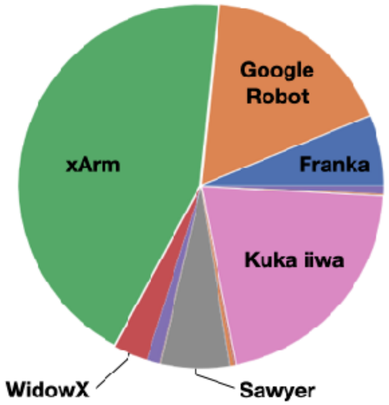

(c) # Trajectories per Embodiment

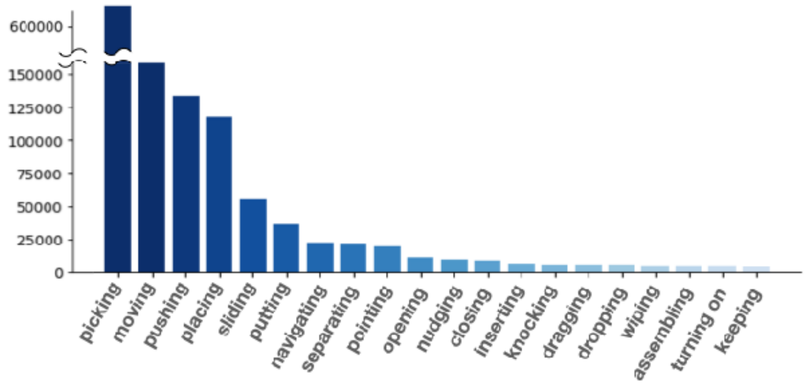

(d) Common Dataset Skills

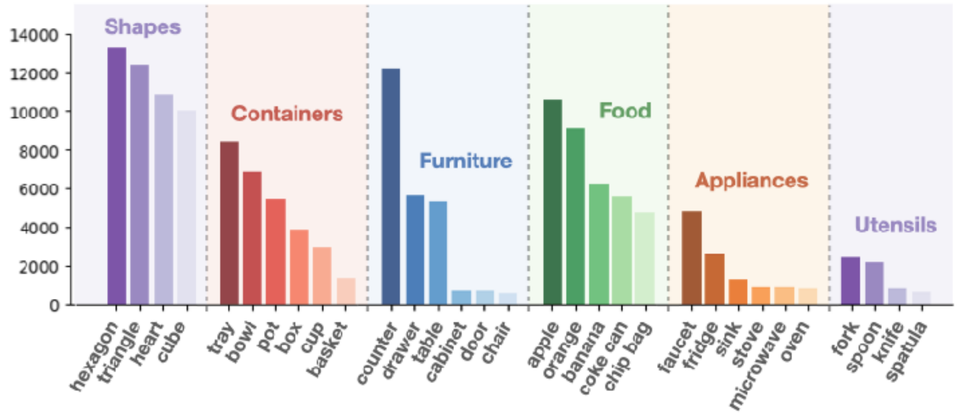

(e) Common Dataset Objects

Open-X Embodiment Dataset

Big Models

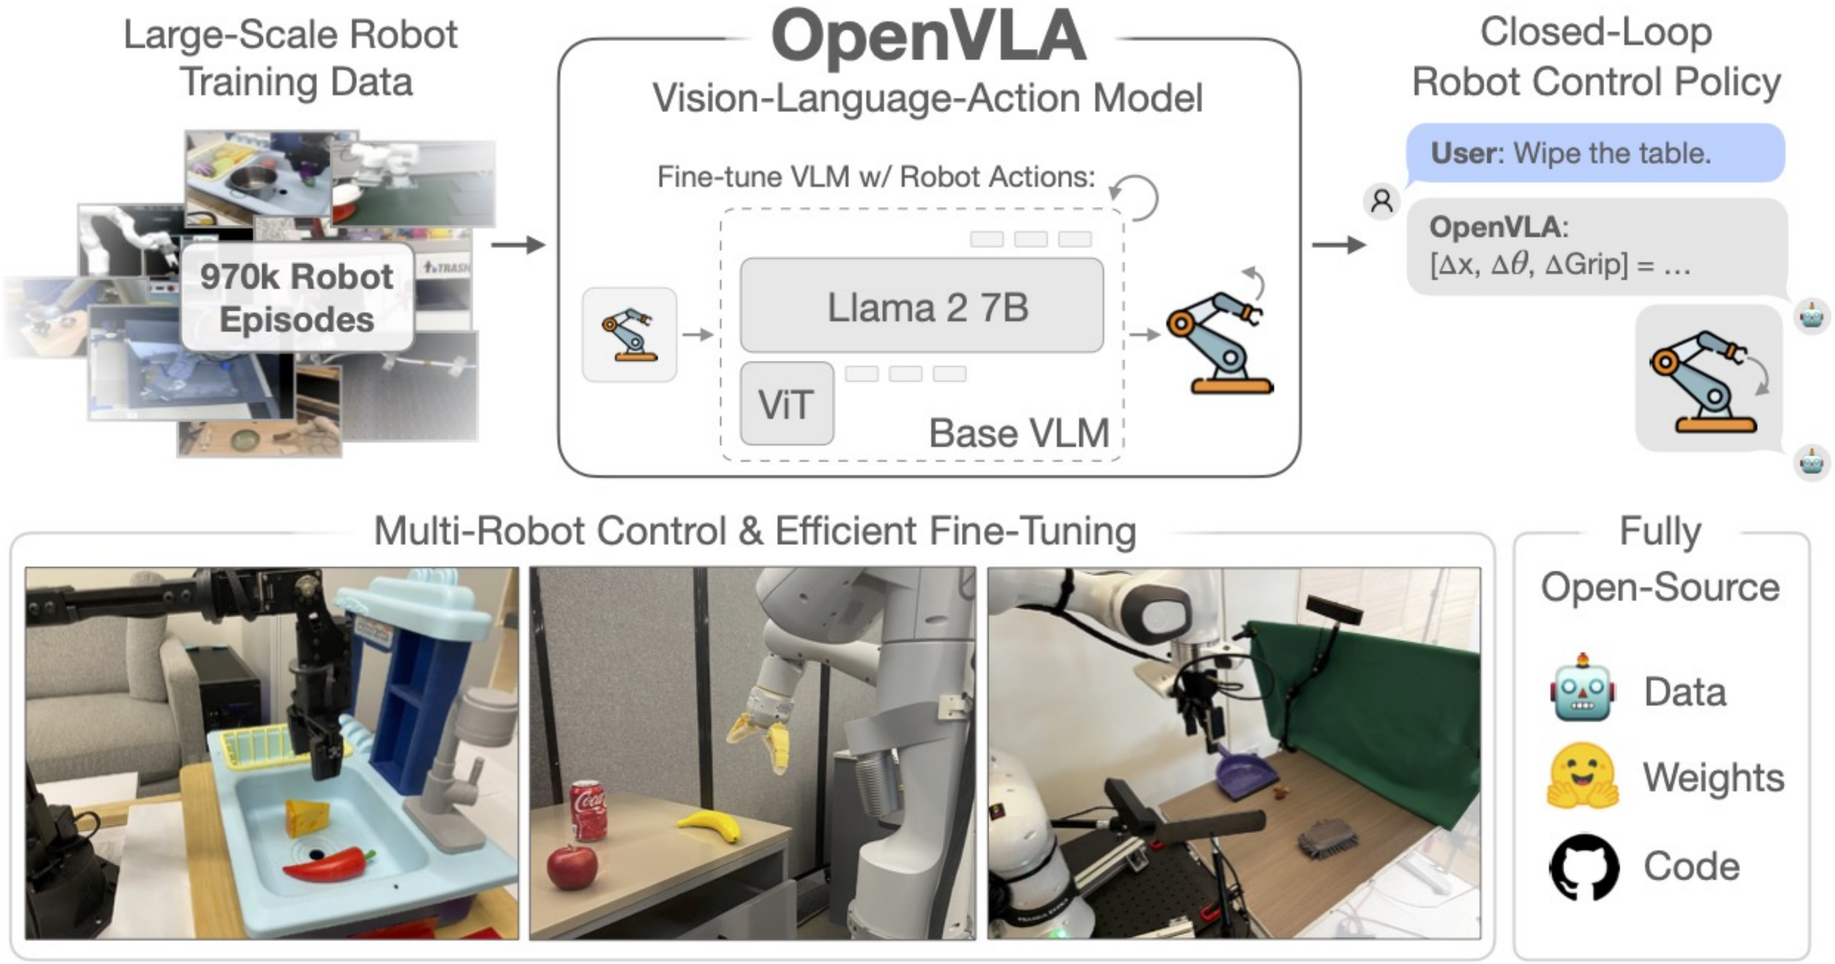

https://openvla.github.io/

Activity!

Train home robot apprentice to help grandma!

Demonstrations,
Language

Interactive
feedback

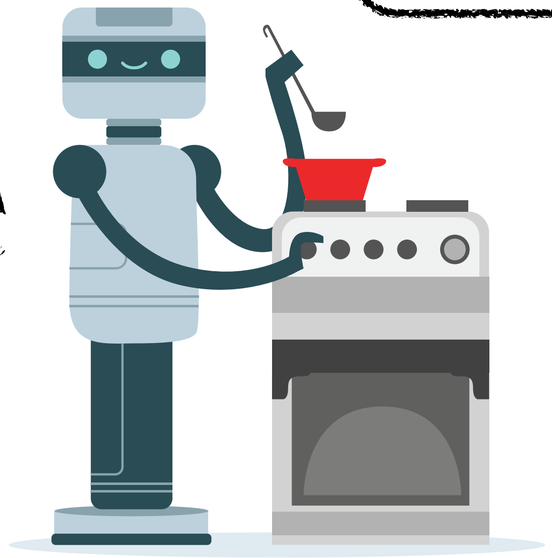

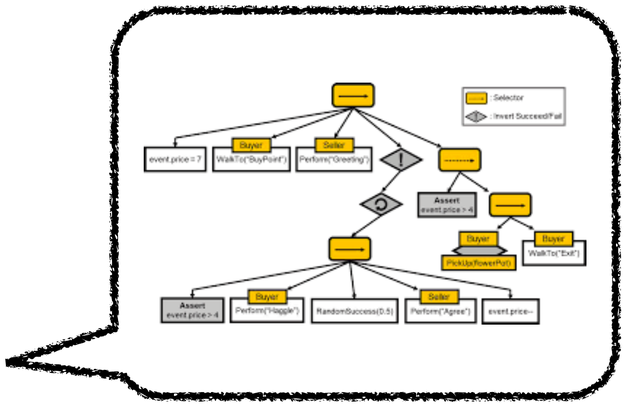

Home Robots in our lab!

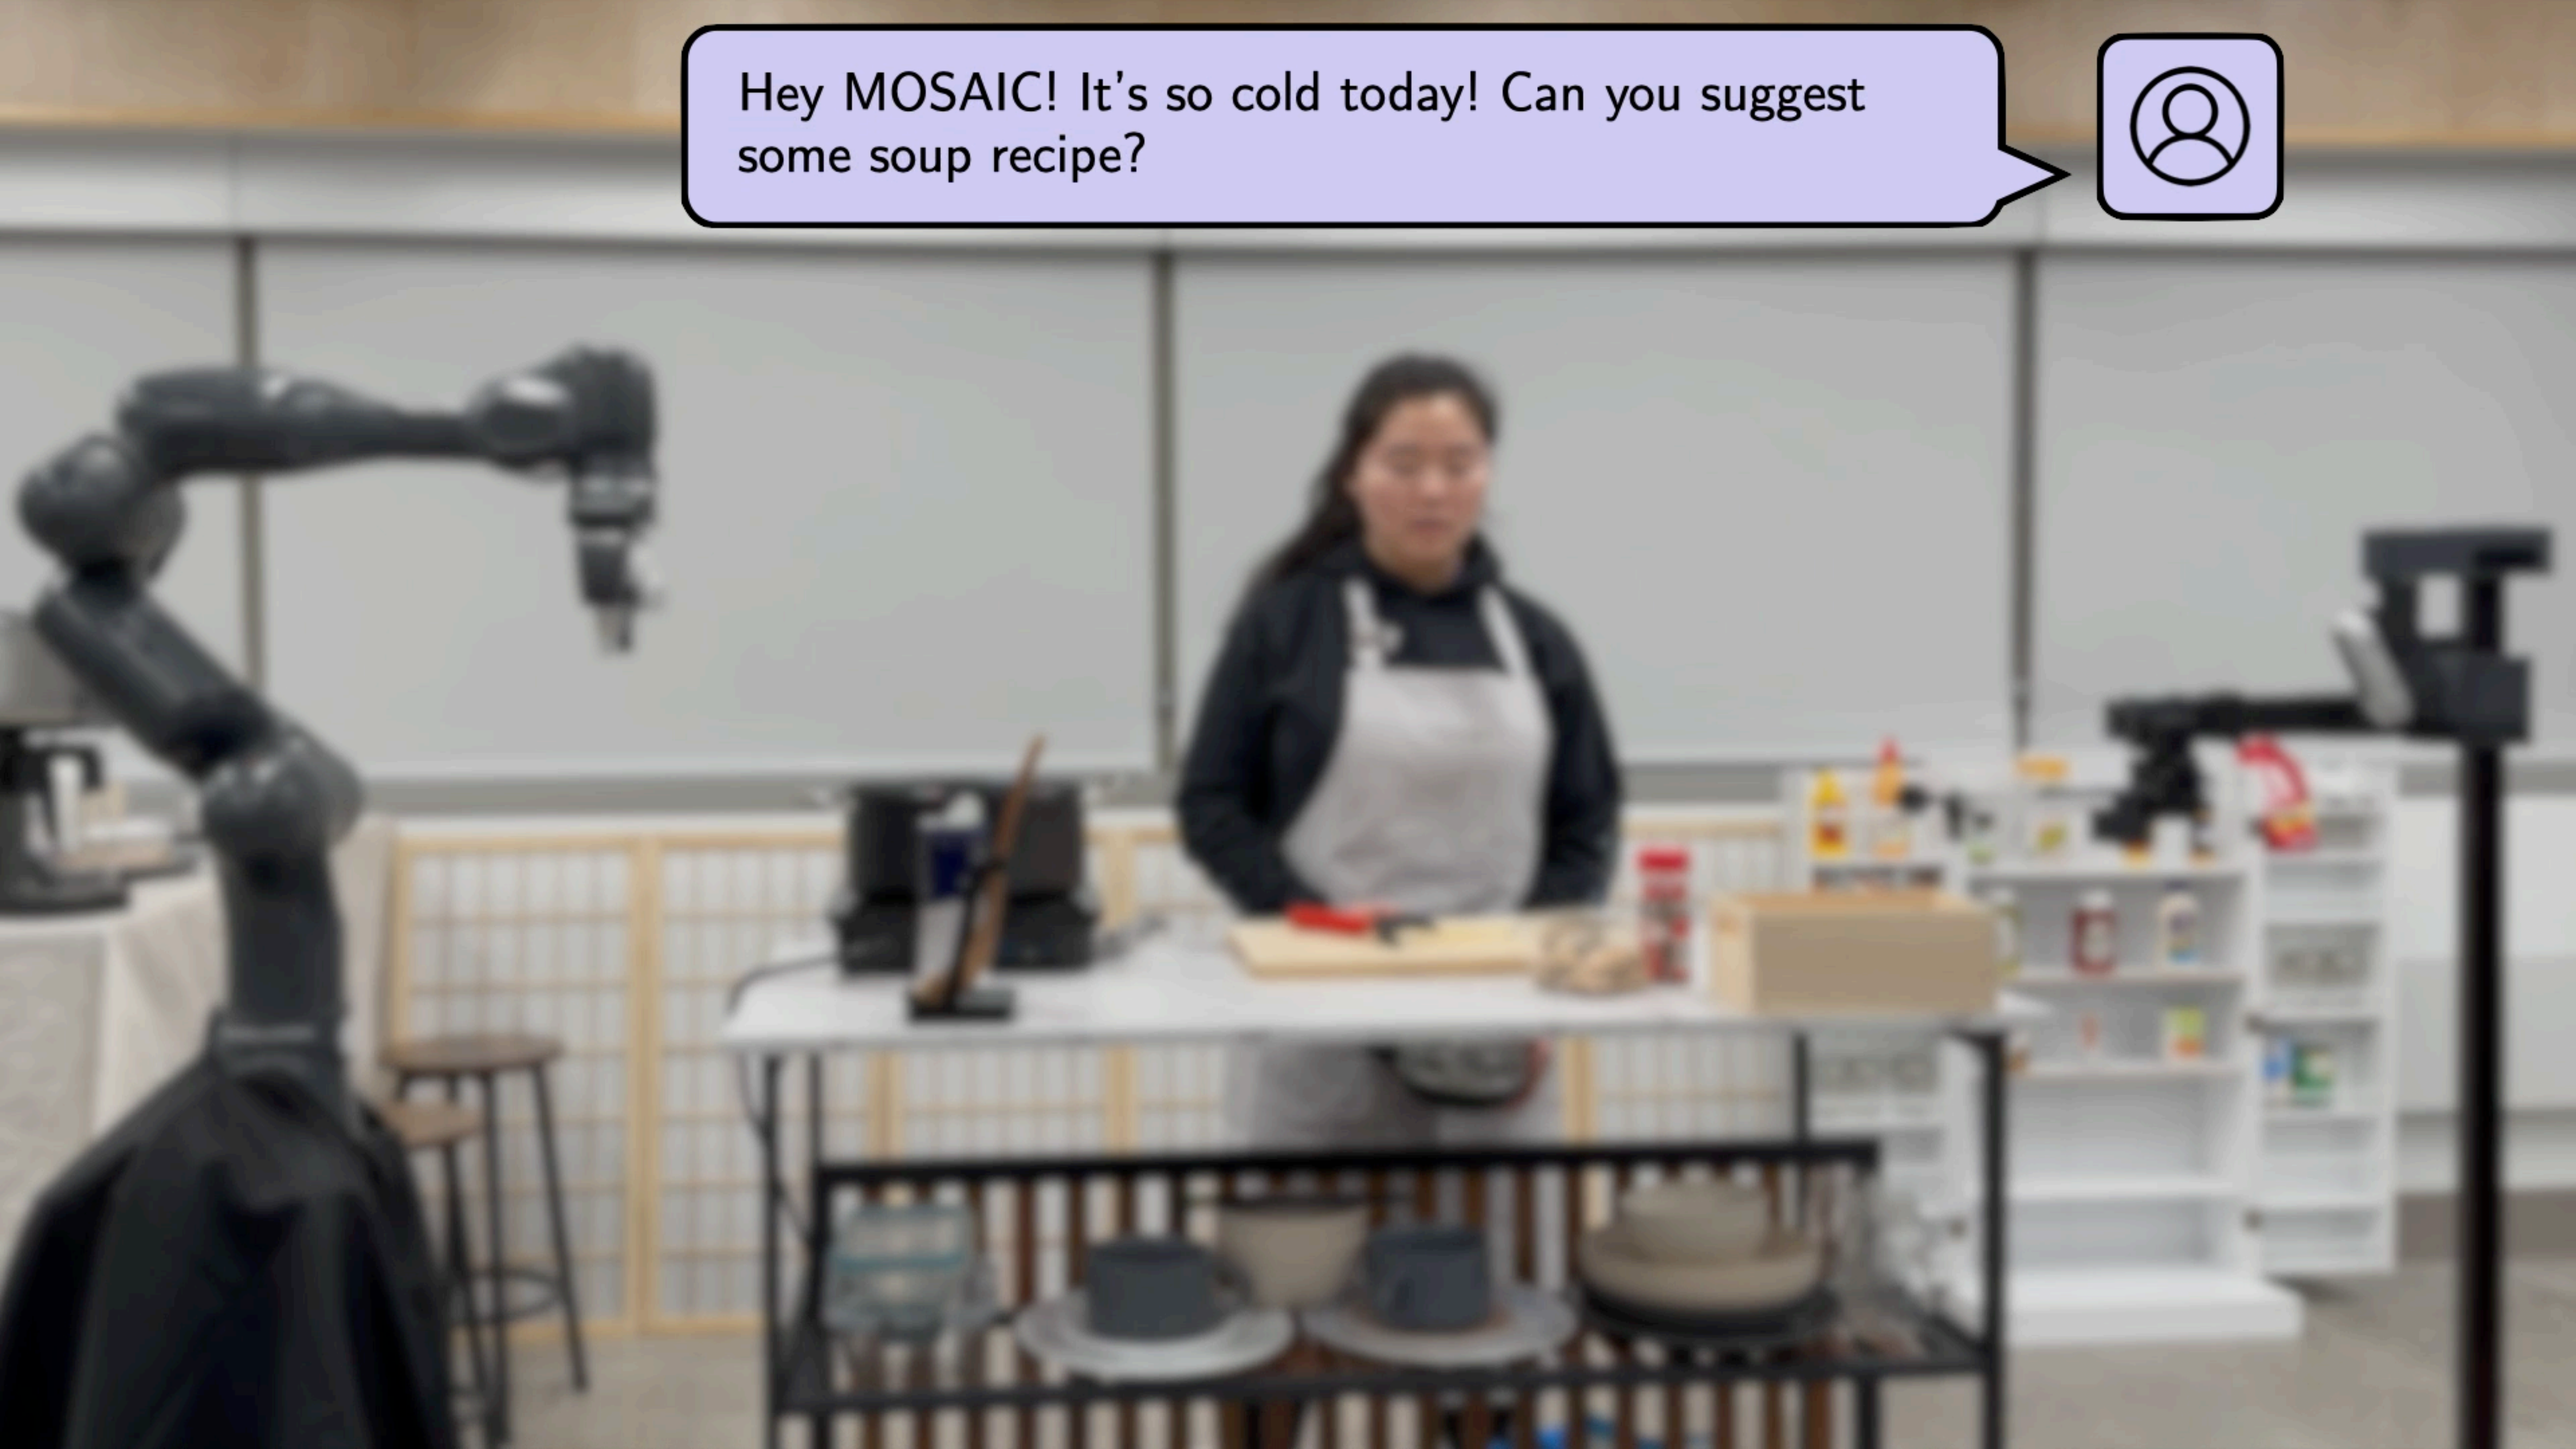
Hey MOSAIC! It's so cold today! Can you suggest some soup recipe?

