CS 4756/5756: Robot Learning

Sanjiban Choudhury

PORTAL

Cornell Bowers CIS Computer Science

WHAT A TIME TO BE ALIVE!

2023 was a big year for Machine Learning


GPT-4

Reinforcement Learning from Human Feedback (RLHF)

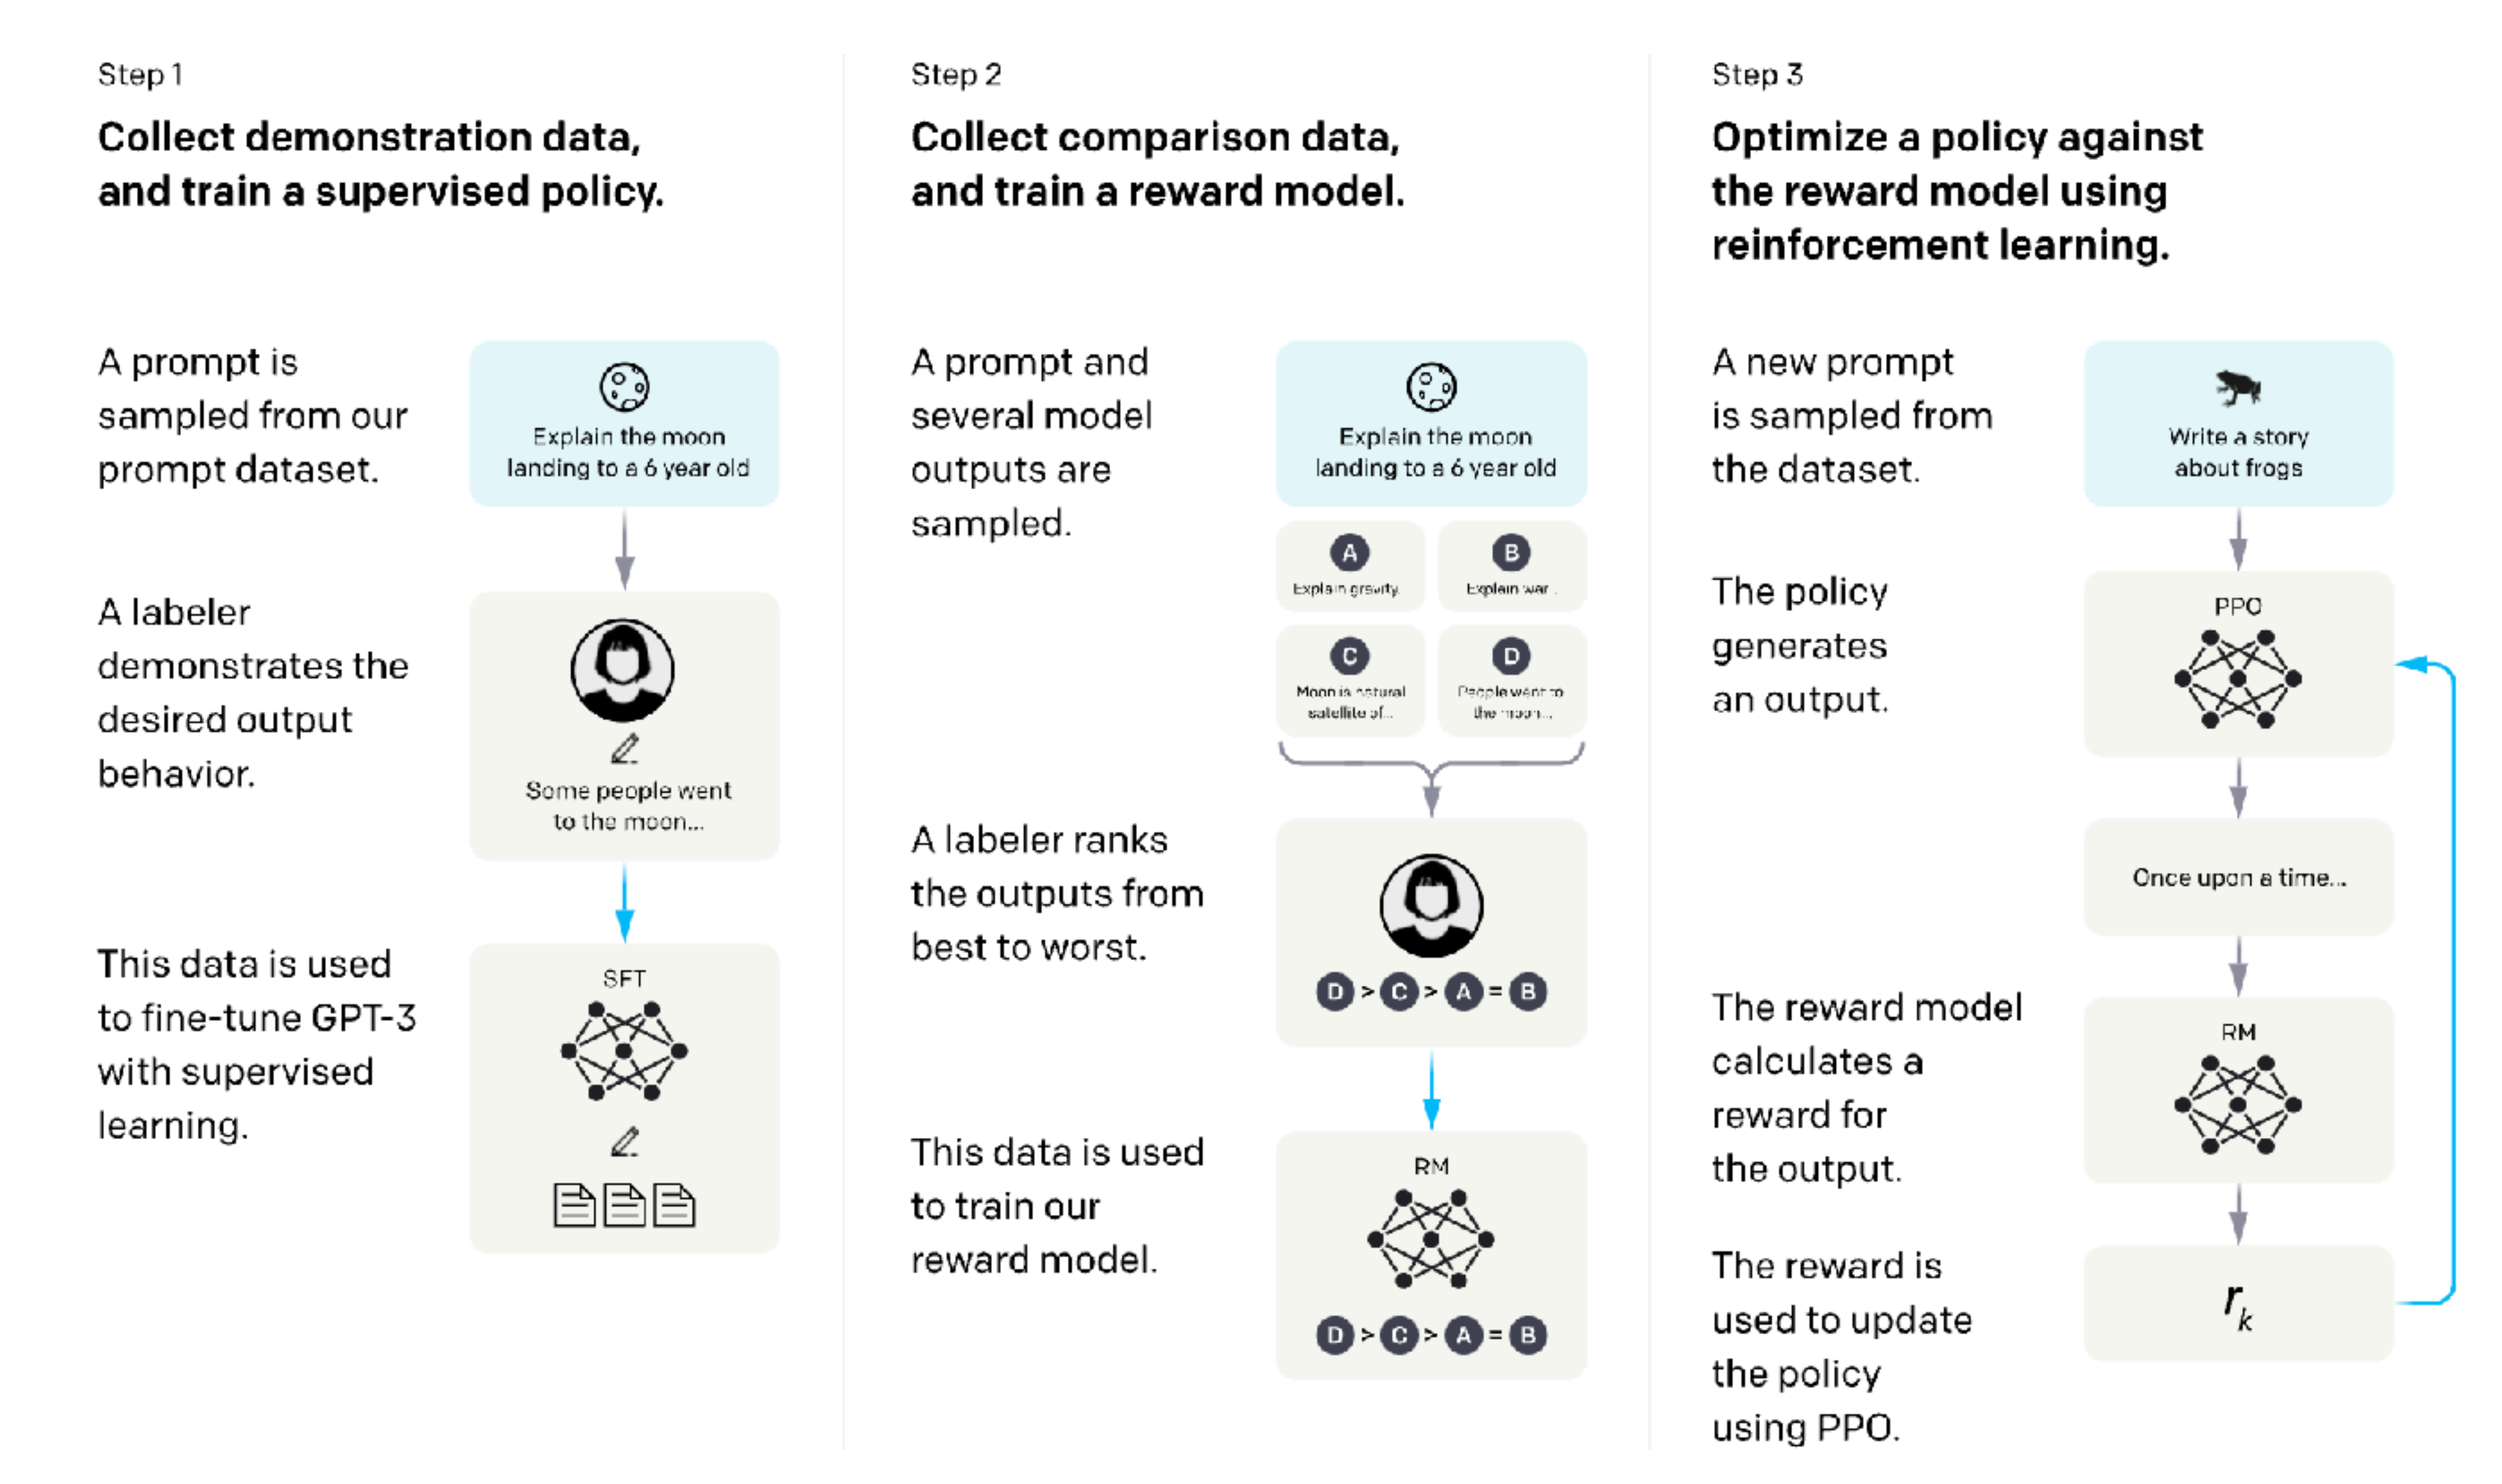

Open-source fine-tunable models

LLAMA

Alpaca

Vicuna

Mistral

Parameter Efficient Fine Tuning (PEFT)

