Introduction

What is learning?

What is learning?

- **Question.** How do human learn?

What is learning?

- **Question.** How do human learn?
- Toddlers like to call the names of many things
  - Learning "concepts" — associations of visual and linguistic signals

What is learning?

- Toddlers learn concepts from many **samples**
  - Images of dogs, and images of non-dogs




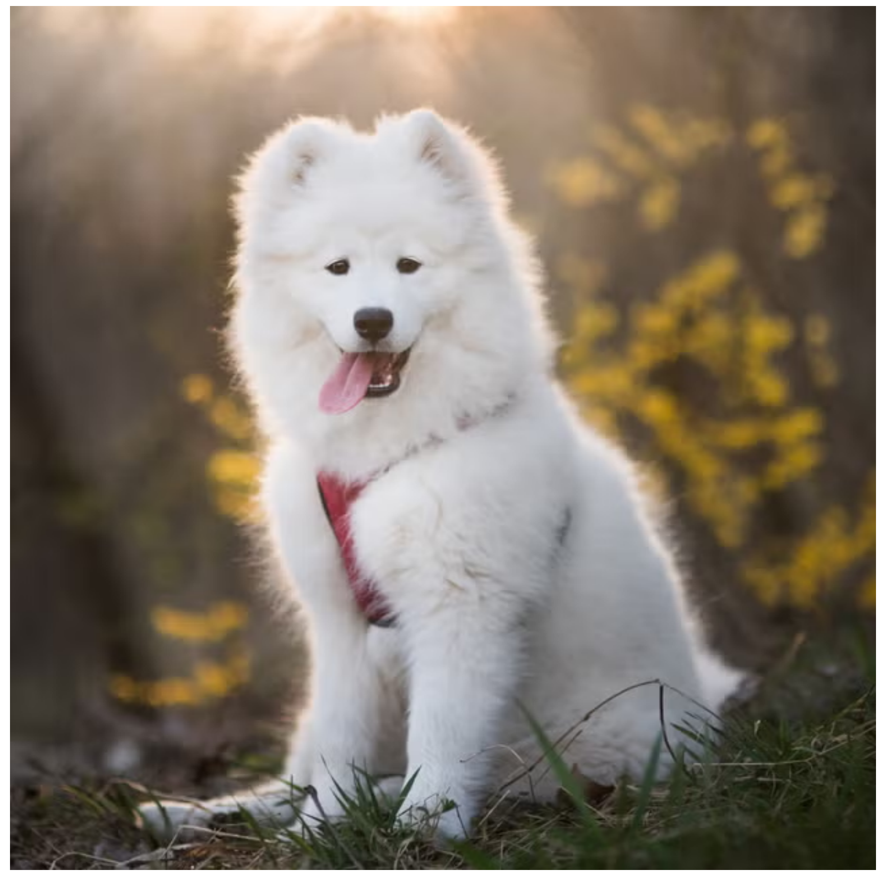


What is learning?

- Importantly, the learned knowledge should be **generalizable**
  - Applicable to the data that have not been observed yet

What is learning?

- **Why?** We simply can't **memorize** everything
  - Limited data: Cannot observe everything
  - Limited memory: Cannot remember everything
  - Limited compute: "Recalling" requires loading and comparing with all the data we have seen and remembered, which is extremely slow

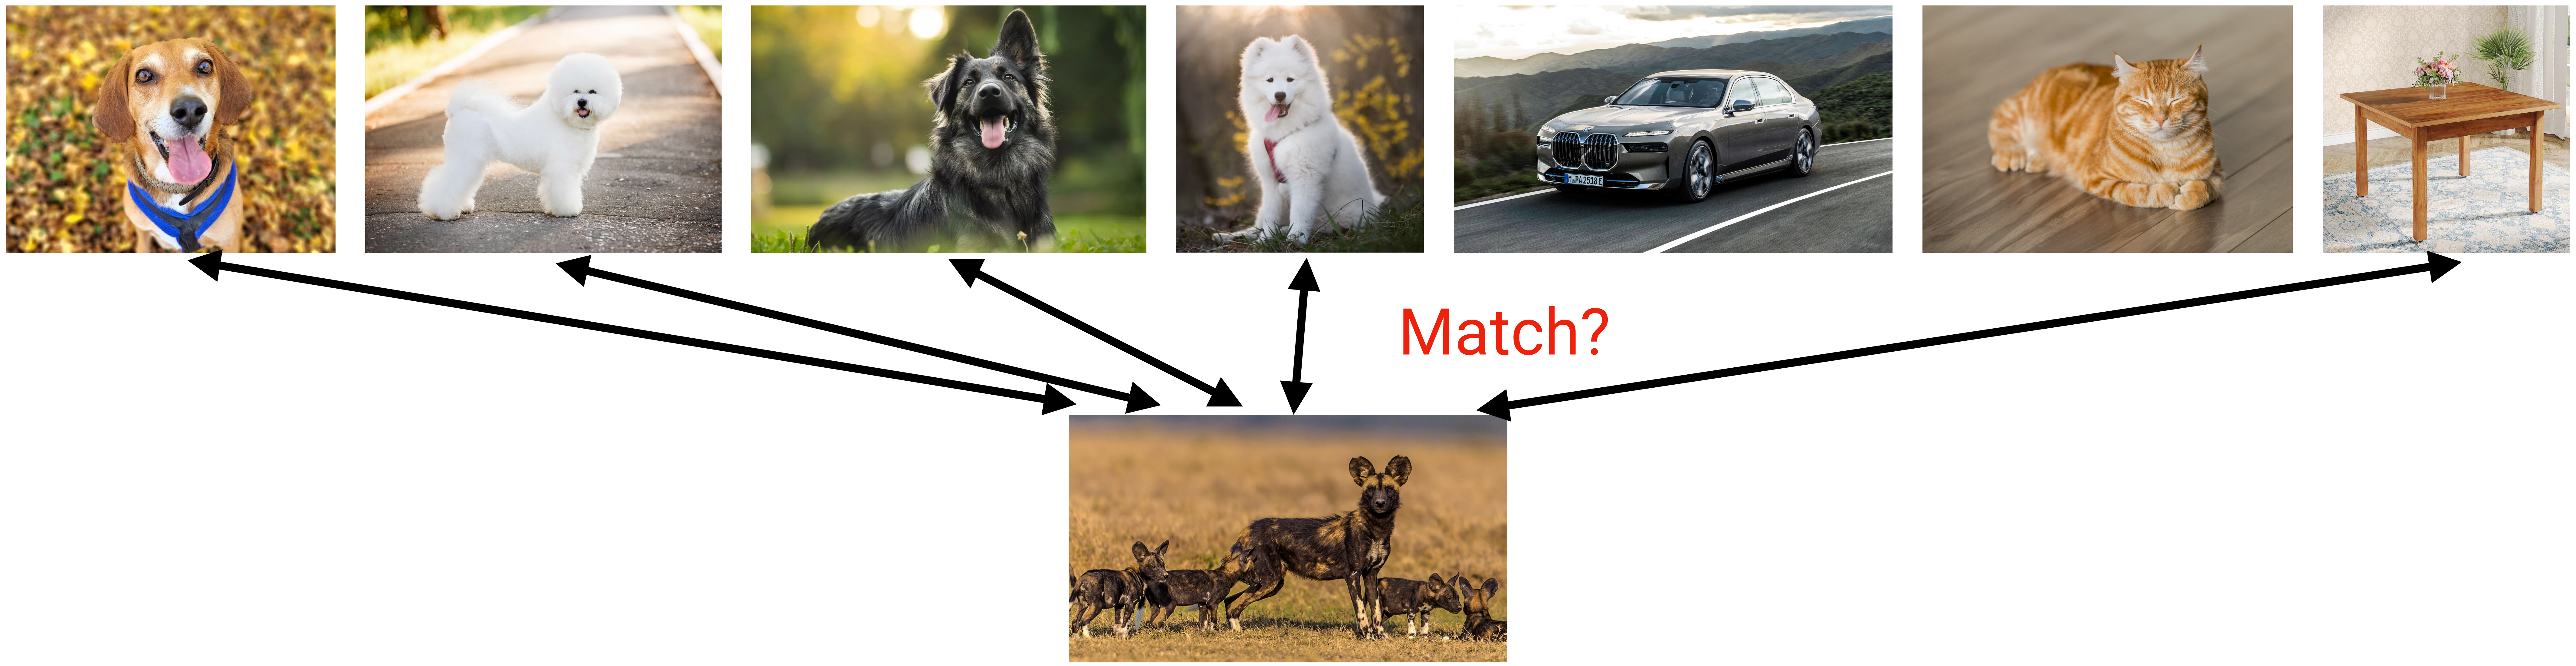

What is learning?

- To generalize, we want to find **patterns** from the observed samples
    - Simple yet widely applicable rule

What is learning?

- Based on the patterns learned, human **behaves**:
  - Predictions, Decisions, Actions

What is learning?