Question

What is main challenge in
apply machine learning to home robots?

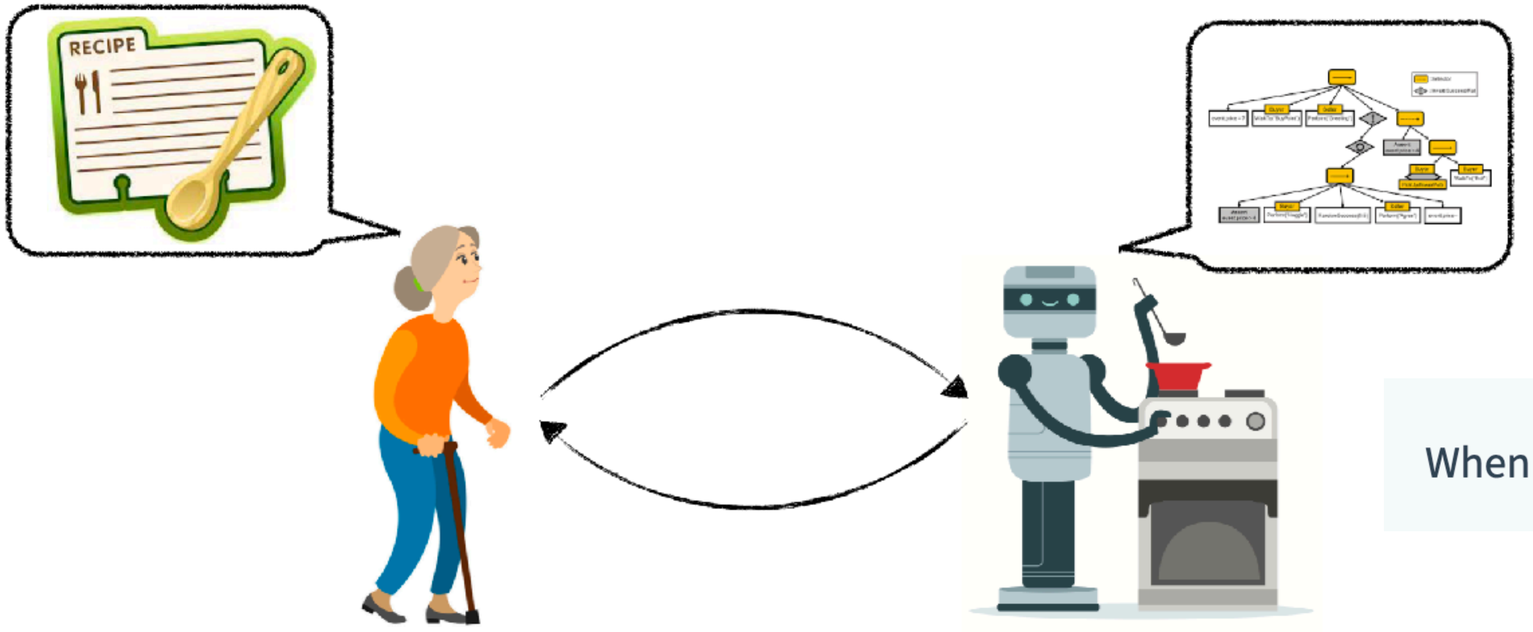

When poll is active respond at **PollEv.com/sc2582**

What is special about robot learning?

$$\min_{\theta} \; \mathbb{E}_{x,y} \, \ell(y, \theta(x))$$

x is a sequence of inputs, y is a sequence of outputs, θ is a model

What is special about robot learning?

$$\min_{\theta} \mathbb{E}_{x,y} \ell(y, \theta(x))$$

x is a sequence of inputs, y is a sequence of outputs, θ is a model

Transformers are pretty standard choice for the model

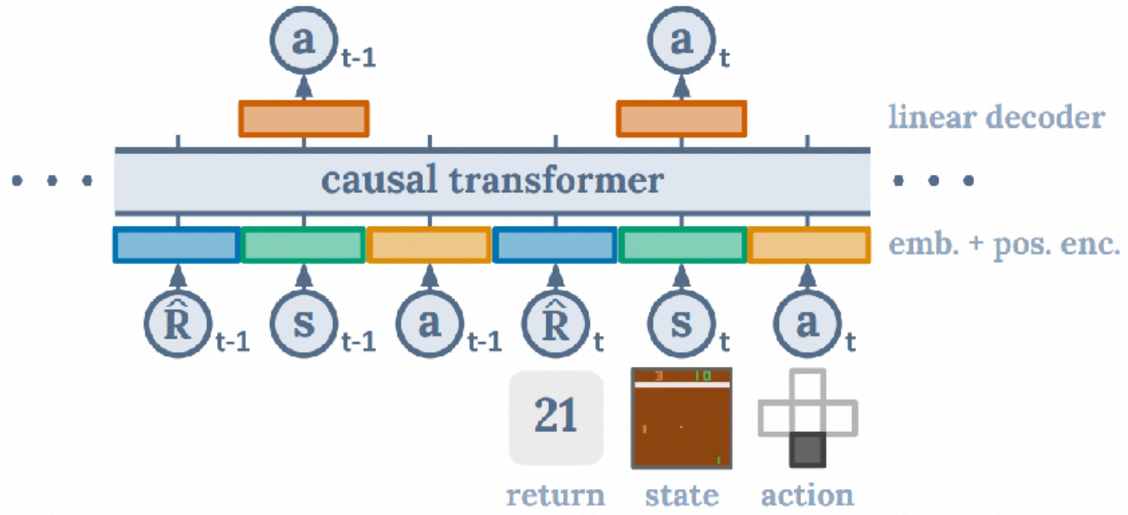

What is special about robot learning?

$$\min_{\theta} \mathbb{E}_{x,y} \ell(y, \theta(x))$$

x is a sequence of inputs, y is a sequence of outputs, θ is a model

Problem 1: How do we gather the right data?

What is special about robot learning?

$$\min_{\theta} \; \mathbb{E}_{x,y} \, \ell(y, \theta(x))$$

x is a sequence of inputs, y is a sequence of outputs, θ is a model

Problem 2: How do we choose the right loss?

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

Office hours:
Tues 11:30 – 1:30pm
Gates 413B

Build robots that can learn from humans!

Undergrad

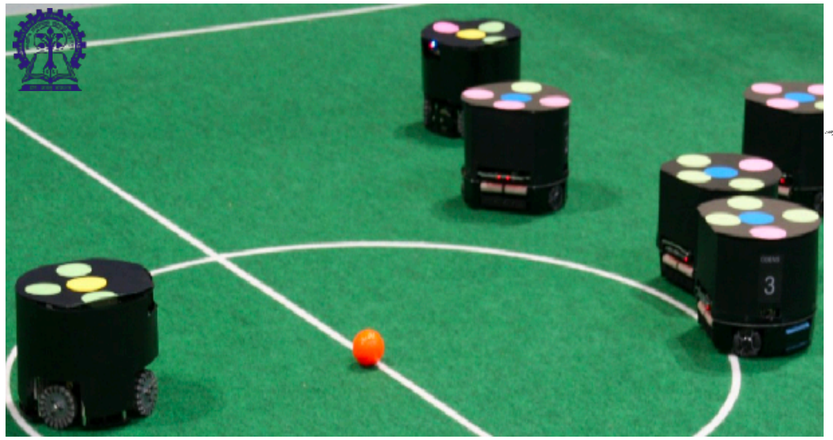

PhD

PostDoc

Research Engineer

We are PoRTaL

(People and Robots, Teaching and Learning)

PORTAL

https://portal.cs.cornell.edu/

Belonging

Some news!

We are expecting a baby this fall (October 22nd!)

We are super excited (and nervous!)

Some news!

We are expecting a baby this fall (October 22nd!)

Since I don't get an official parental leave, I will need YOUR help

I will teach actively up until the due date.
After, I will have my brilliant colleagues come in for guest lectures.
These are amazing researchers from different ML fields.

I need YOUR help to make these classes engaging and give a good impression.

Let’s get started!

Self-driving

A brief history of self-driving

One of the first self-driving car drove from Pittsburgh to Sandiego with 2800 miles of autonomy. Which year did this happen?

1995

CMU Navlab Minivan

Pittsburgh -> San Diego,
2800 miles of autonomy

(... but really only lane-keeping)