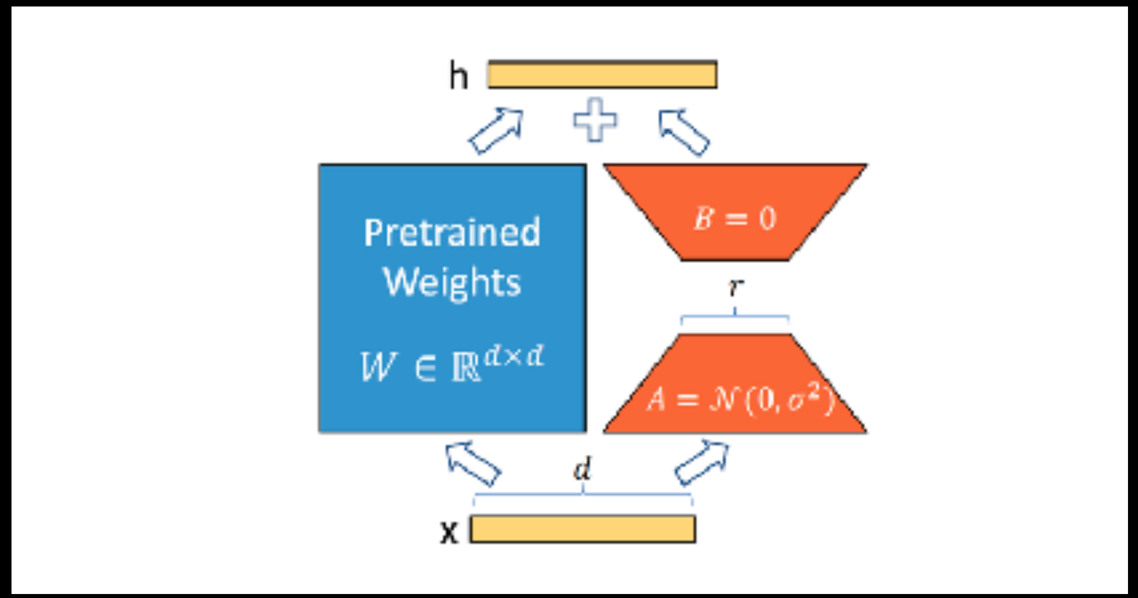

LORA

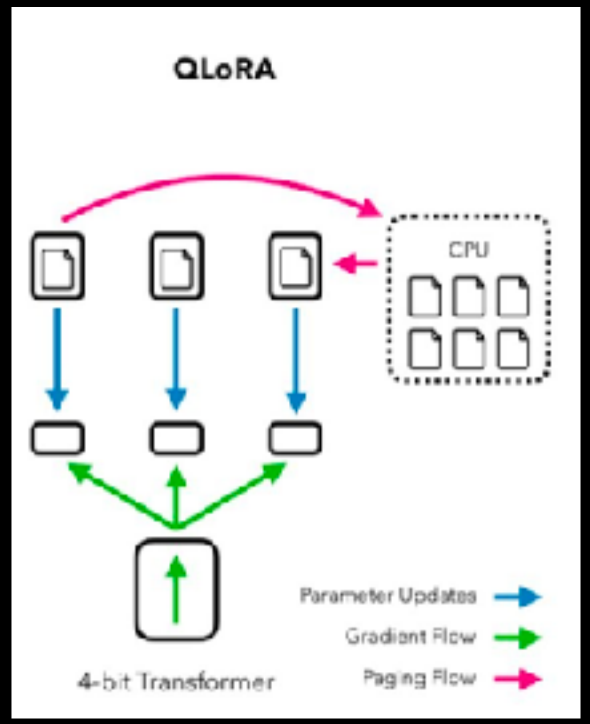

QLORA

Multi-modal models

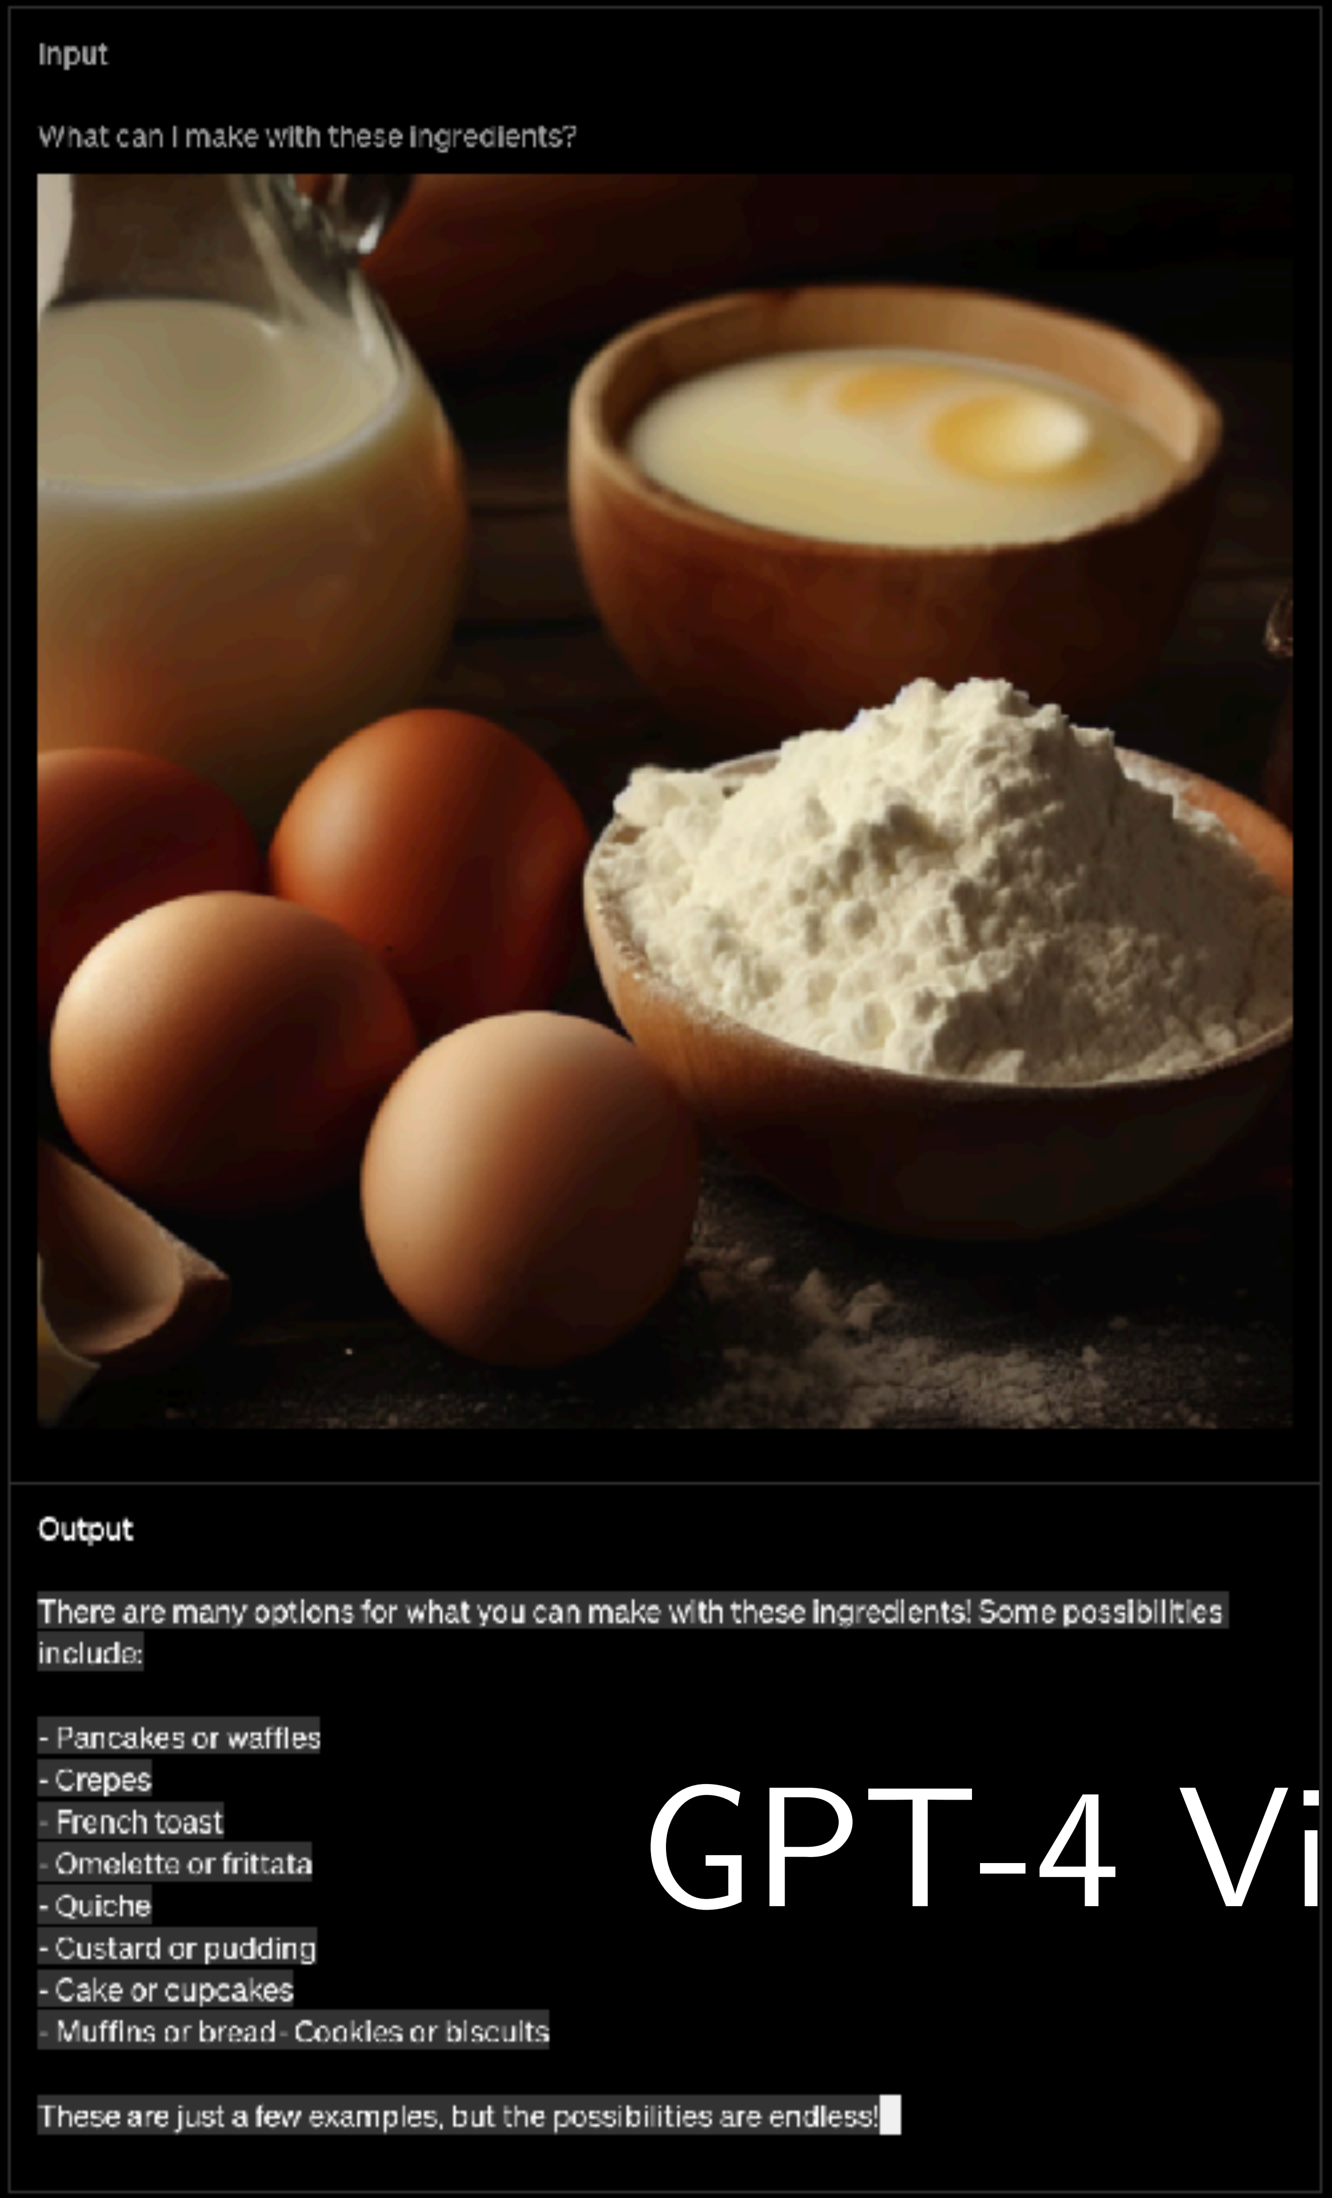

GPT-4 Vision

Video Diffusion Models

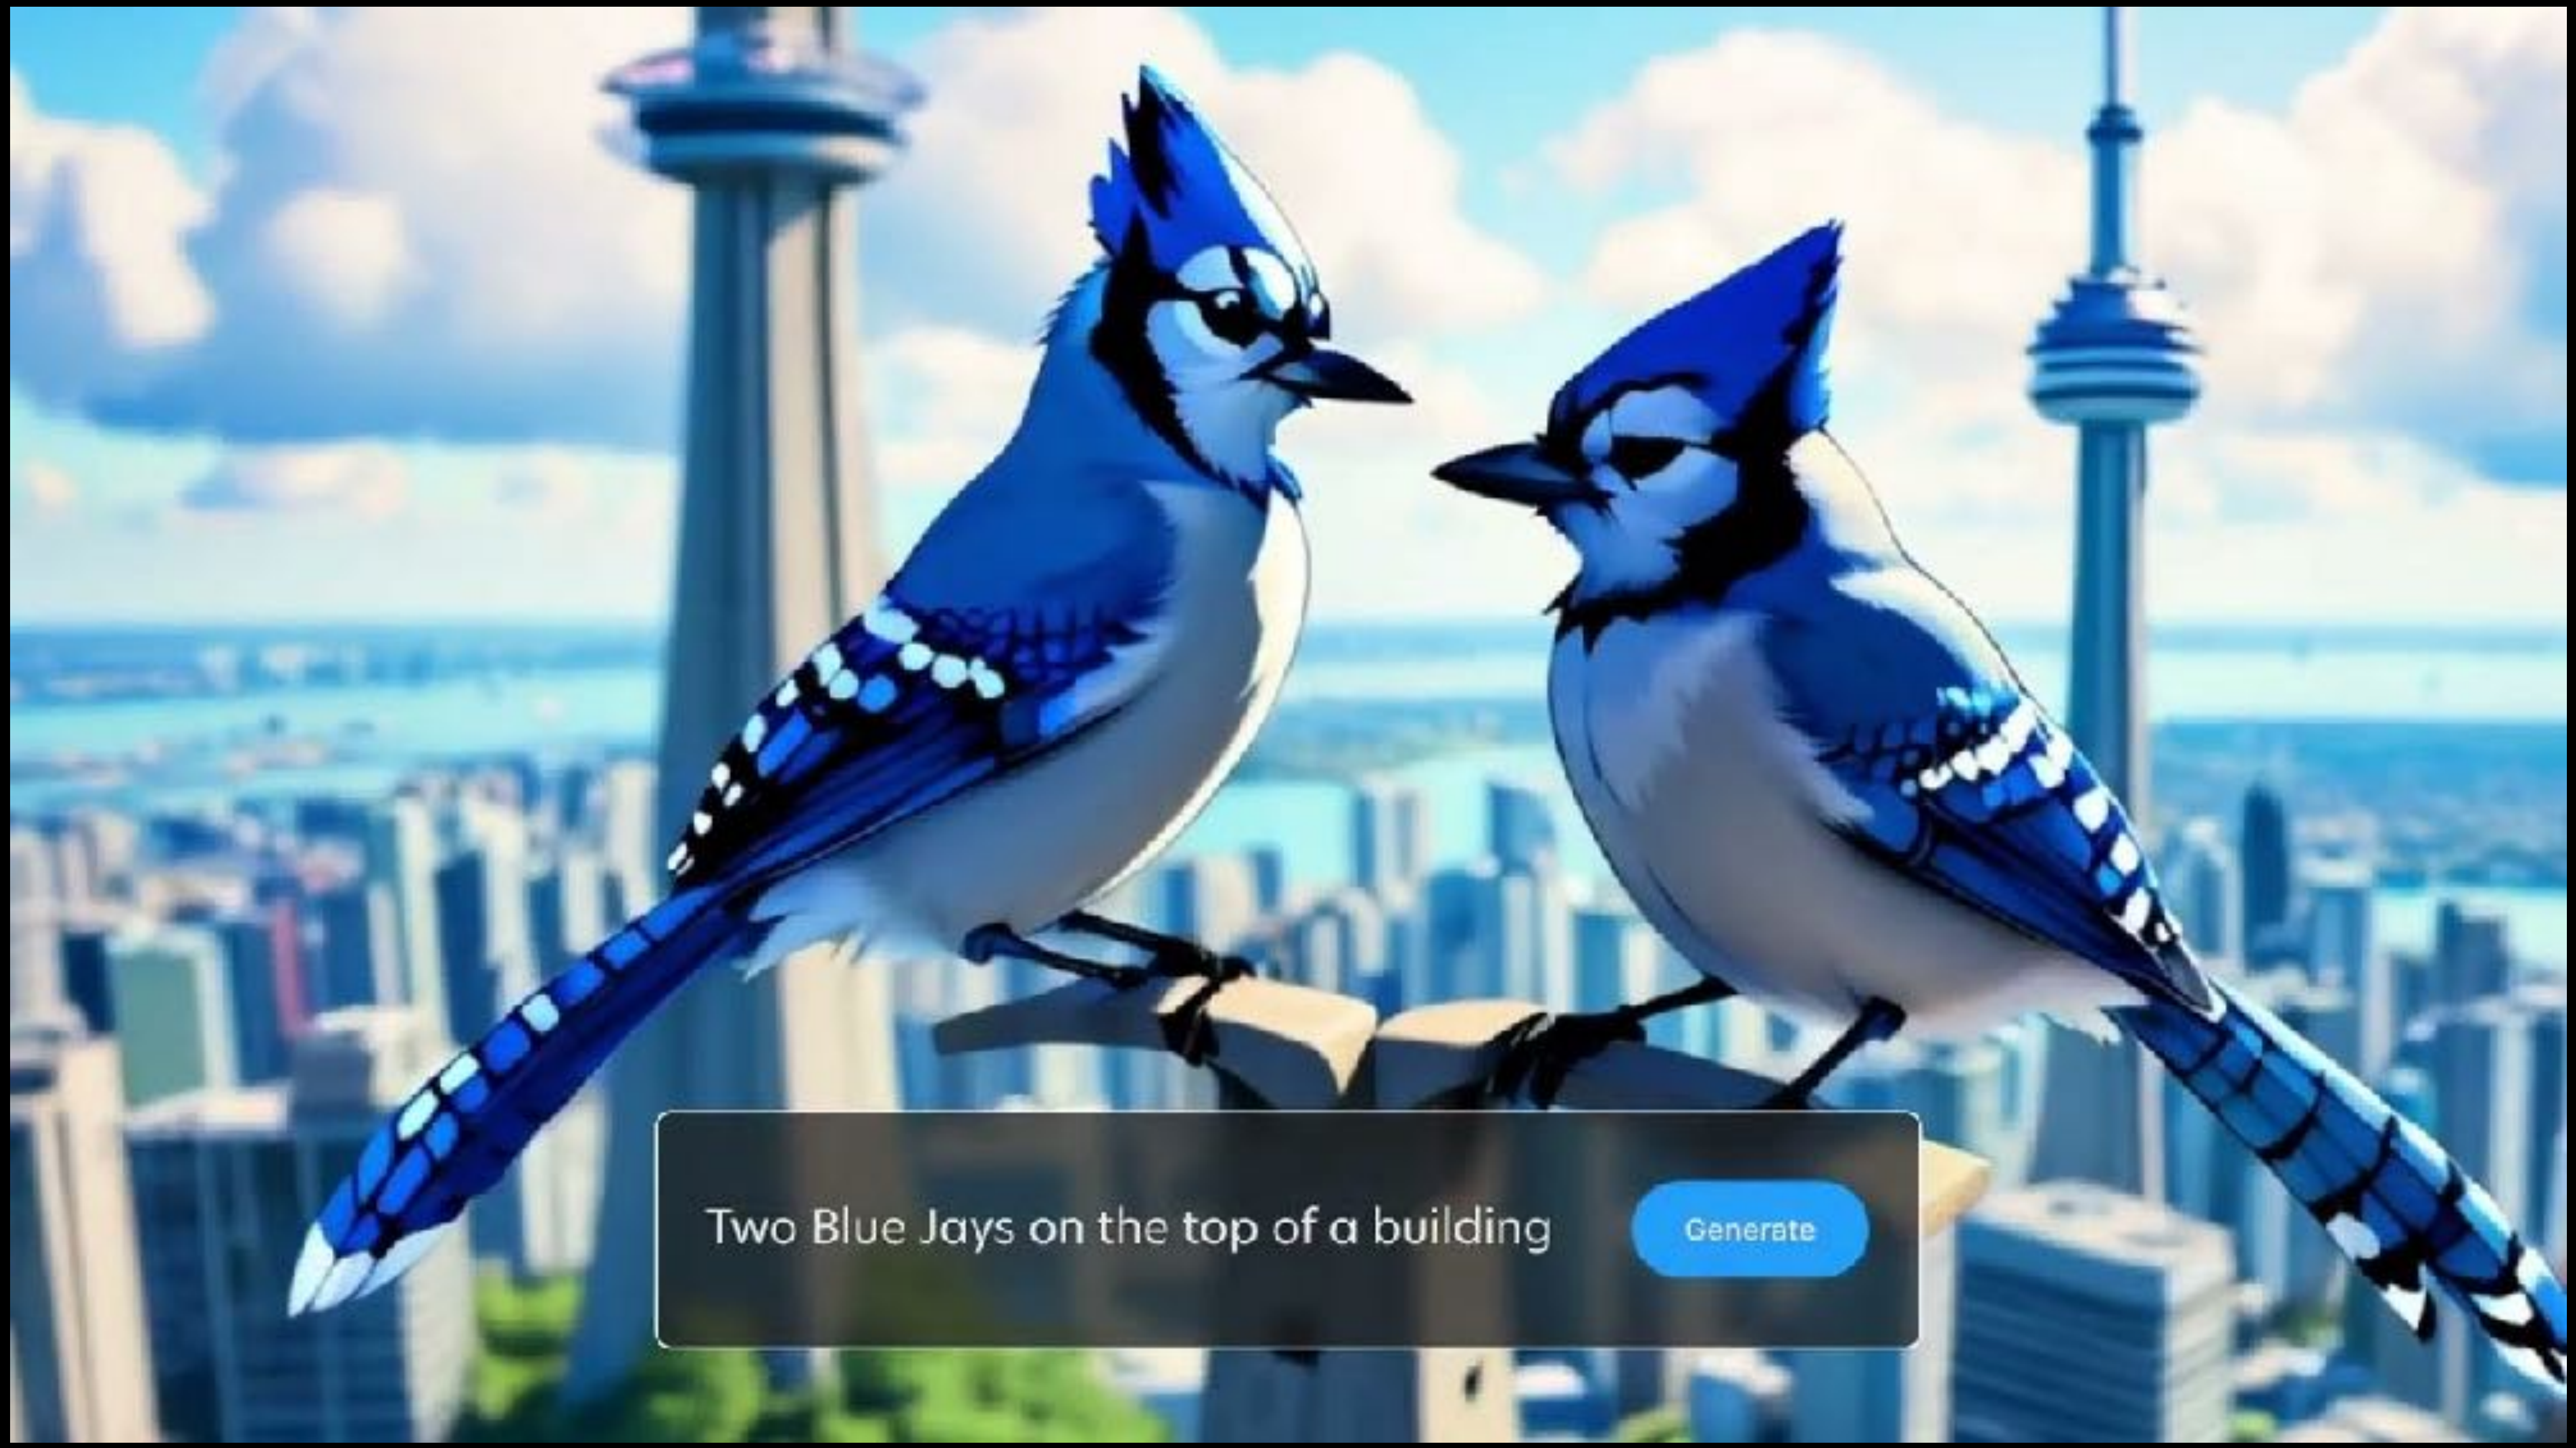

Where are the robots?