- Summing up, learning is the process of extracting and utilizing the patterns from the samples

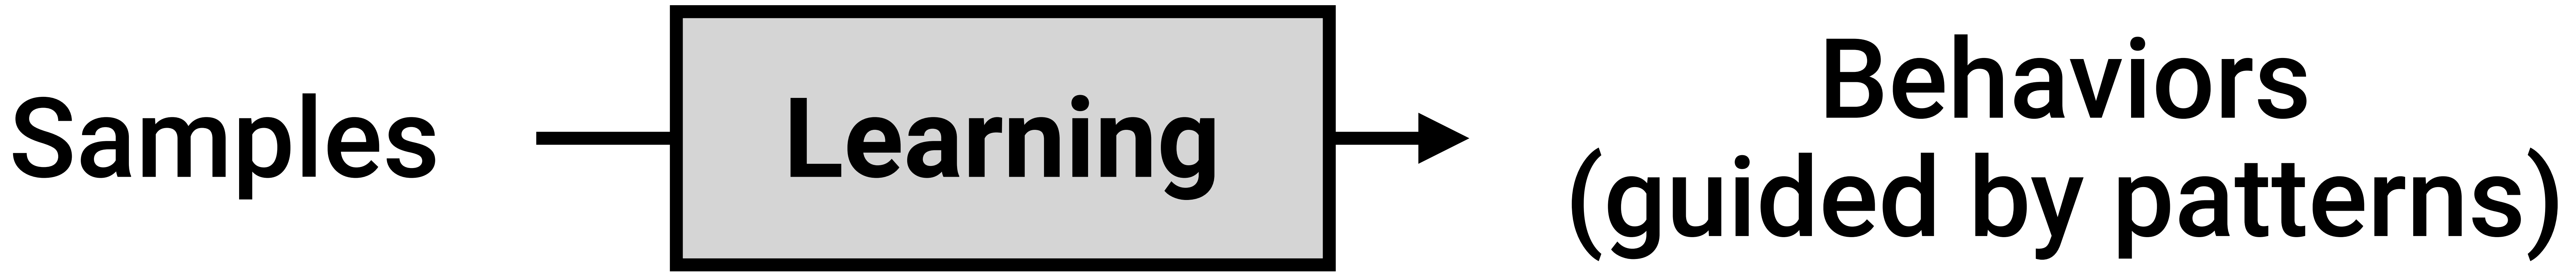

What is learning?

- Summing up, learning is the process of extracting and utilizing the patterns from the samples

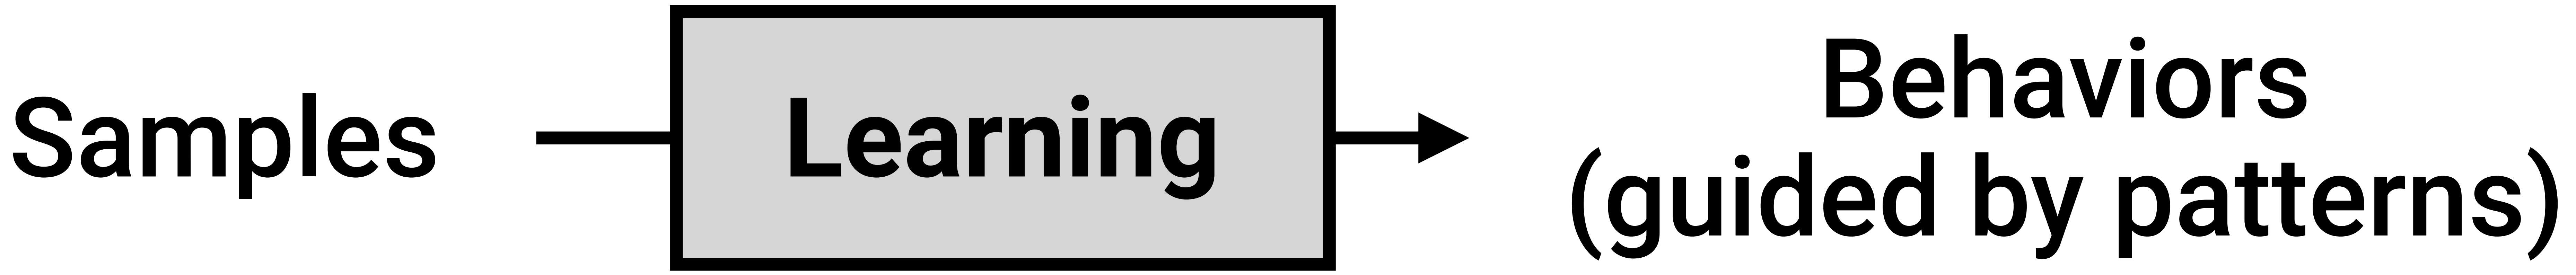

- **Machine Learning.** Letting a machine do this

Machine learning

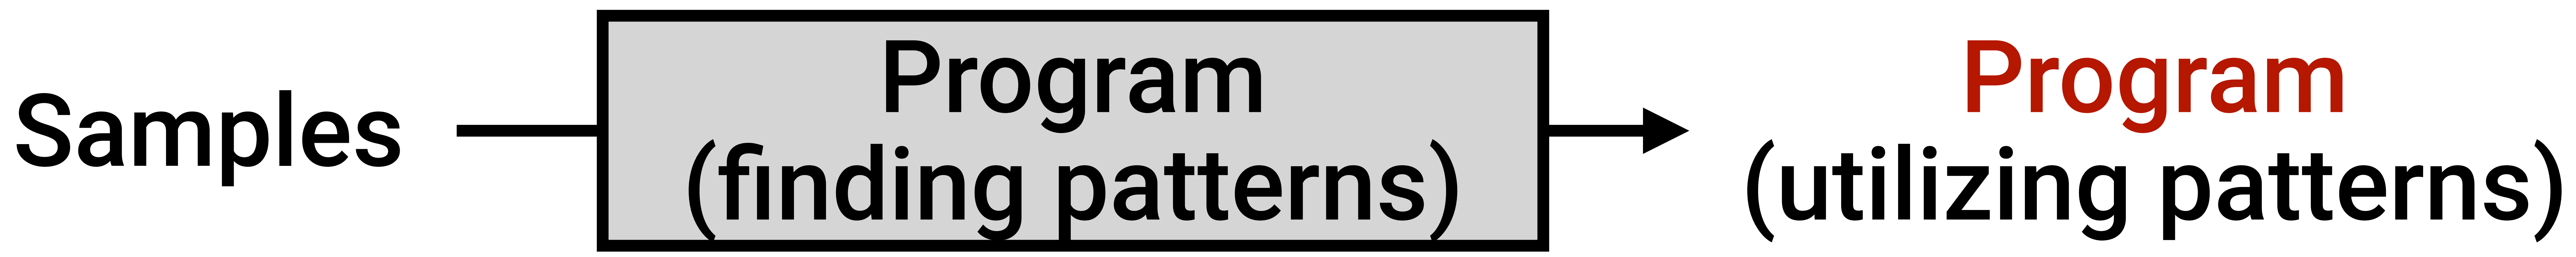

- Two programs in action:
  - **One** utilizes patterns
    - The program is called "**model**" (or "predictor" or "hypothesis")
    - Running the program is called "**inference**" (or "prediction")

Machine learning

- Two programs in action:
  - **One** utilizes patterns
    - The program is called "**model**" (or "predictor" or "hypothesis")
    - Running the program is called "**inference**" (or "prediction")
  - **Another** finds patterns from samples
    - The program is called "**learning algorithm**"
    - Running the program is called "**training**"

Machine learning

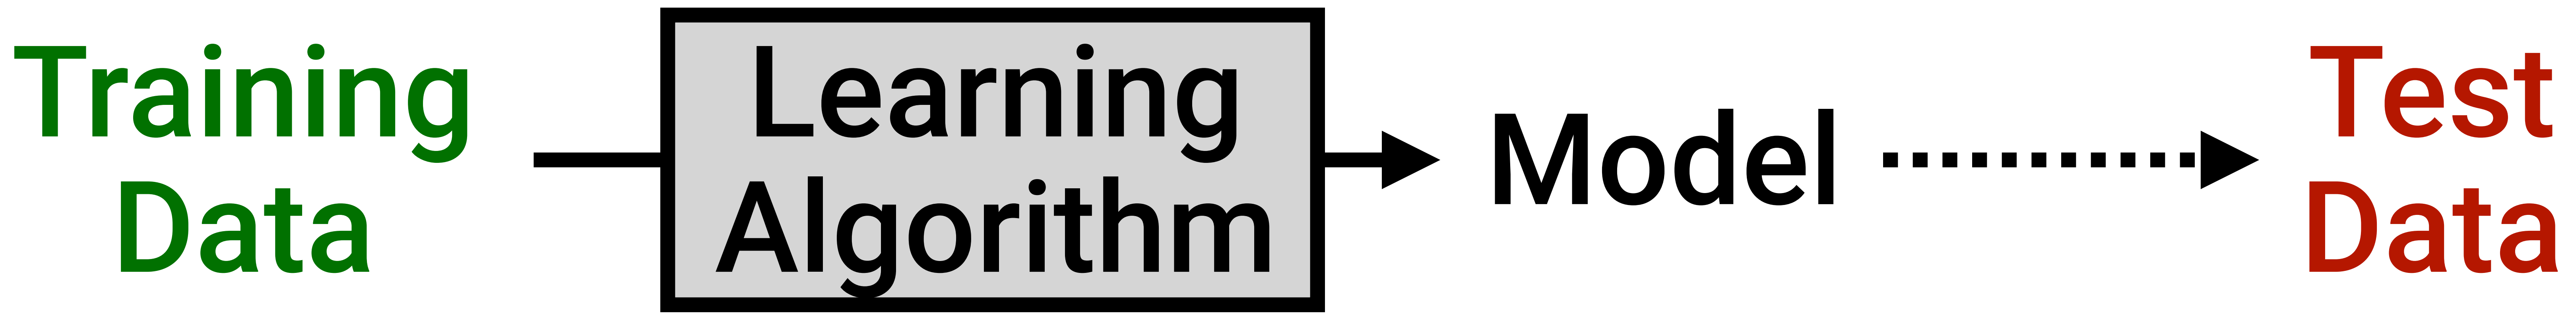

- The data that learning algorithm sees is called **training data**
  - The samples from which the patterns are discovered
- The data that model sees is called **test data**
  - Never observed during the training
  - The ultimate goal is to do well on this data!