2005

Stanford's Stanley

Wins the first
DARPA Grand Challenge
beating both of CMU's cars

*(Tested full, driverless
autonomy ... but all in a desert)*

2007

CMU's BOSS

Wins the DARPA
Urban Challenge

*(Urban setting, interaction with
other cars, traffic rules)*

Self-driving 1.0

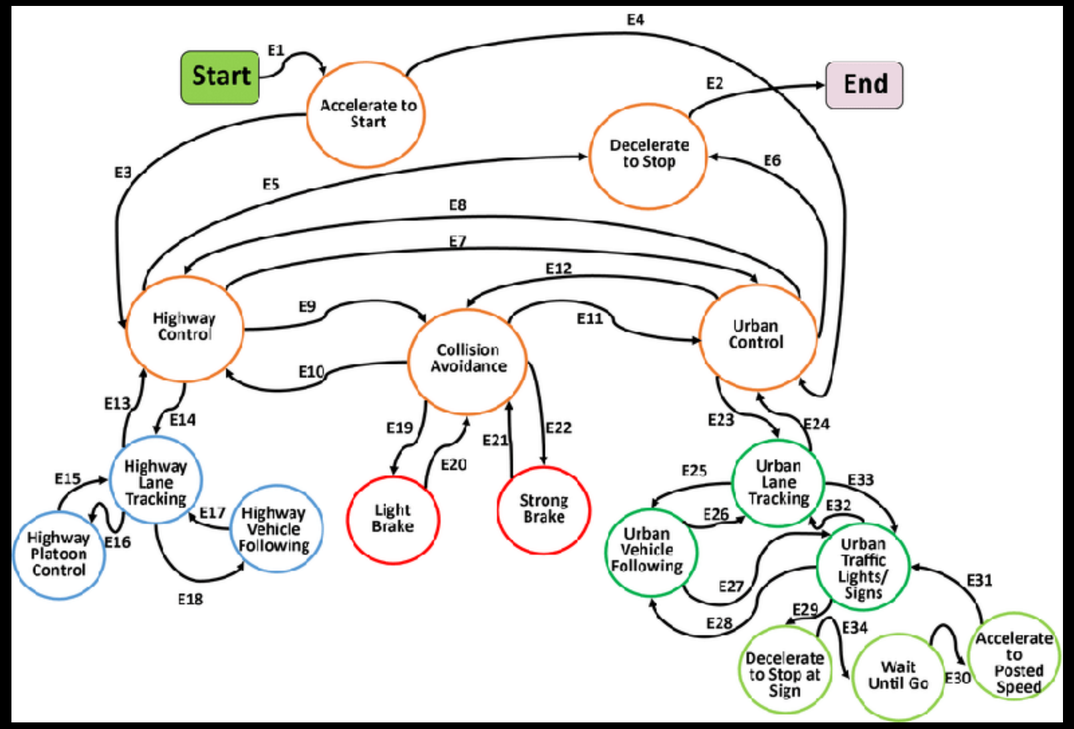

Hand-engineered rules of driving

Limited use of machine learning

Software that fundamentally could not scale over time

Self-driving 2.0

Software 2.0

Andrej Karpathy · Follow

9 min read · Nov 11, 2017

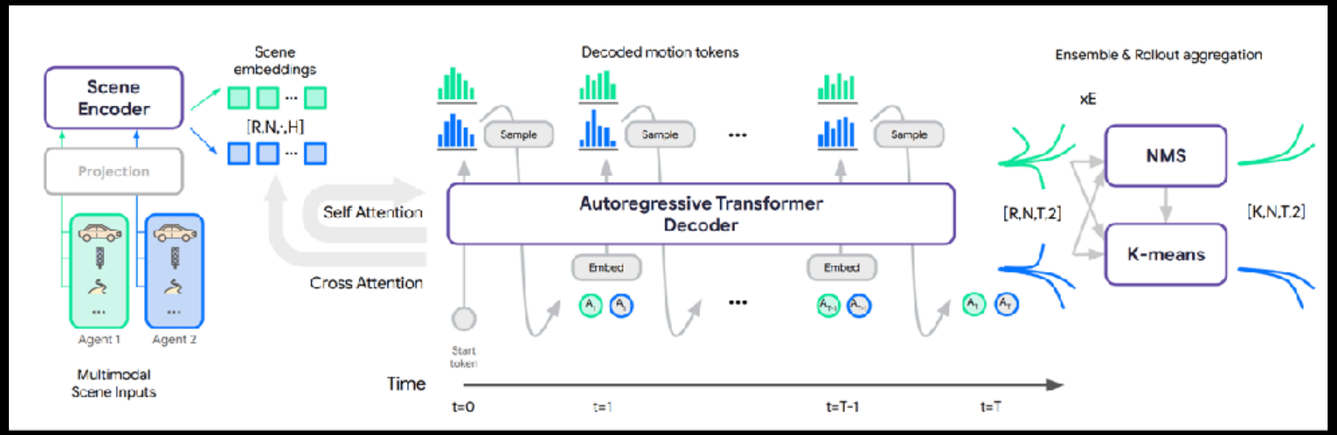

Design software from the ground-up to be learnable

Scalable pipelines that turn data into tests

Learning not just for perception, but also decision making

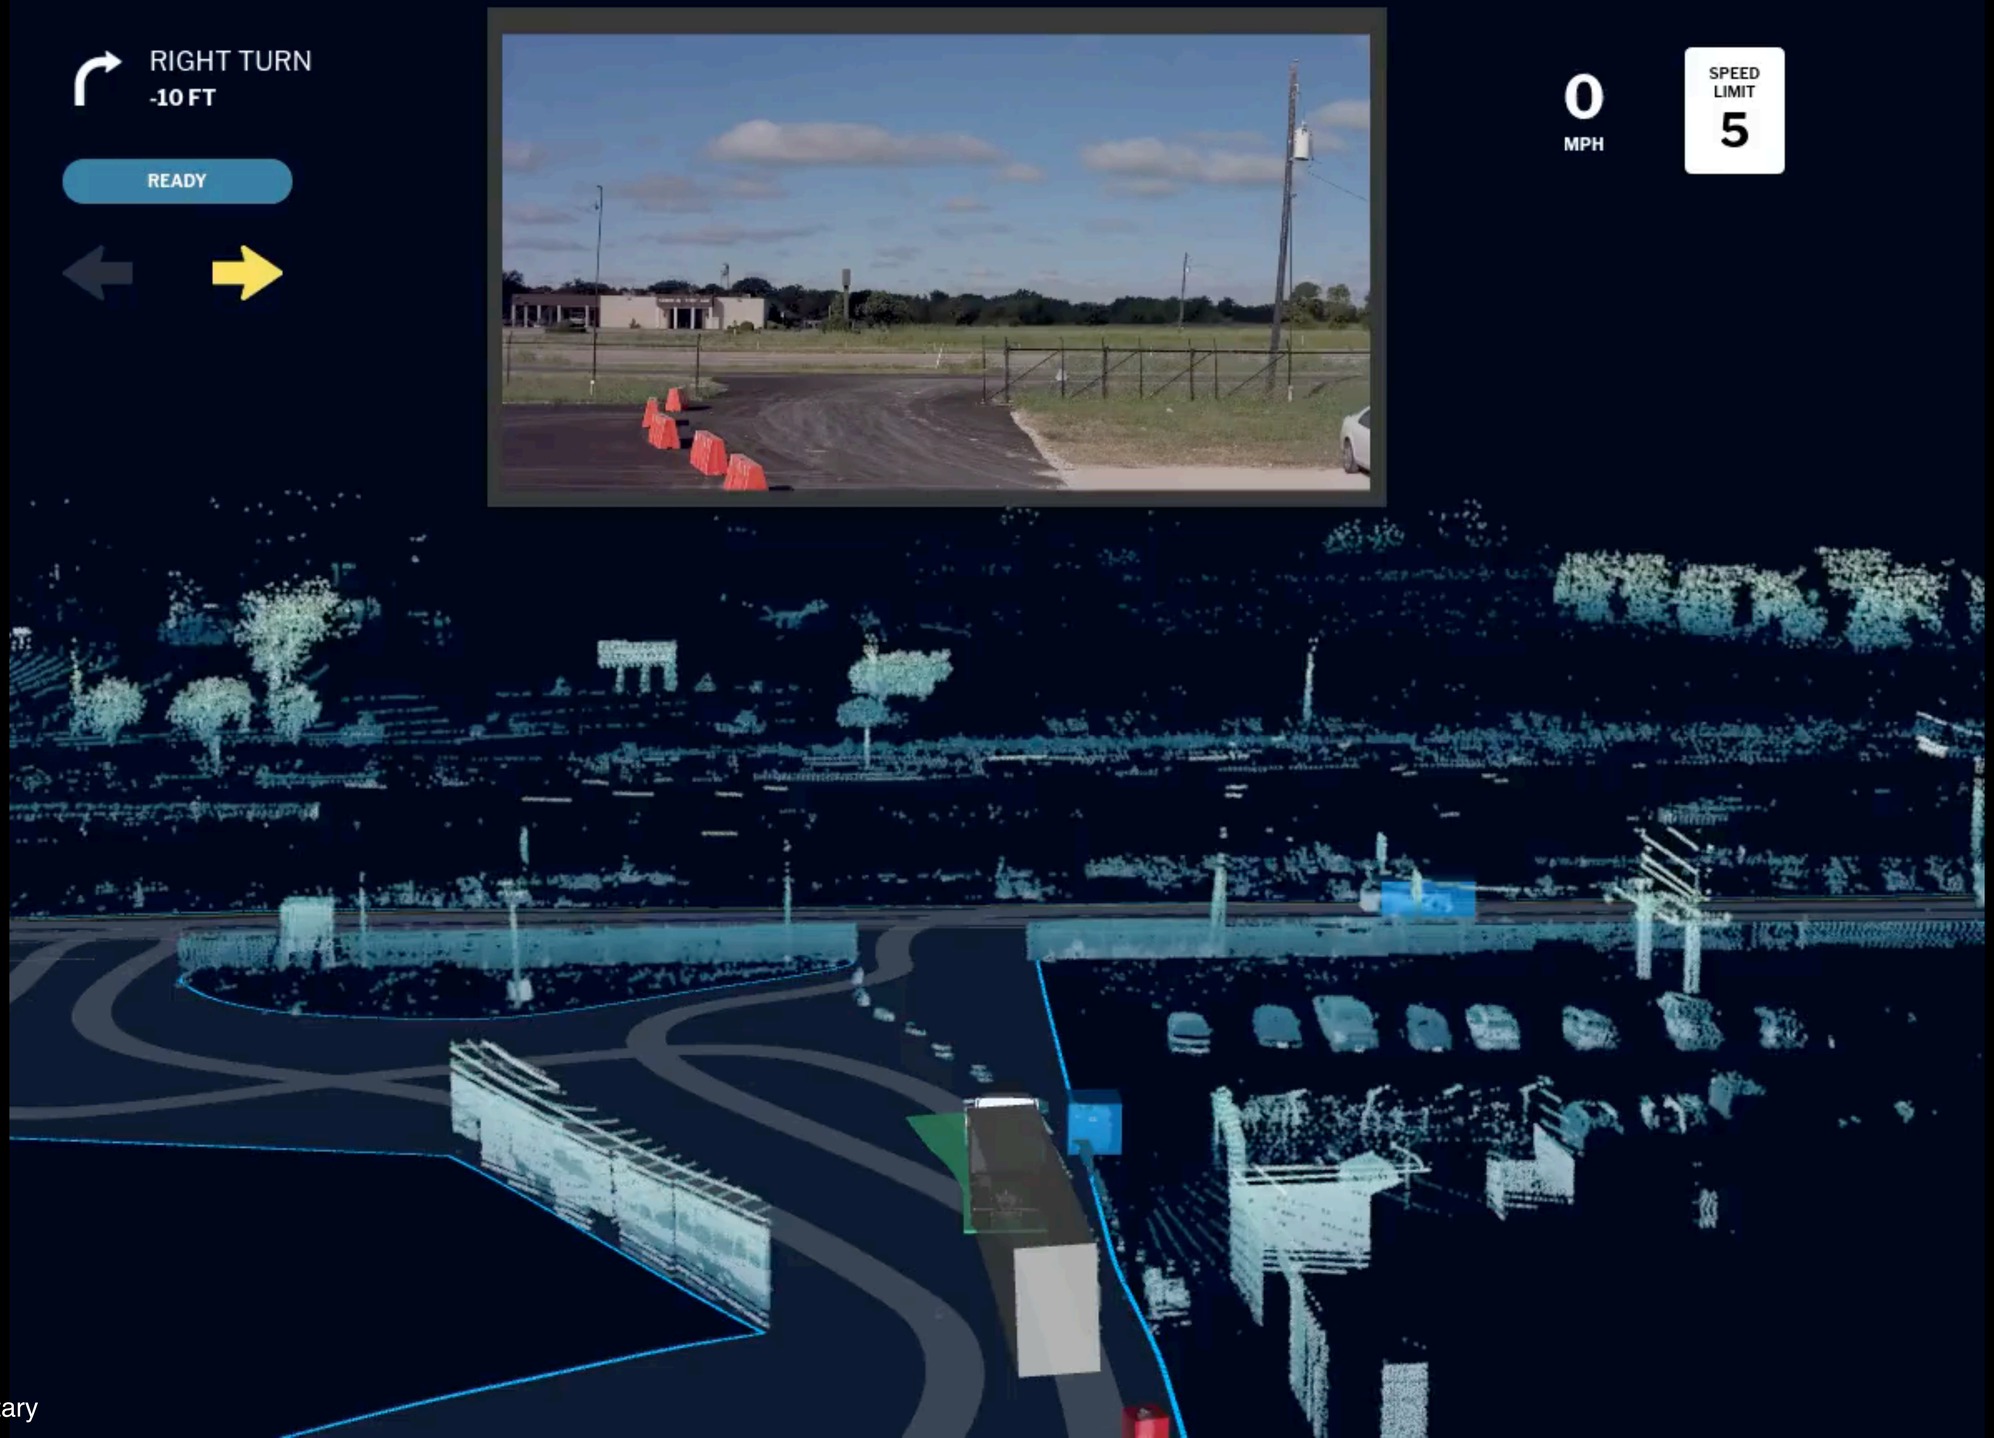

A grim state of affairs

Tesla Recalls Autopilot Software in 2 Million Vehicles

Federal regulators pressed the automaker to make updates to ensure drivers are paying attention while using Autopilot, a system that can steer, accelerate and brake on its own.

Cruise Stops All Driverless Taxi Operations in the United States

The move comes just two days after California regulators told the company to take its autonomously driven cars off the road.

While machine learning is very powerful,
getting it to do the right thing in all possible situations has been hard

Even when it makes a mistake, it's hard to know why

The BILLION dollar question

Is this a scaling issue? Should we 10x our data, have bigger models?

Do we need new ways to teach our self-driving cars?
(Think of language models before and after RLHF)

Do we need more powerful simulators and have self-driving cars evolve via natural selection?

Do we need new policies for safety and interpretability?

Let’s dive a bit deeper

Lidar

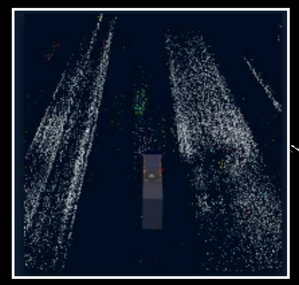

Radar

Camera

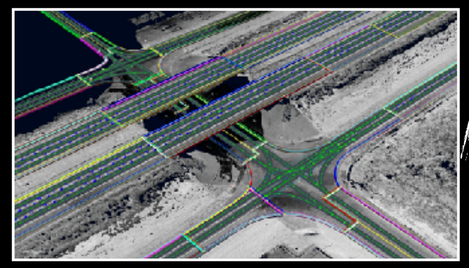

Maps

State

Perception

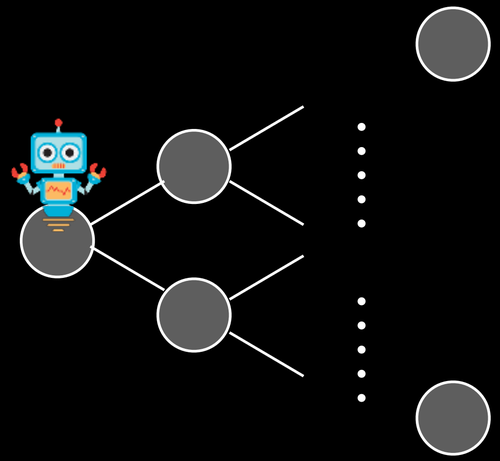

Prediction

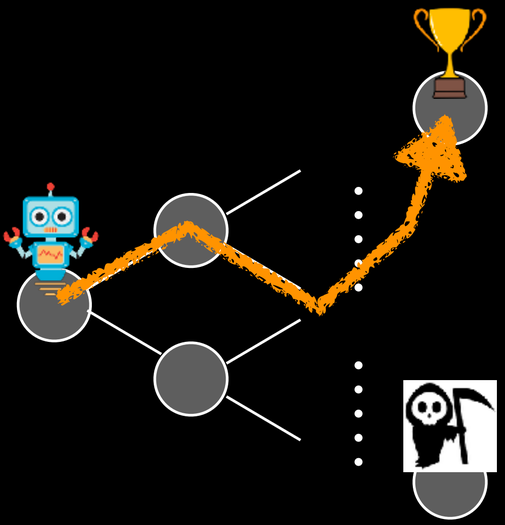

Decision Making

Lidar

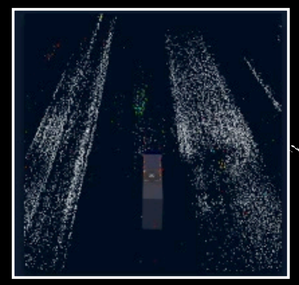

Radar

Camera

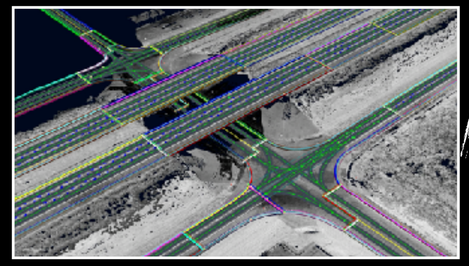

Maps

Learning!!