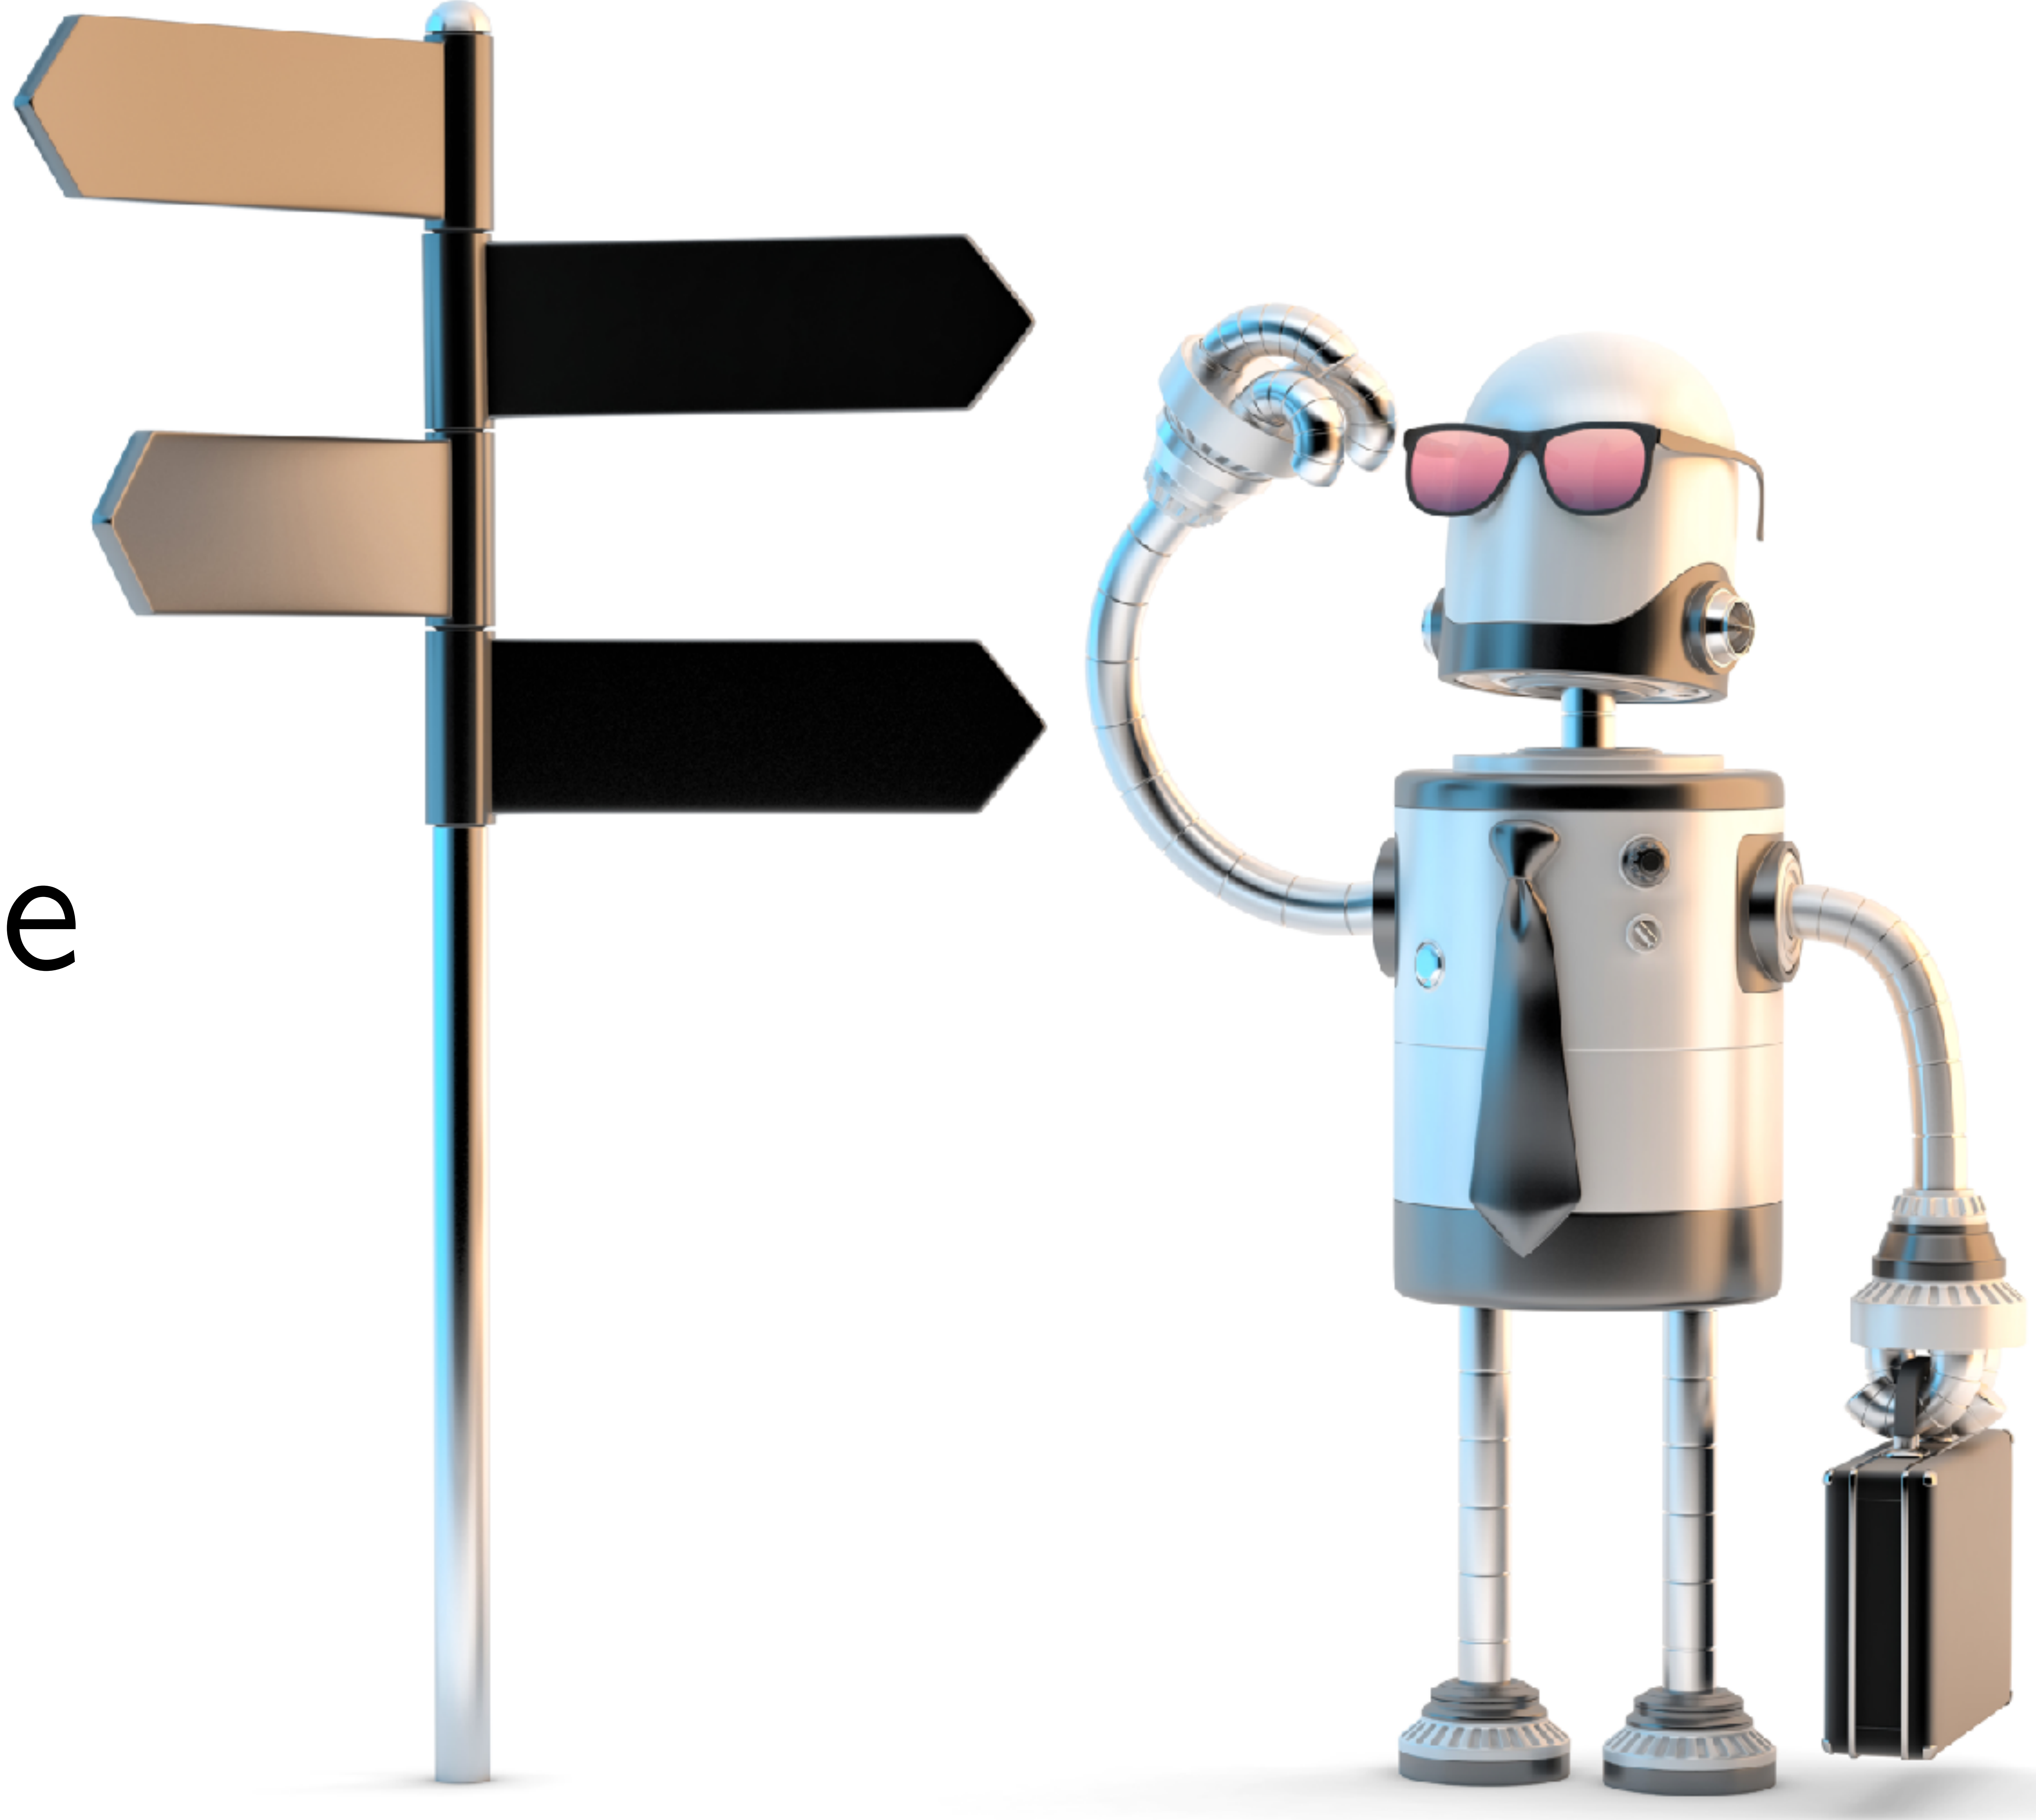

Rise of the Humanoids

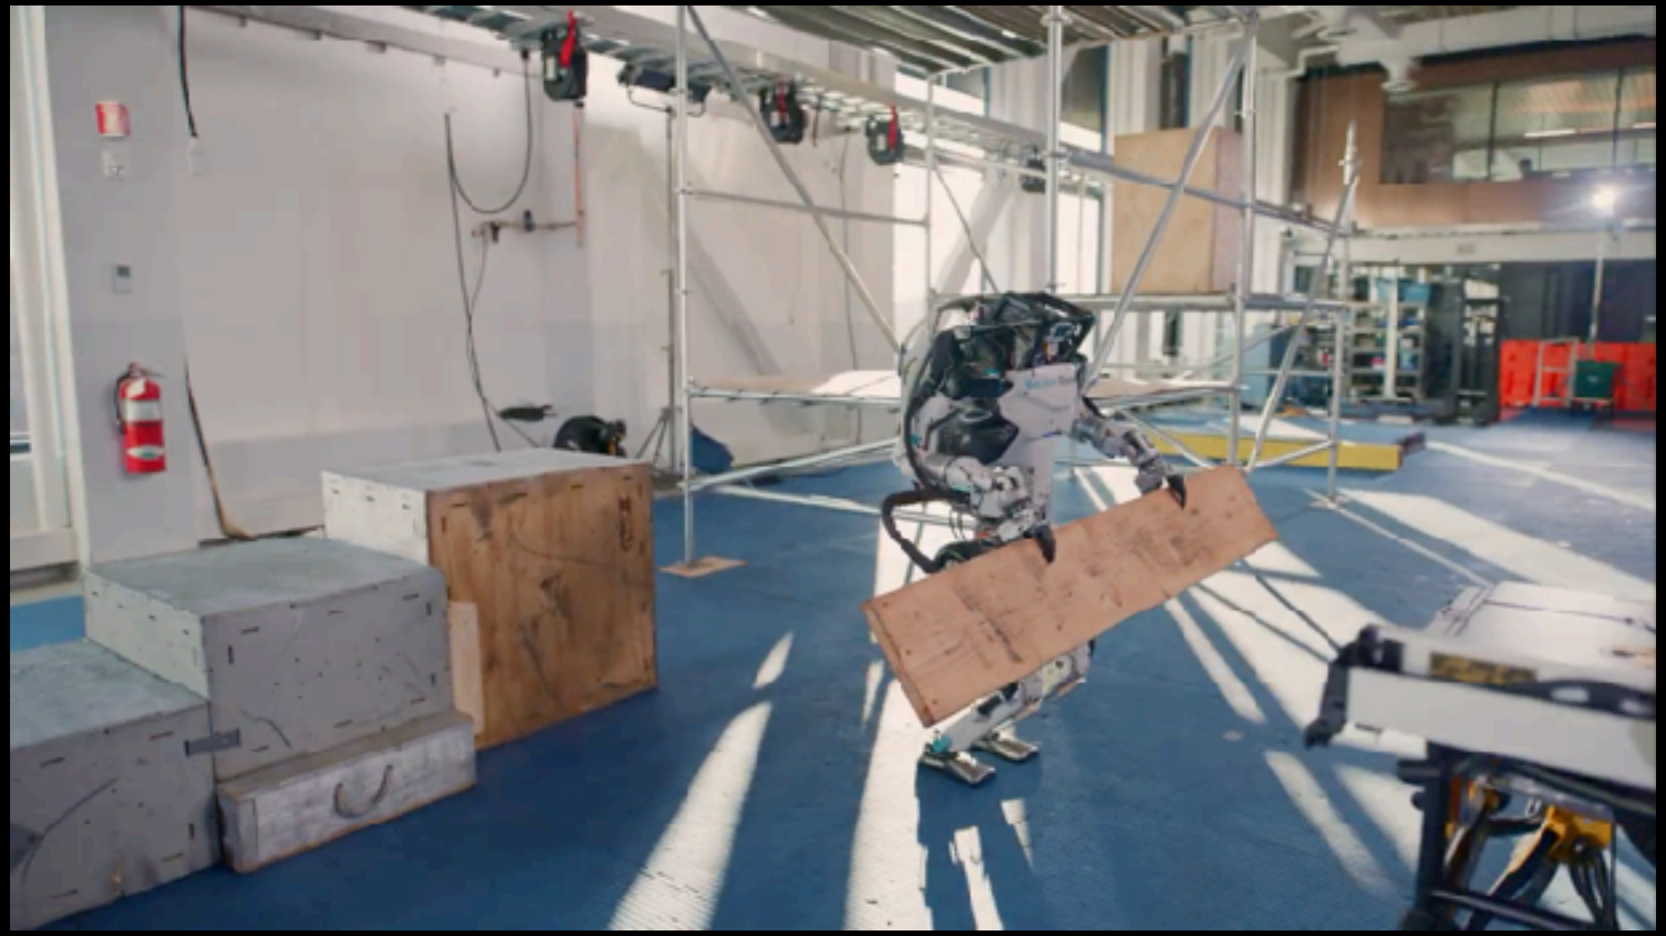

Boston Dynamics

Tesla

Agility Robotics

Figure AI

Quadrupeds going strong

Boston Dynamics SPOT

Unitree GO 2

ANYmal

Self-driving continues driverless runs

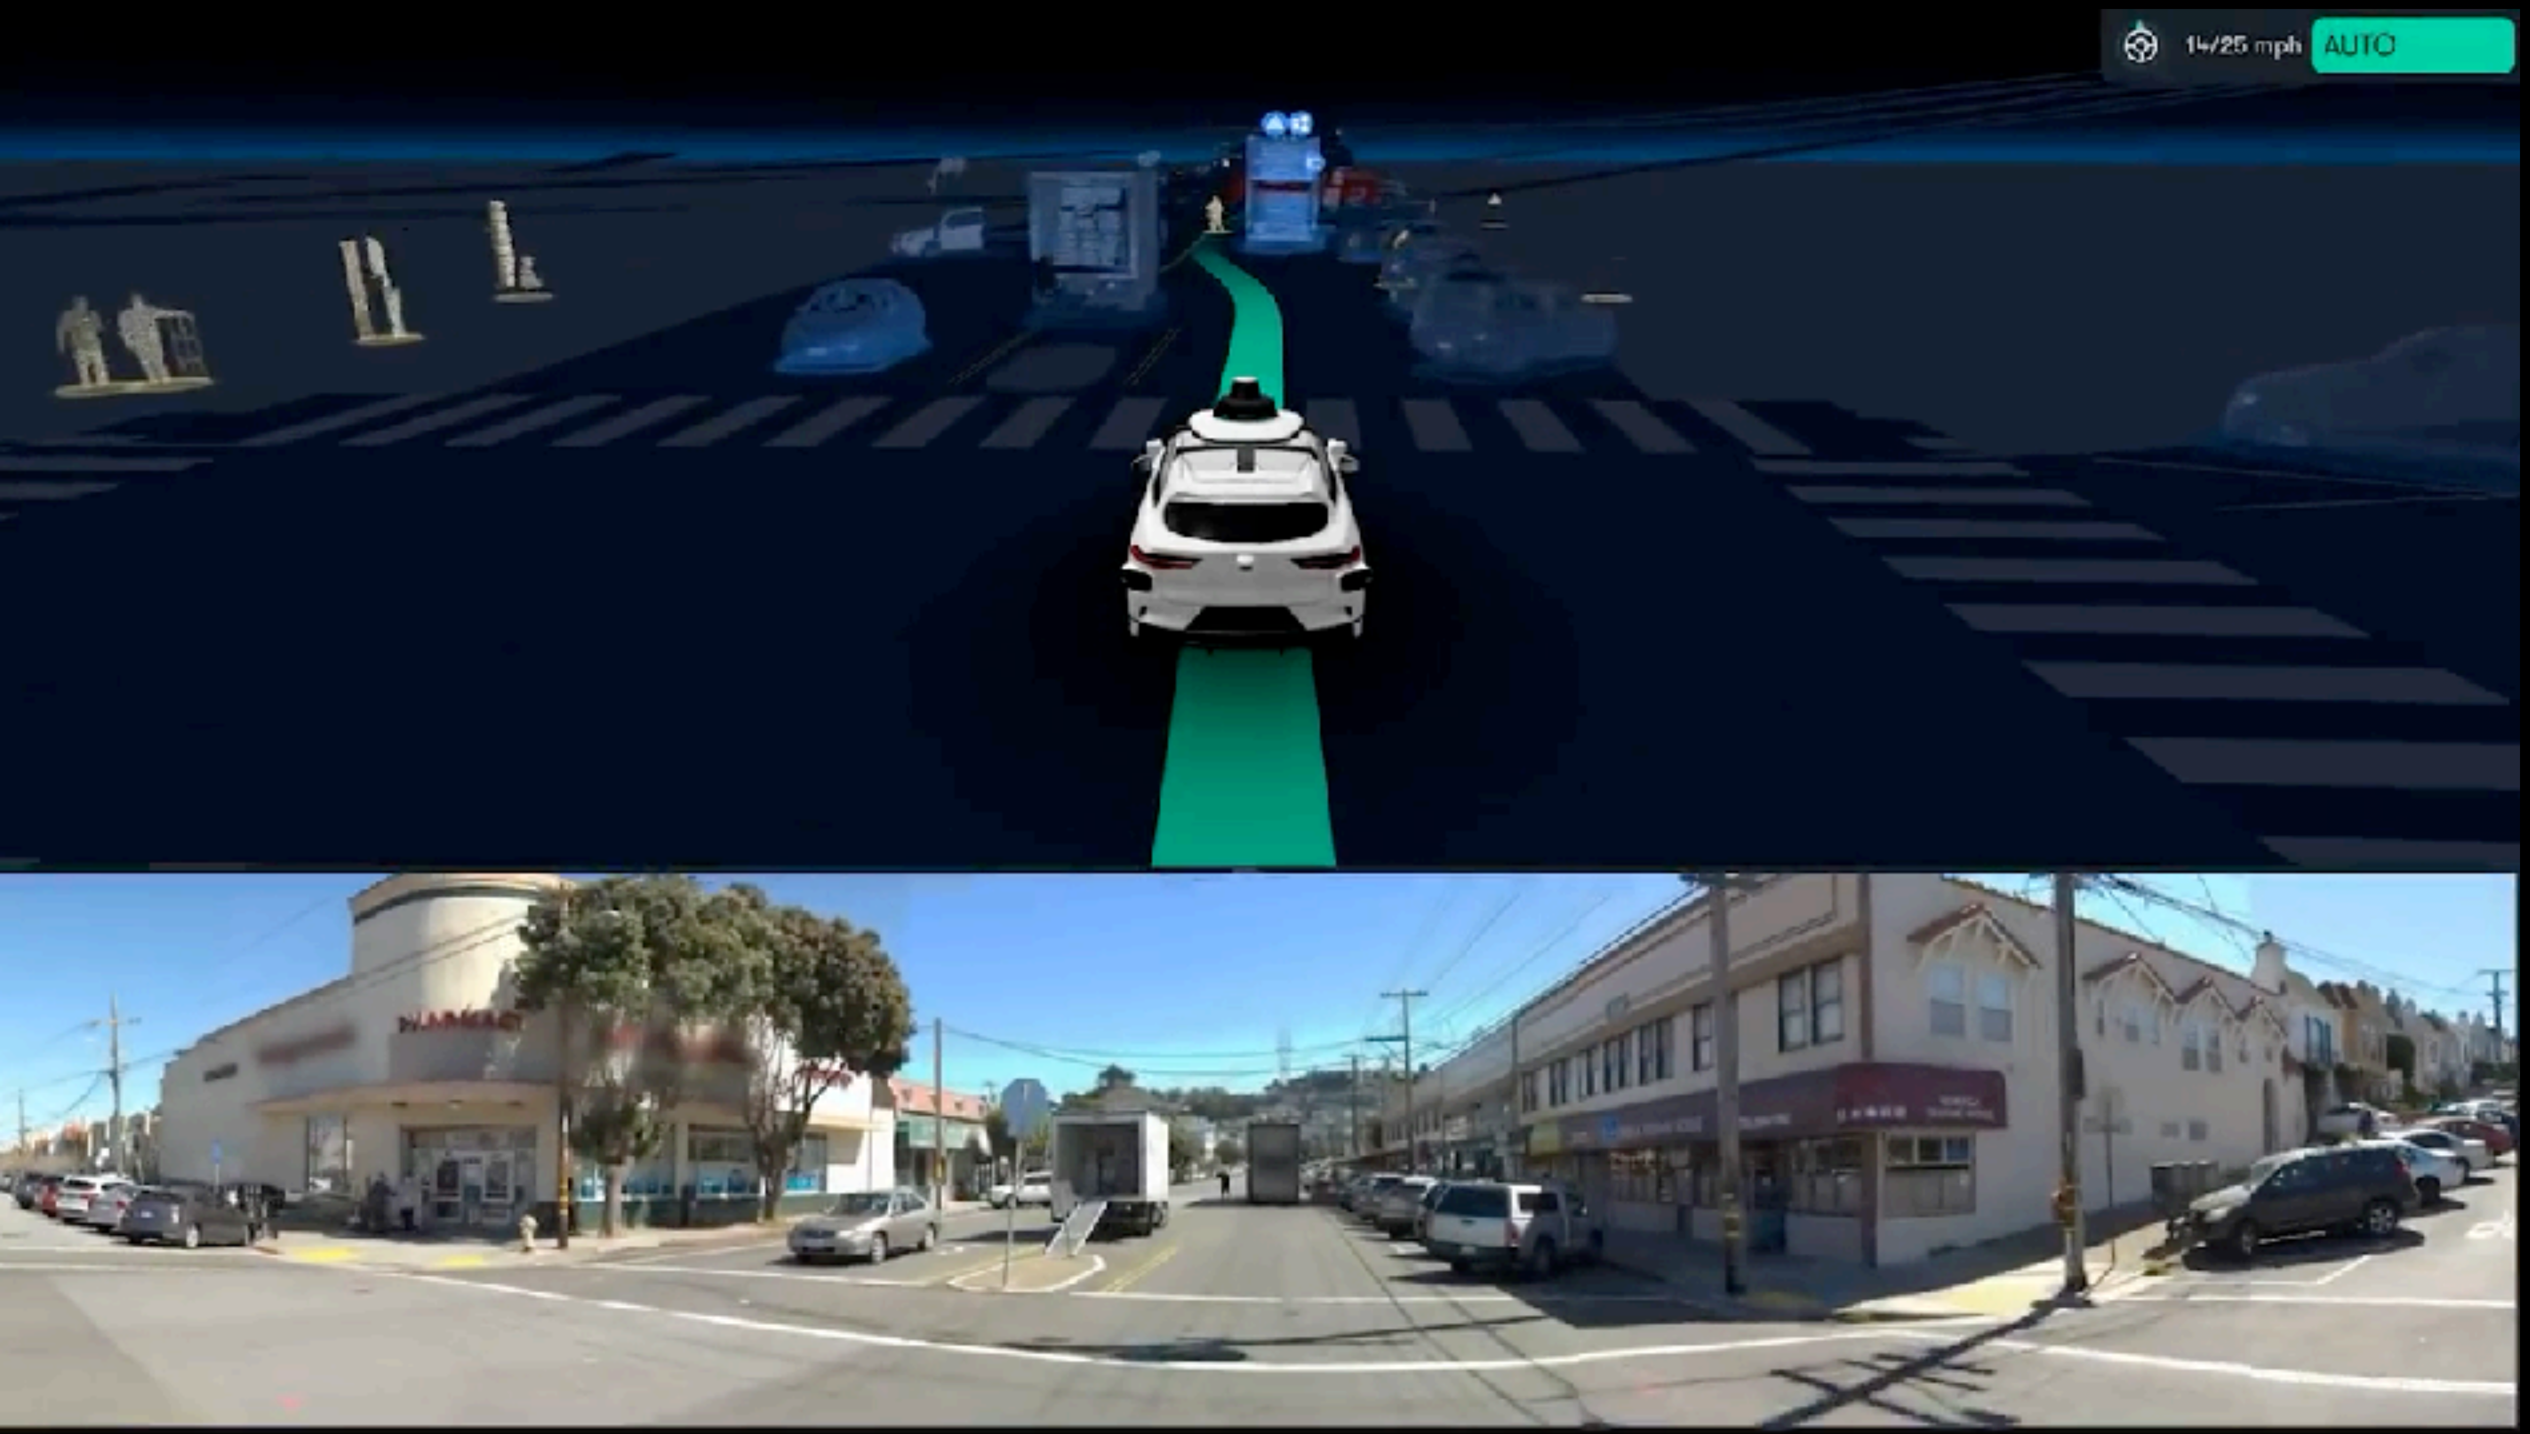

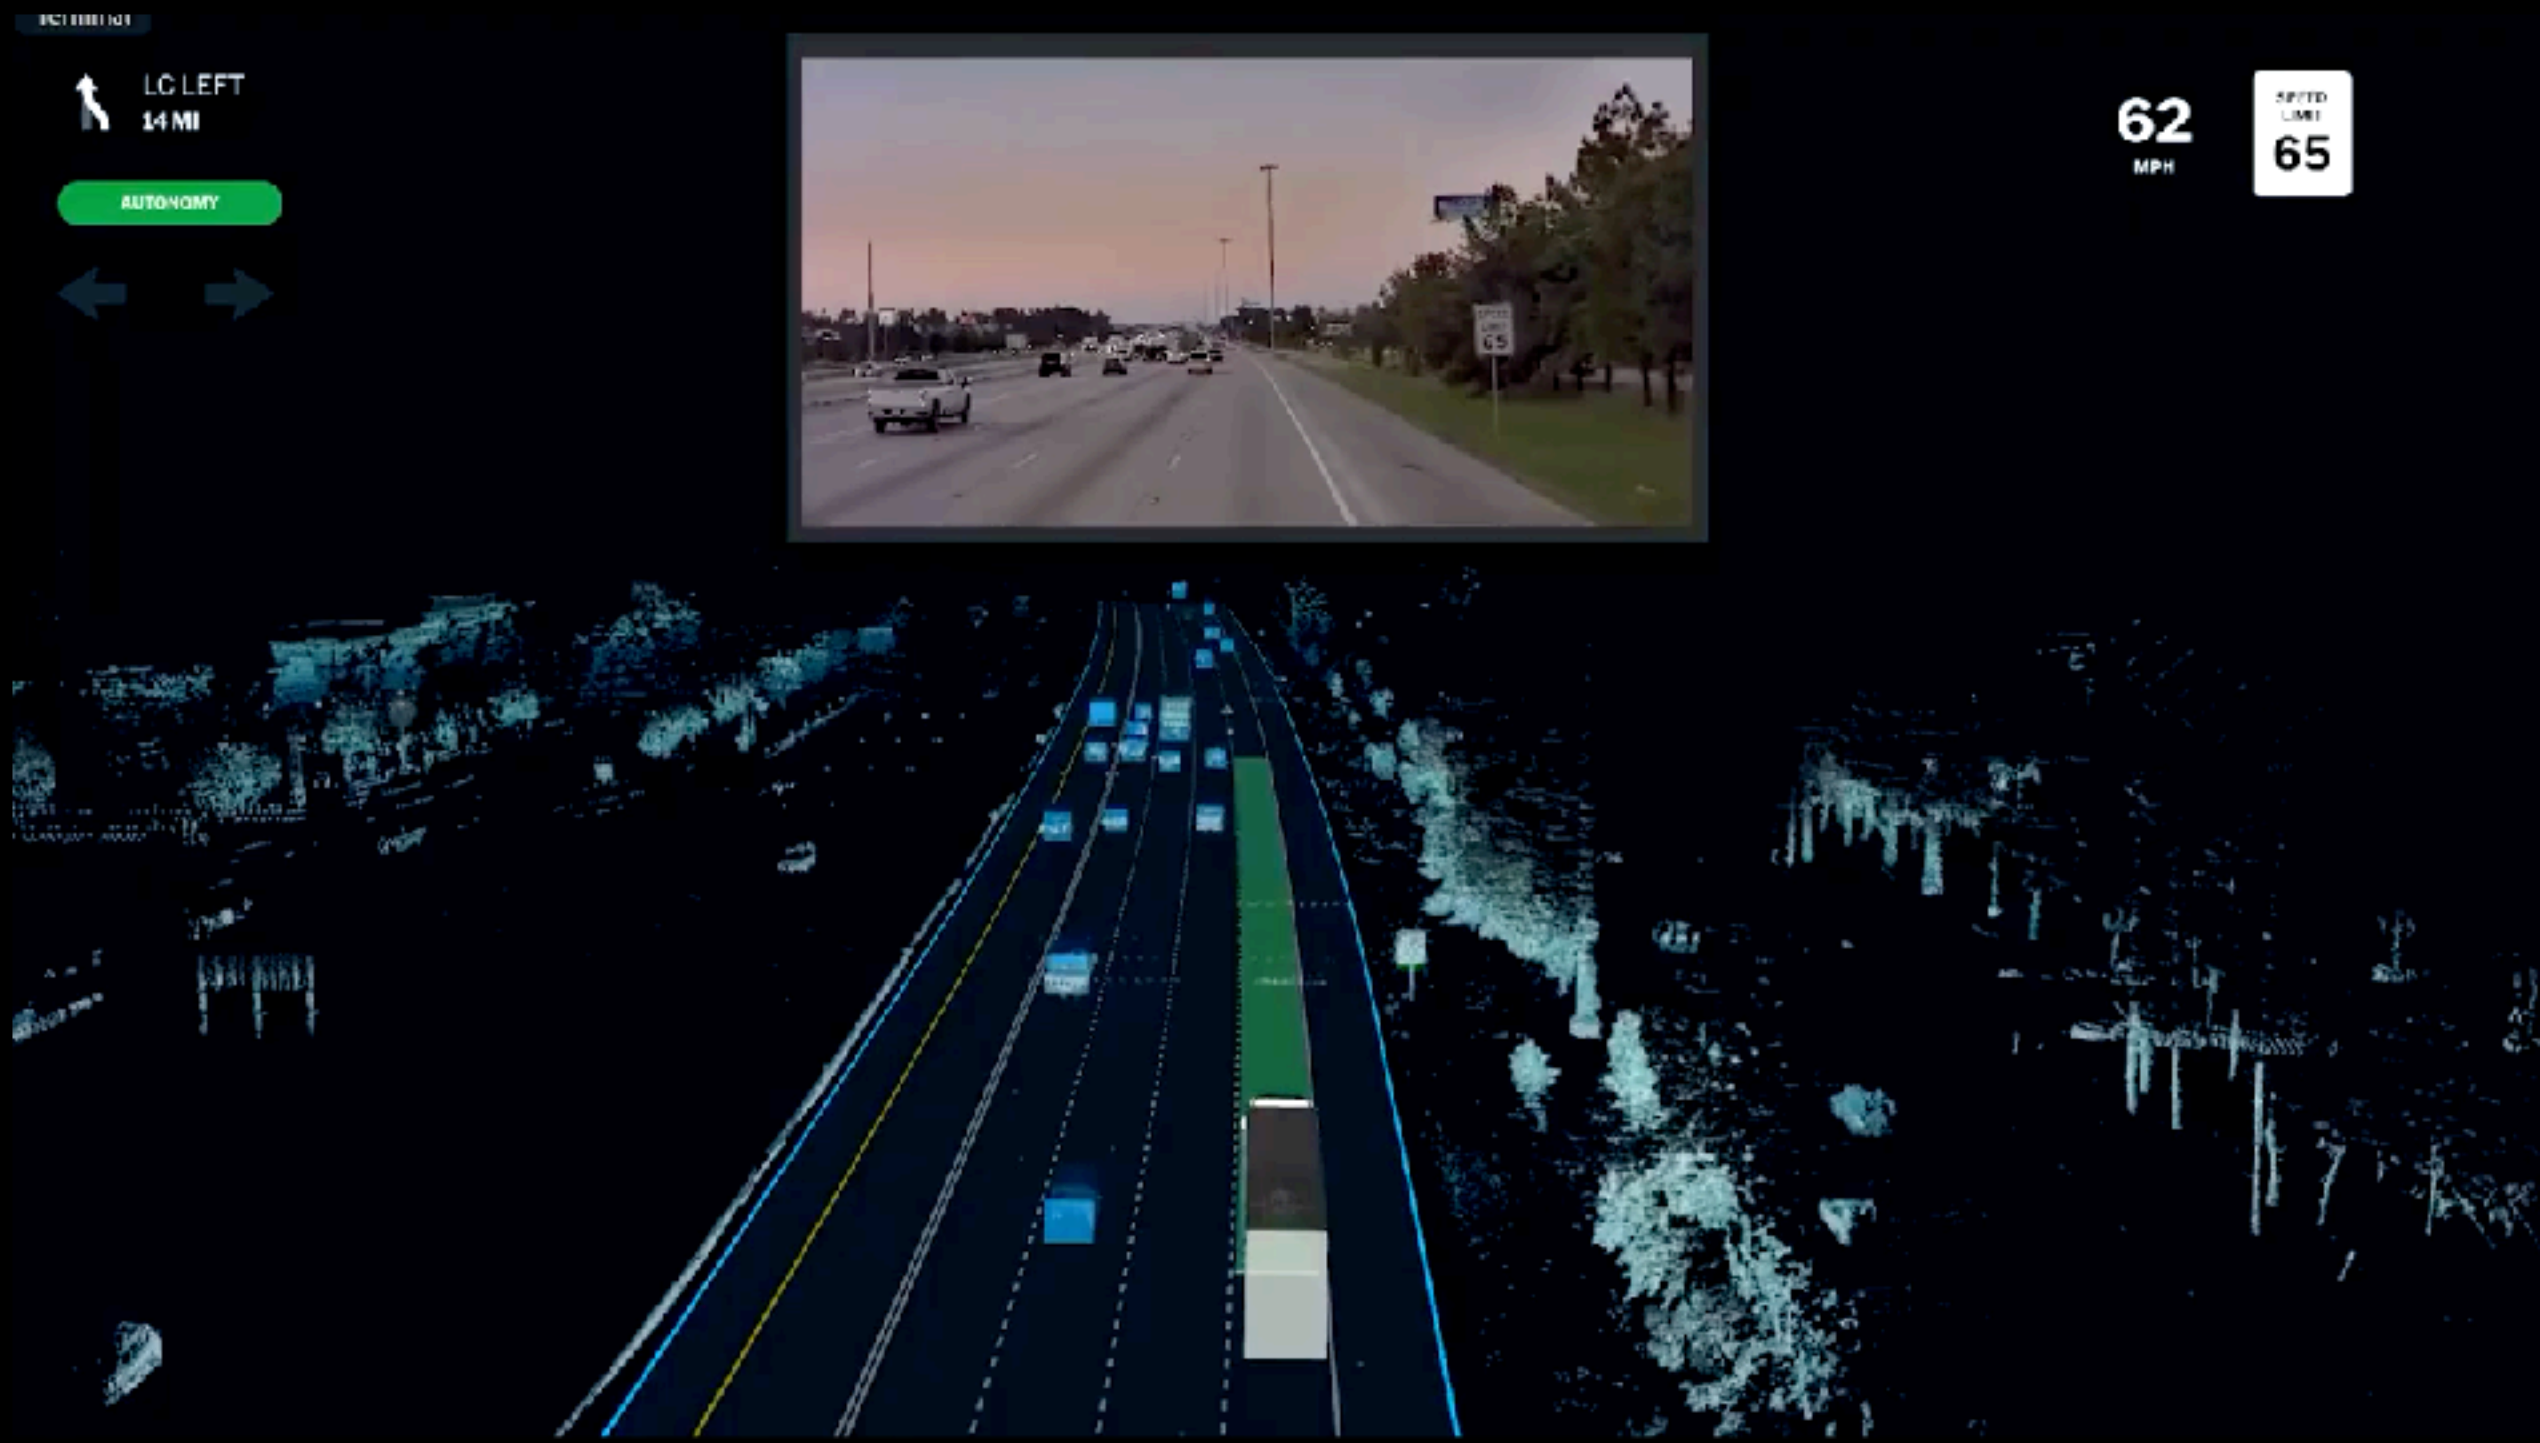

WAYMO

Aurora

But ...

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

Why?

The Problem: Real world is complex!

Why can't we just collect data and use machine learning to solve this problem?

Activity!

Build home robot apprentice to help grandma!