Why machine learning?
(instead of human learning)

Why machine learning?

- In principle:
  - Machines are better at **utilizing** patterns (inference)
  - Machines are better at **finding** patterns (training)

(Caveat: it takes much effort to make this advantage happen)

Machine for inference

- Human are slow (e.g., high-frequency trading)

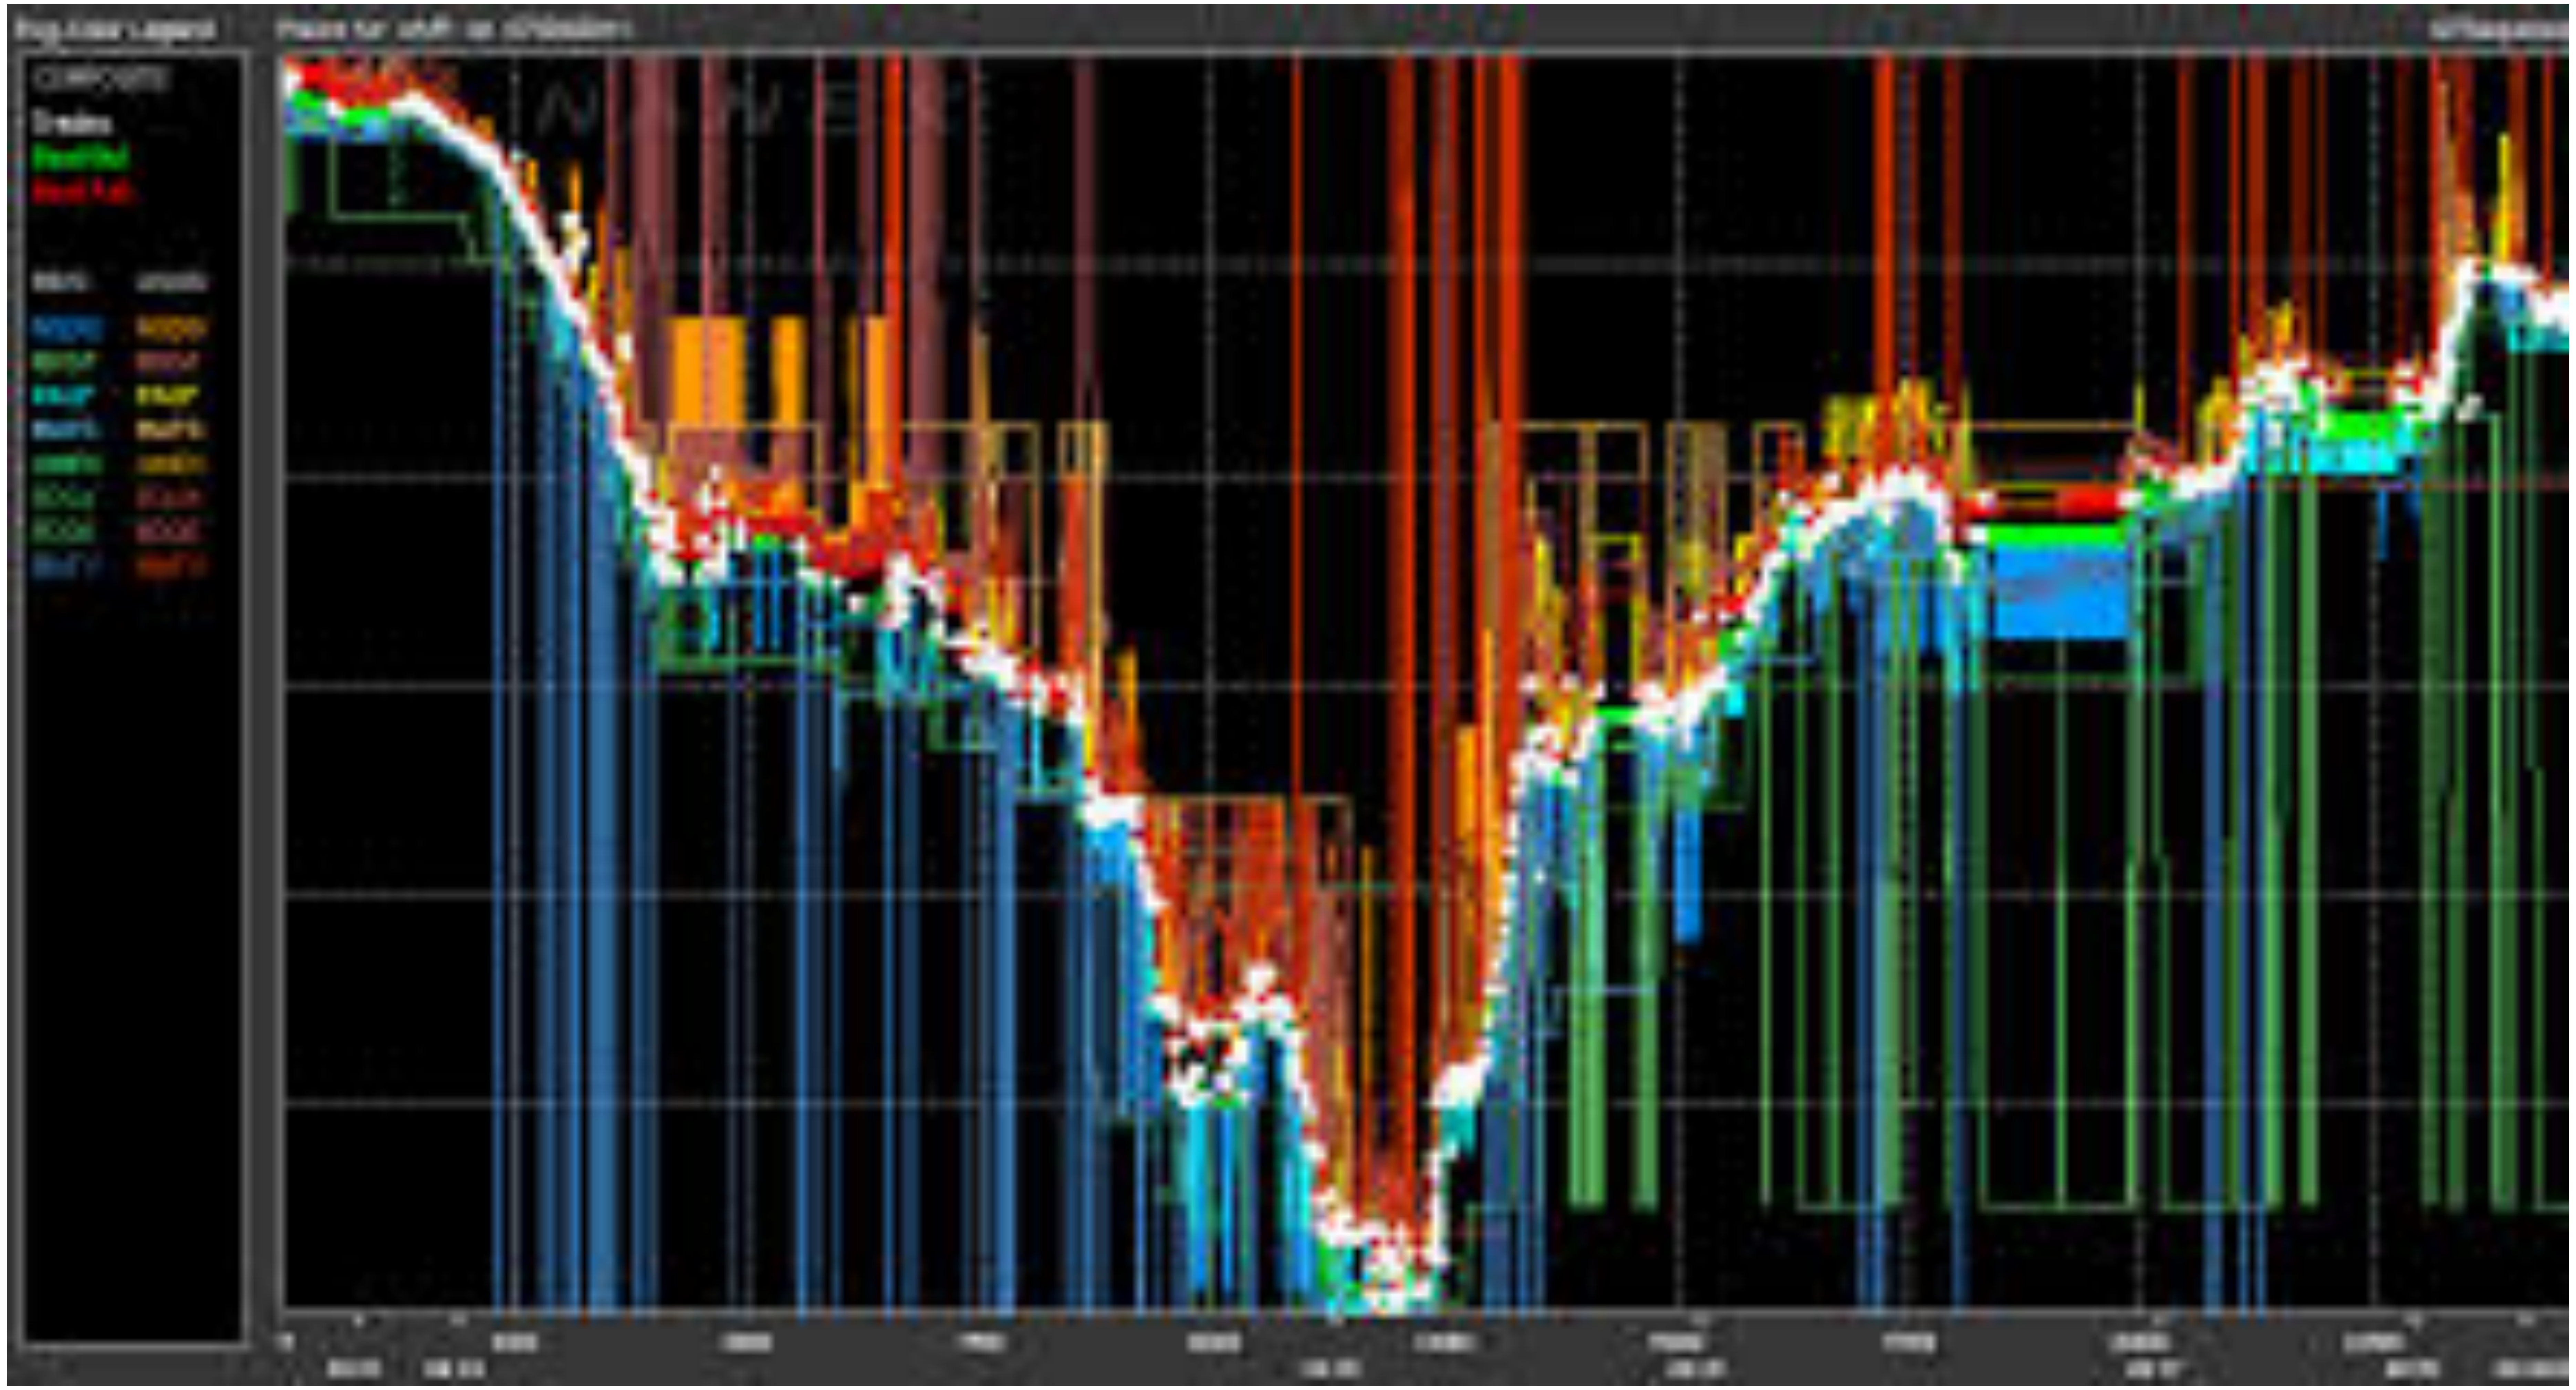

Machine for inference

- Human has limited or inconsistent sensory capabilities (e.g., self-driving)

Machine for inference

- Human are vulnerable (e.g., space mission / explosive disposal)