State

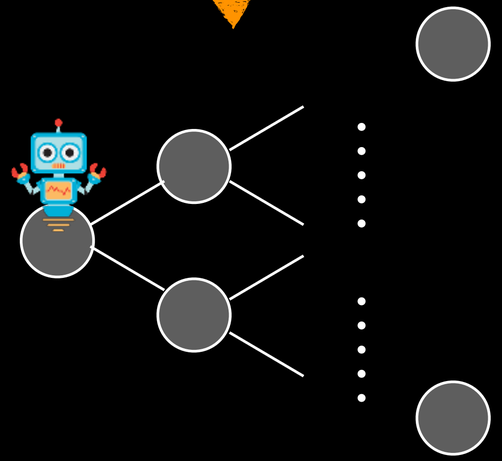

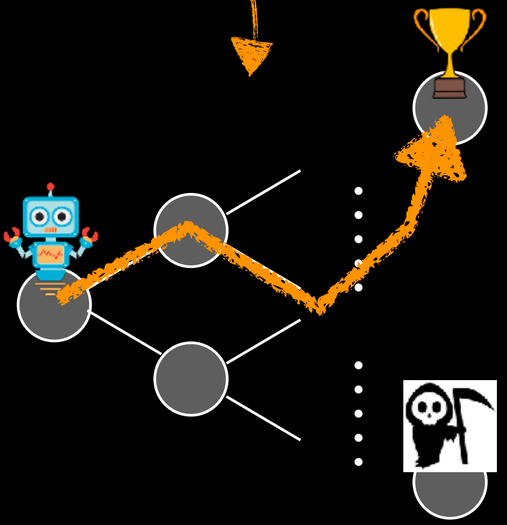

Perception

Prediction

Decision Making

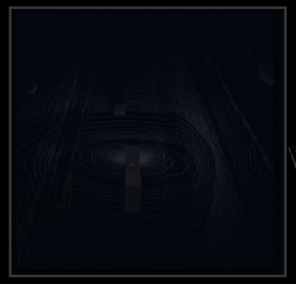

Lidar

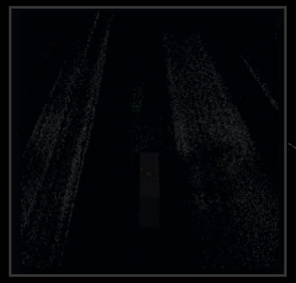

Radar

Camera

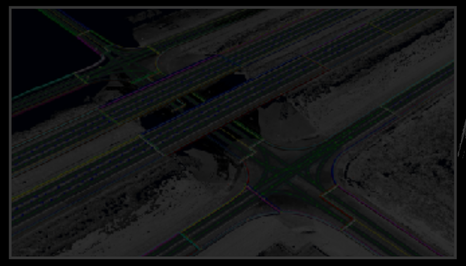

Maps

State

Perception

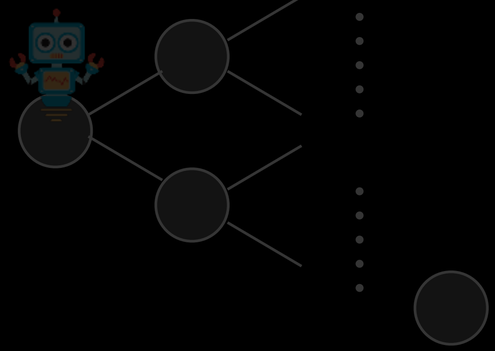

Prediction

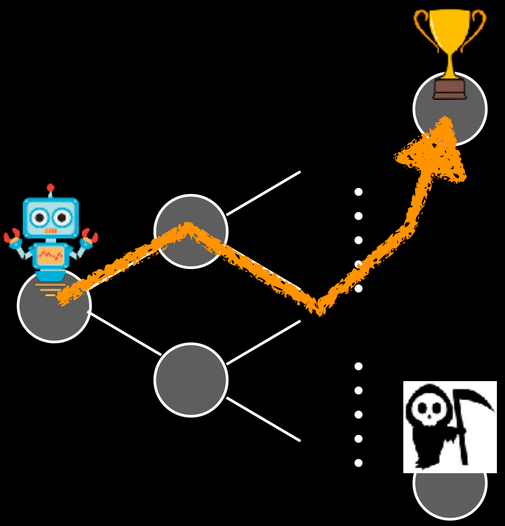

Decision Making

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

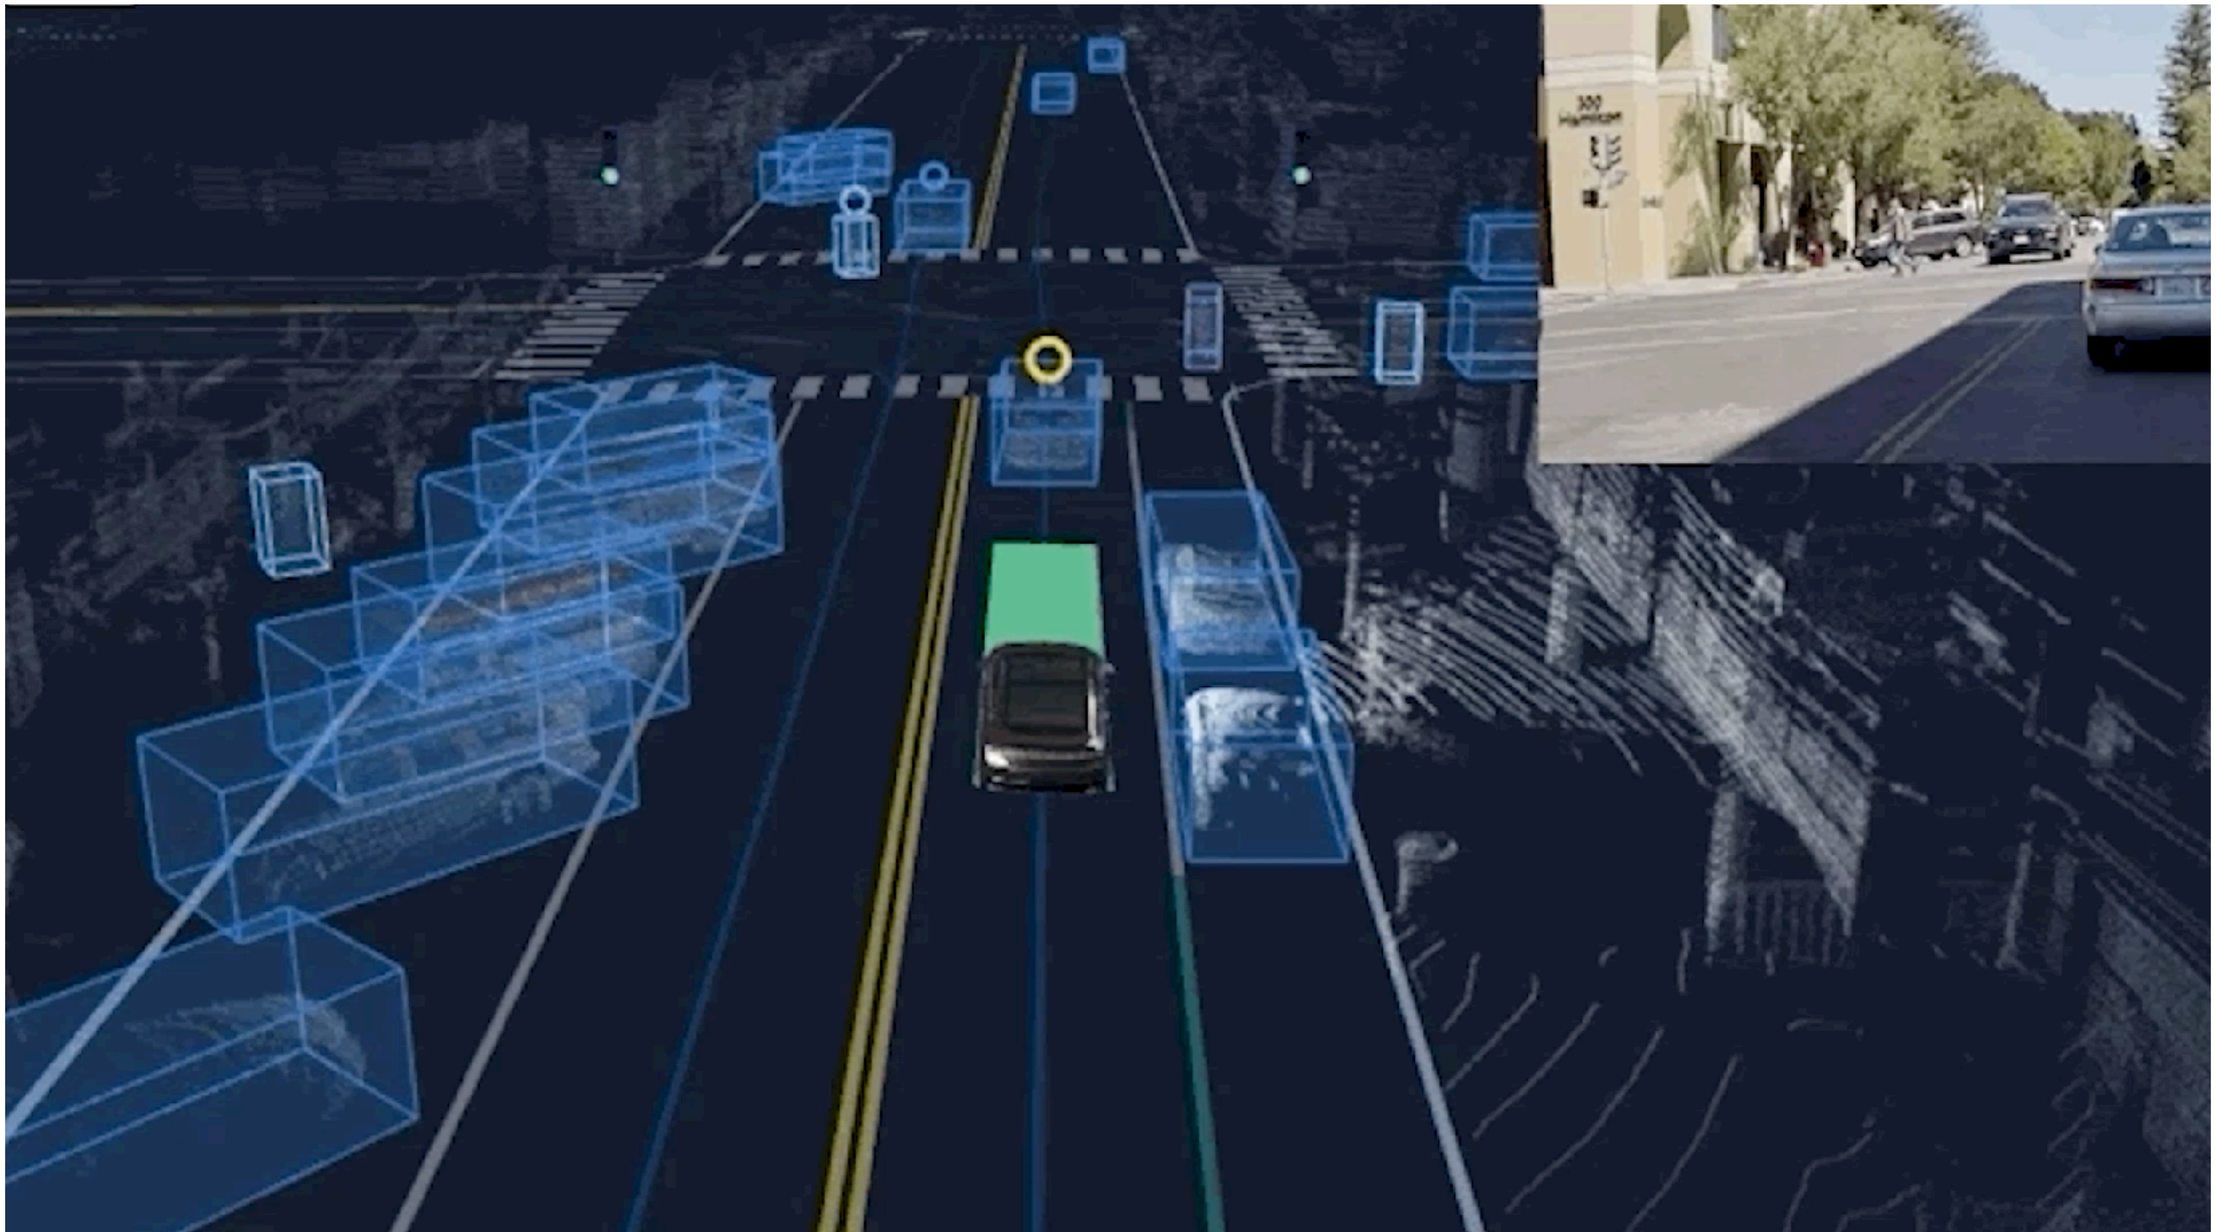

Activity!

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

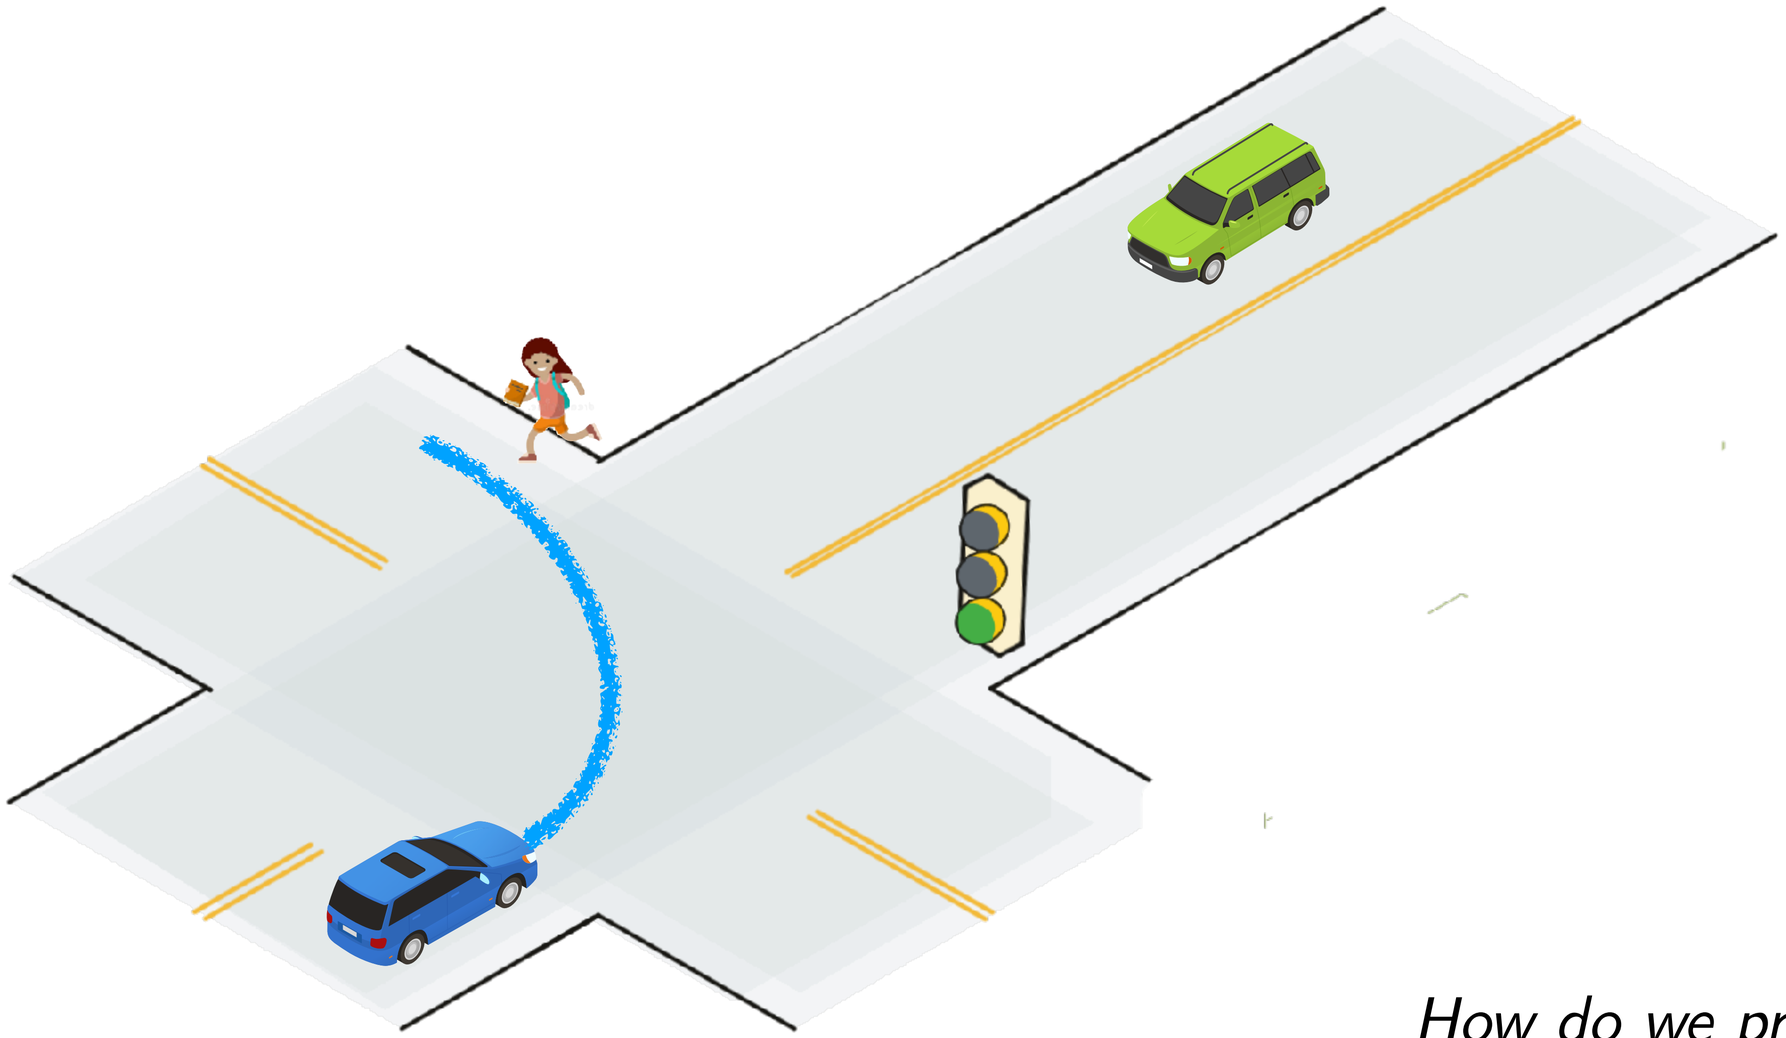

How do we program in these values?