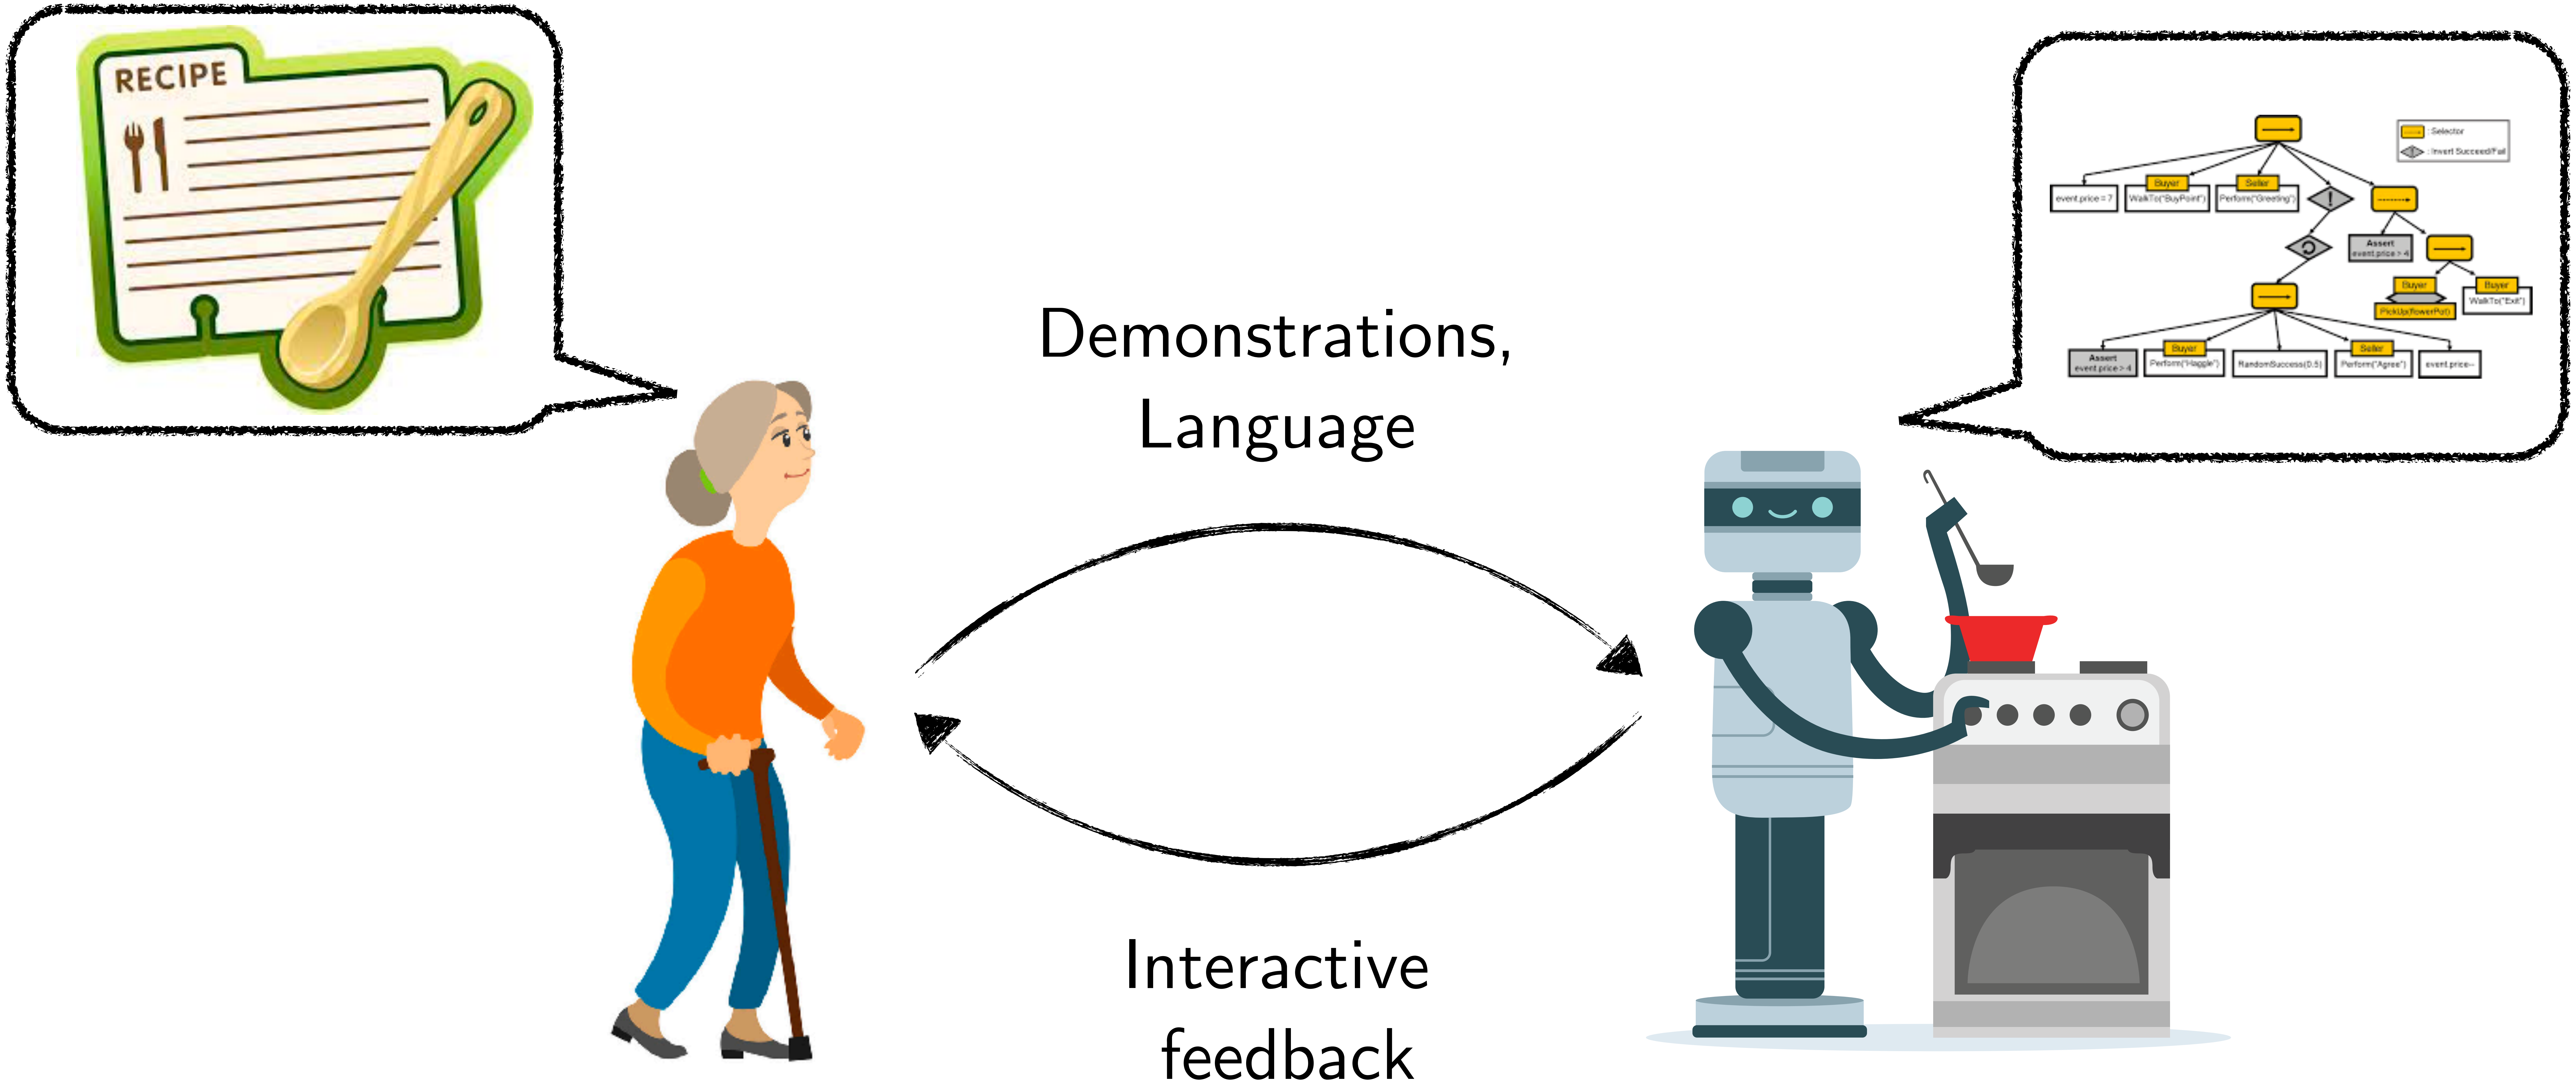

HAL & DORA

Helping Out In the Kitchen

PORTAL

Question

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

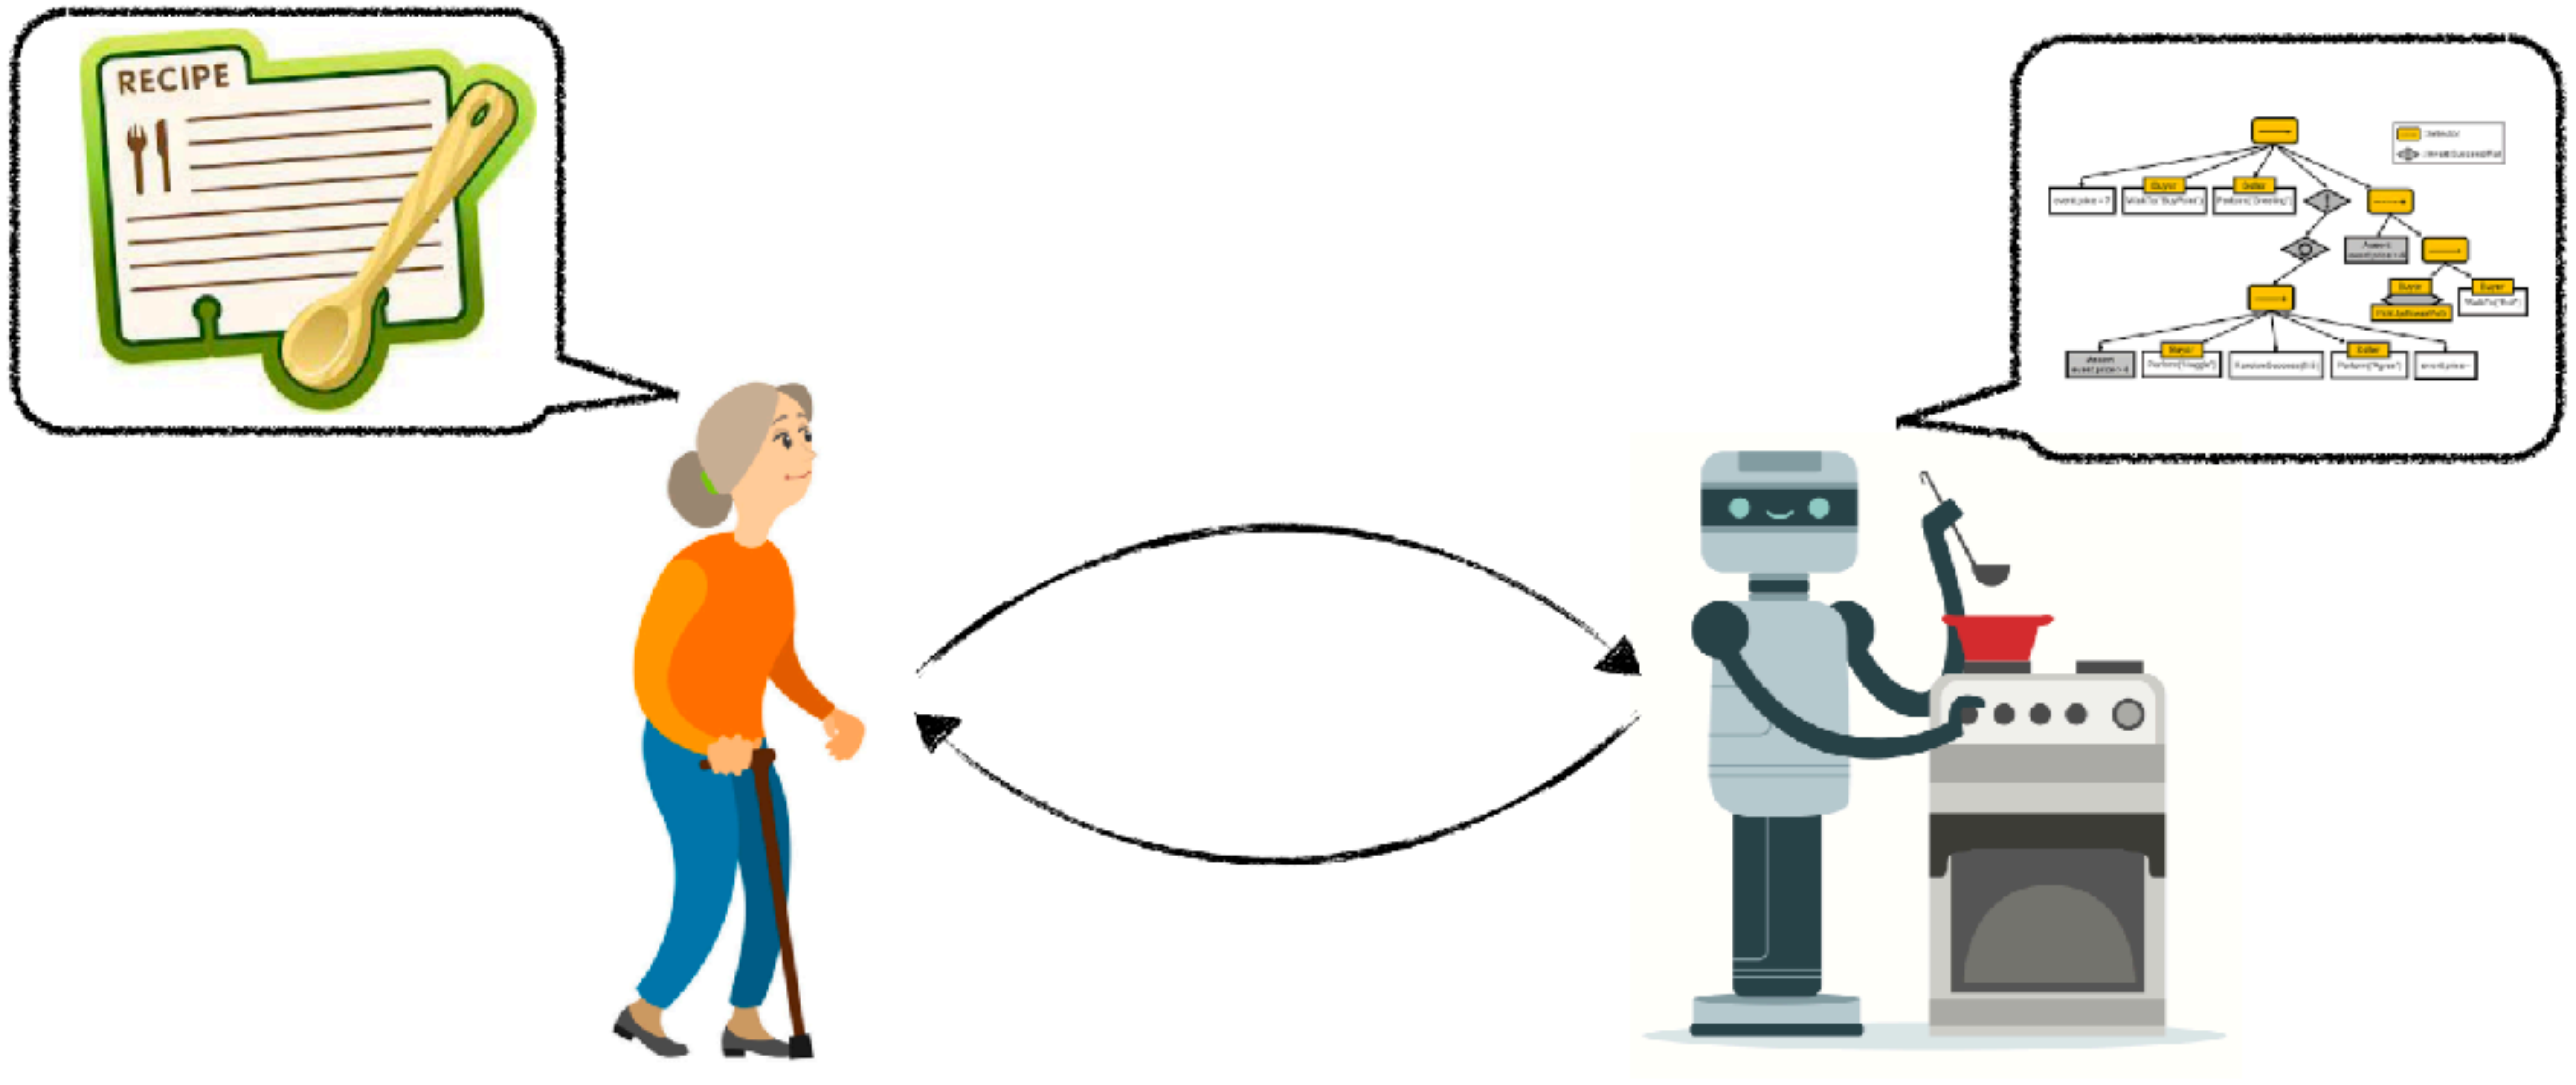

PollEv.com/sc2582

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

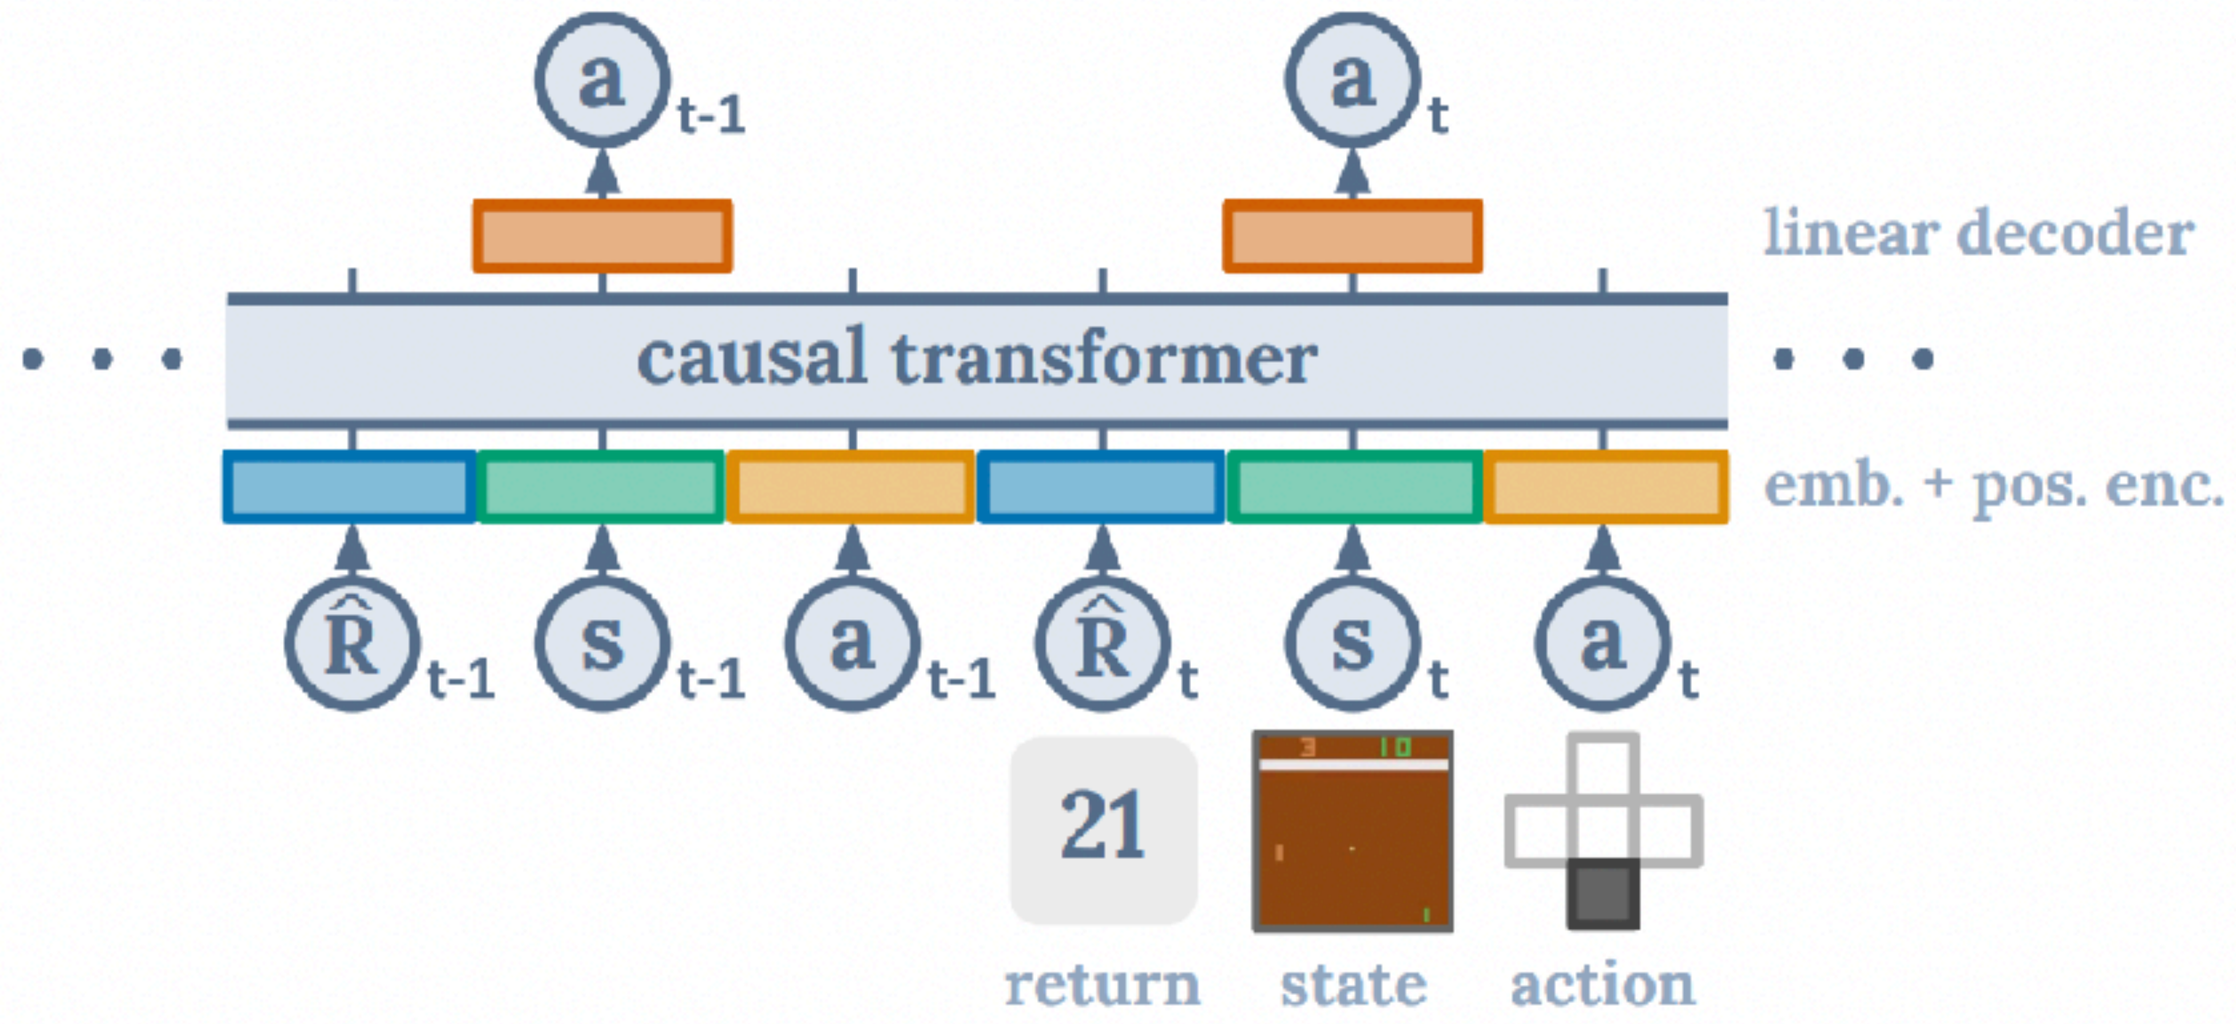

What is special about robot learning?

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

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?

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

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?

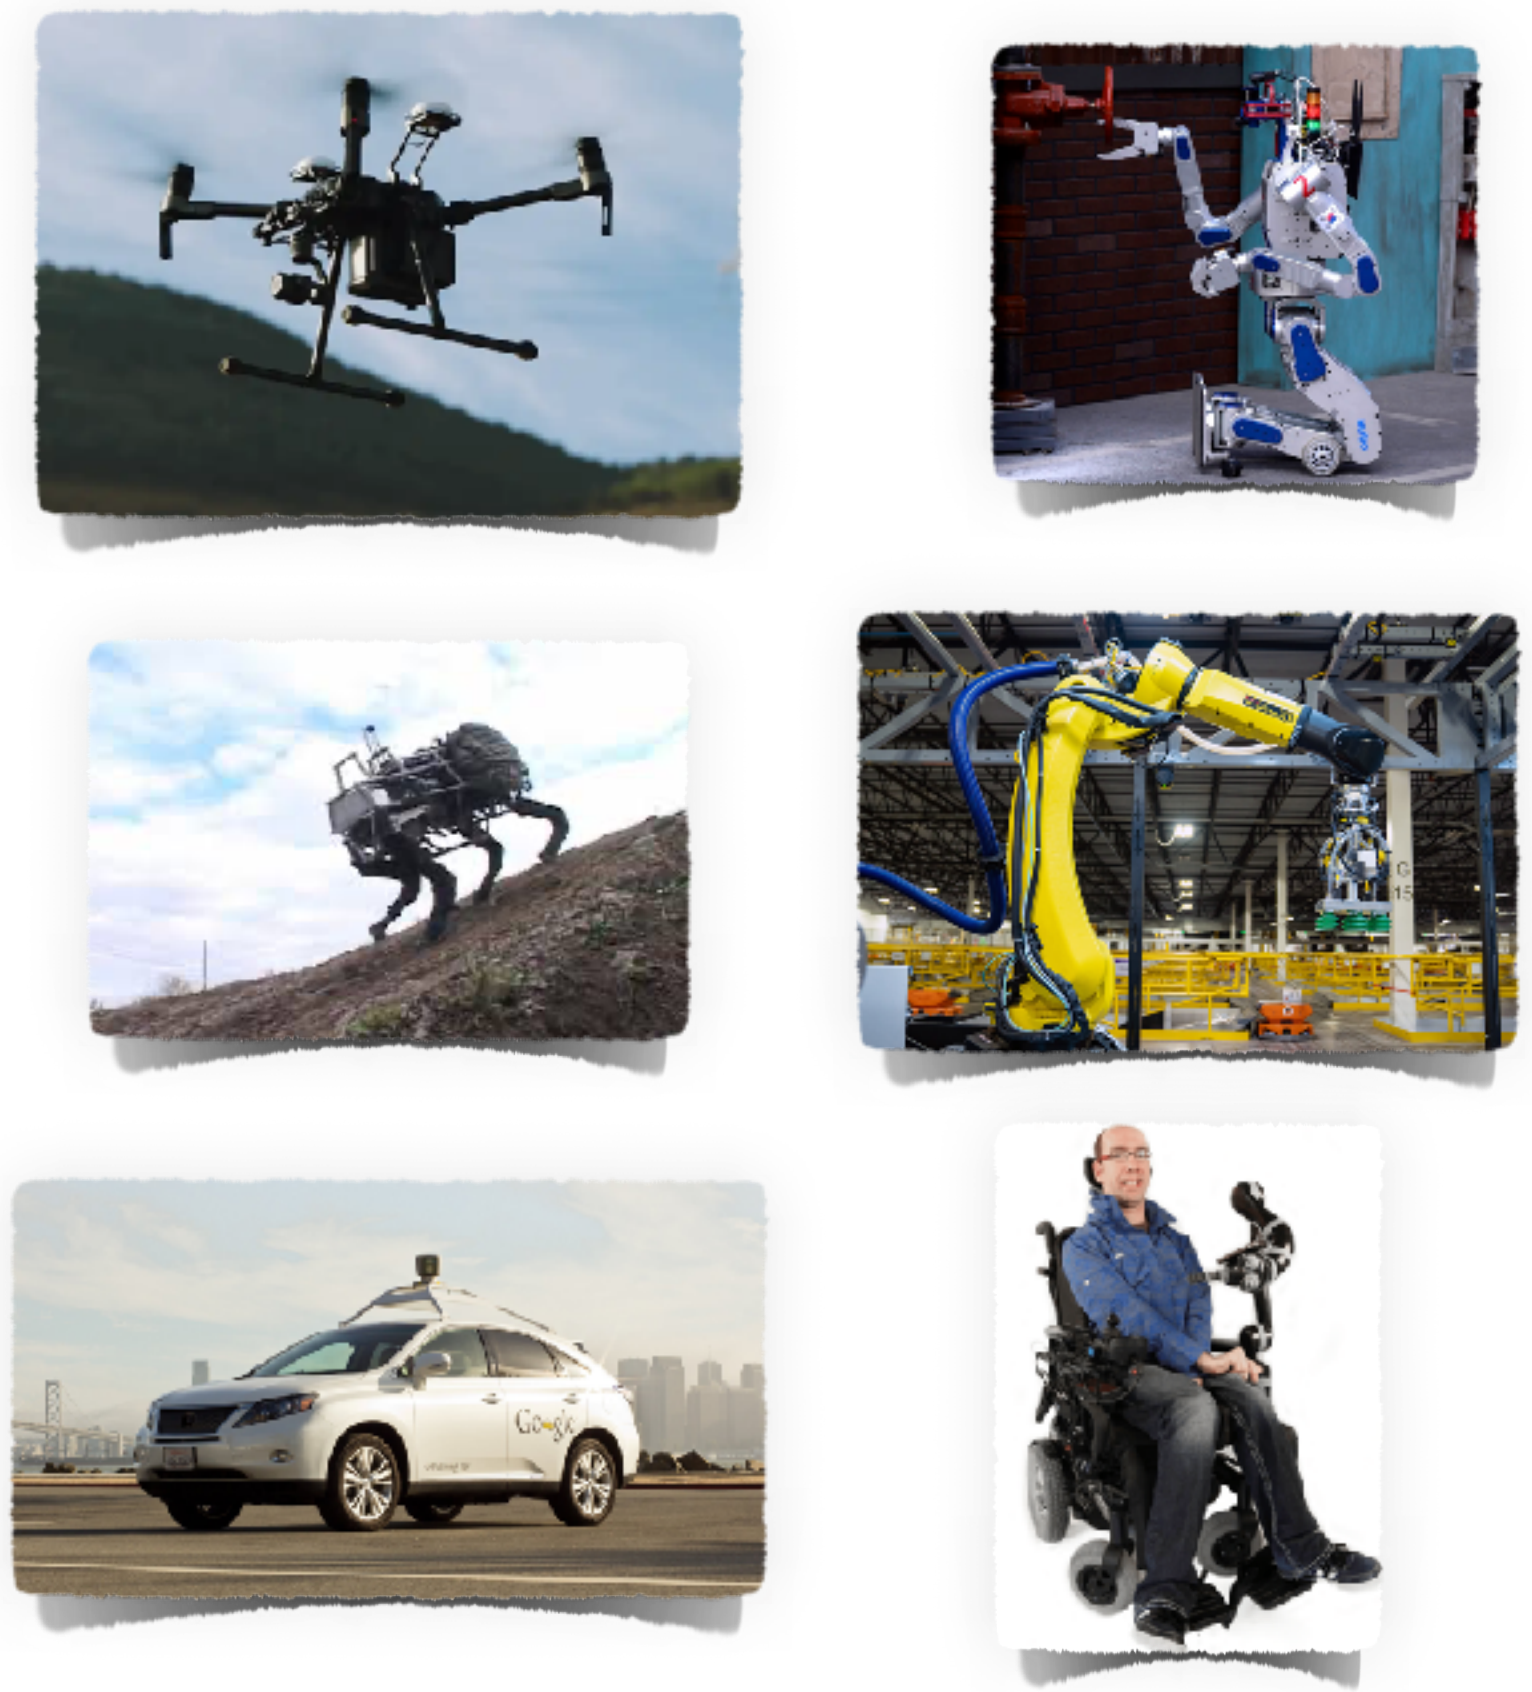

Take *any*
robot application

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

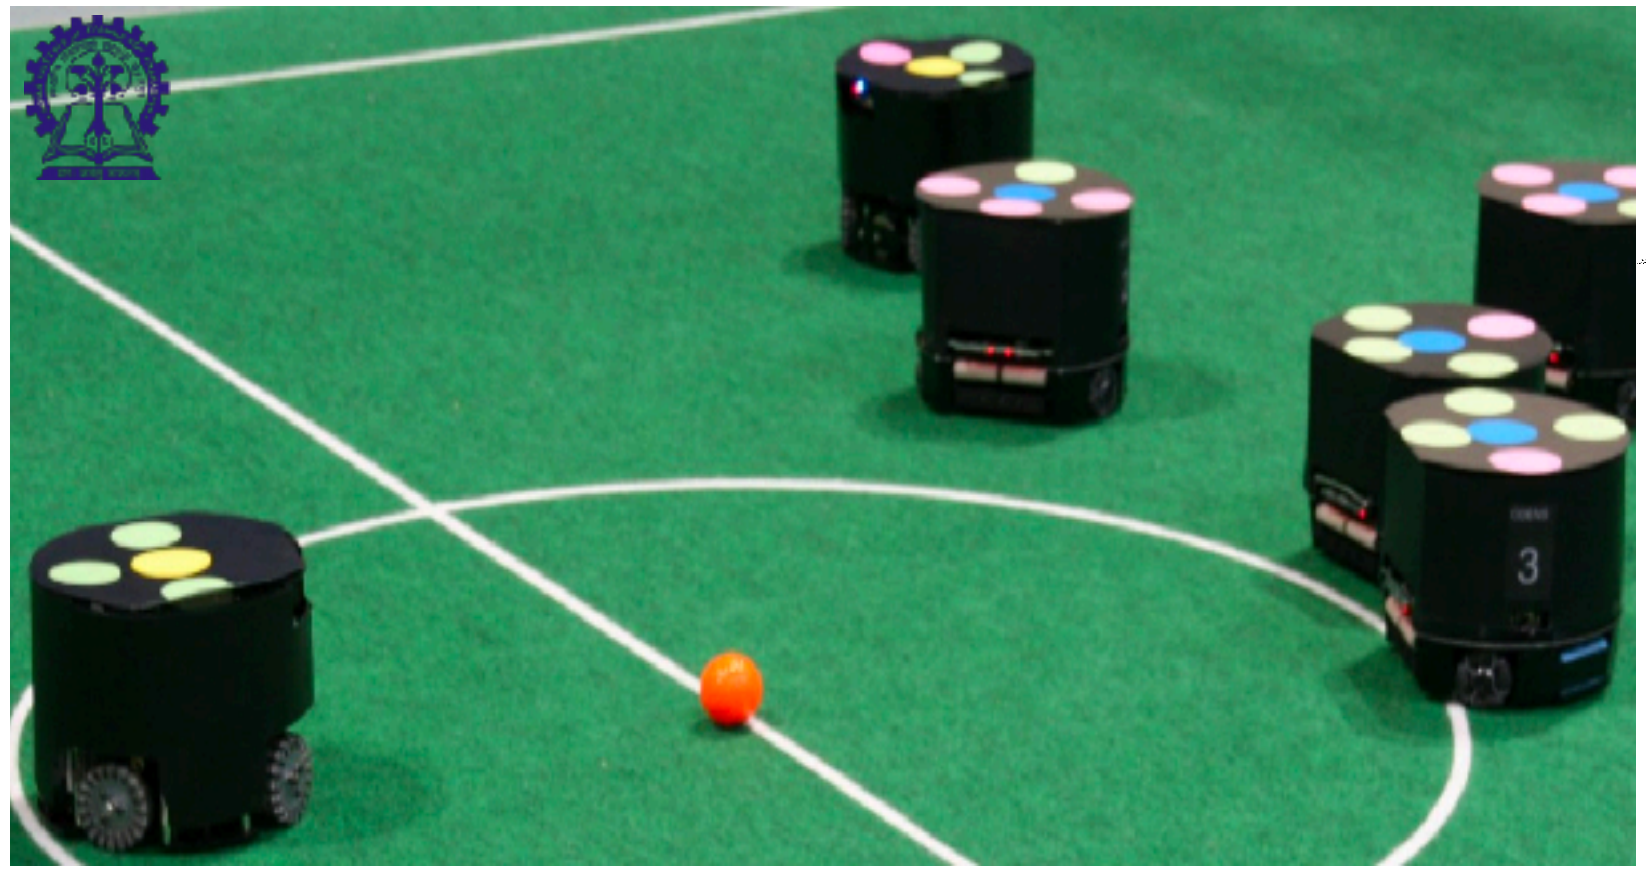

PhD

PostDoc

Research Engineer

We are PoRTaL
(People and Robots, Teaching and Learning)