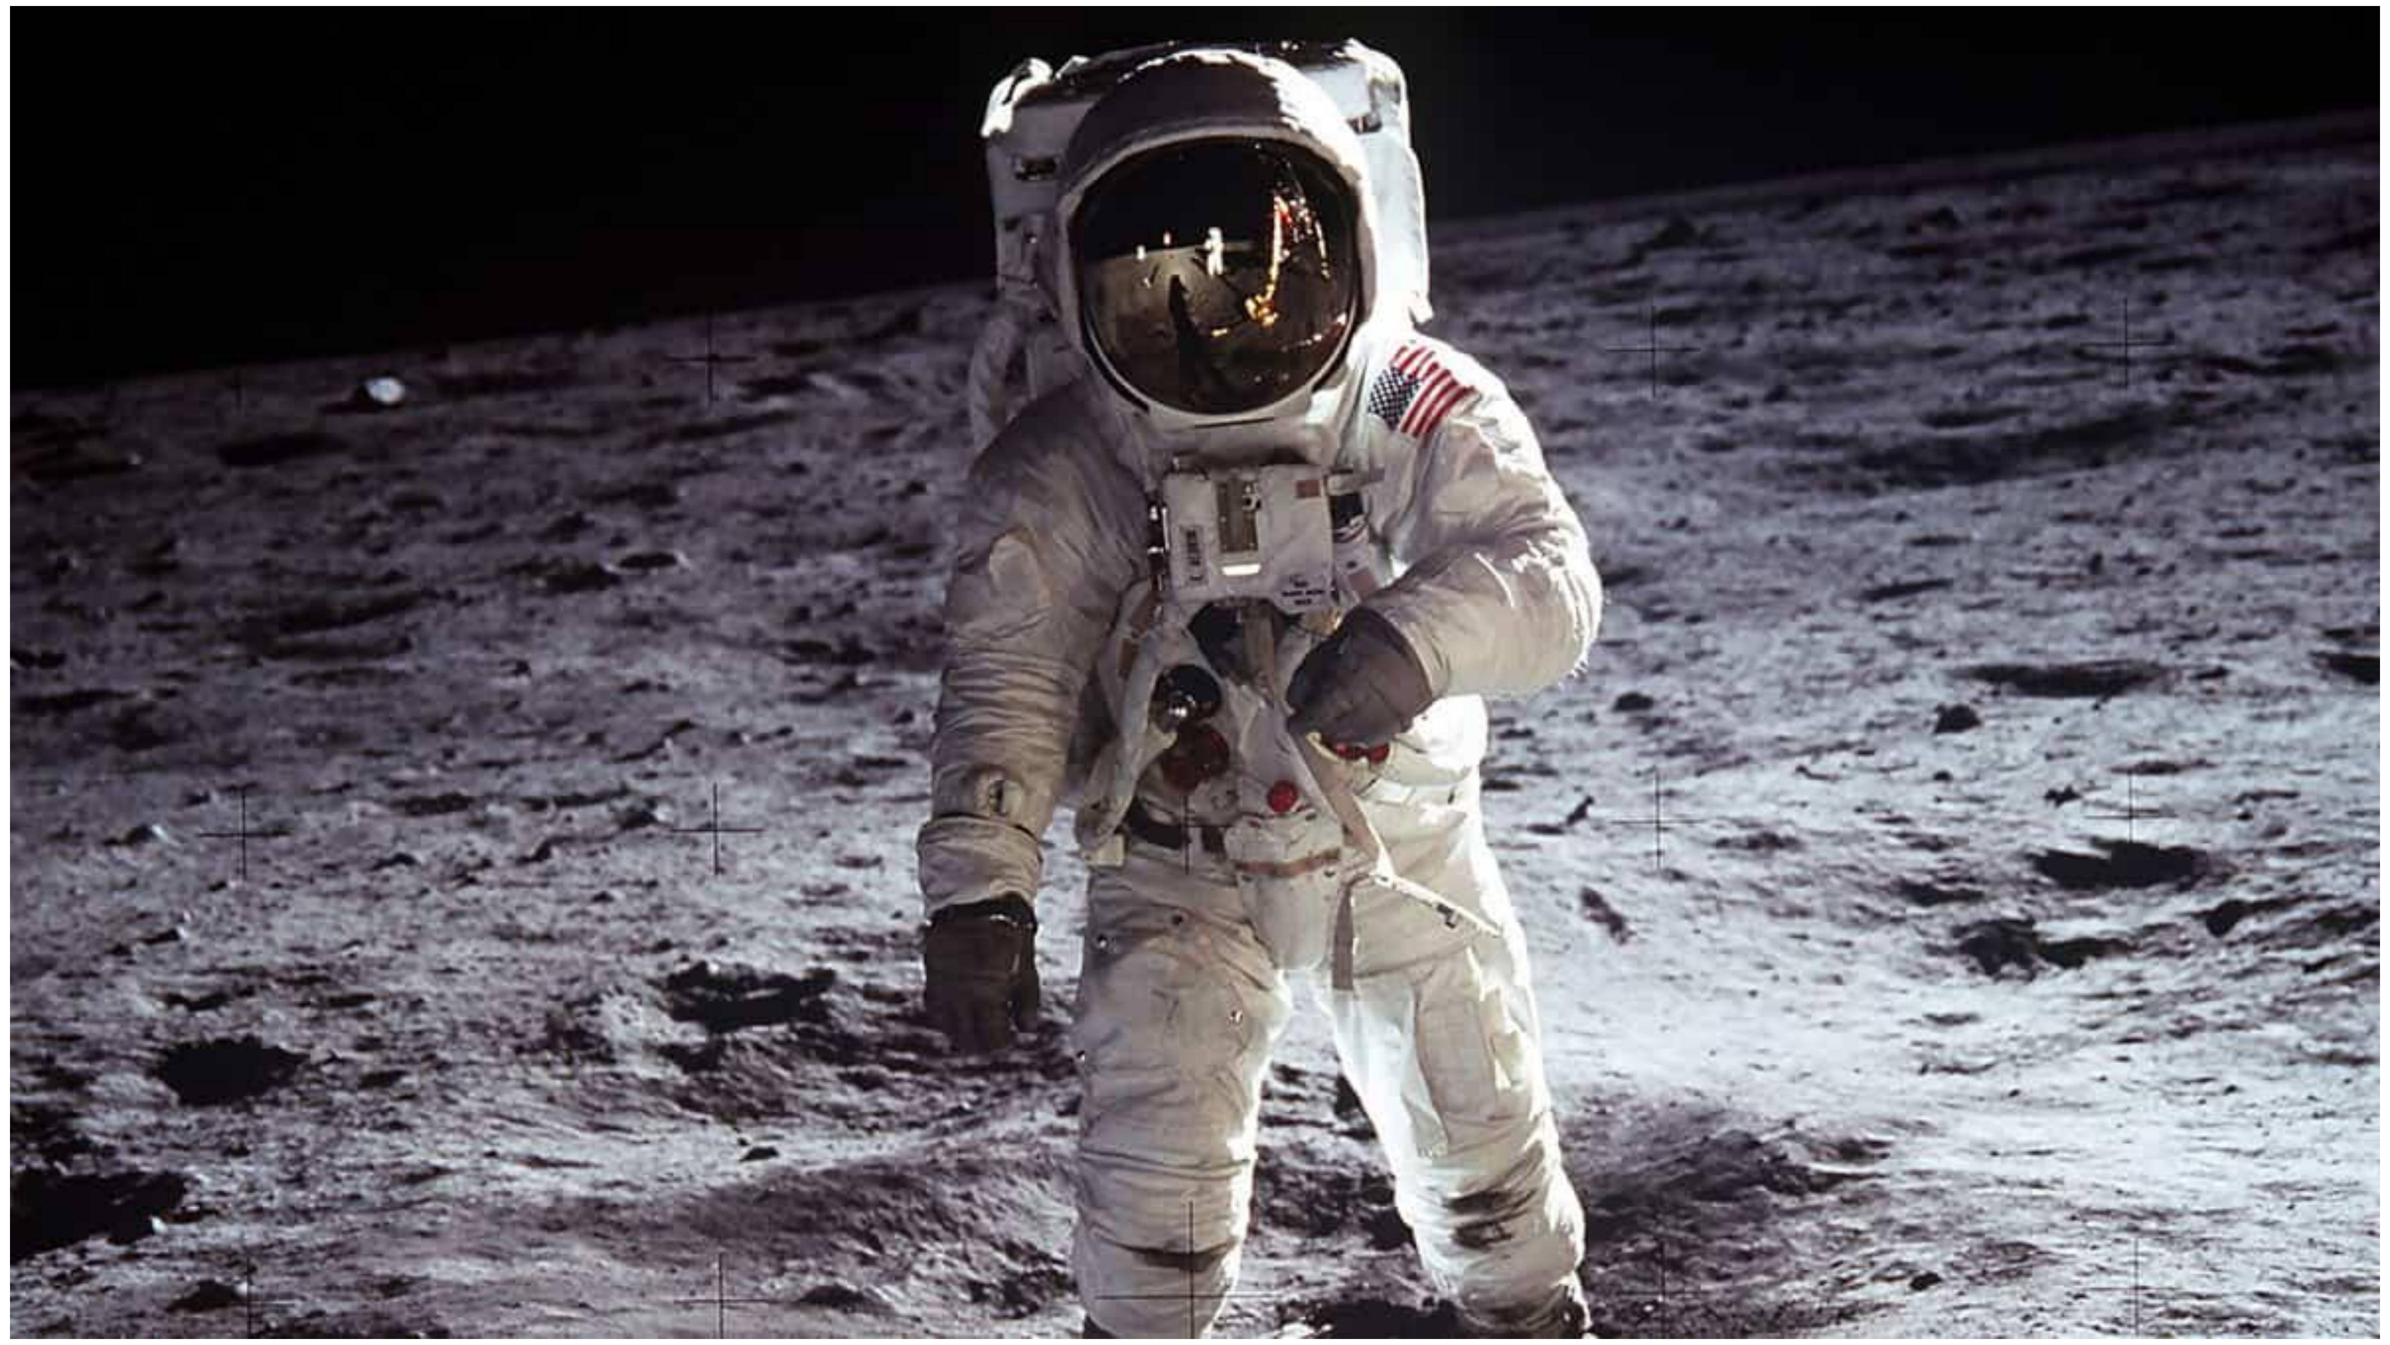

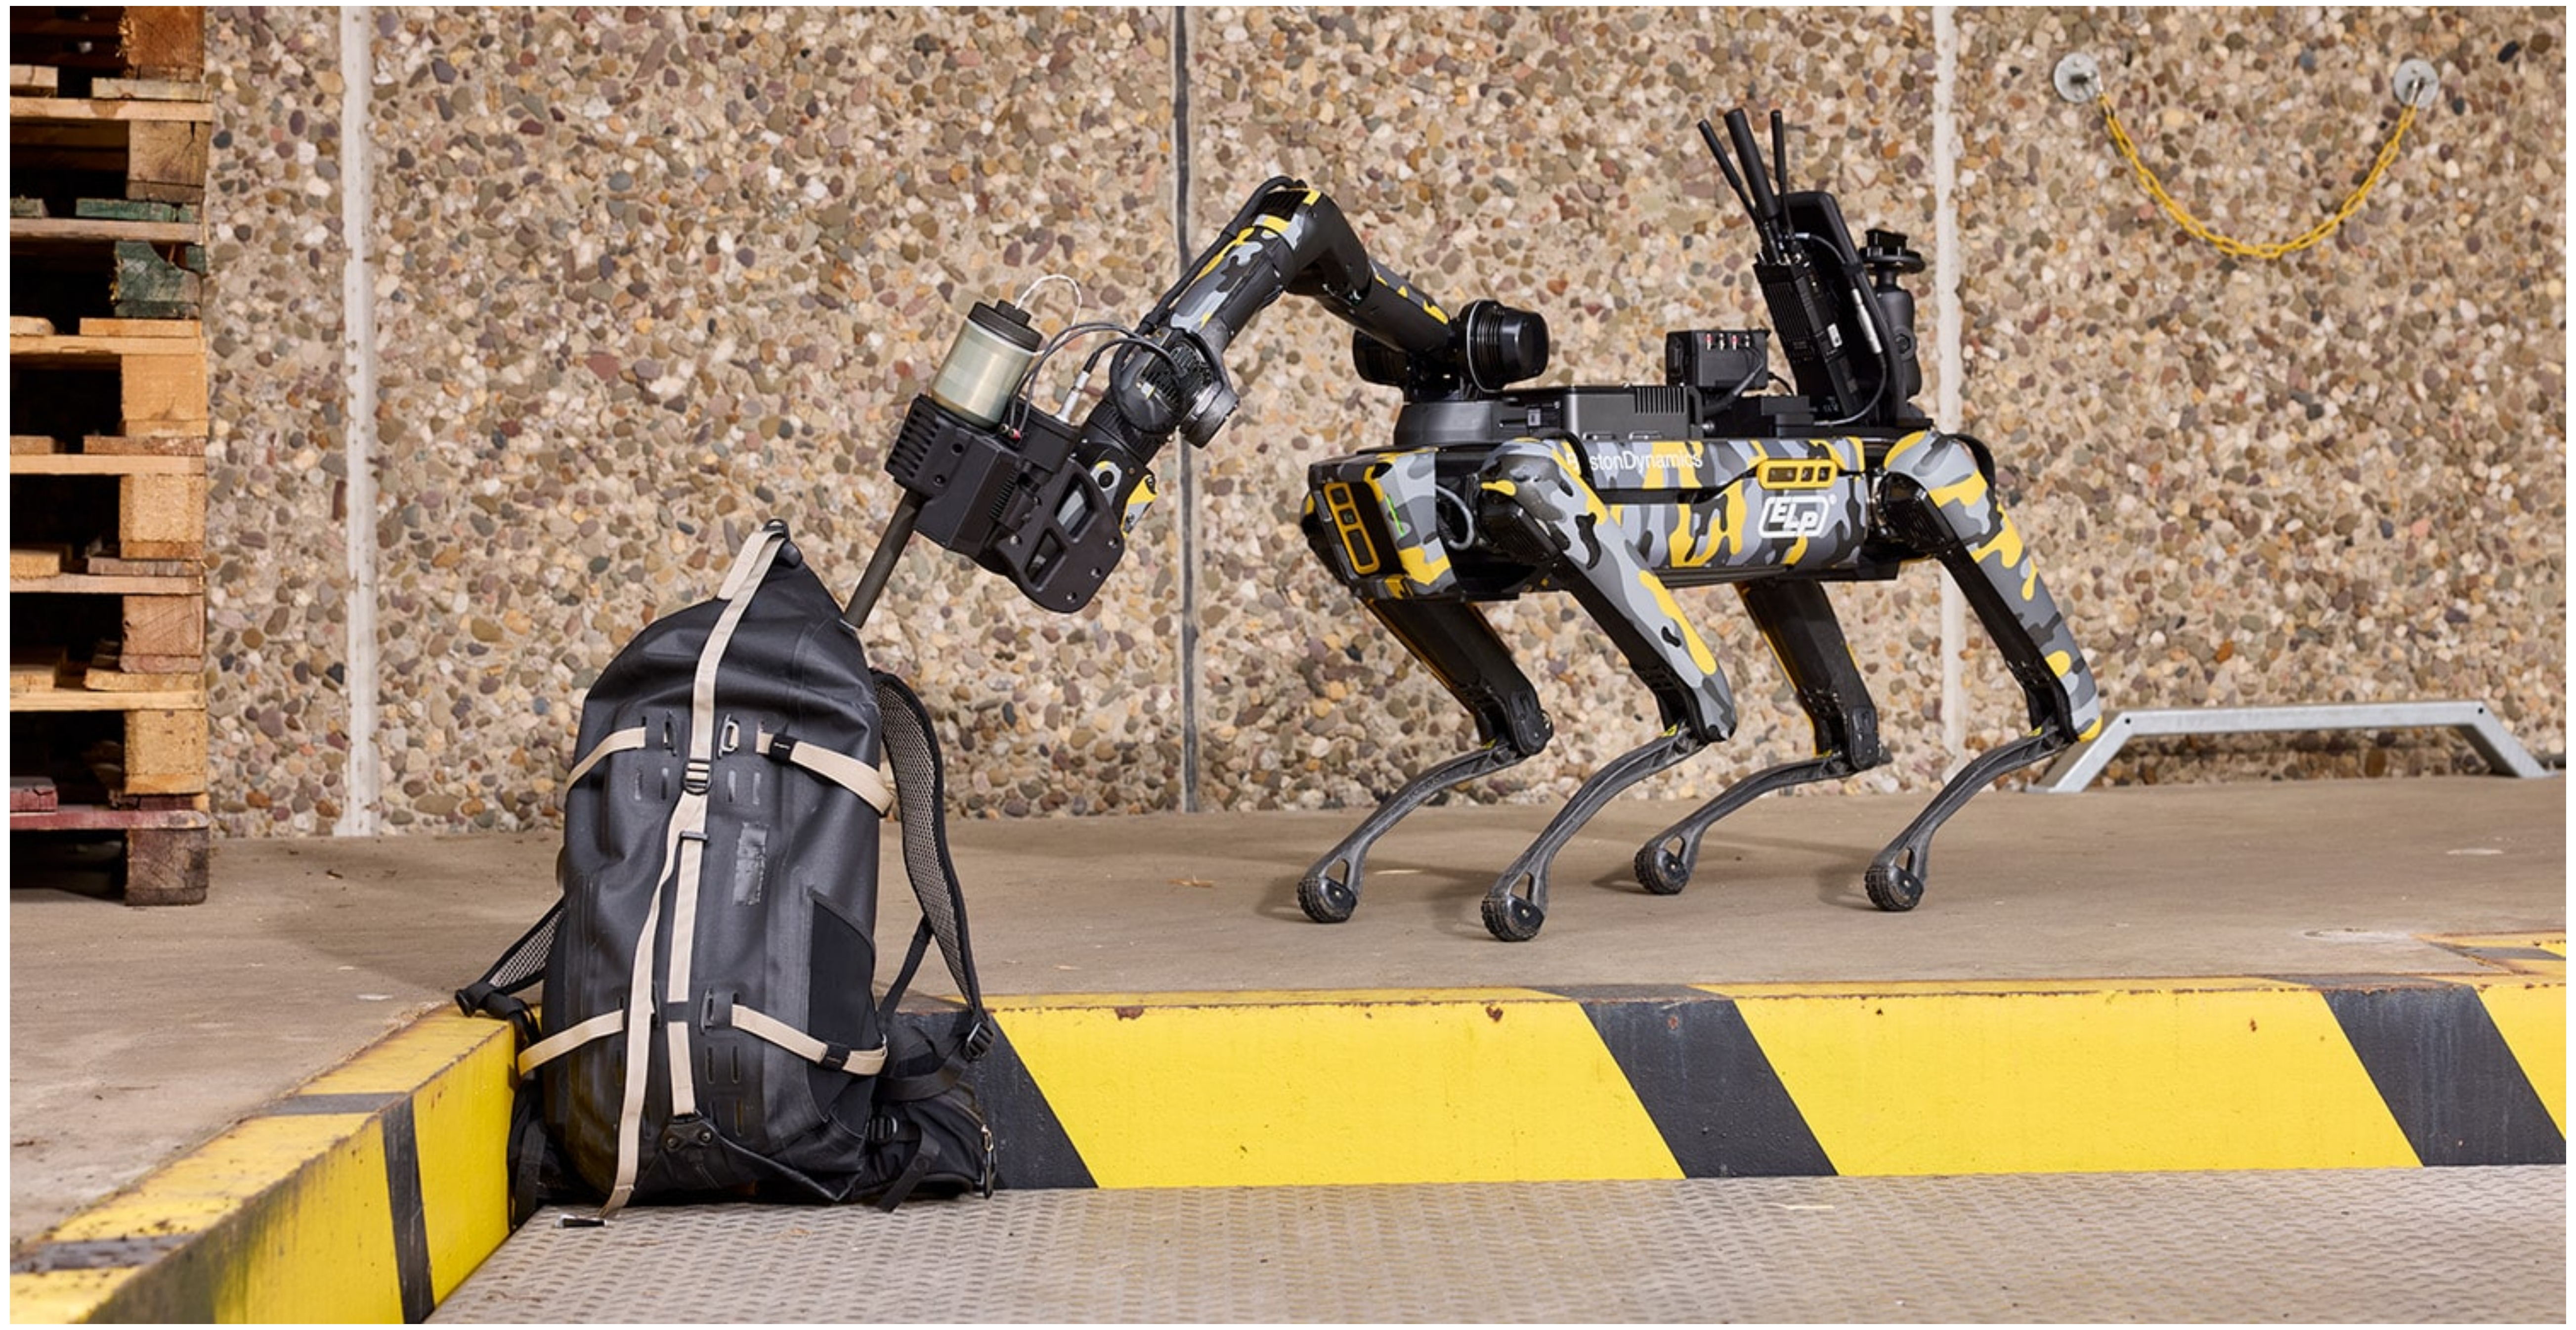

Machine for inference

- Human are expensive (e.g., manufacturing or drawing)