Lesson #1

Values are implicit in human driving!

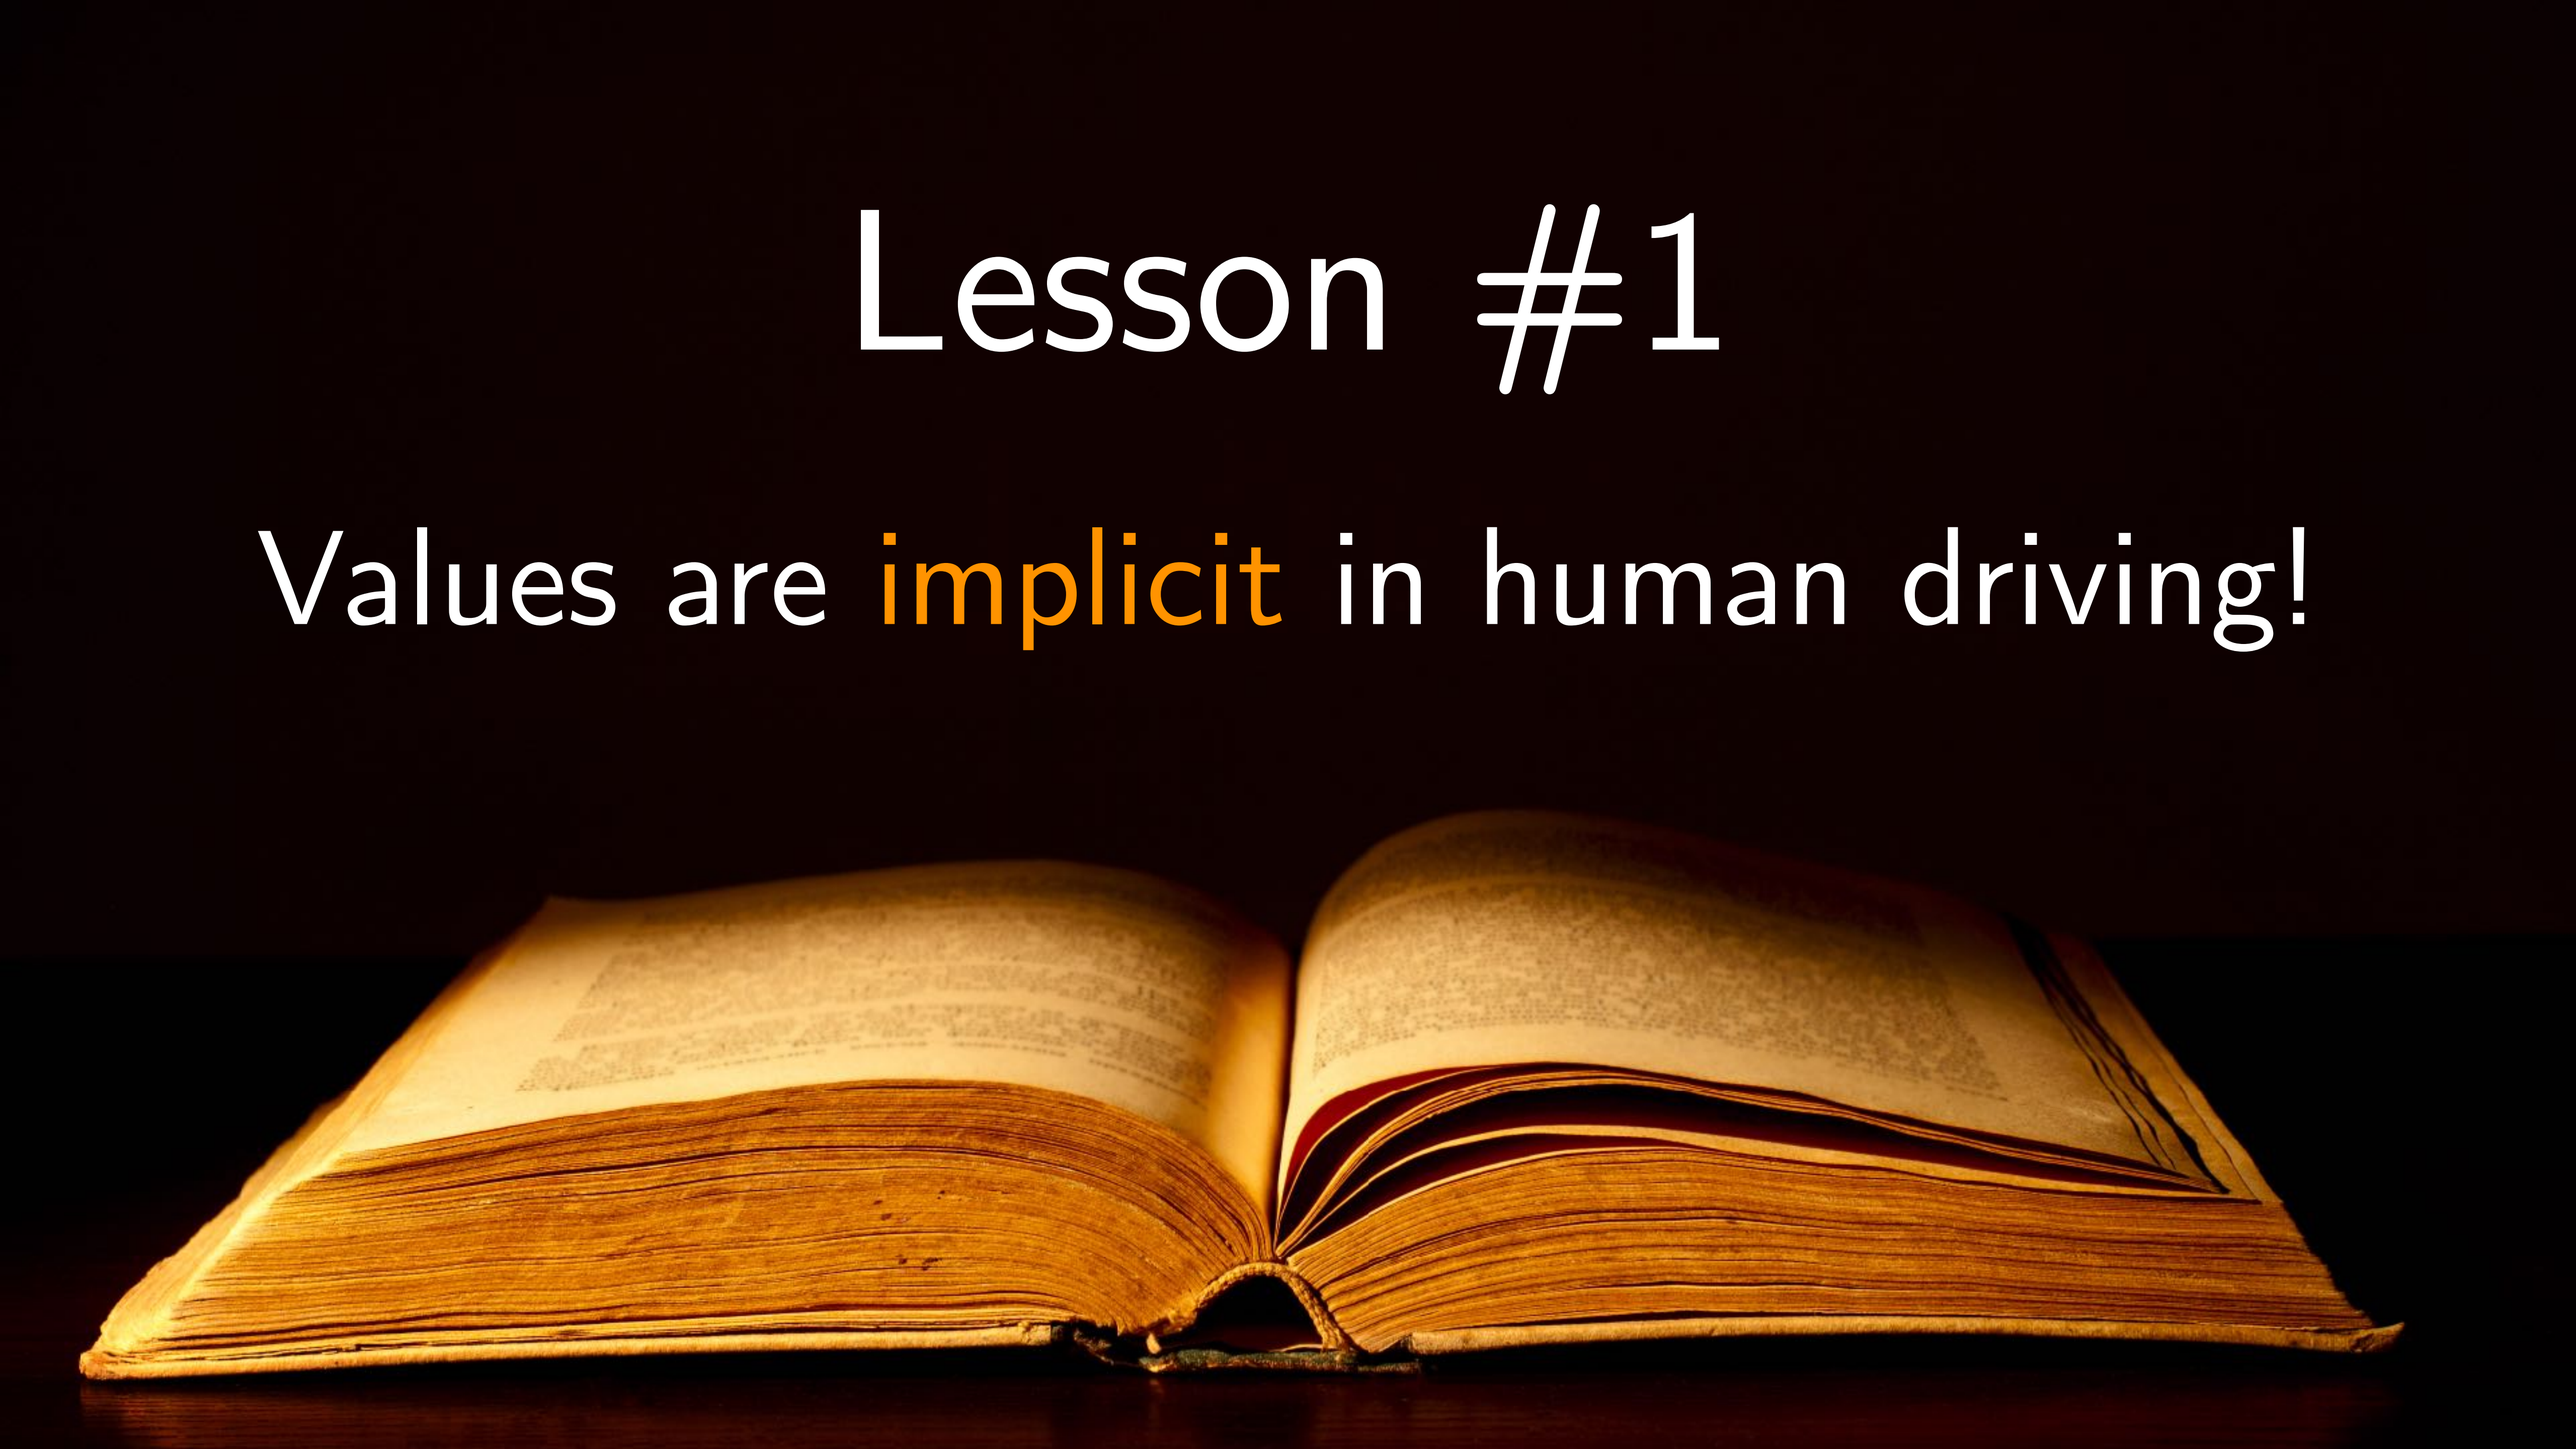

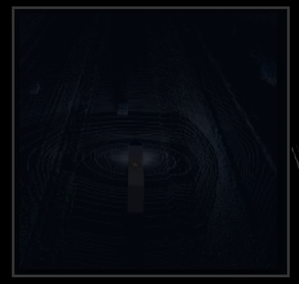

Lidar

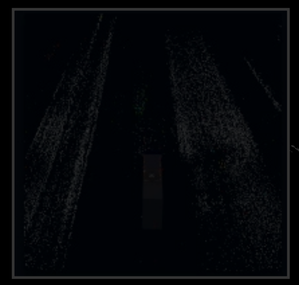

Radar

Camera

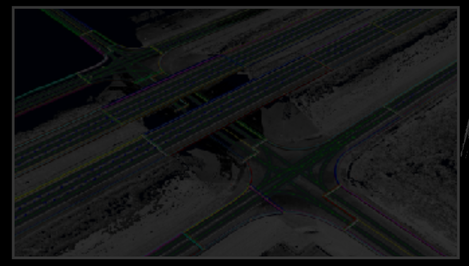

Maps

State

Perception