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

PORTAL

Belonging

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

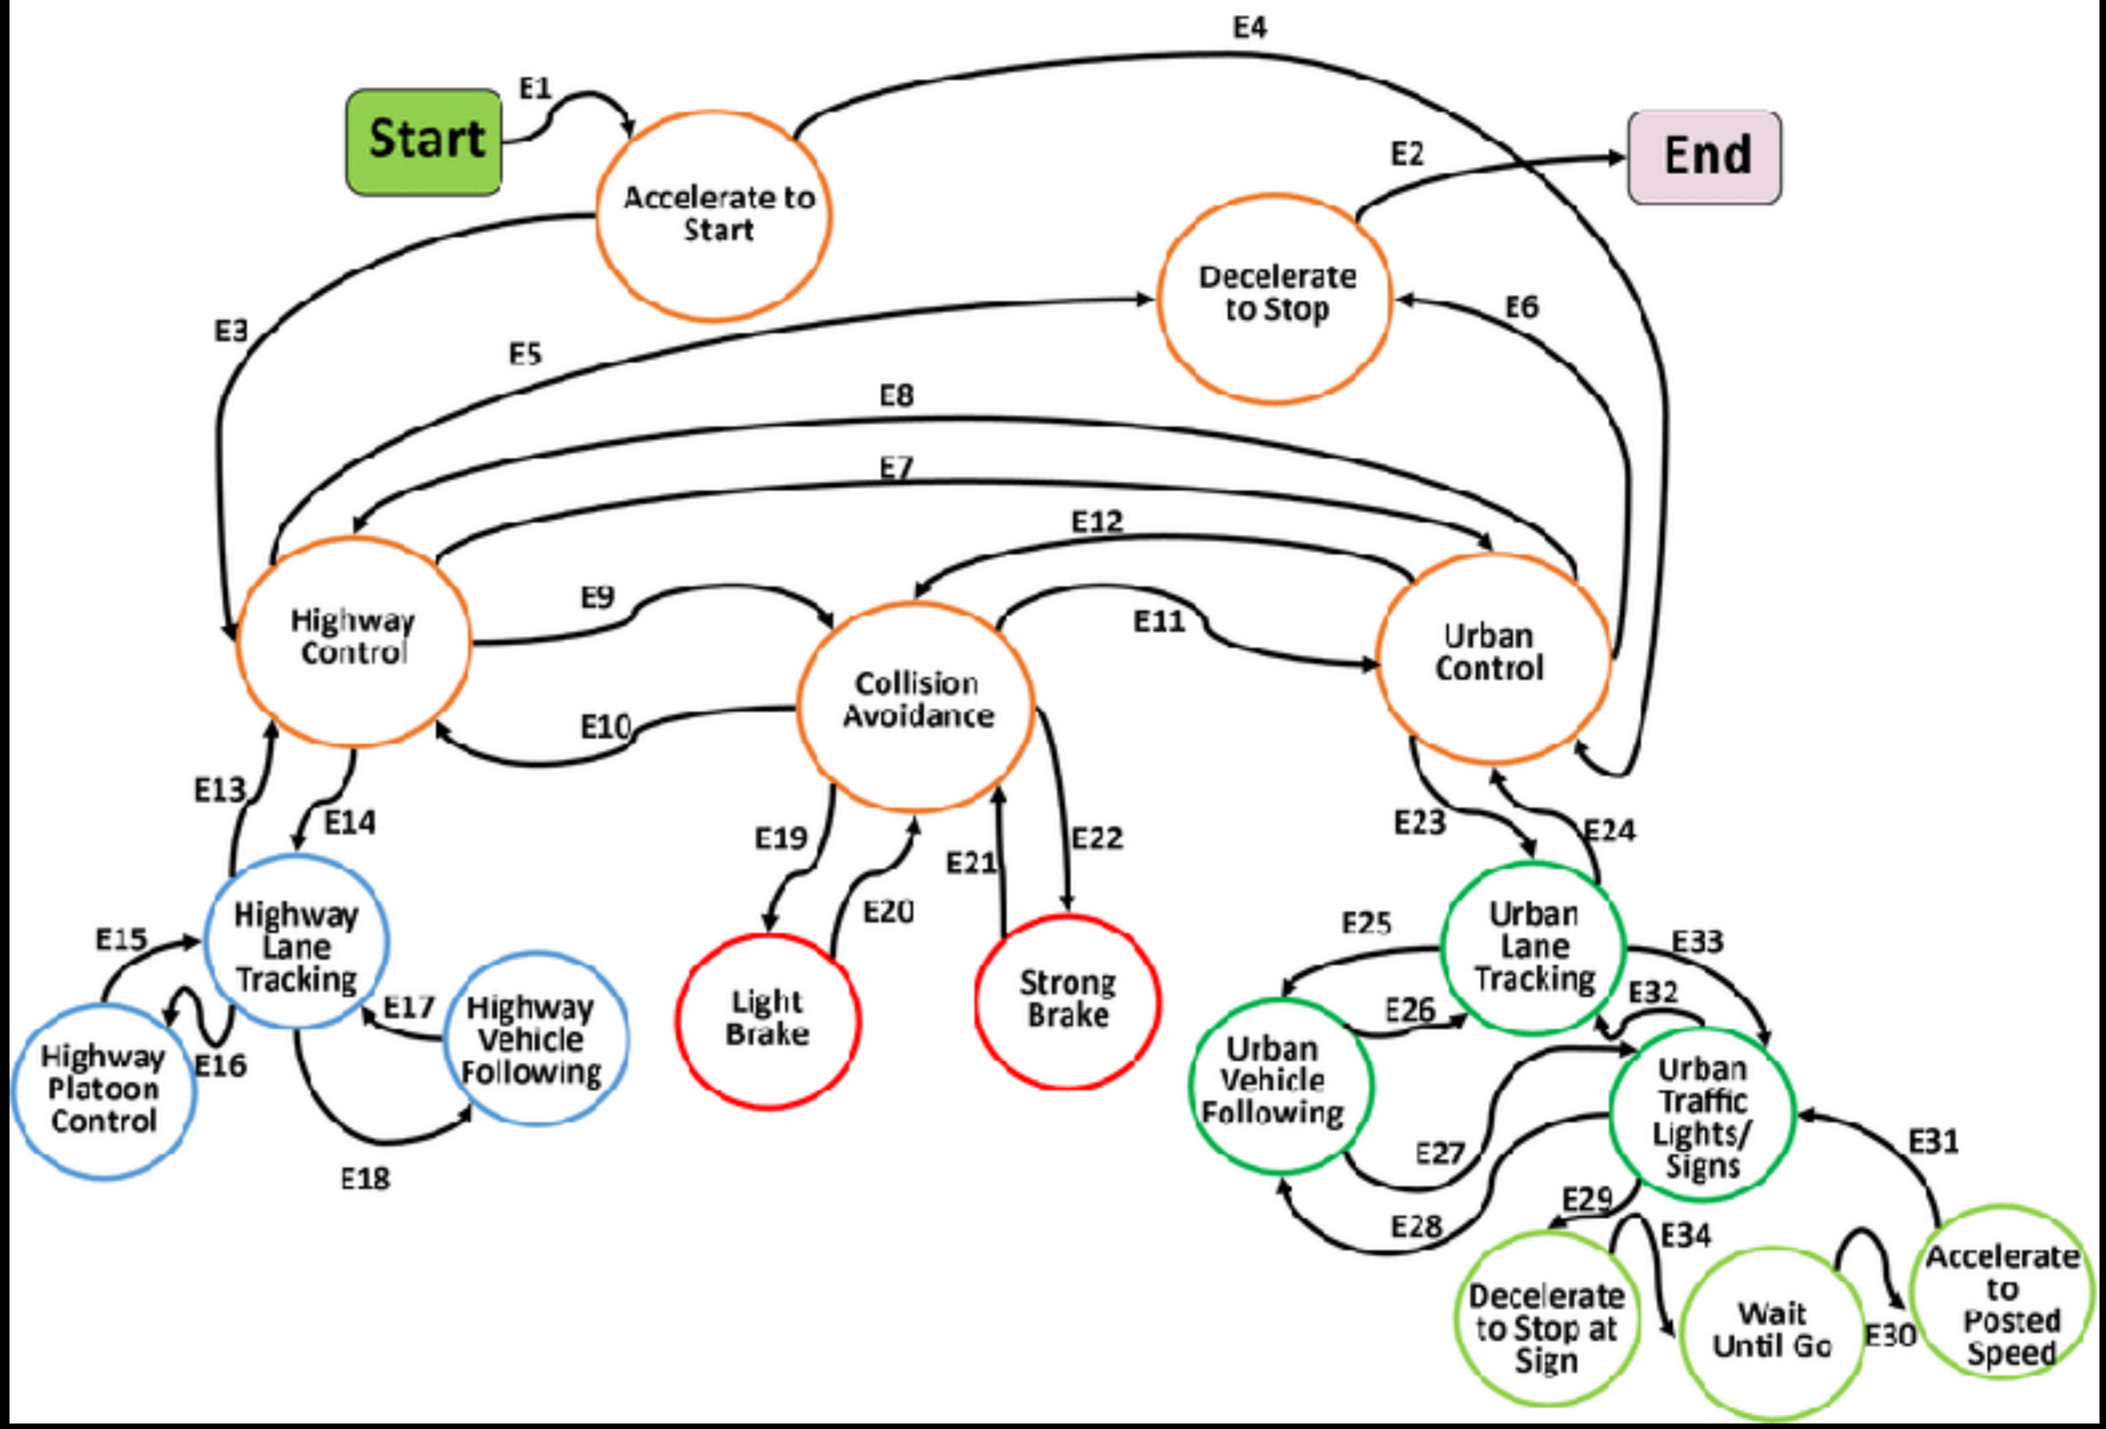

Hand-engineered rules of driving

Limited use of machine learning

Software that fundamentally could not scale over time

Self-driving 2.0

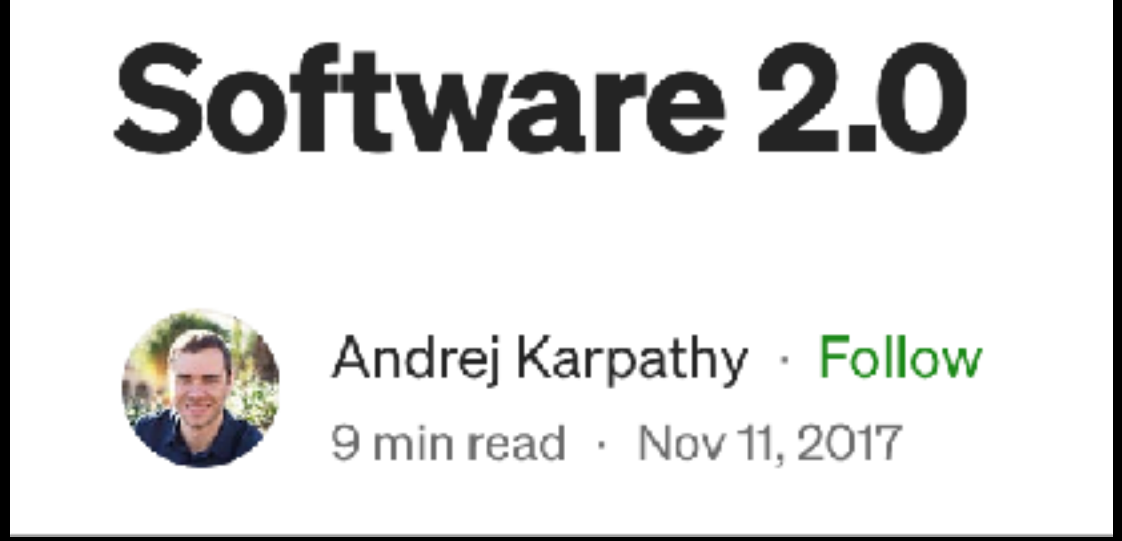

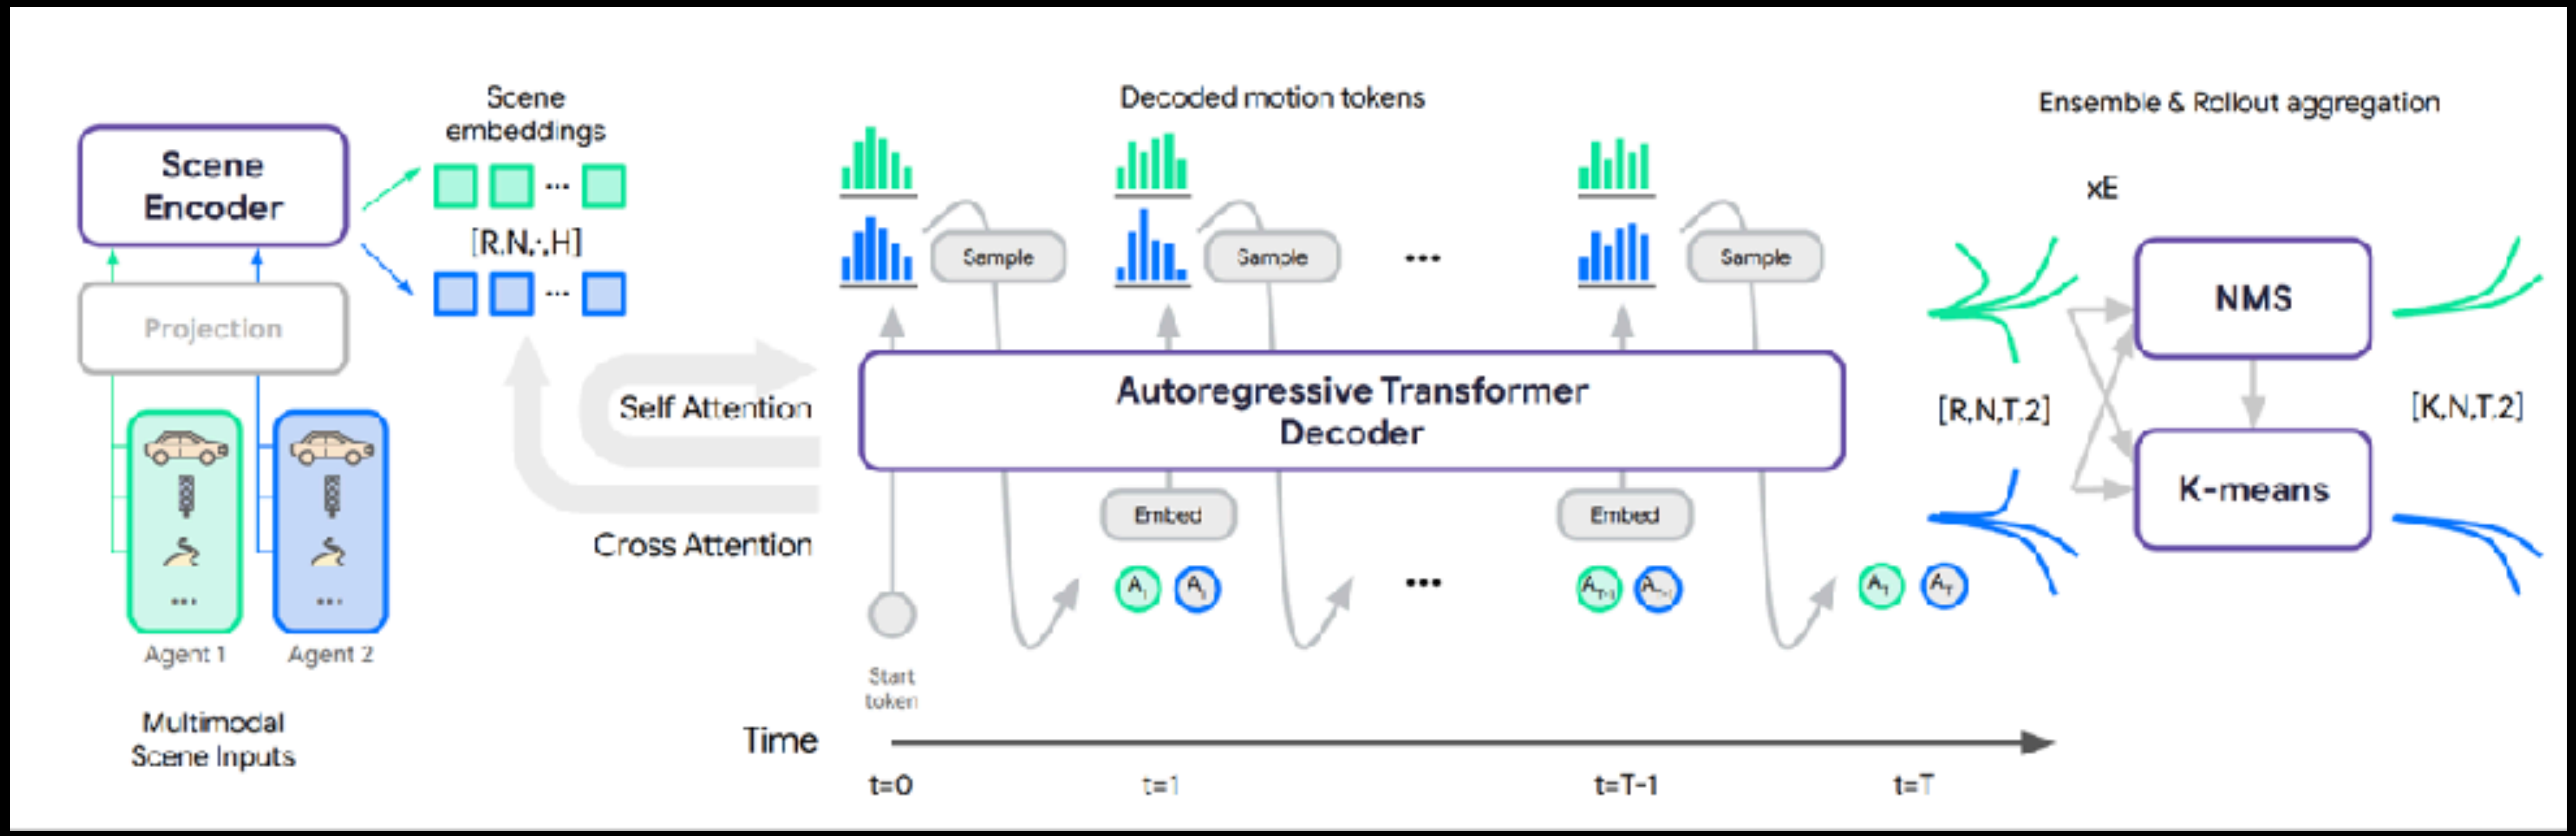

Design software from the ground-up to be learnable

Scalable pipelines that turn data into tests

Learning not just for perception, but also decision making

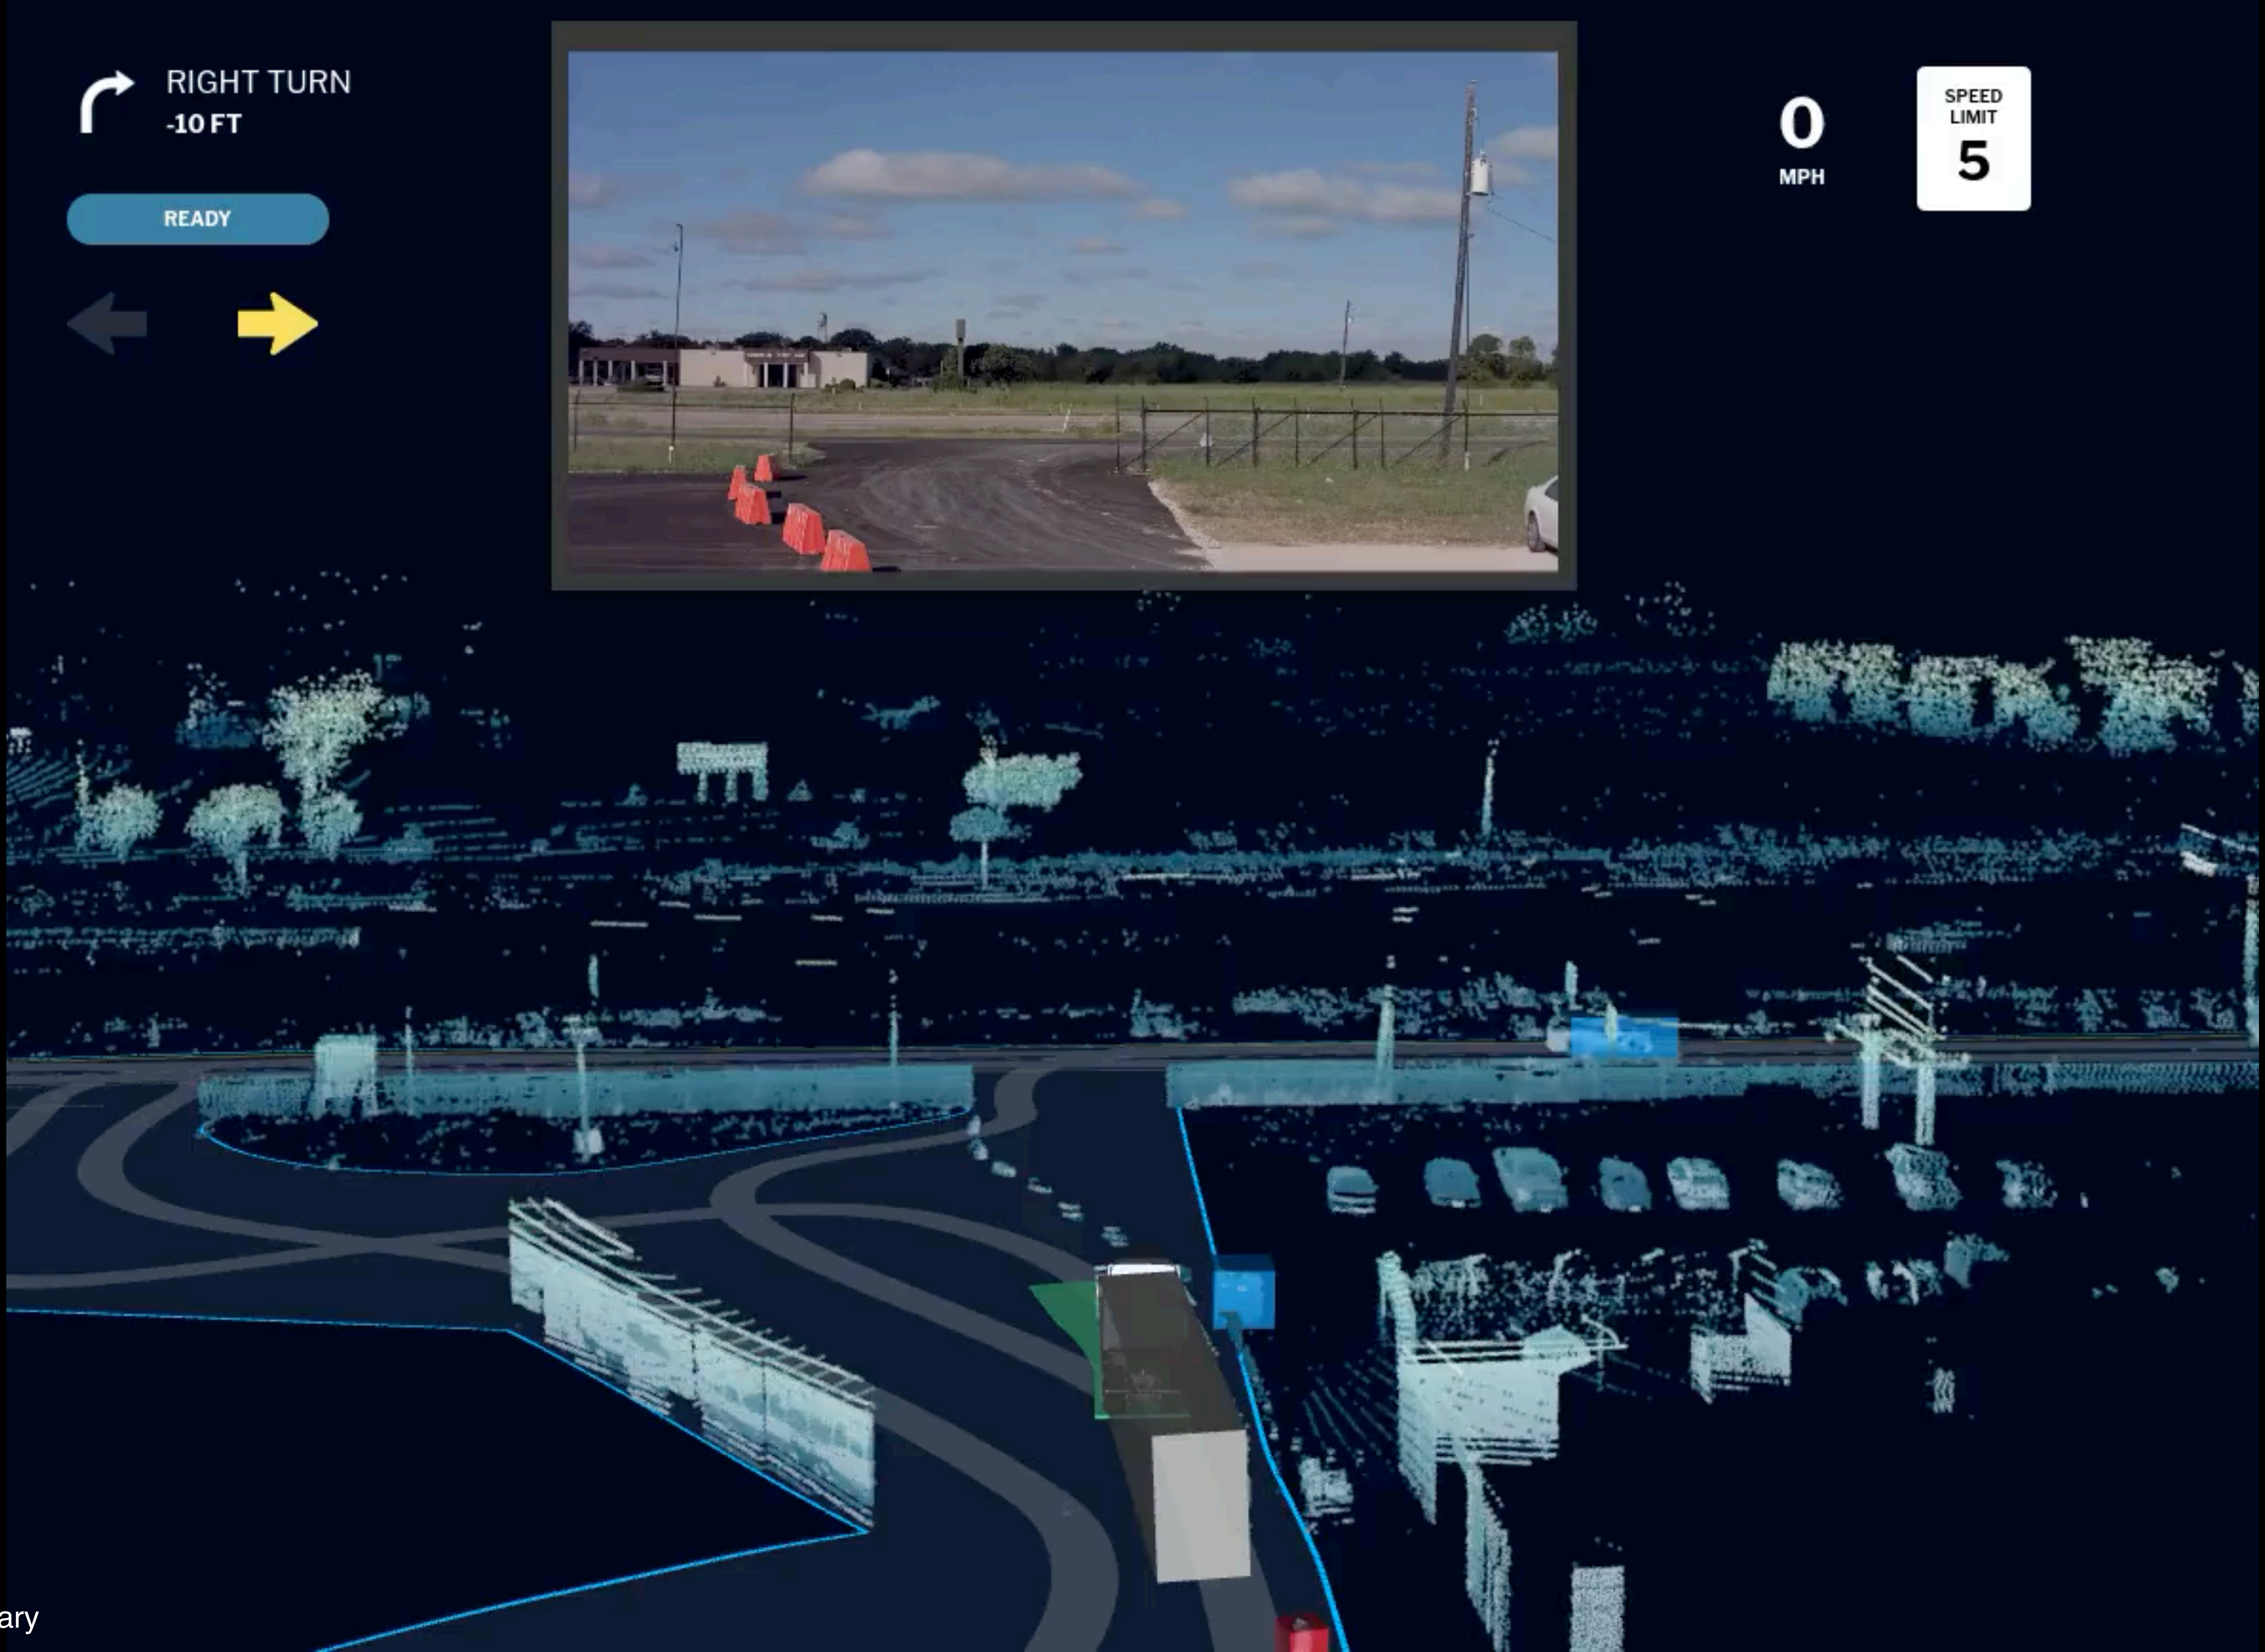
RIGHT TURN
-10 FT
READY
0
MPH
SPEED LIMIT
5

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

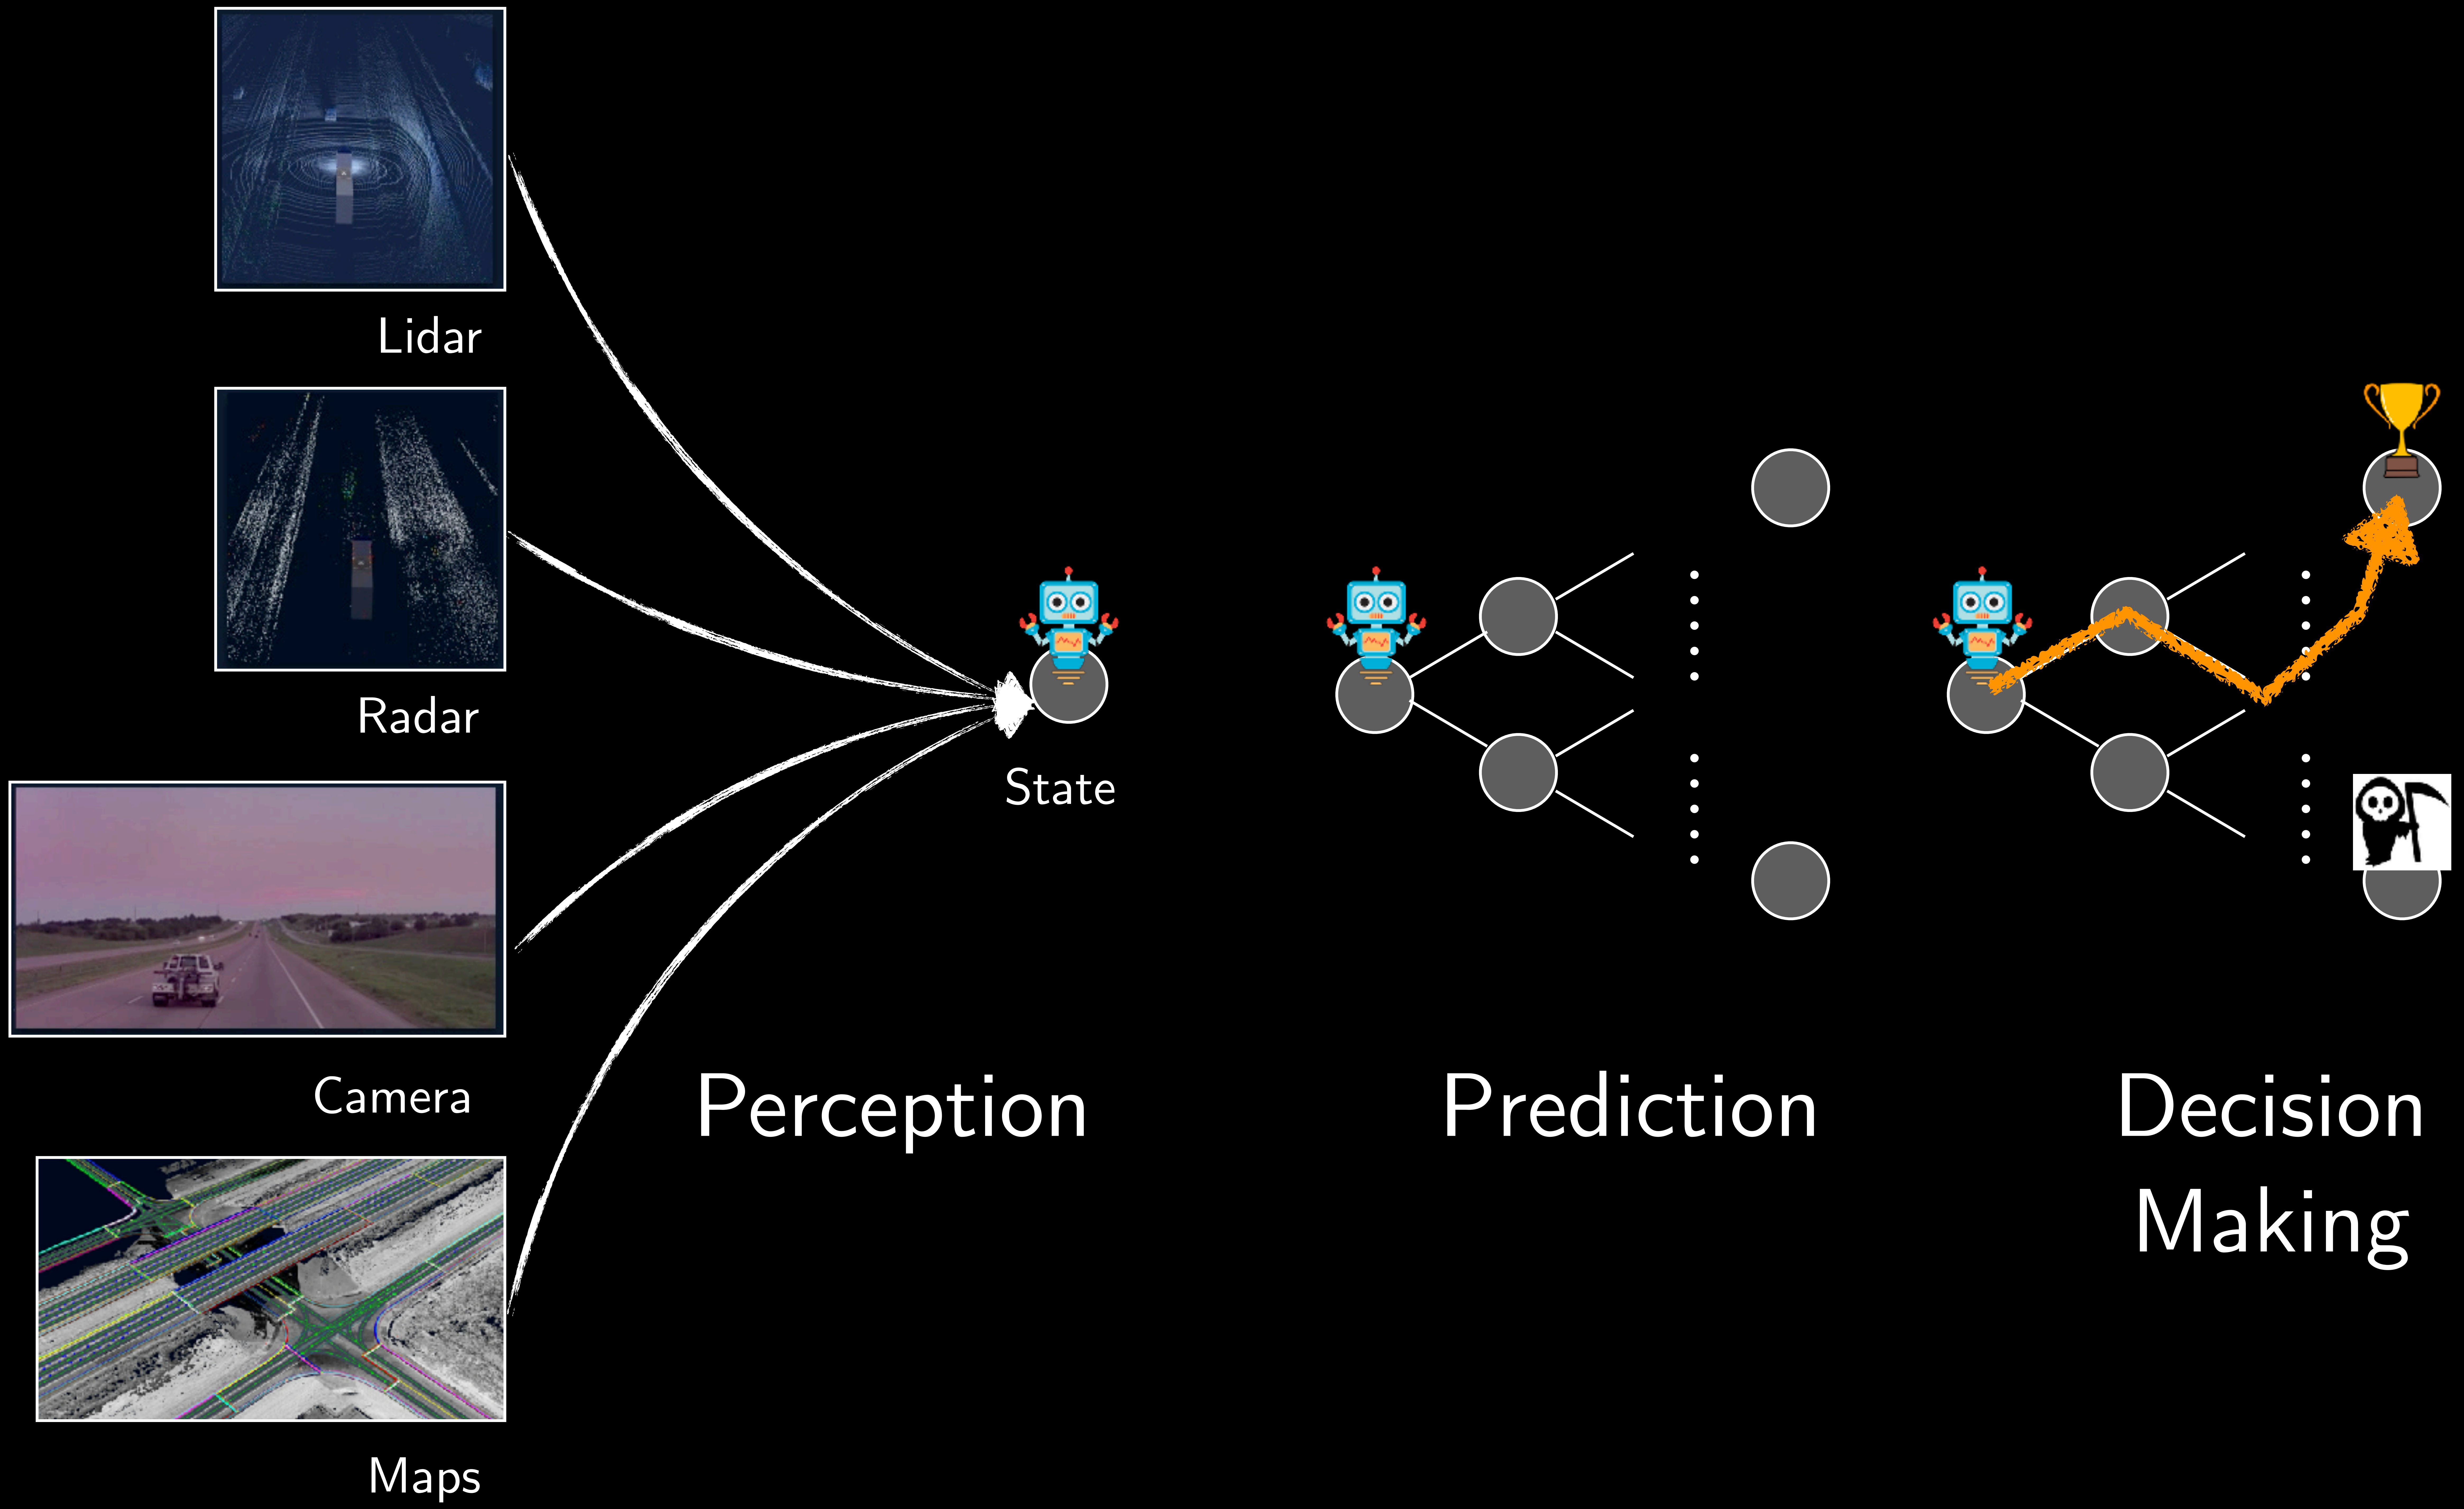
Lidar
Radar
Camera
Maps
State
Perception
Prediction
Decision Making