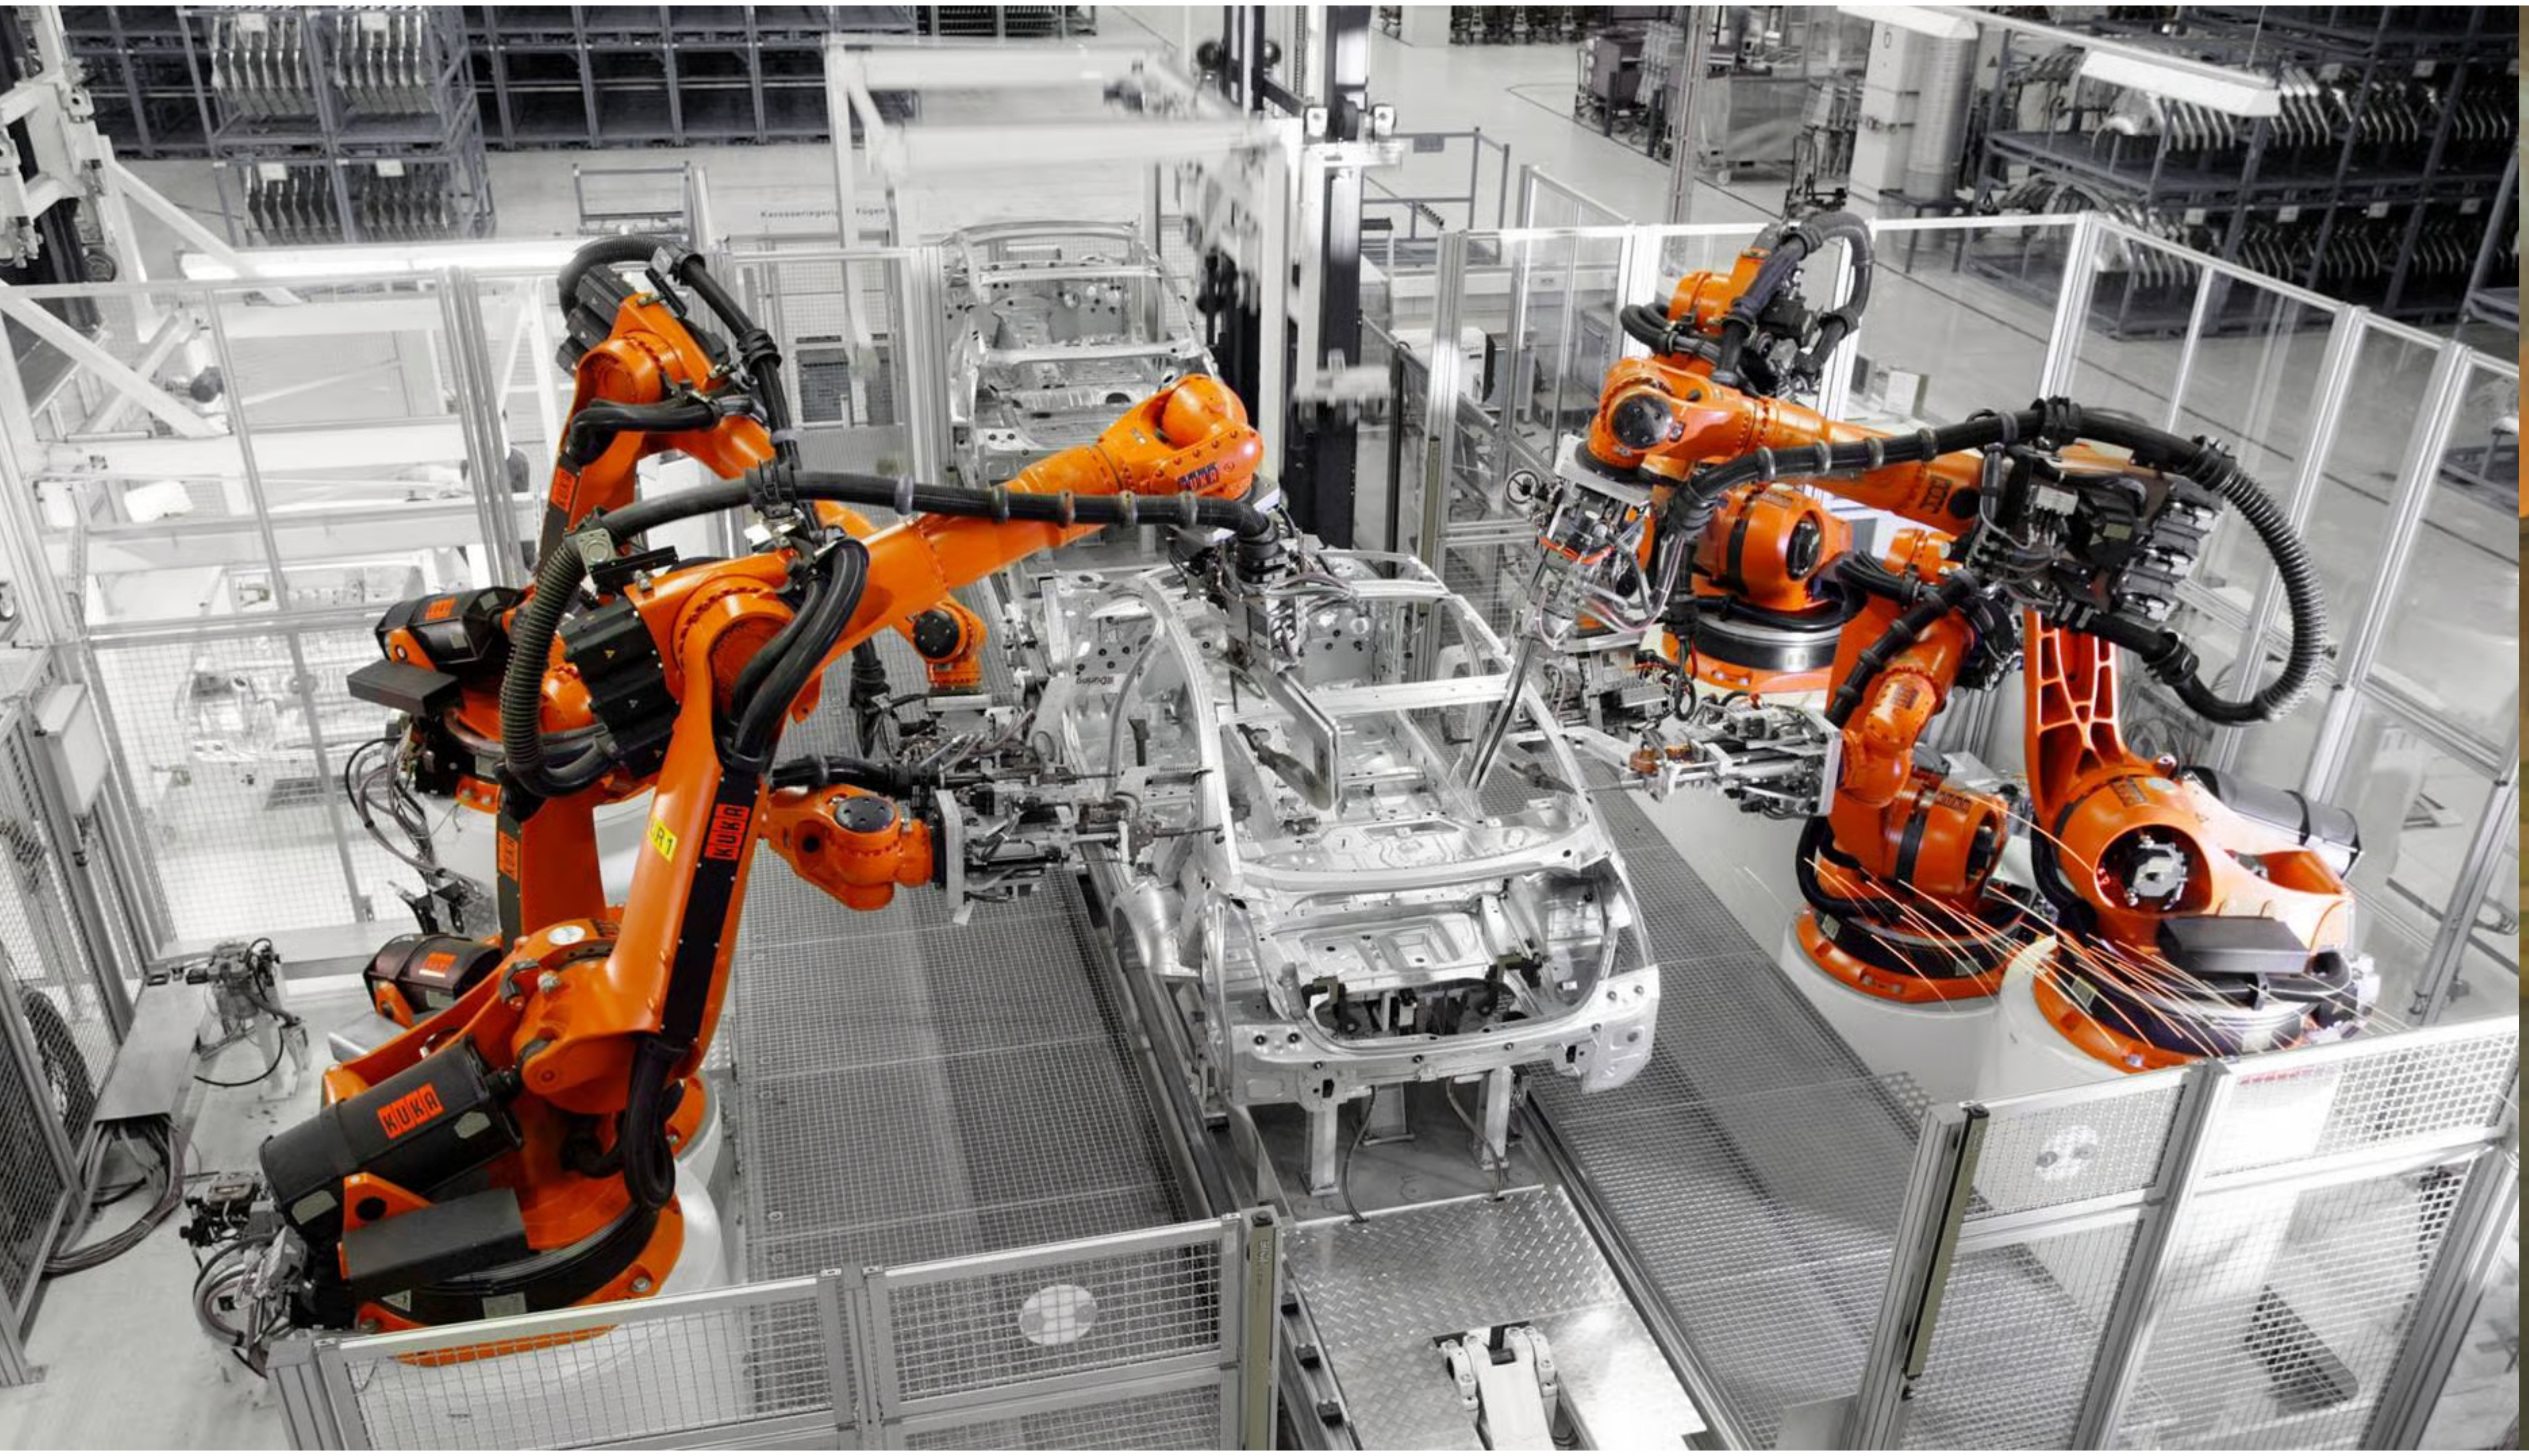

Machine for training

- Dataset is too big to be learned by a human (e.g., machine translation)

Machine for training

- Human are often guided by prejudice (e.g., AlphaGo)