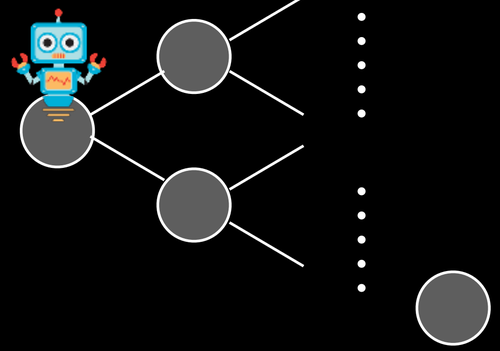

Prediction

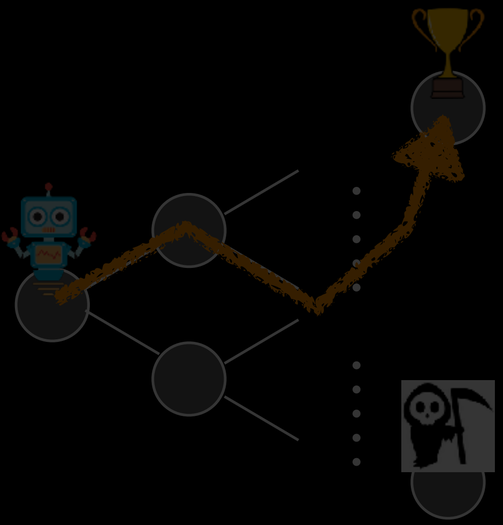

Decision Making

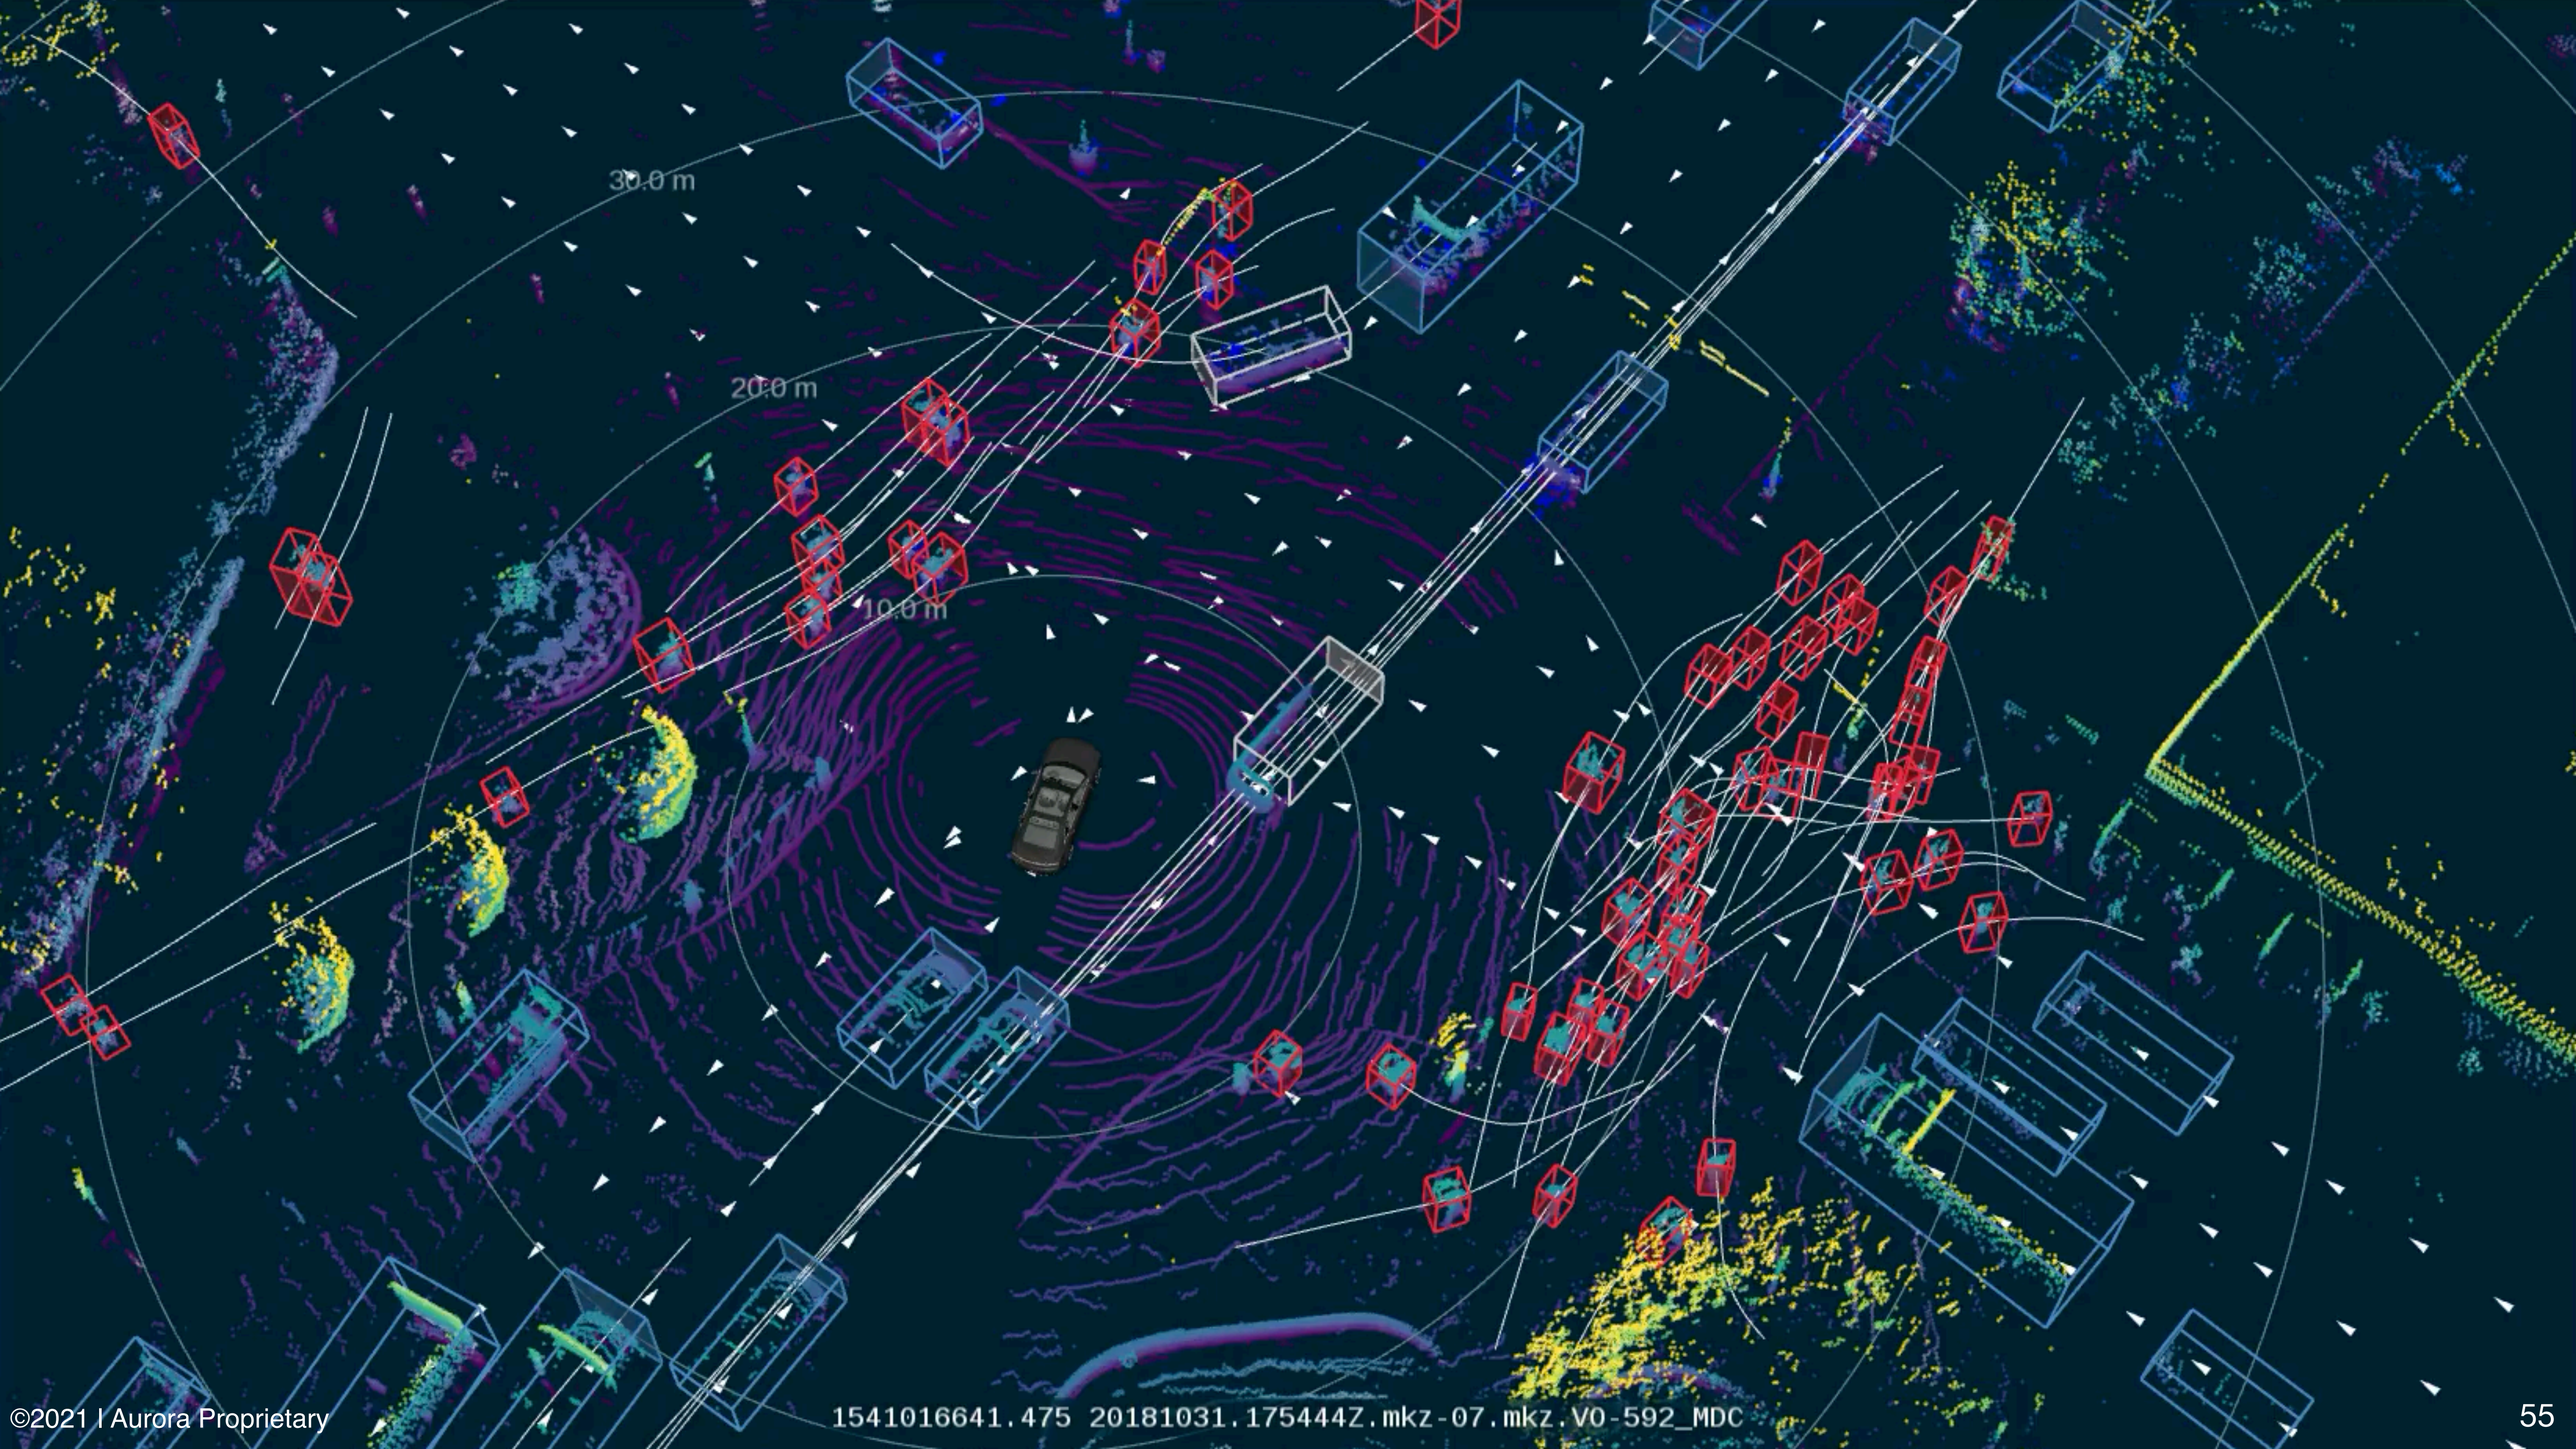

Activity: How can we predict pedestrian motion?

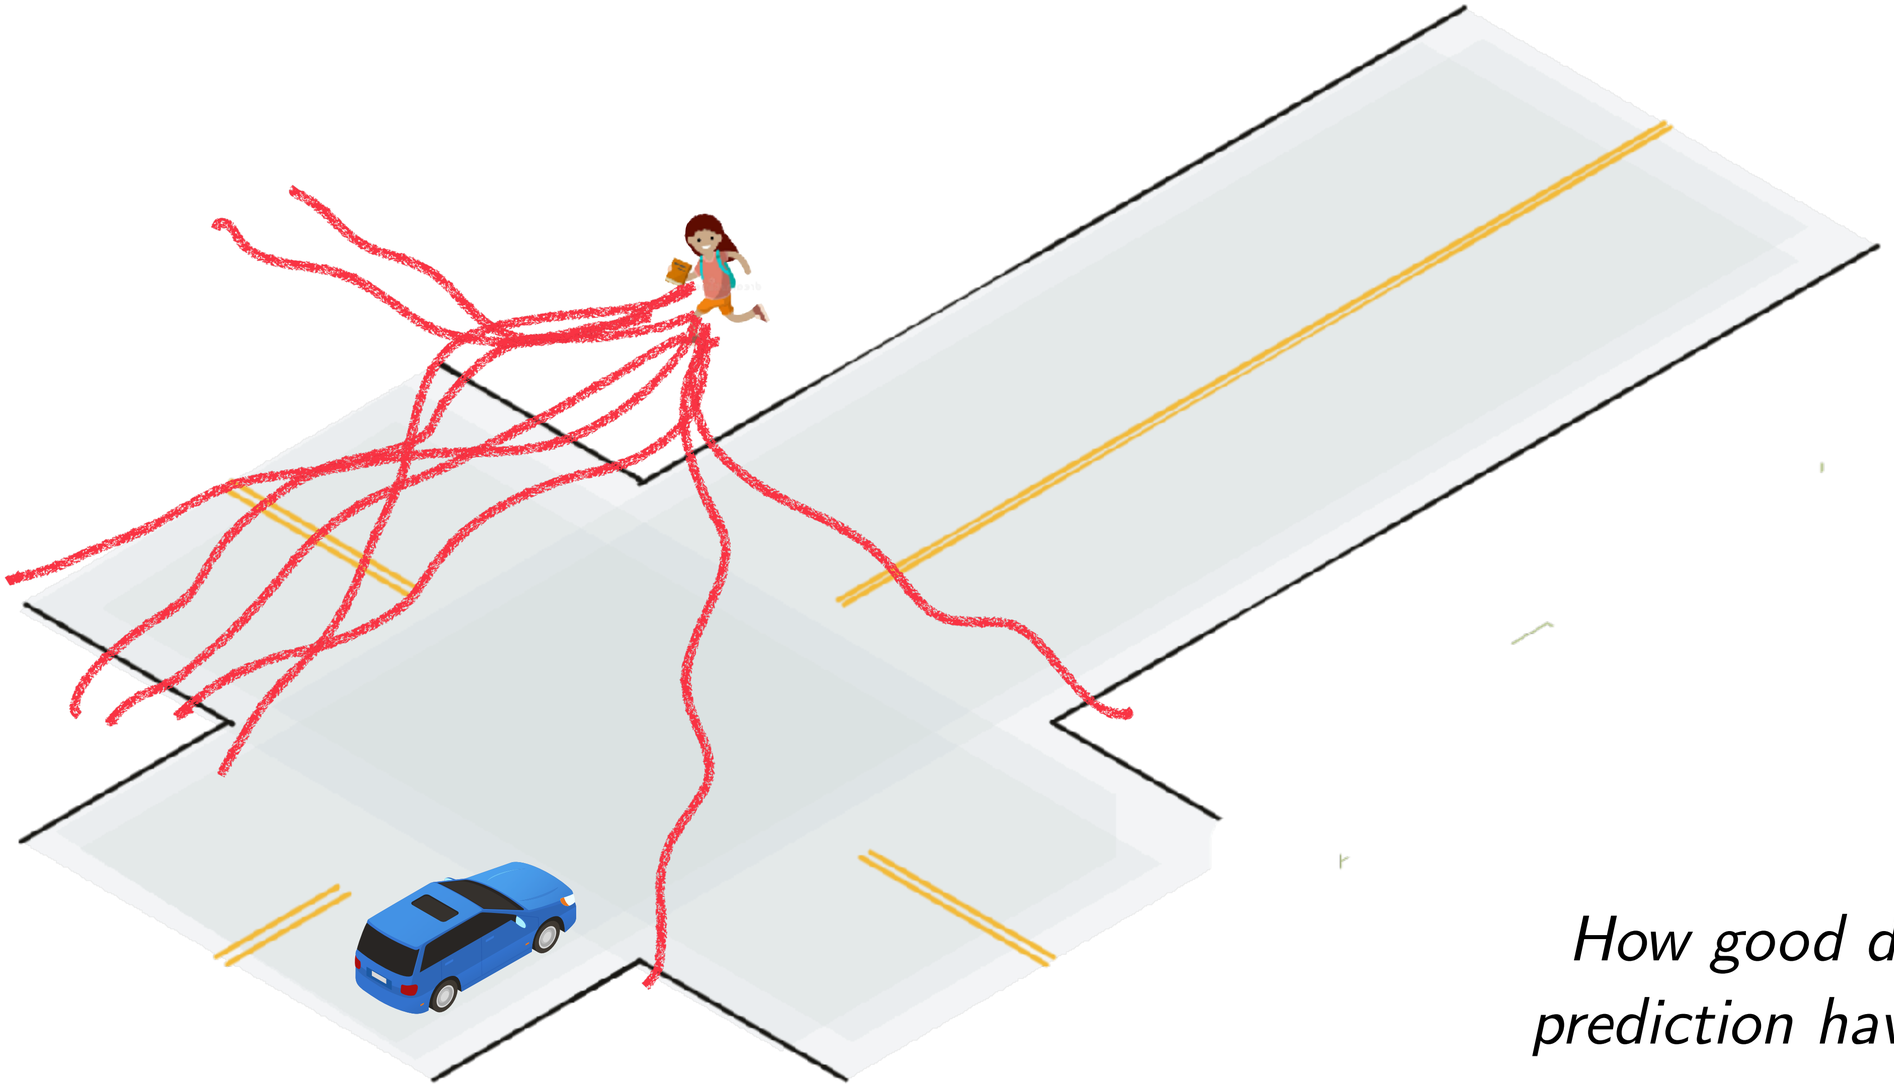

How good does the prediction have to be?