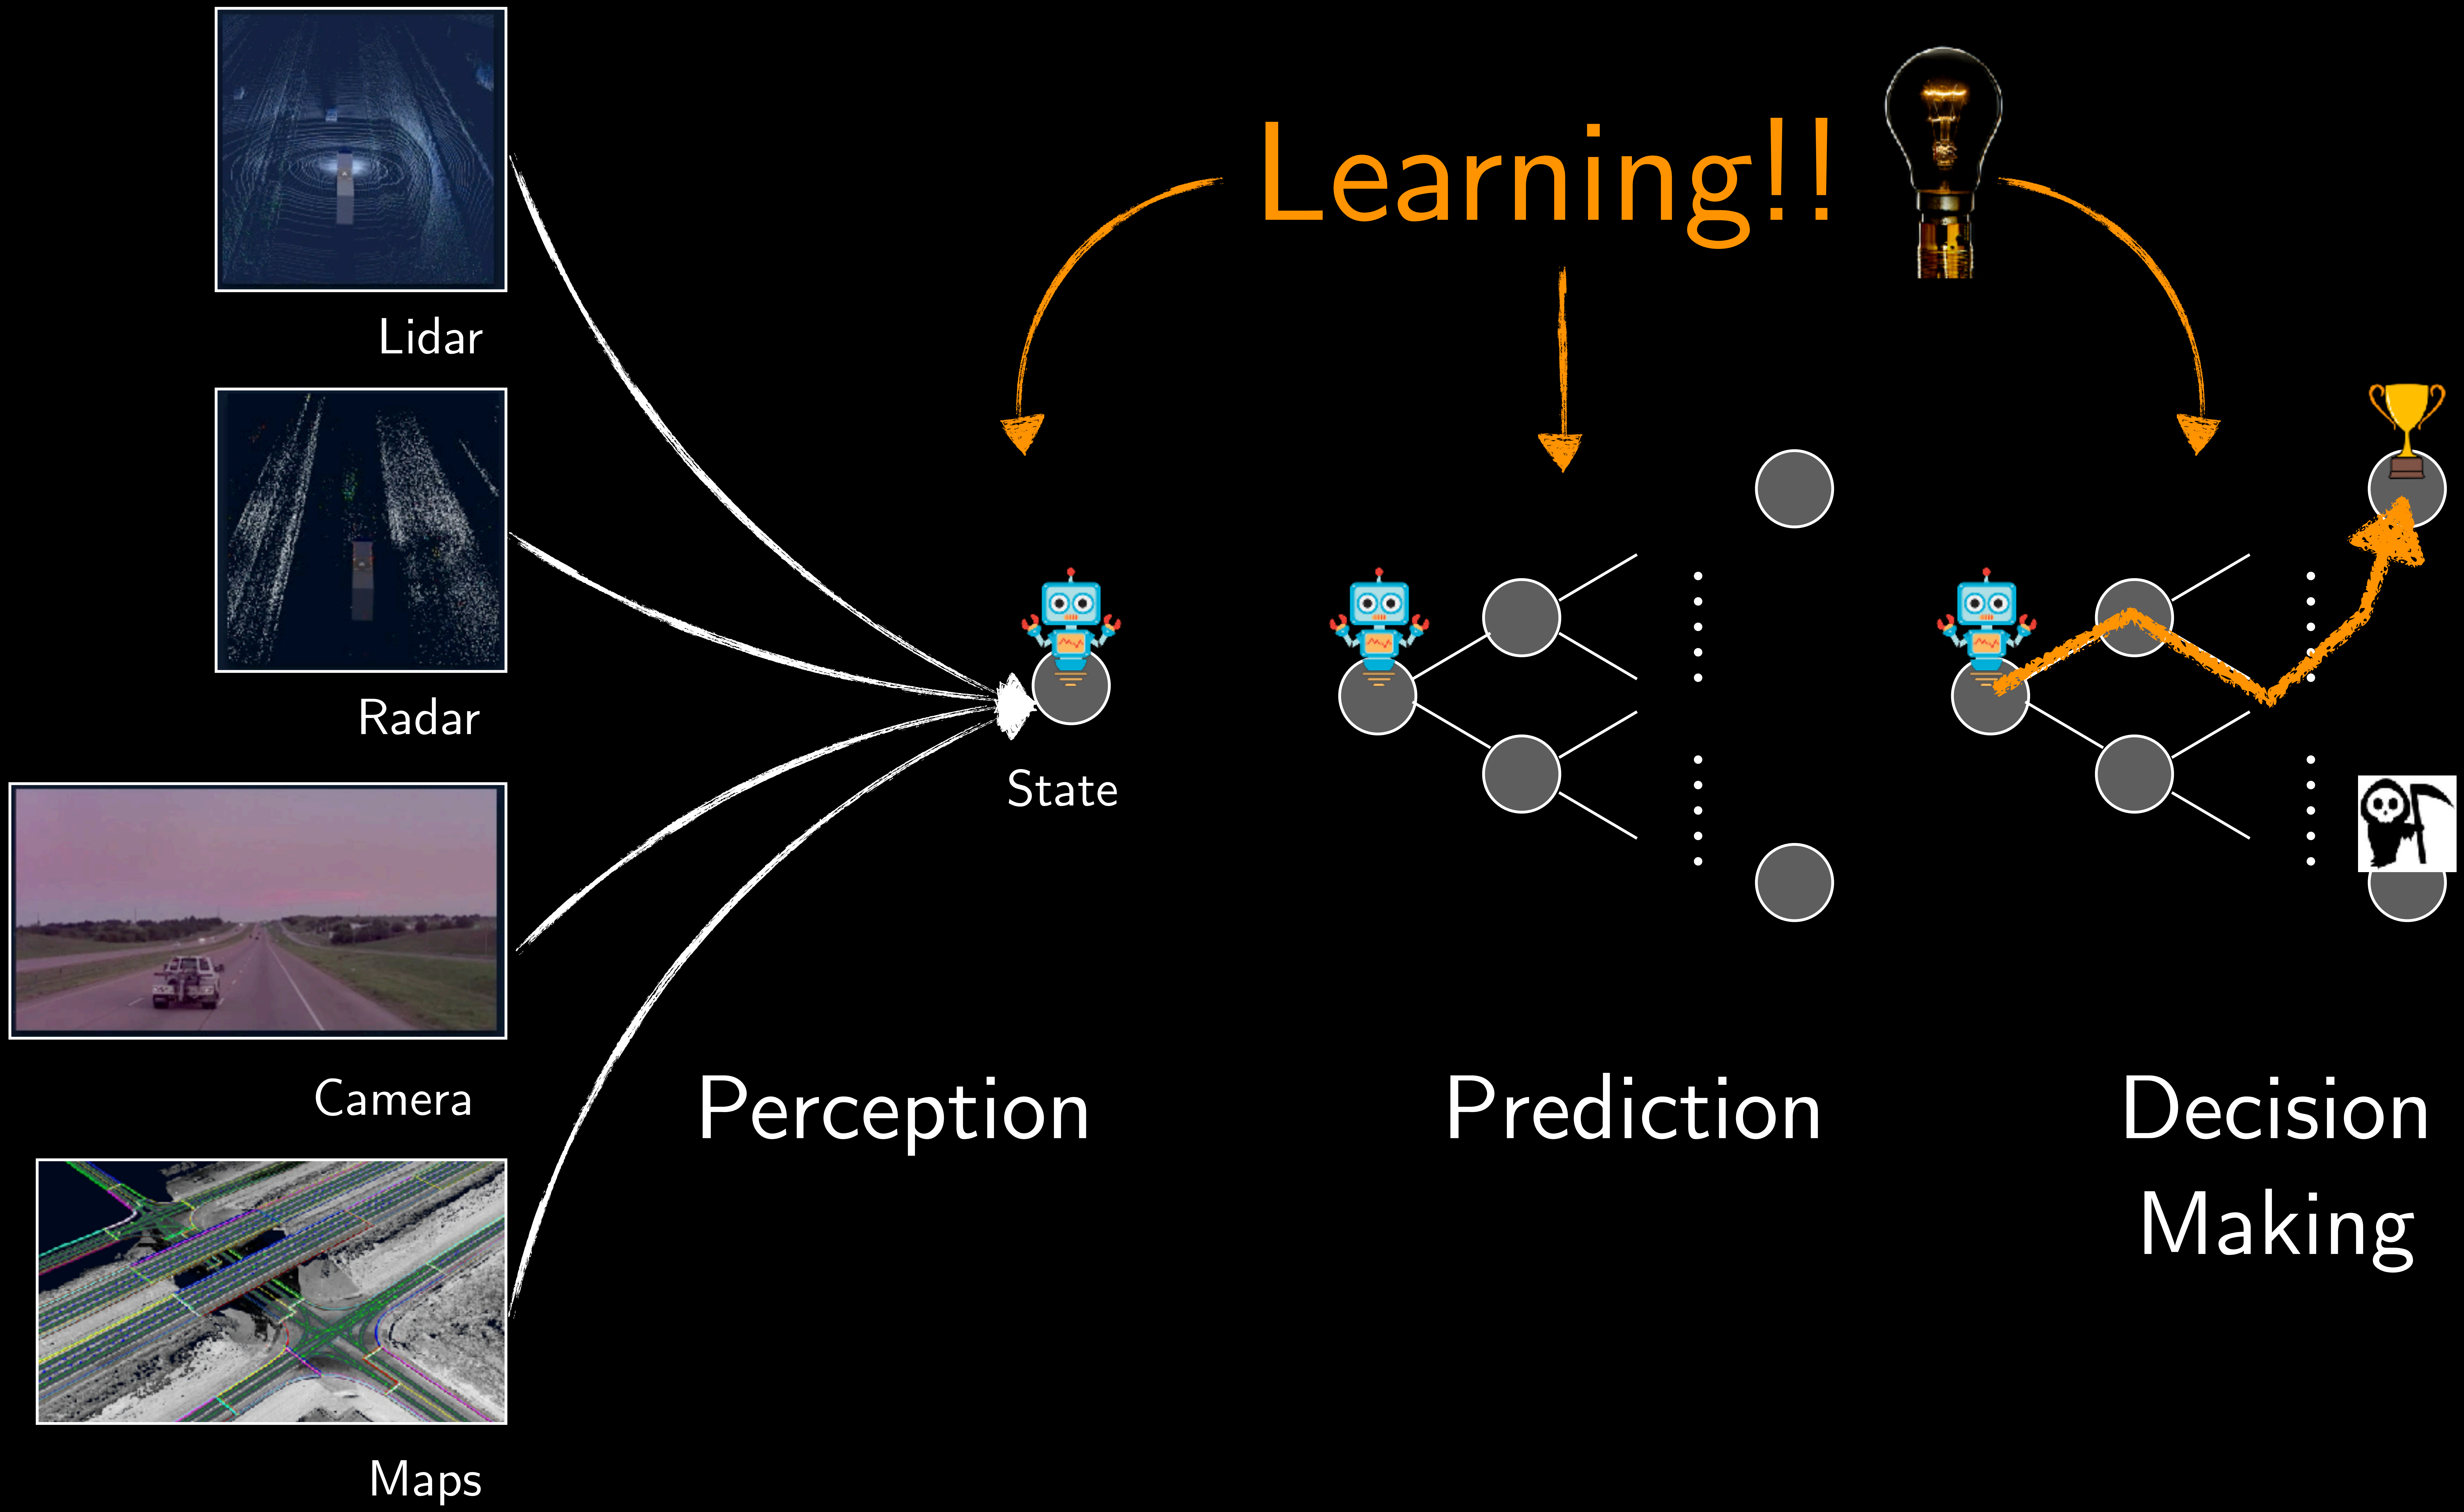
Learning!!
Lidar
Radar
Camera
Maps
State
Perception
Prediction
Decision
Making

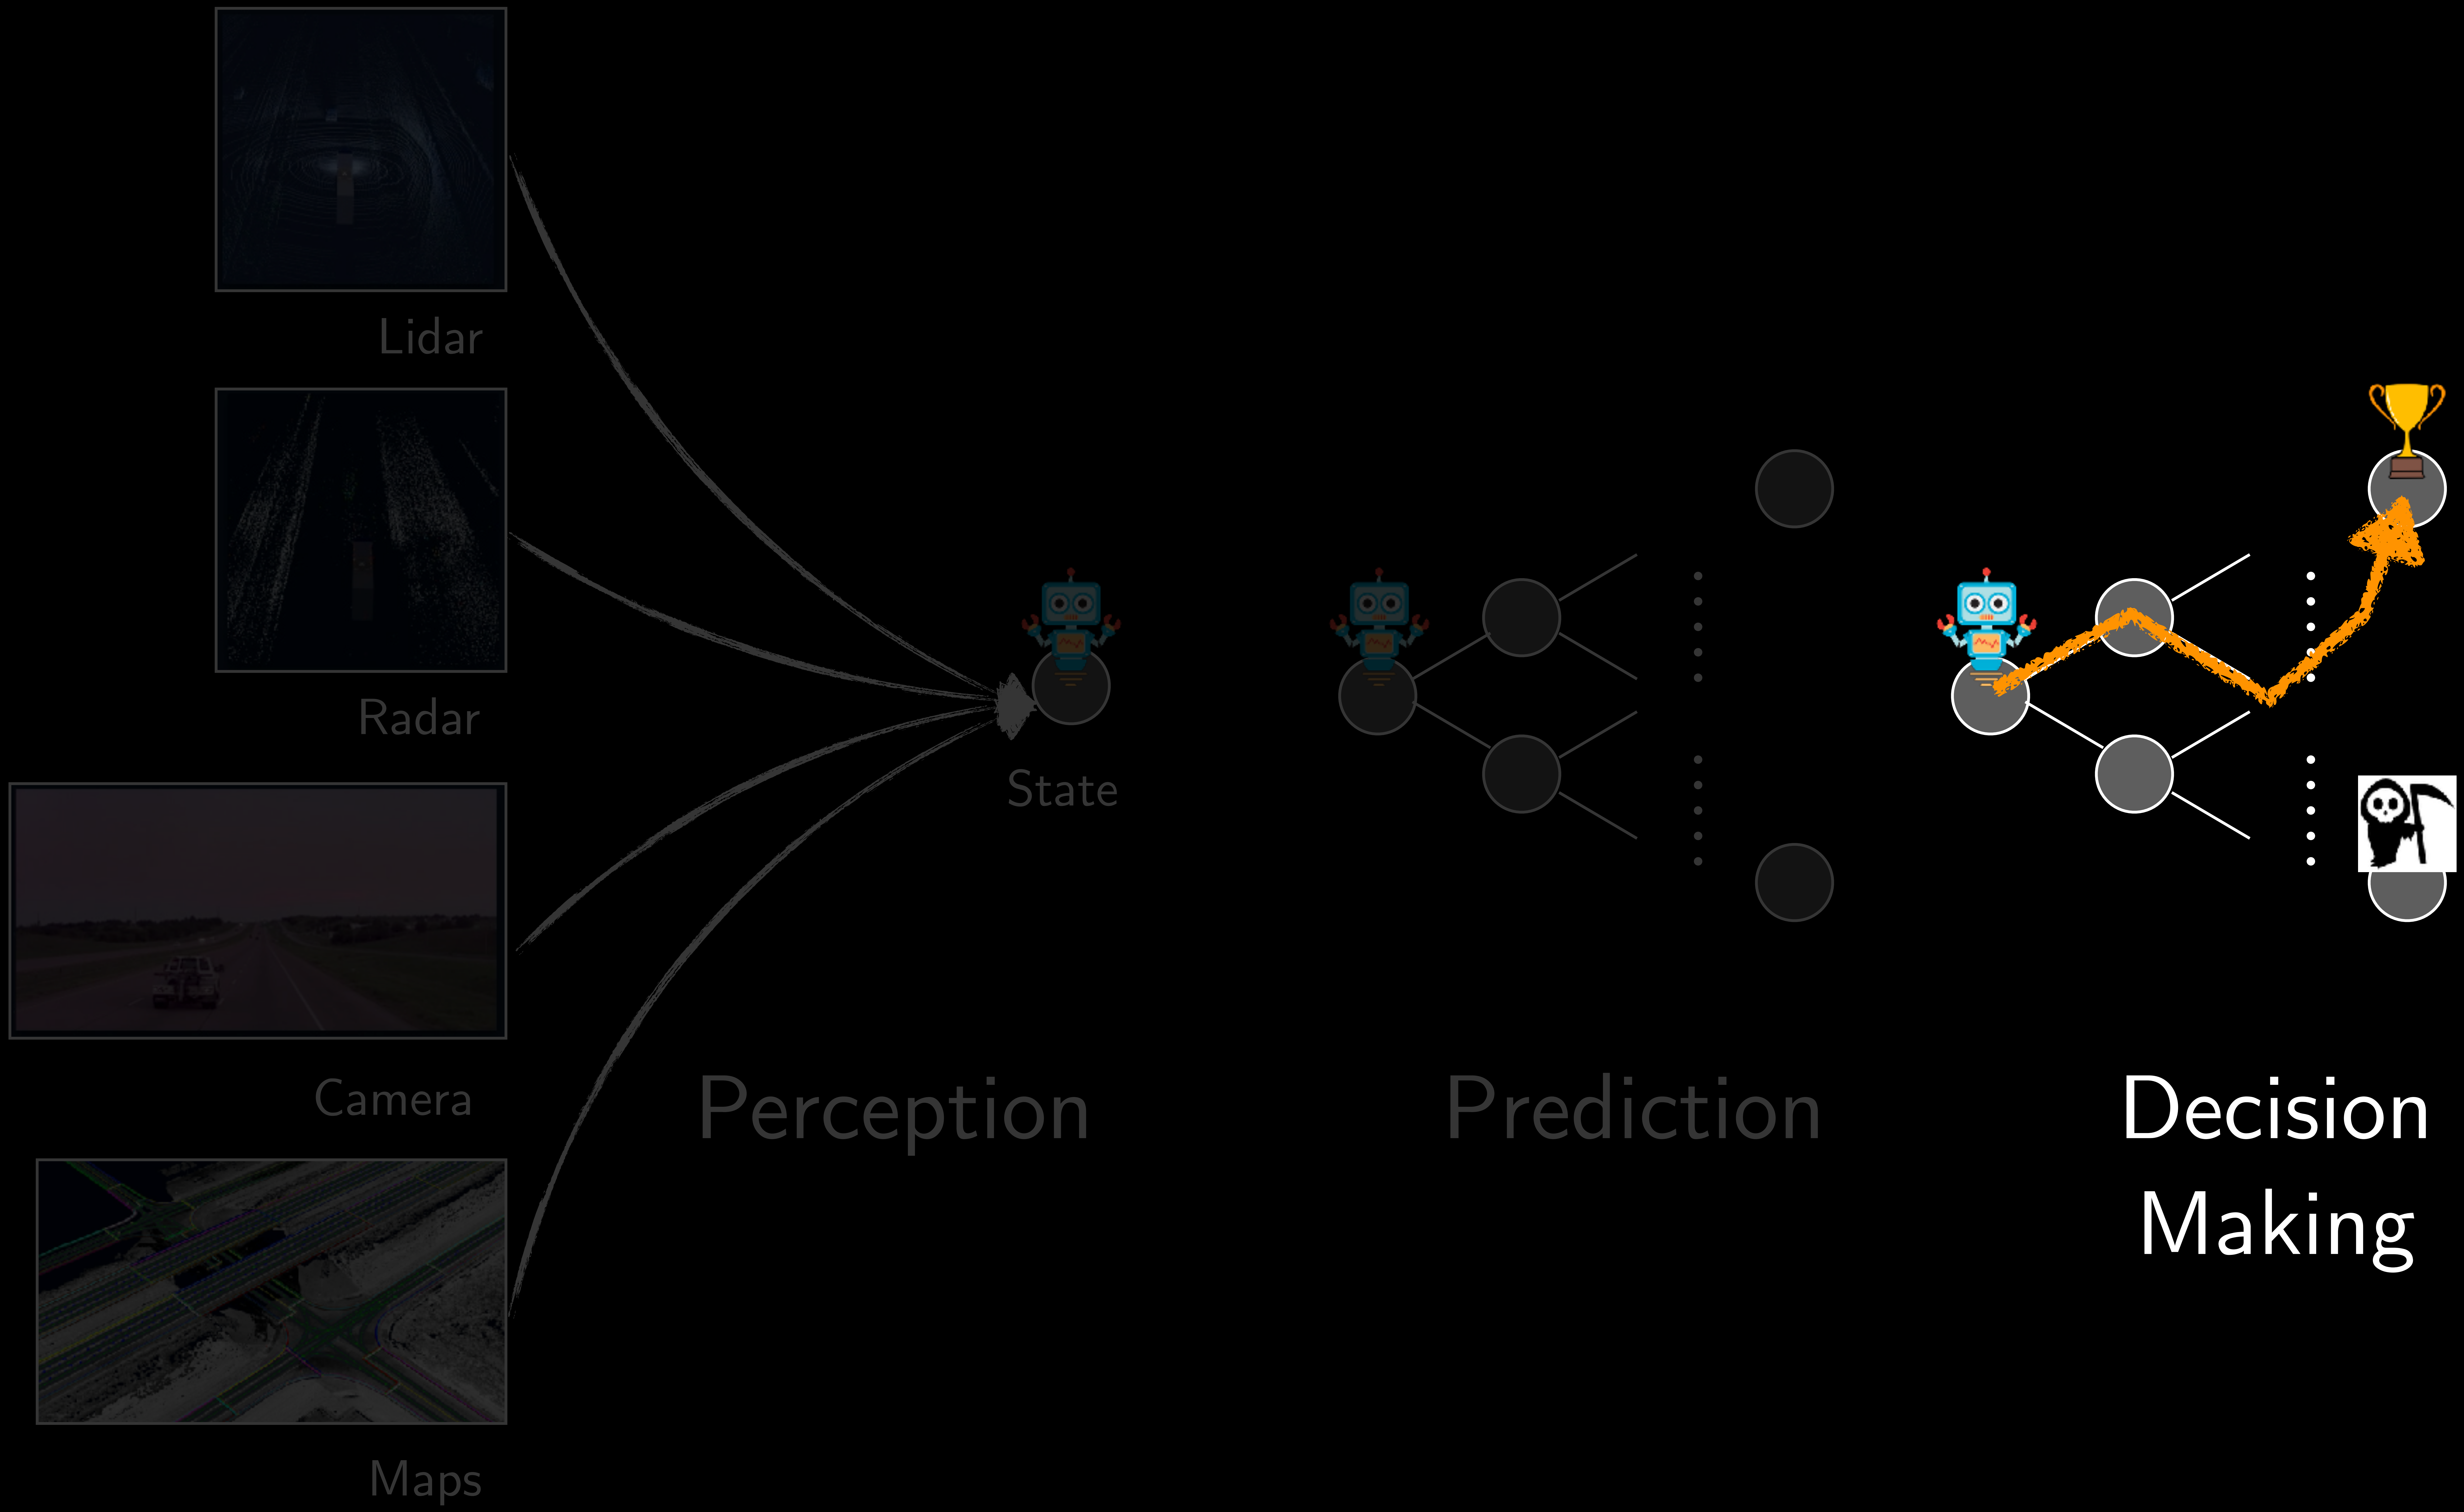
Lidar
Radar
Camera
Maps
State
Perception
Prediction
Decision Making

Activity: What is "good" behavior in a left turn?

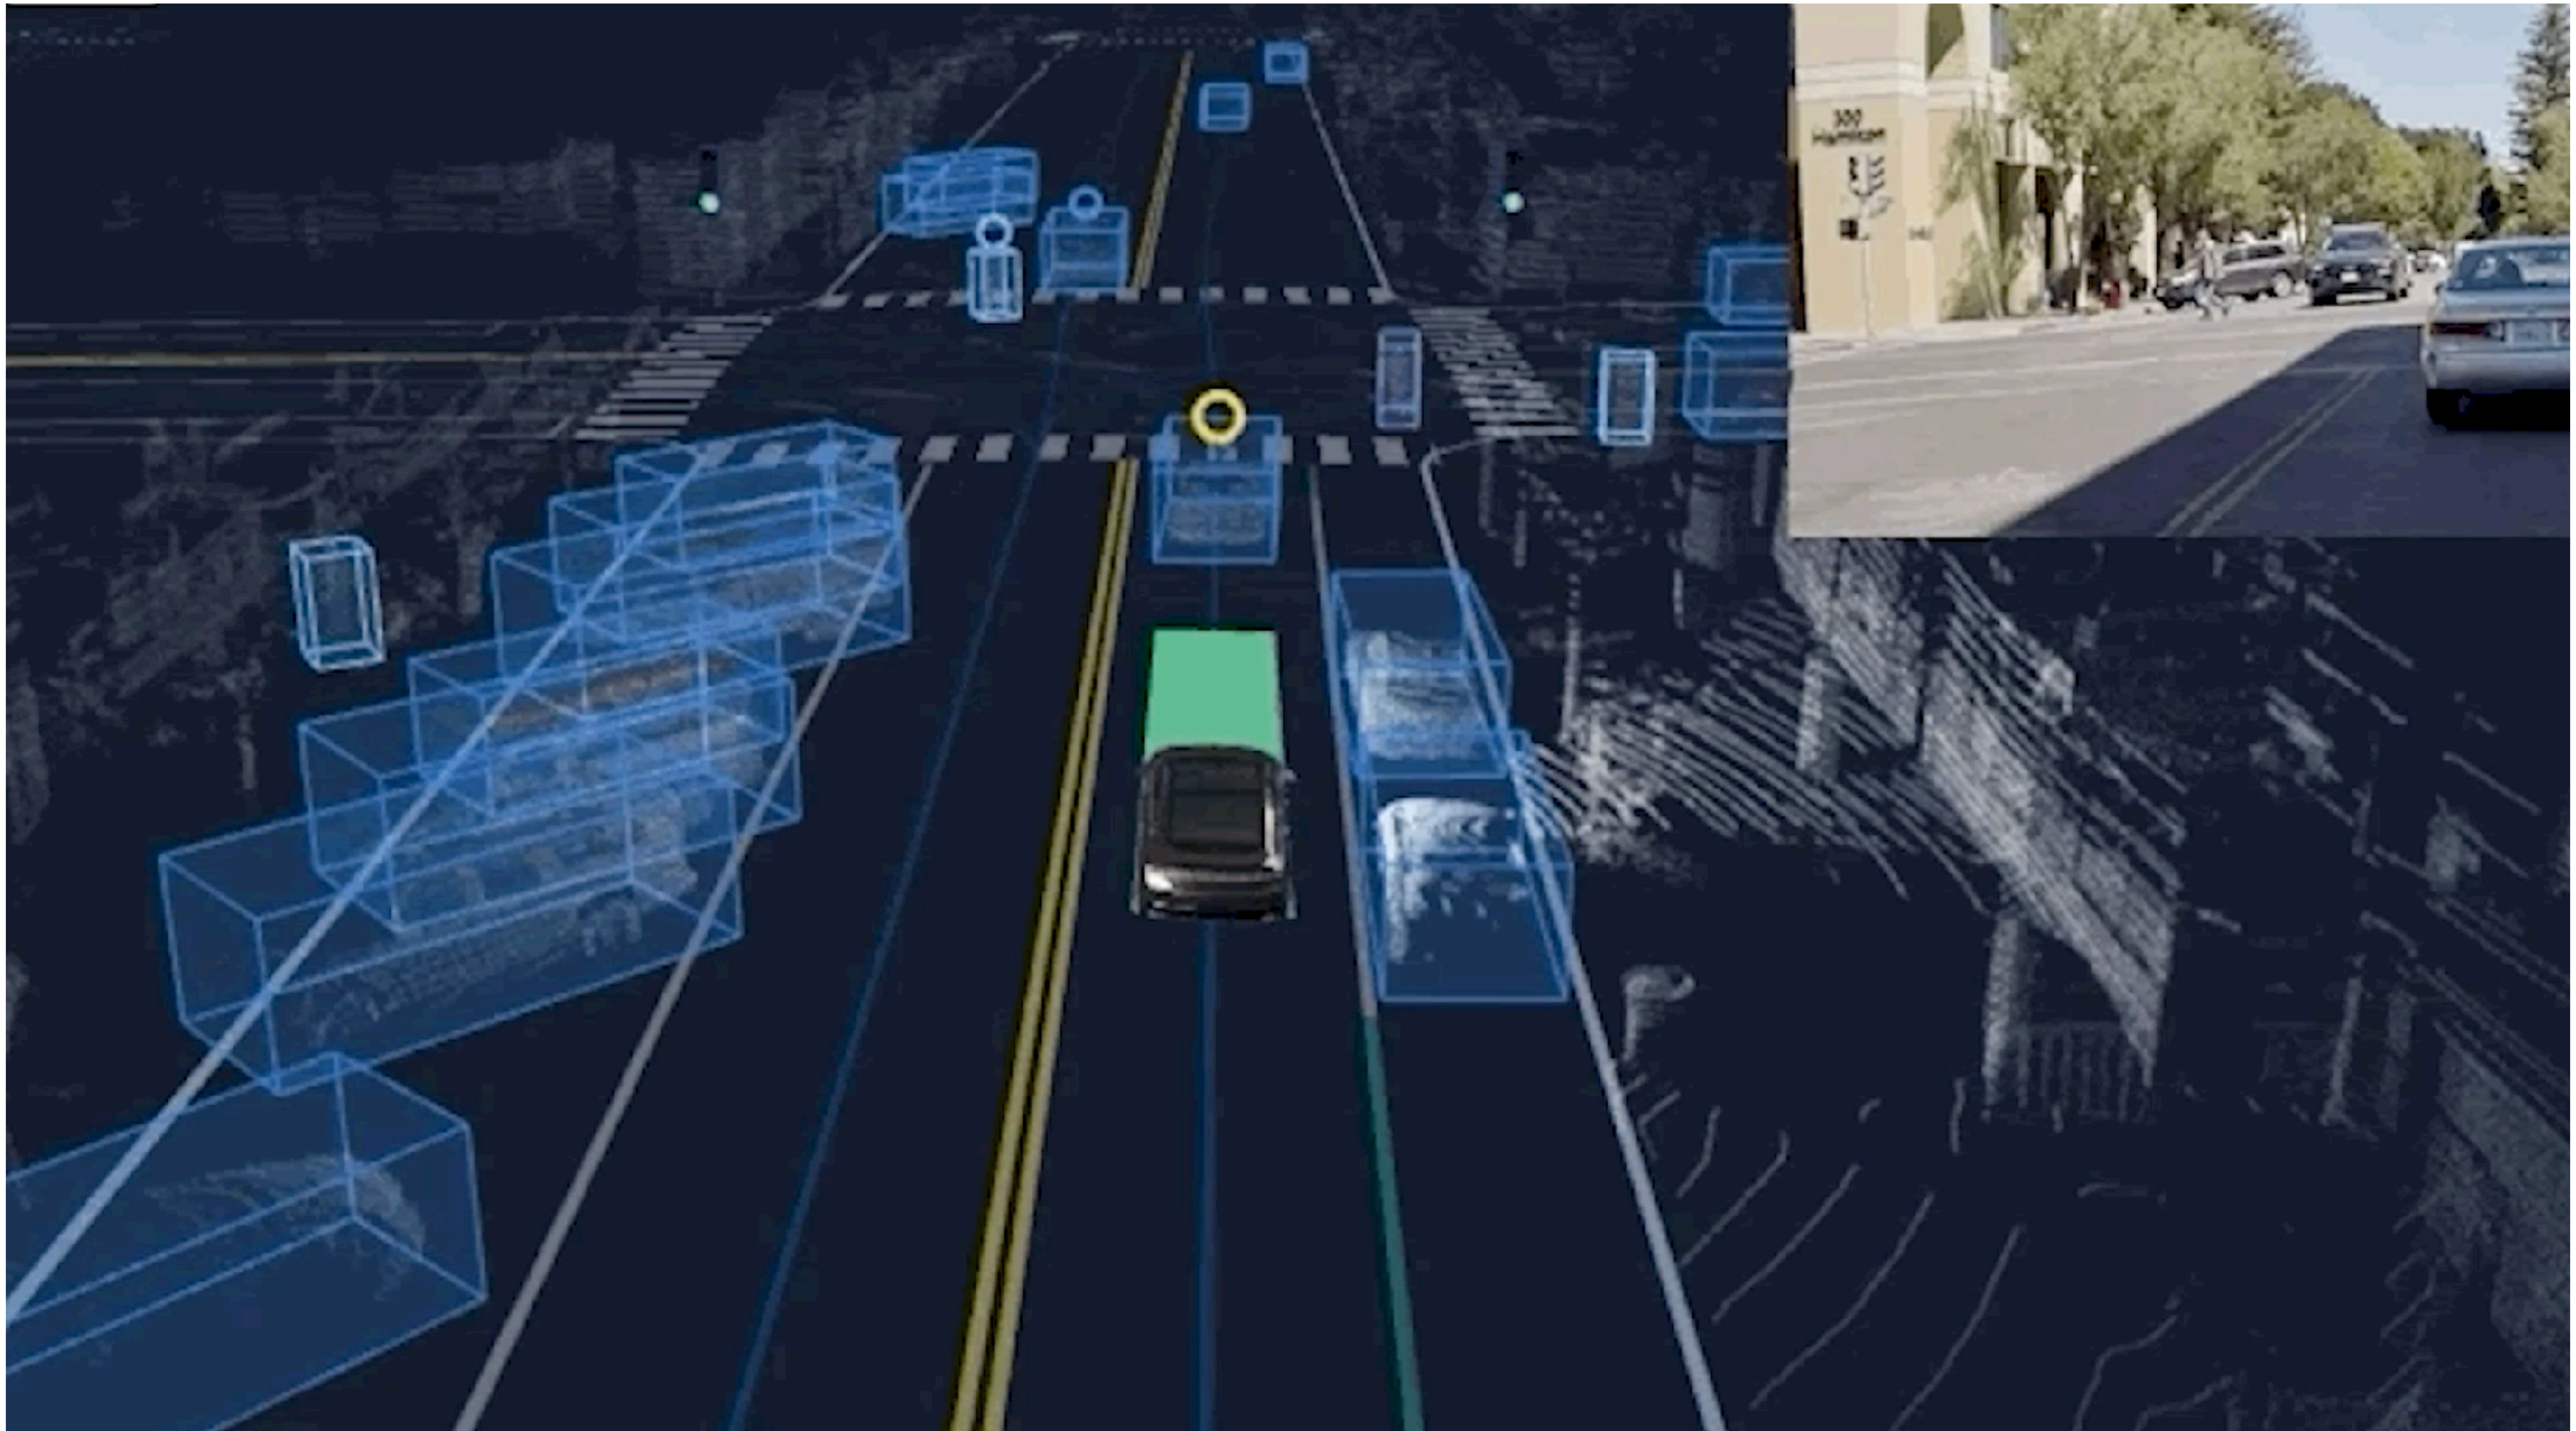

Activity!

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

How do we program in these values?

Lesson #1

Values are implicit in human driving!

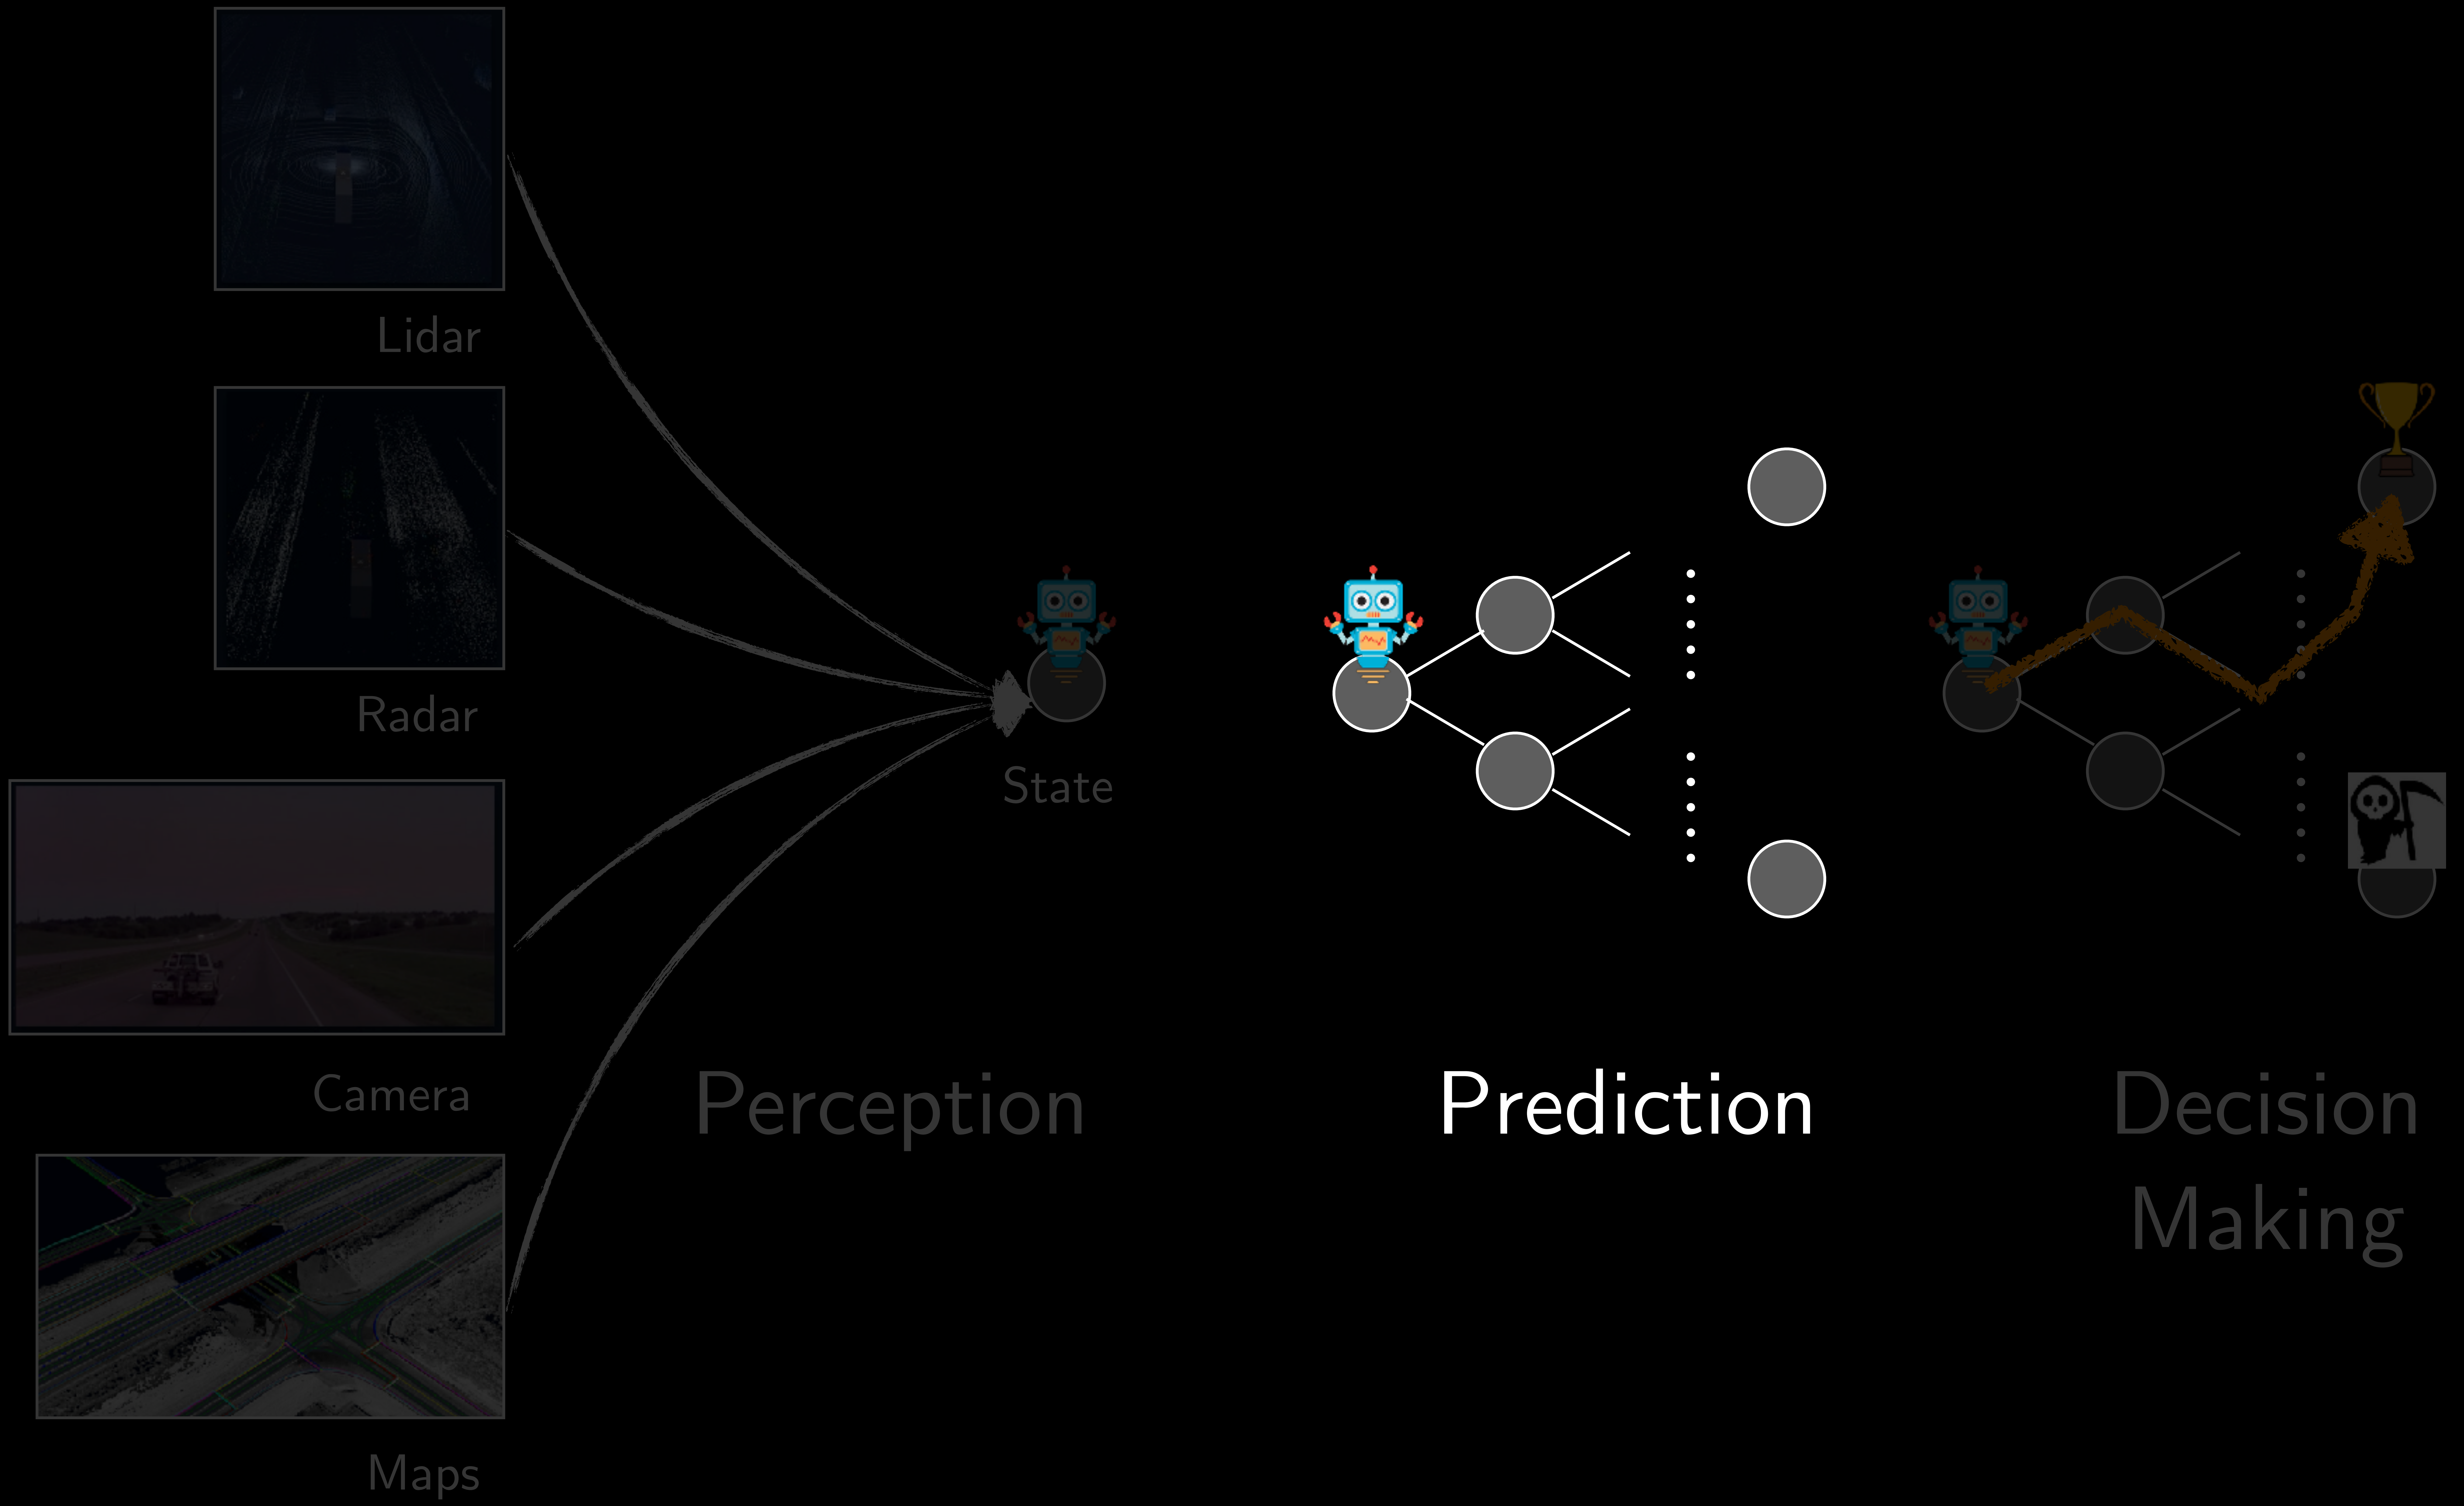
Lidar
Radar
Camera
Maps
State
Perception
Prediction
Decision Making

30.0 m
20.0 m
10.0 m
1541016641.475 20181031.175444Z.mkz-07.mkz.V0-592_MDC

Activity: How can we predict pedestrian motion?

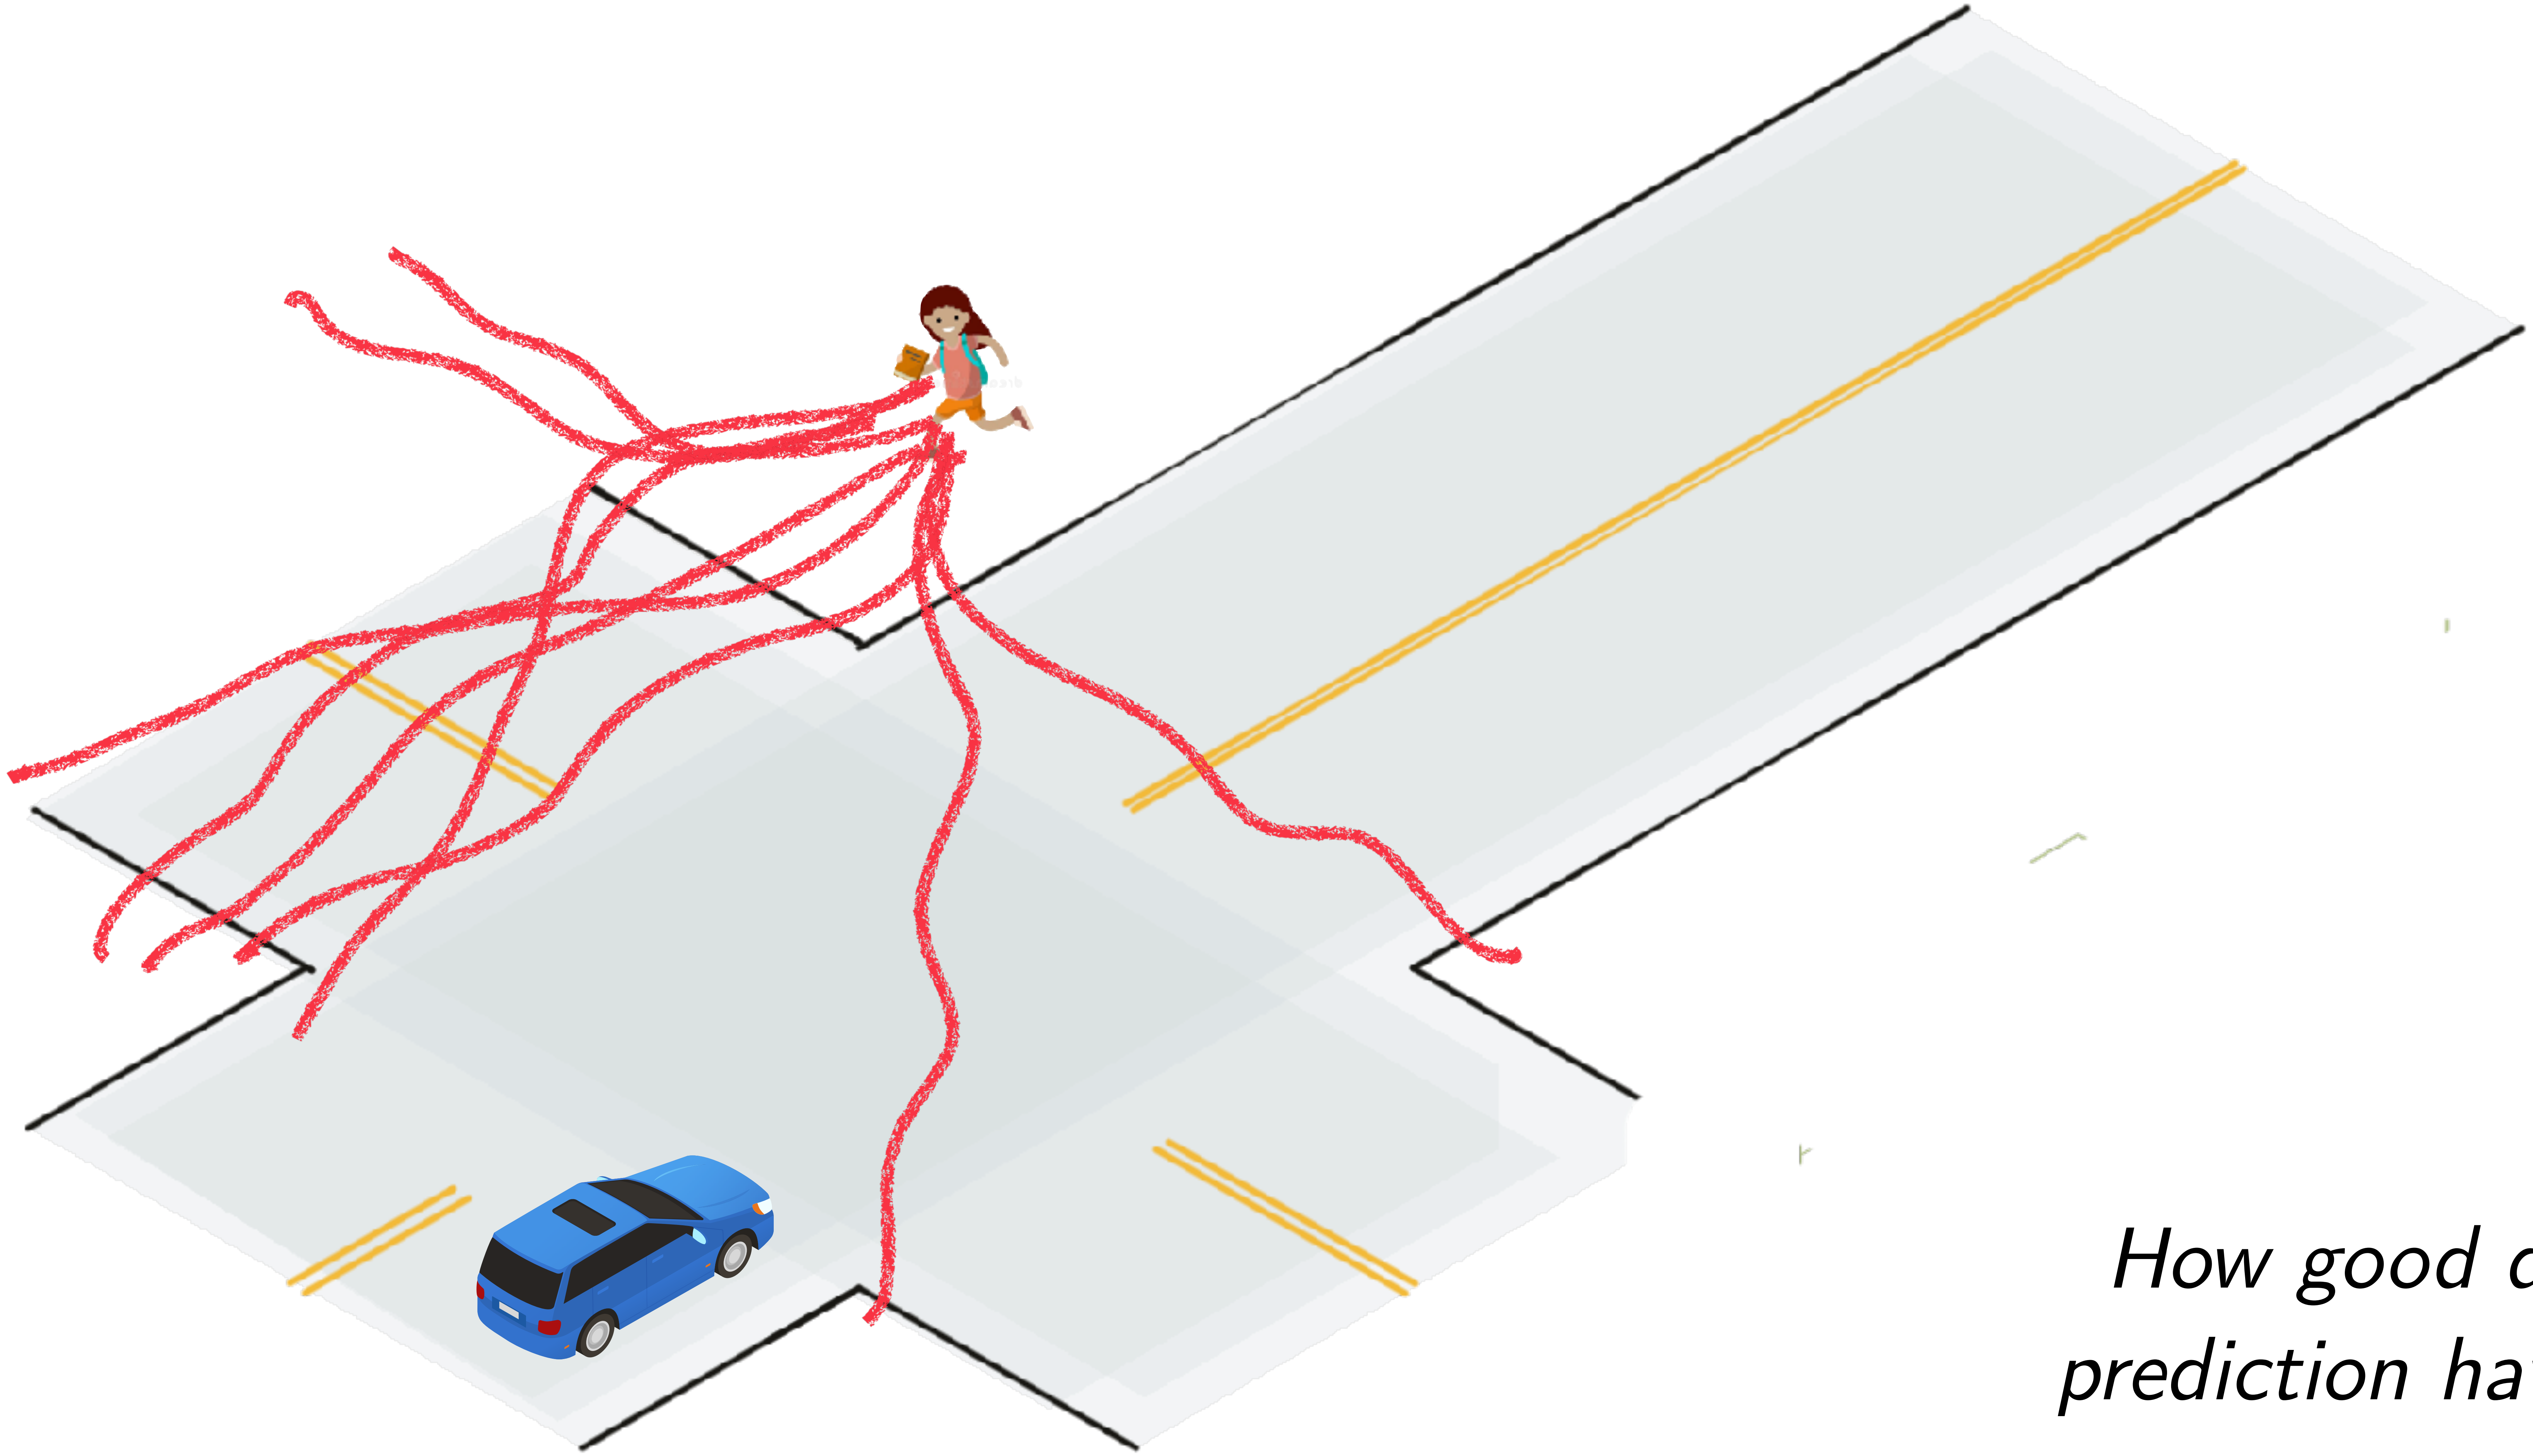

How good does the prediction have to be?