Lesson #2

Models are useful fictions

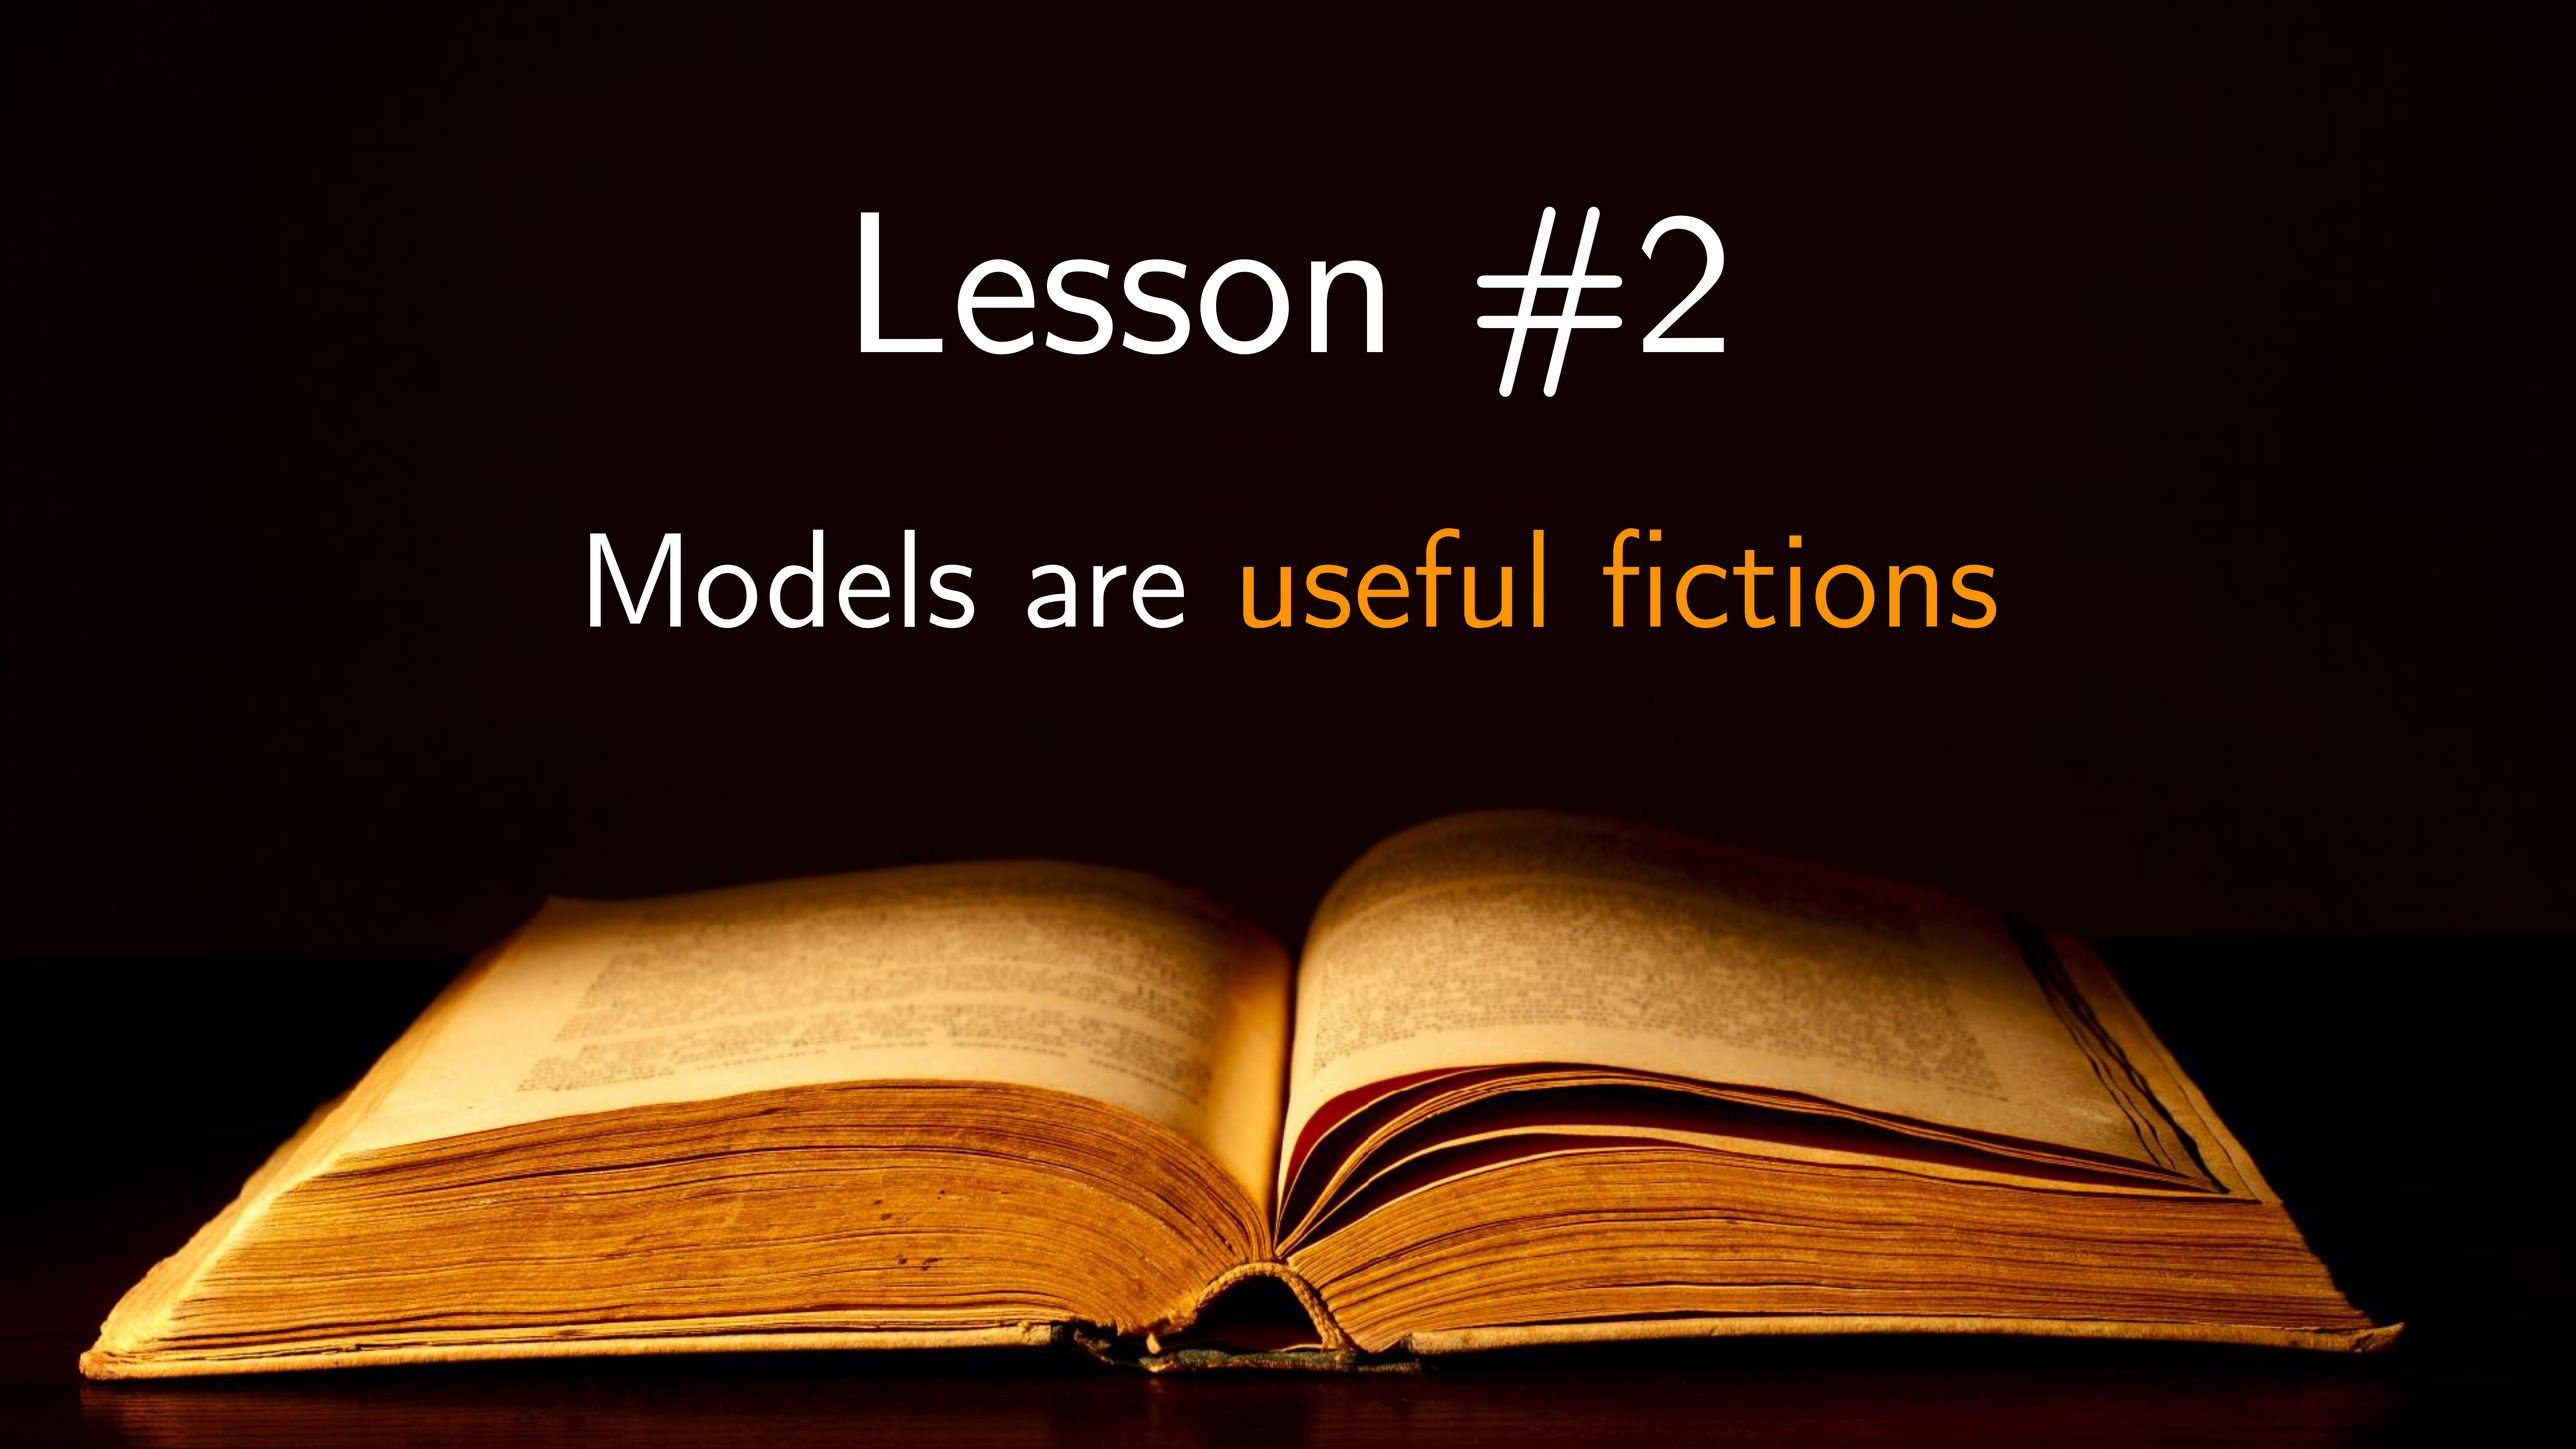

Lidar

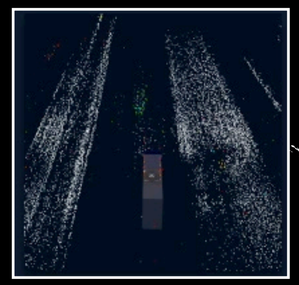

Radar

Camera

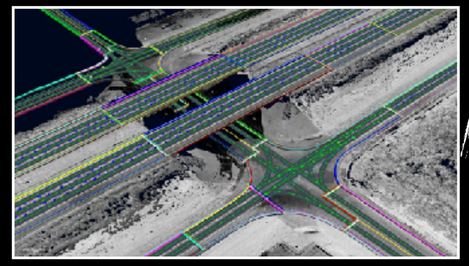

Maps

State

Perception

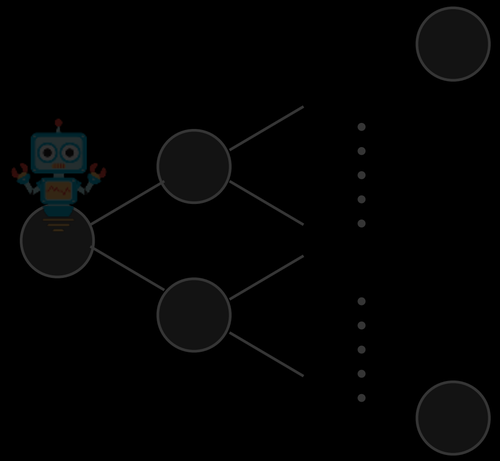

Prediction

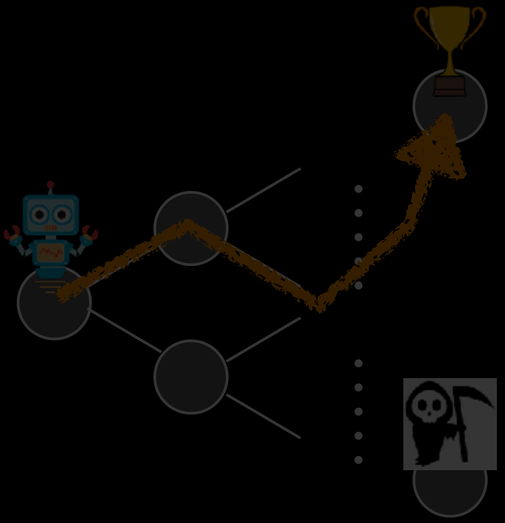

Decision Making

No one sensor tells the whole story!

LiDAR

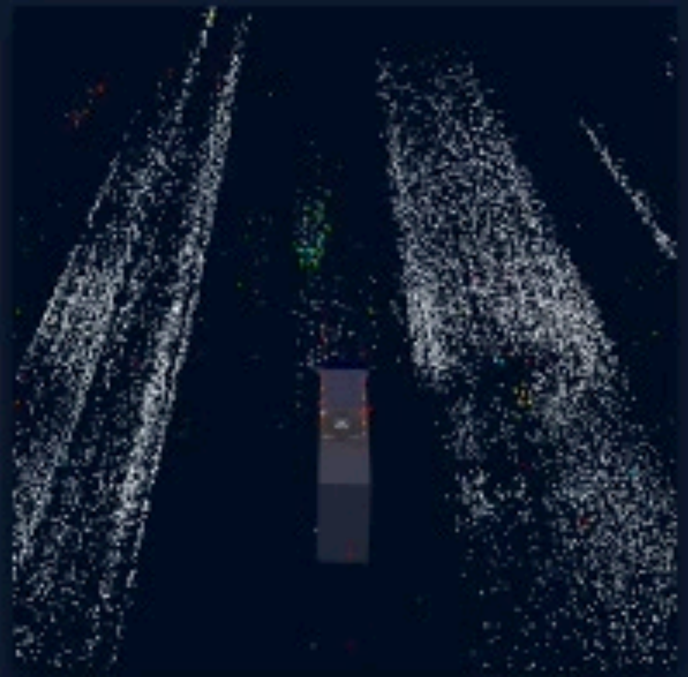

Radar

Camera

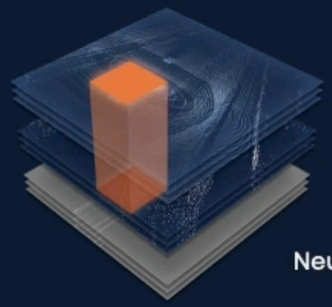

Neural Convolution Engine

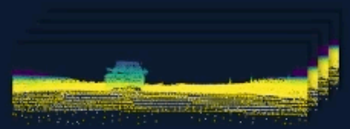

Range Conv. Engine

Image Fusion Engine

Euclidian Ray
Scatter Engine

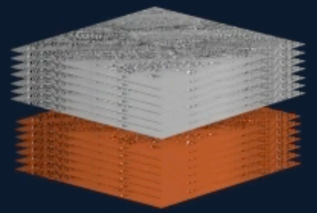

Neural Convolution
Engine

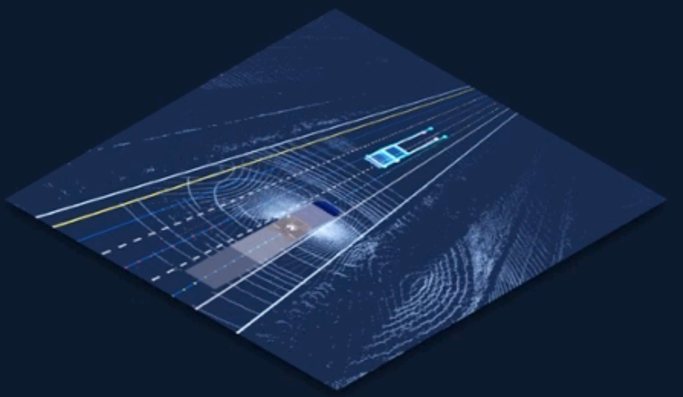

Lesson #3

Solve for the state

that explains all observations

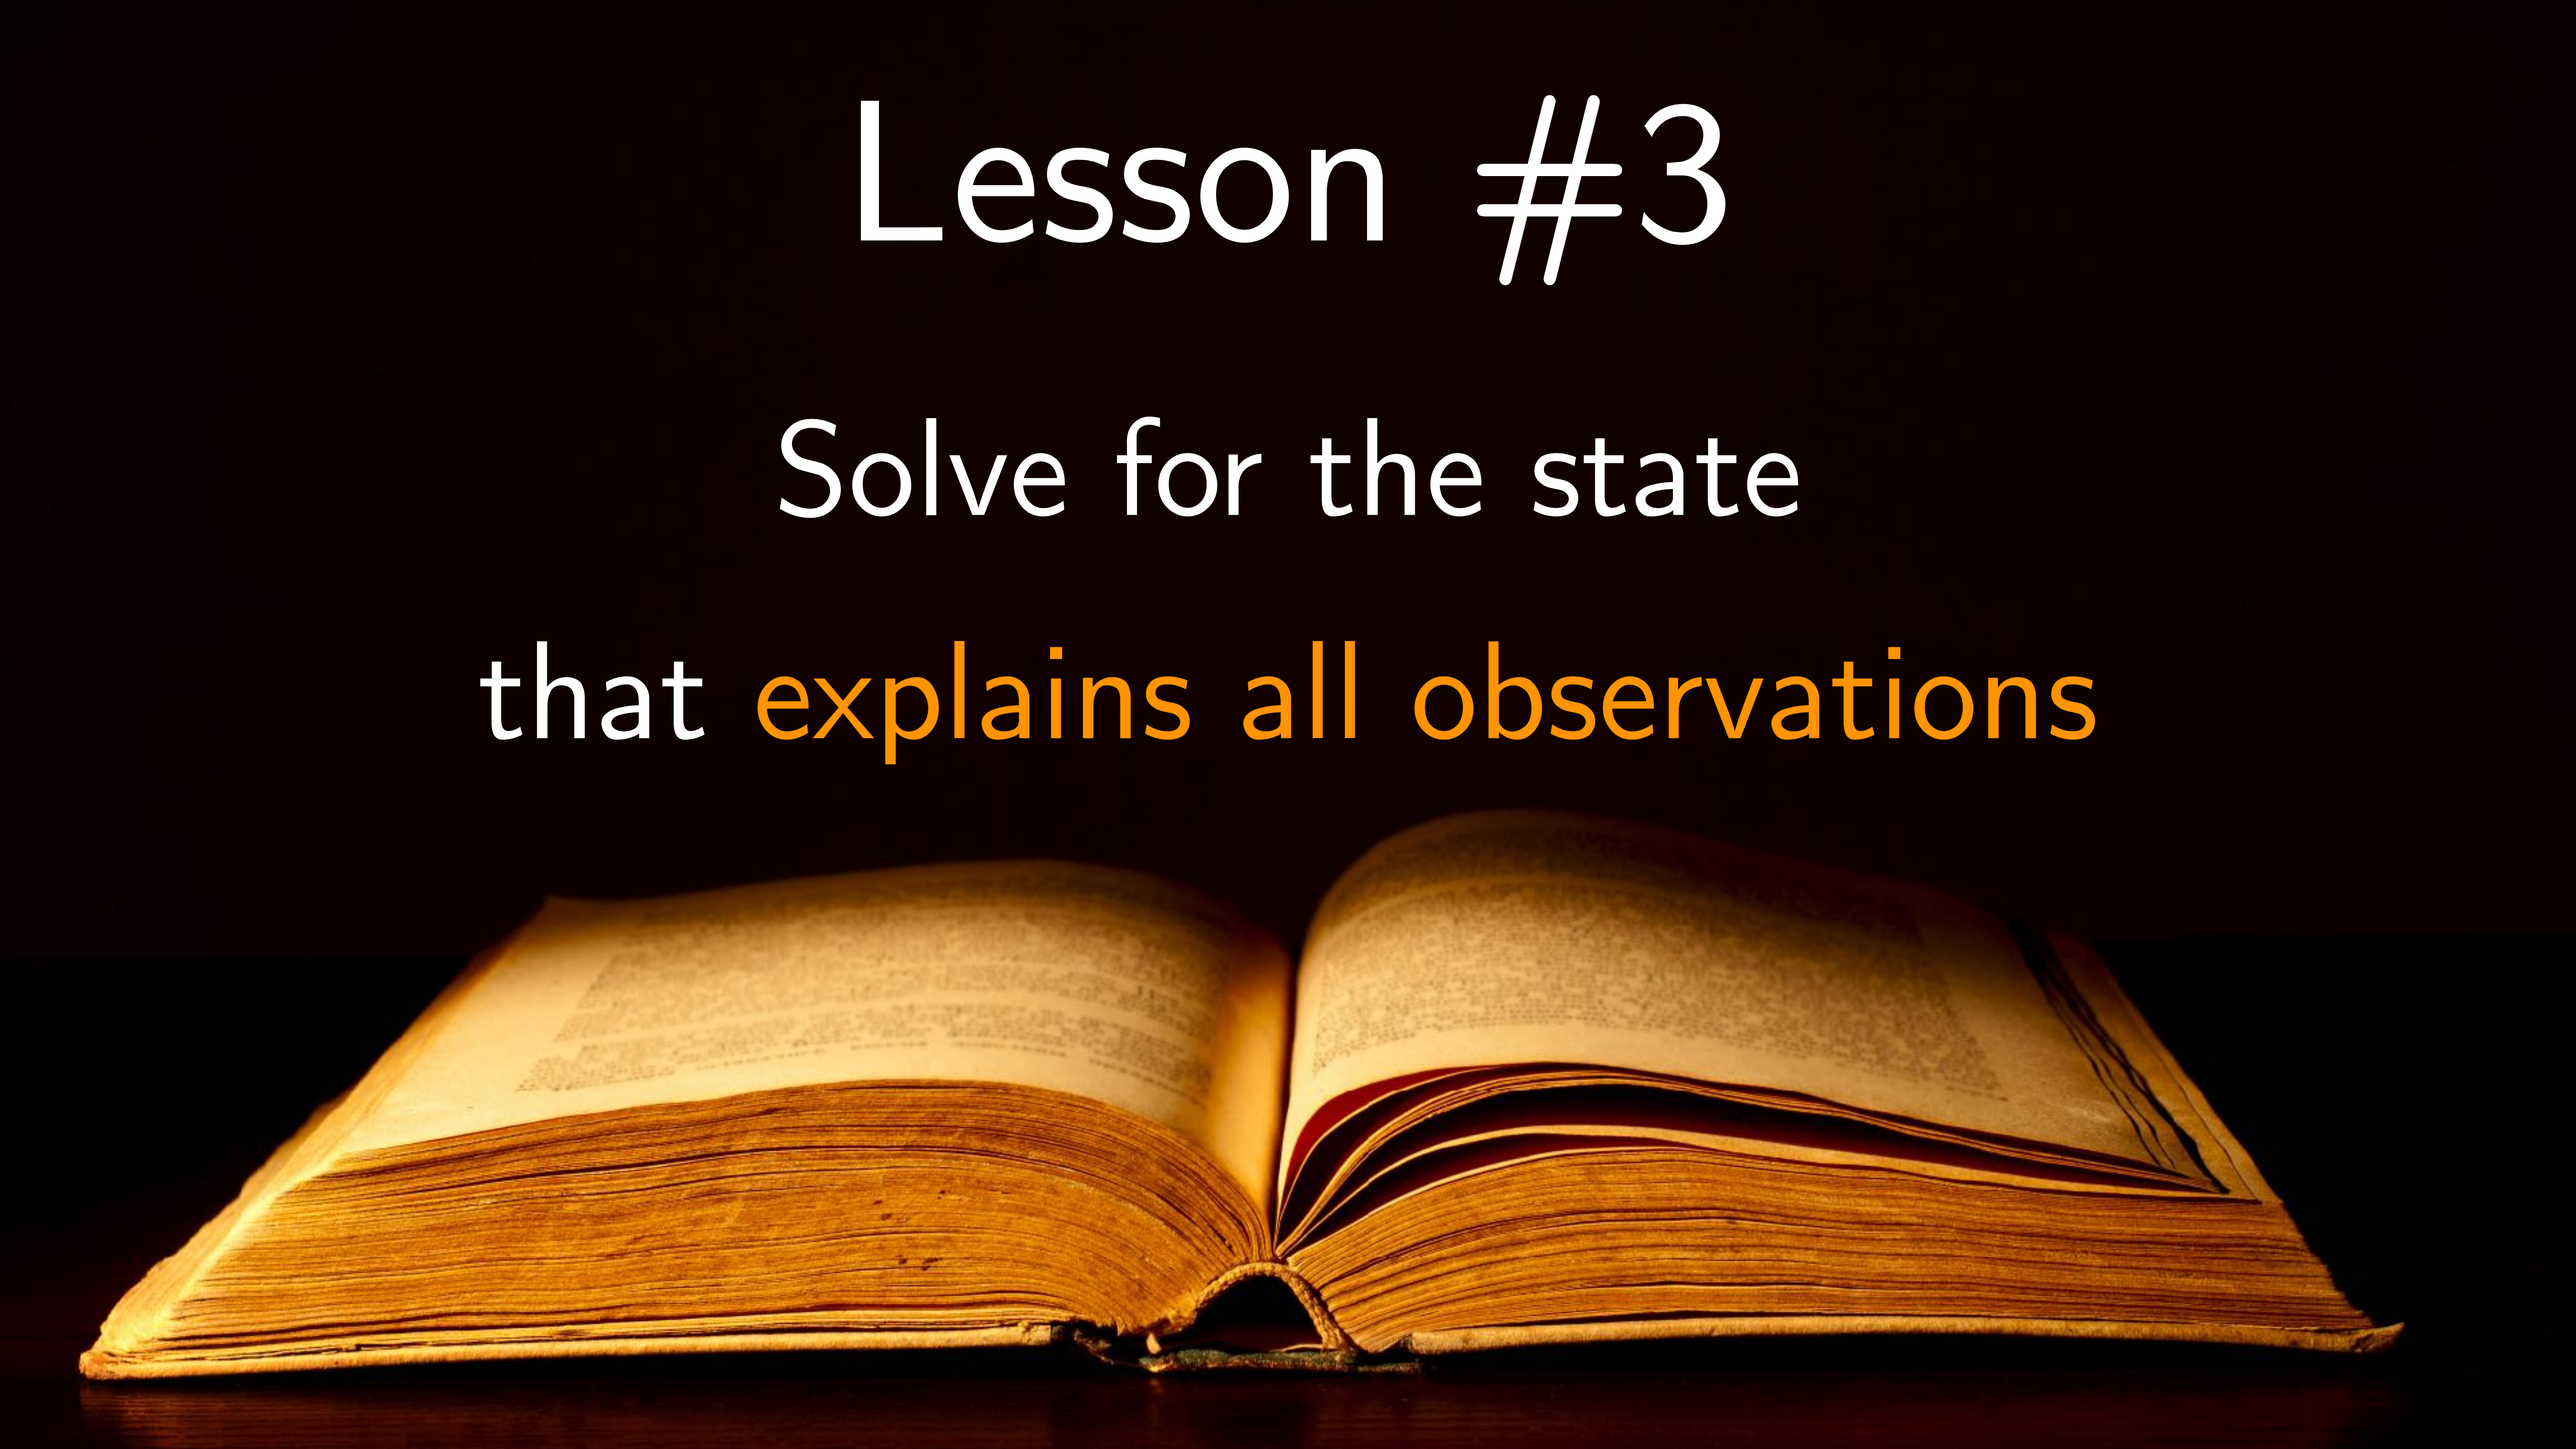

The journey ahead!

Frontiers

World Models
& Forecasting

Multi-modal
Models

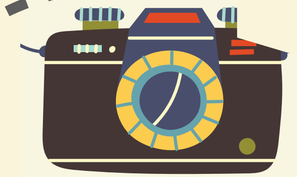

Visual
Representation

Reinforcement
Learning

Imitation
Learning

Planning
& Control

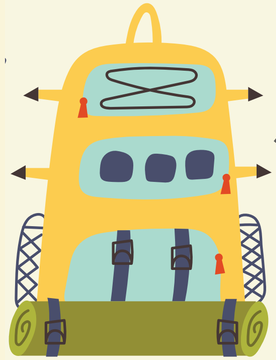

Fundamentals

Logistics

Website is the ONE true hub

https://www.cs.cornell.edu/courses/cs4756/2024fa/

Course Book

Modern Adaptive Control and Reinforcement Learning (MACRL)

Drew Bagnell, Byron Boots, Sanjiban Choudhury

https://macrl-book.github.io/

Pre-reading and Resources

Date	Lecture	Preread	Resources
	Fundamentals		
08/27/24	Introduction to Robot Learning		The Bitter Lesson
08/29/24	Robots as Markov Decision Problems	MACRL Ch. 1	Dan Klein slides I

Please look at the pre-reading before coming to lecture!

Resources are for *after* the lecture if you want to go deeper into a concept.

6 Assignments [50%]

A0: Intro assignment

A1, A3: Written assignment

A2, A4, A5: Programming assignment

Assignments will be based off of concepts / exercises from class!

In-class Prelim [20%]

Use written assignments as a reference

Use course book (pre-reading chapters) as a reference

Final Project [20%]

This is your chance to be creative and apply concepts to solve some robot learning problems!

See this doc for ideas.

We, unfortunately, do not have GPUs to offer, so choose projects wisely that you can run on your machines. Talk to TAs!

The best projects are simple ideas that convey insight!

Participation [10%]

Participate in class polls and exercises!

Graduate Version (CS5756)

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

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.

Generative AI

The work you do consists of writing code and natural language descriptions.

To some extent, the new crop of "generative AI" (GAI) tools can do both of these things for you.

However, **we require that the vast majority of the intellectual work must be originated by you,** not by GAI. You may use GAI to look up helper functions, or to proofread your text, but clearly document how you used it.

Generative AI

In this class, for every assignment and final project, you can choose between two options:

Option 1: Avoid all GAI tools. Disable GitHub Copilot in your editor, do not ask chatbots any questions related to the assignment, etc. If you choose this option, you have nothing more to do.

Option 2: Use GAI tools with caution and include a one-paragraph description of everything you used them for along with your writeup. This paragraph must:

1. Link to exactly which tools you used and describe how you used each of them, for which parts of the work.
2. Give at least one concrete example (e.g., generated code or Q&A output) that you think is particularly illustrative of the "help" you got from the tool.
3. Describe any times when the tool was unhelpful, especially if it was wrong in a particularly hilarious way.
4. Conclude with your current opinion about the strengths and weaknesses of the tools you used for real-world compiler implementation.

Remember that you can pick whether to use GAI tools for every assignment, so using them on one set of tasks doesn't mean you have to keep using them forever.

The Crew

Sanjiban Choudhury

Instructor

sanjibanc@cornell.edu

Nathaniel Chin

Teaching Assistant

nlc62@cornell.edu

Atiksh Bhardwaj

Teaching Assistant

ab2635@cornell.edu

Zach Garcia

Teaching Assistant

zag7@cornell.edu

Zubin Bhaumik

Teaching Assistant

jir429@cornell.edu

Giorgi Berndt

Teaching Assistant

gb449@cornell.edu

Ved Sriraman

Teaching Assistant

vs346@cornell.edu

Riley Coogan

Teaching Assistant

rmc329@cornell.edu

Teaching Assistant

abc256@cornell.edu

Qian Meng

Teaching Assistant

qm34@cornell.edu

Zach Garcia

Teaching Assistant

zag7@cornell.edu

Ved Sriraman

Teaching Assistant

vs346@cornell.edu

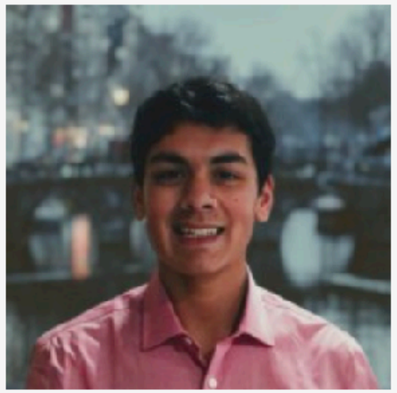

Zubin Bhaumik

Teaching Assistant

jlr429@cornell.edu

Riley Coogan

Teaching Assistant

rmc329@cornell.edu

Giorgi Berndt

Teaching Assistant

gb449@cornell.edu

Wait list info

- Sanjiban cannot do anything about the waitlist / getting into the course
- If you are on the waitlist, you will eventually receive a PIN. Keep checking your emails.
- If you are not on the waitlist, please wait for the waitlist to be cleared before adding yourself on.
- If you are unable to add yourself to the waitlist, reach out to the registrar's office
- Historically, everyone has gotten off of the waitlist. We *hope* this is the case this semester!

TLDR

Checkout course website for all details:

https://www.cs.cornell.edu/courses/cs4756/2024fa/

Checkout pre-reading for next lecture!