Lesson #2

Models are useful fictions

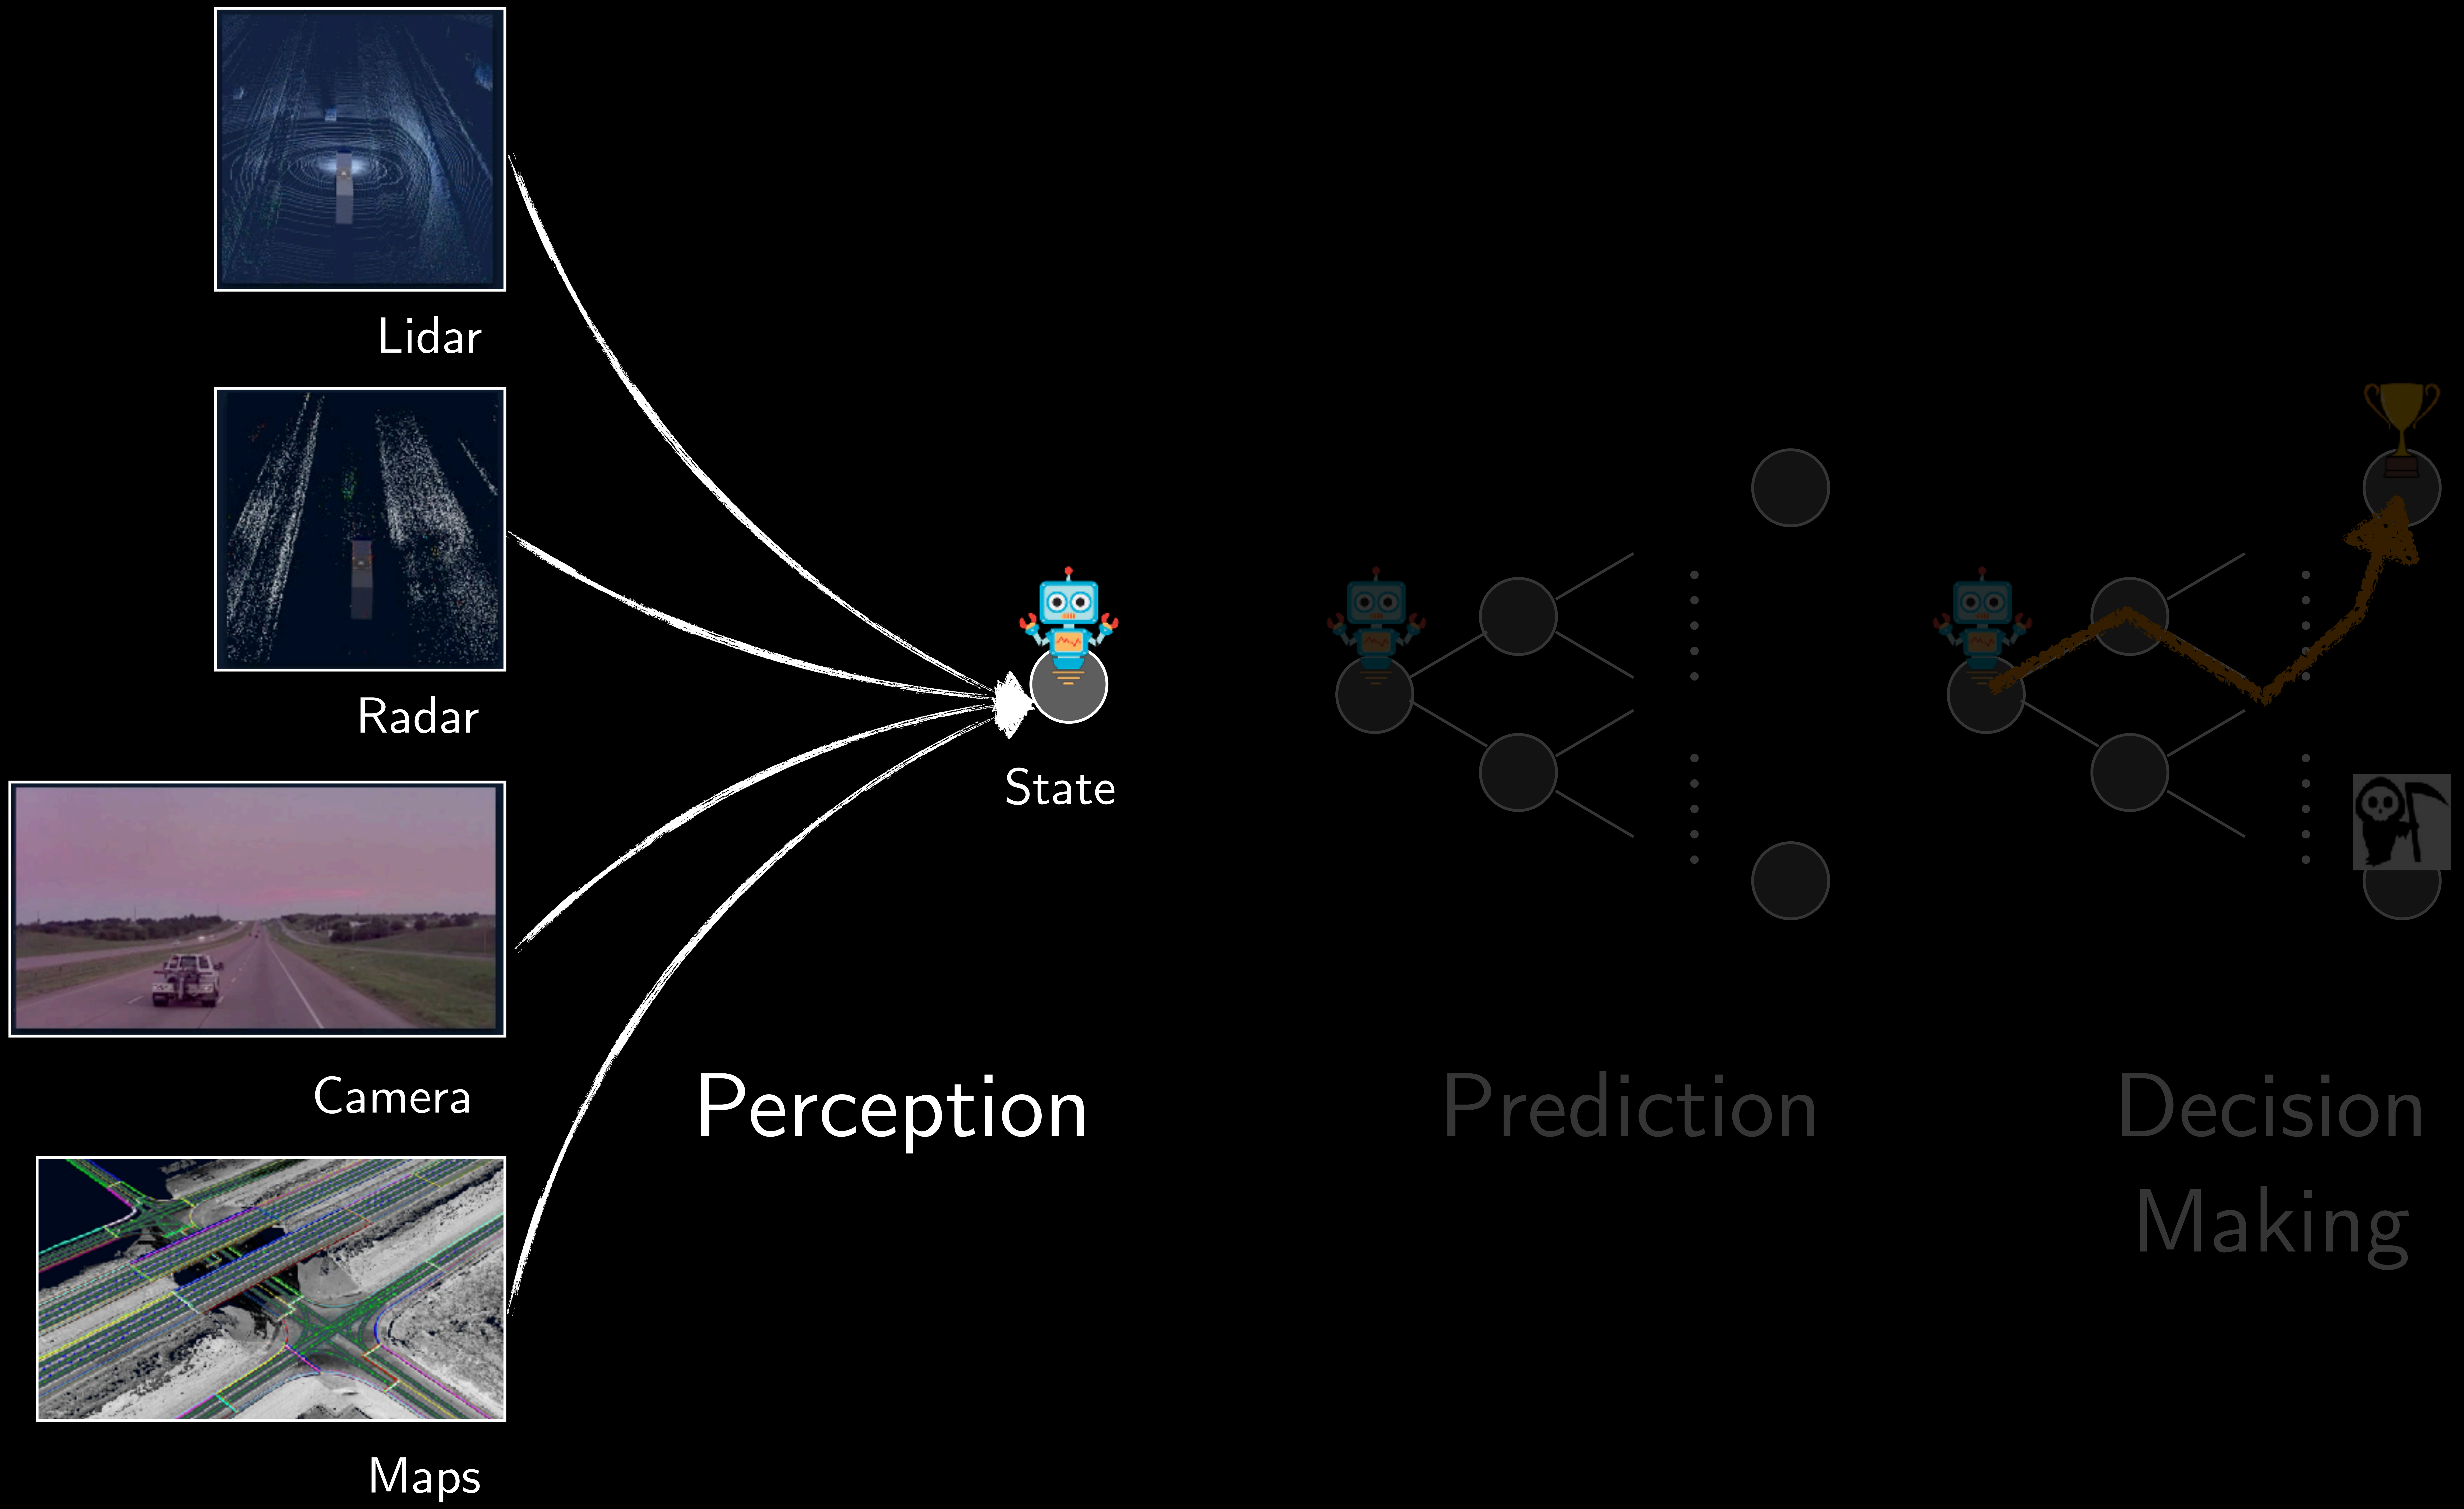
Lidar
Radar
Camera
Maps
State
Perception
Prediction
Decision Making

No one sensor tells the whole story!

LiDAR

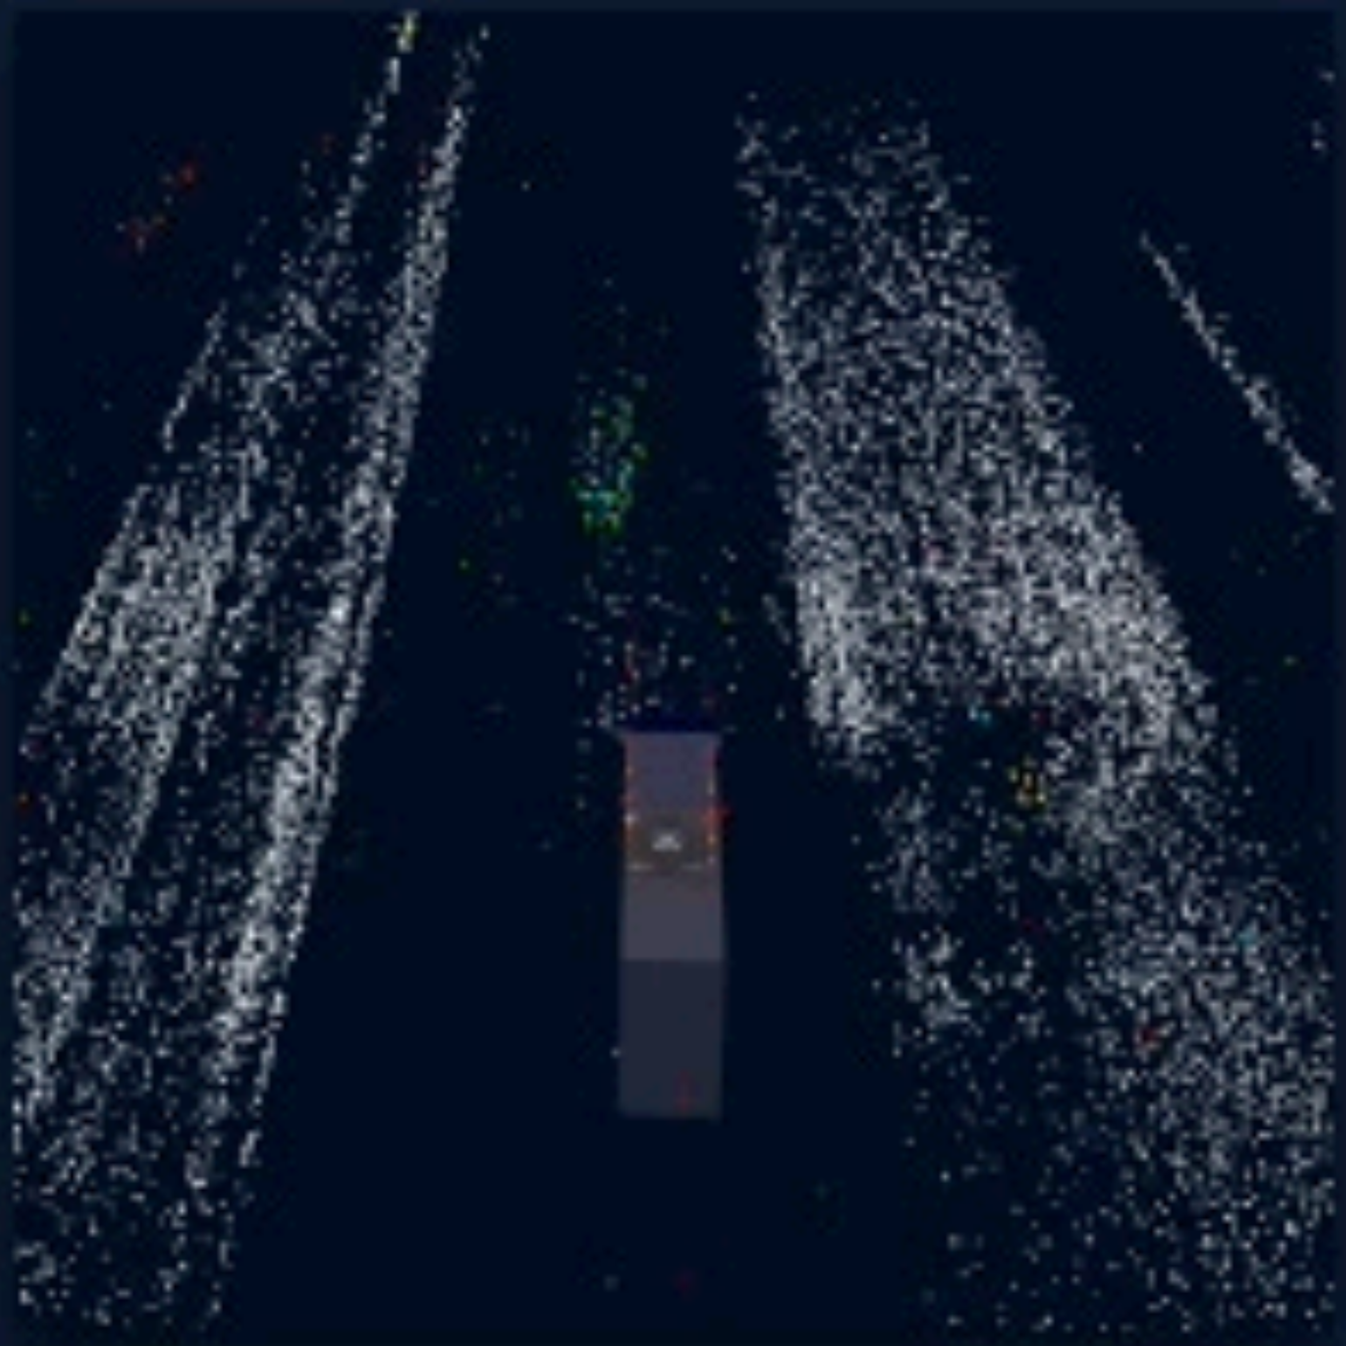

Radar

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!

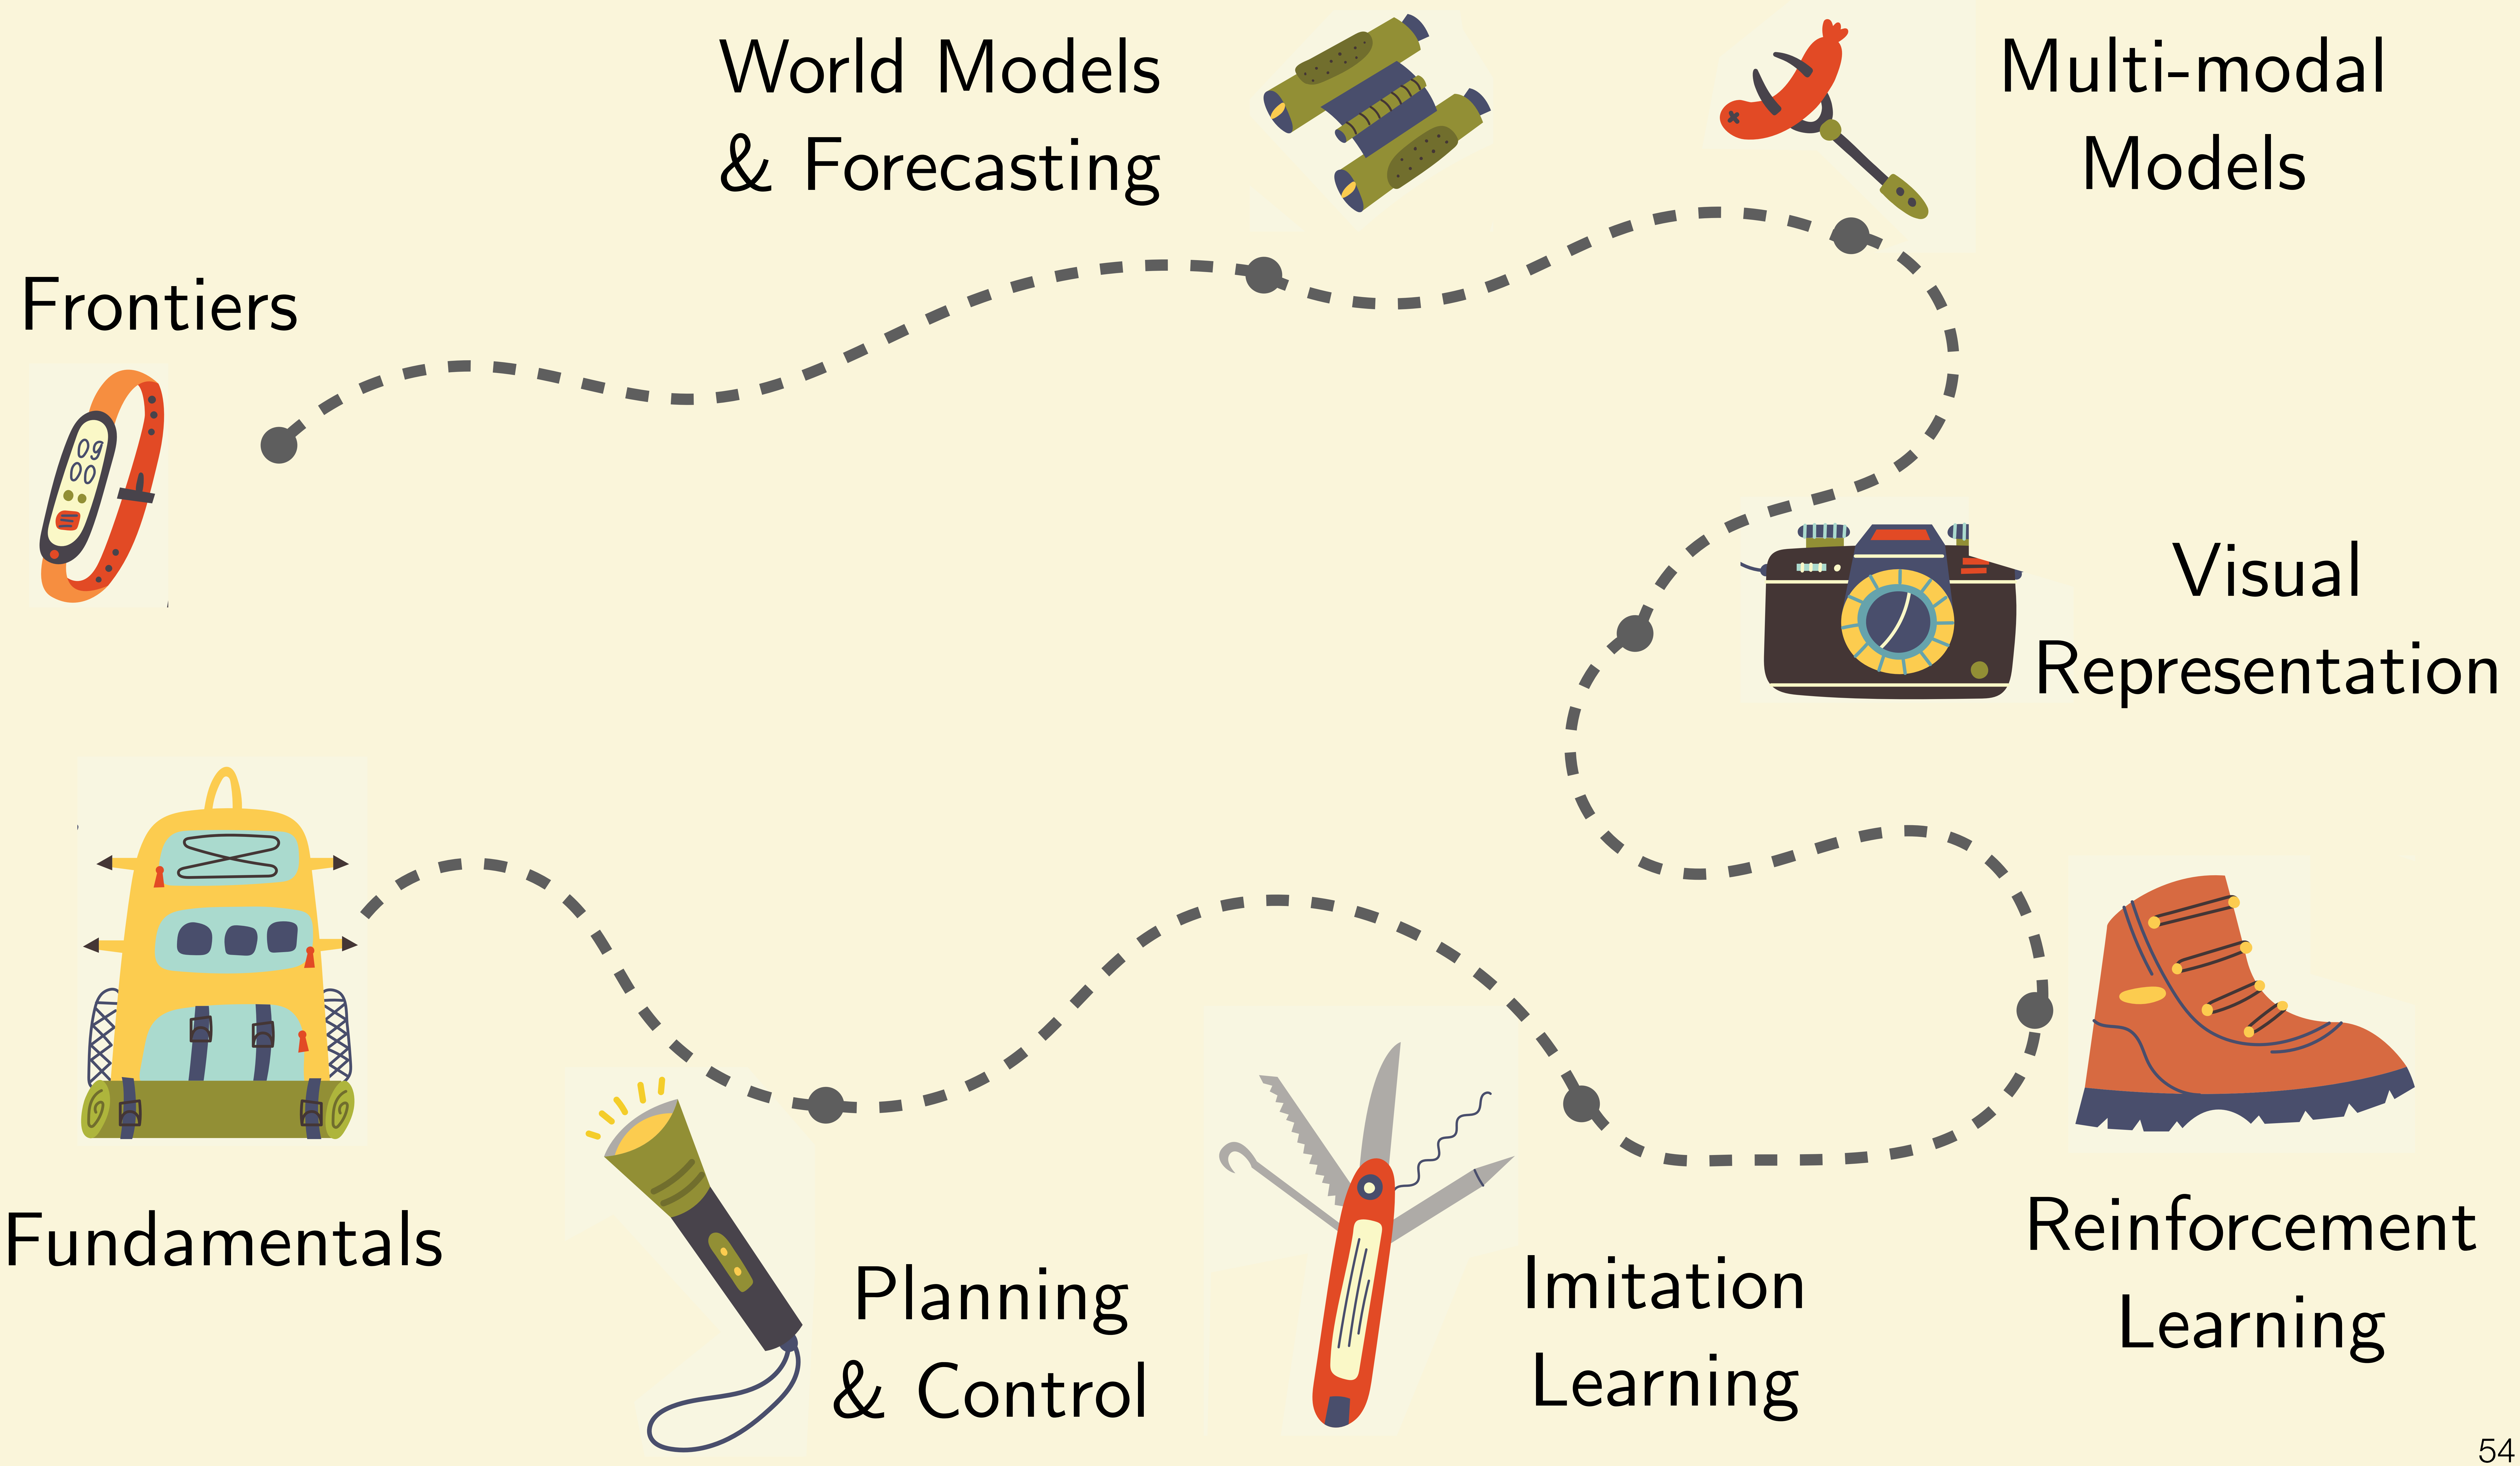
Fundamentals
Planning
& Control
Imitation
Learning
Reinforcement
Learning
Visual
Representation
Multi-modal
Models
World Models
& Forecasting
Frontiers

Logistics

Website is the ONE true hub

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

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		
1/23/24	Introduction to Robot Learning **(Assignment 0 released)**		The Bitter Lesson
1/25/24	Robots as Markov Decision Problems	MACRL Ch. 1	Dan Klein slides I , slides II

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.

7 Assignments [50%]

A0: Intro assignment

A1, A3, A5: Written assignment

A2, A4, A6: 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!

Announcement!

Assignment 0

https://github.com/portal-cornell/cs4756-robot-learning-sp24/tree/main/assignments/A0

Link in website!

Checks familiarity with PyTorch!

Due Thursday 2/1!

Graduate Version (CS5756)

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

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

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

Office Hours:
Tuesday 11:30 - 1:30 pm, Gates 413 B

Gonzalo Gonzalez

Teaching Assistant

gg387@cornell.edu

Office Hours:
Sunday 12:00 - 2:00 pm, Rhodes 404

"My new year's resolution is to cook everyday"

Lisa Asriev

Teaching Assistant

laa97@cornell.edu

Office Hours:
Friday 11:00 am - 1:00 pm, Rhodes 404

Adam Cahall

Teaching Assistant

abc256@cornell.edu

Office Hours:
Thursday 1:15 - 3:15 pm, Rhodes 404

"I enjoy ice skating in my free time, and I used to compete nationally as a synchronized figure skater."

"I enjoy playing Ultimate Frisbee."

Angela Chao

Teaching Assistant

ac2323@cornell.edu

Office Hours:
Wednesday 11:00 am - 1:00 pm, Rhodes 404

Sometimes I like to play shows I know too well as background noise when doing work.

Prithwish Dan

Teaching Assistant

pd337@cornell.edu

Office Hours:
Tuesday 2:00 - 4:00 pm, Rhodes 404

I enjoy chinese yo-yoing and dancing

Vaishnavi Gupta

Teaching Assistant

vg222@cornell.edu

Office Hours:
Monday 9:00 - 11:00 am, Rhodes 404

Bryant Har

Teaching Assistant

bjh254@cornell.edu

Office Hours:
Monday 11:00 am - 12:00 pm,
Friday 10:00 - 11:00 am Zoom

I enjoy reading books about games and hobbies I don't have.

Juntao Ren

Teaching Assistant

jlr429@cornell.edu

Office Hours:
Saturday 9:00 - 11:00 am, Rhodes 404

My two favorite authors are probably John Steinbeck and Stephen King.

TLDR

A0 is due 1/2!

Checkout pre-reading for next lecture!

Checkout course website for all details:

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