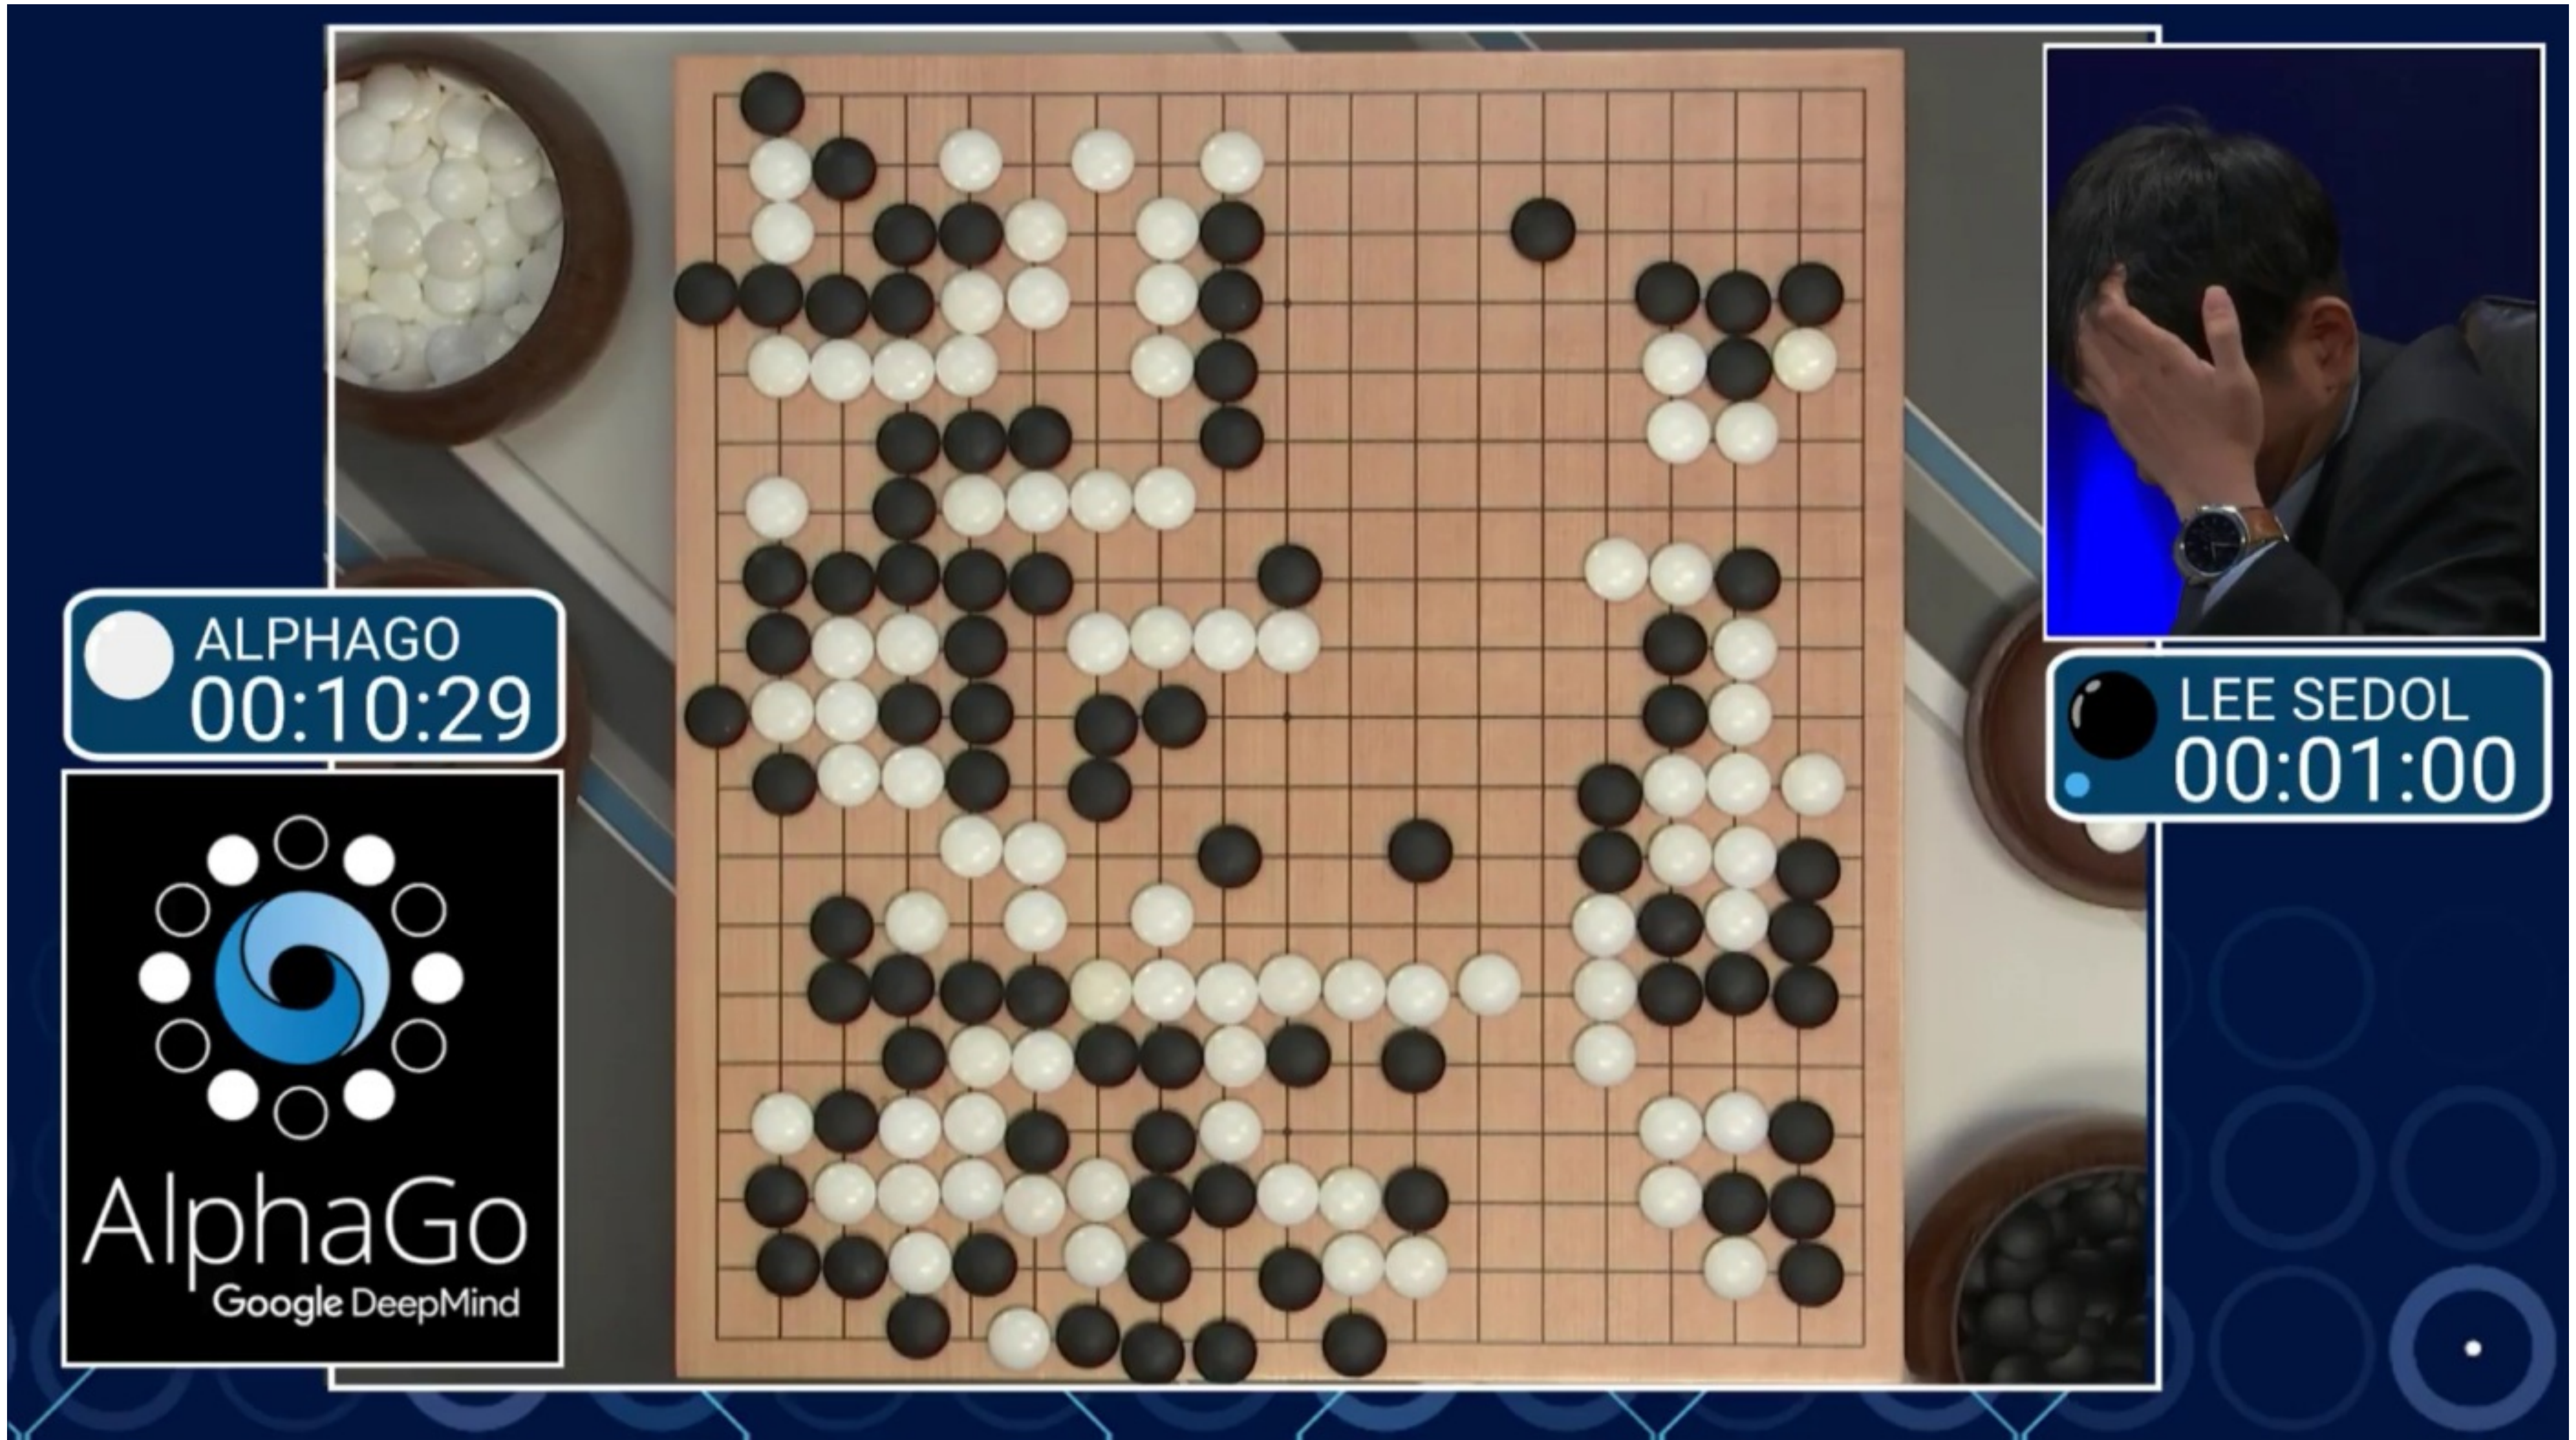

Machine for training

- Even when human can find better patterns, difficult to write it as a code:
  - How would you code up a “dog classifier?”

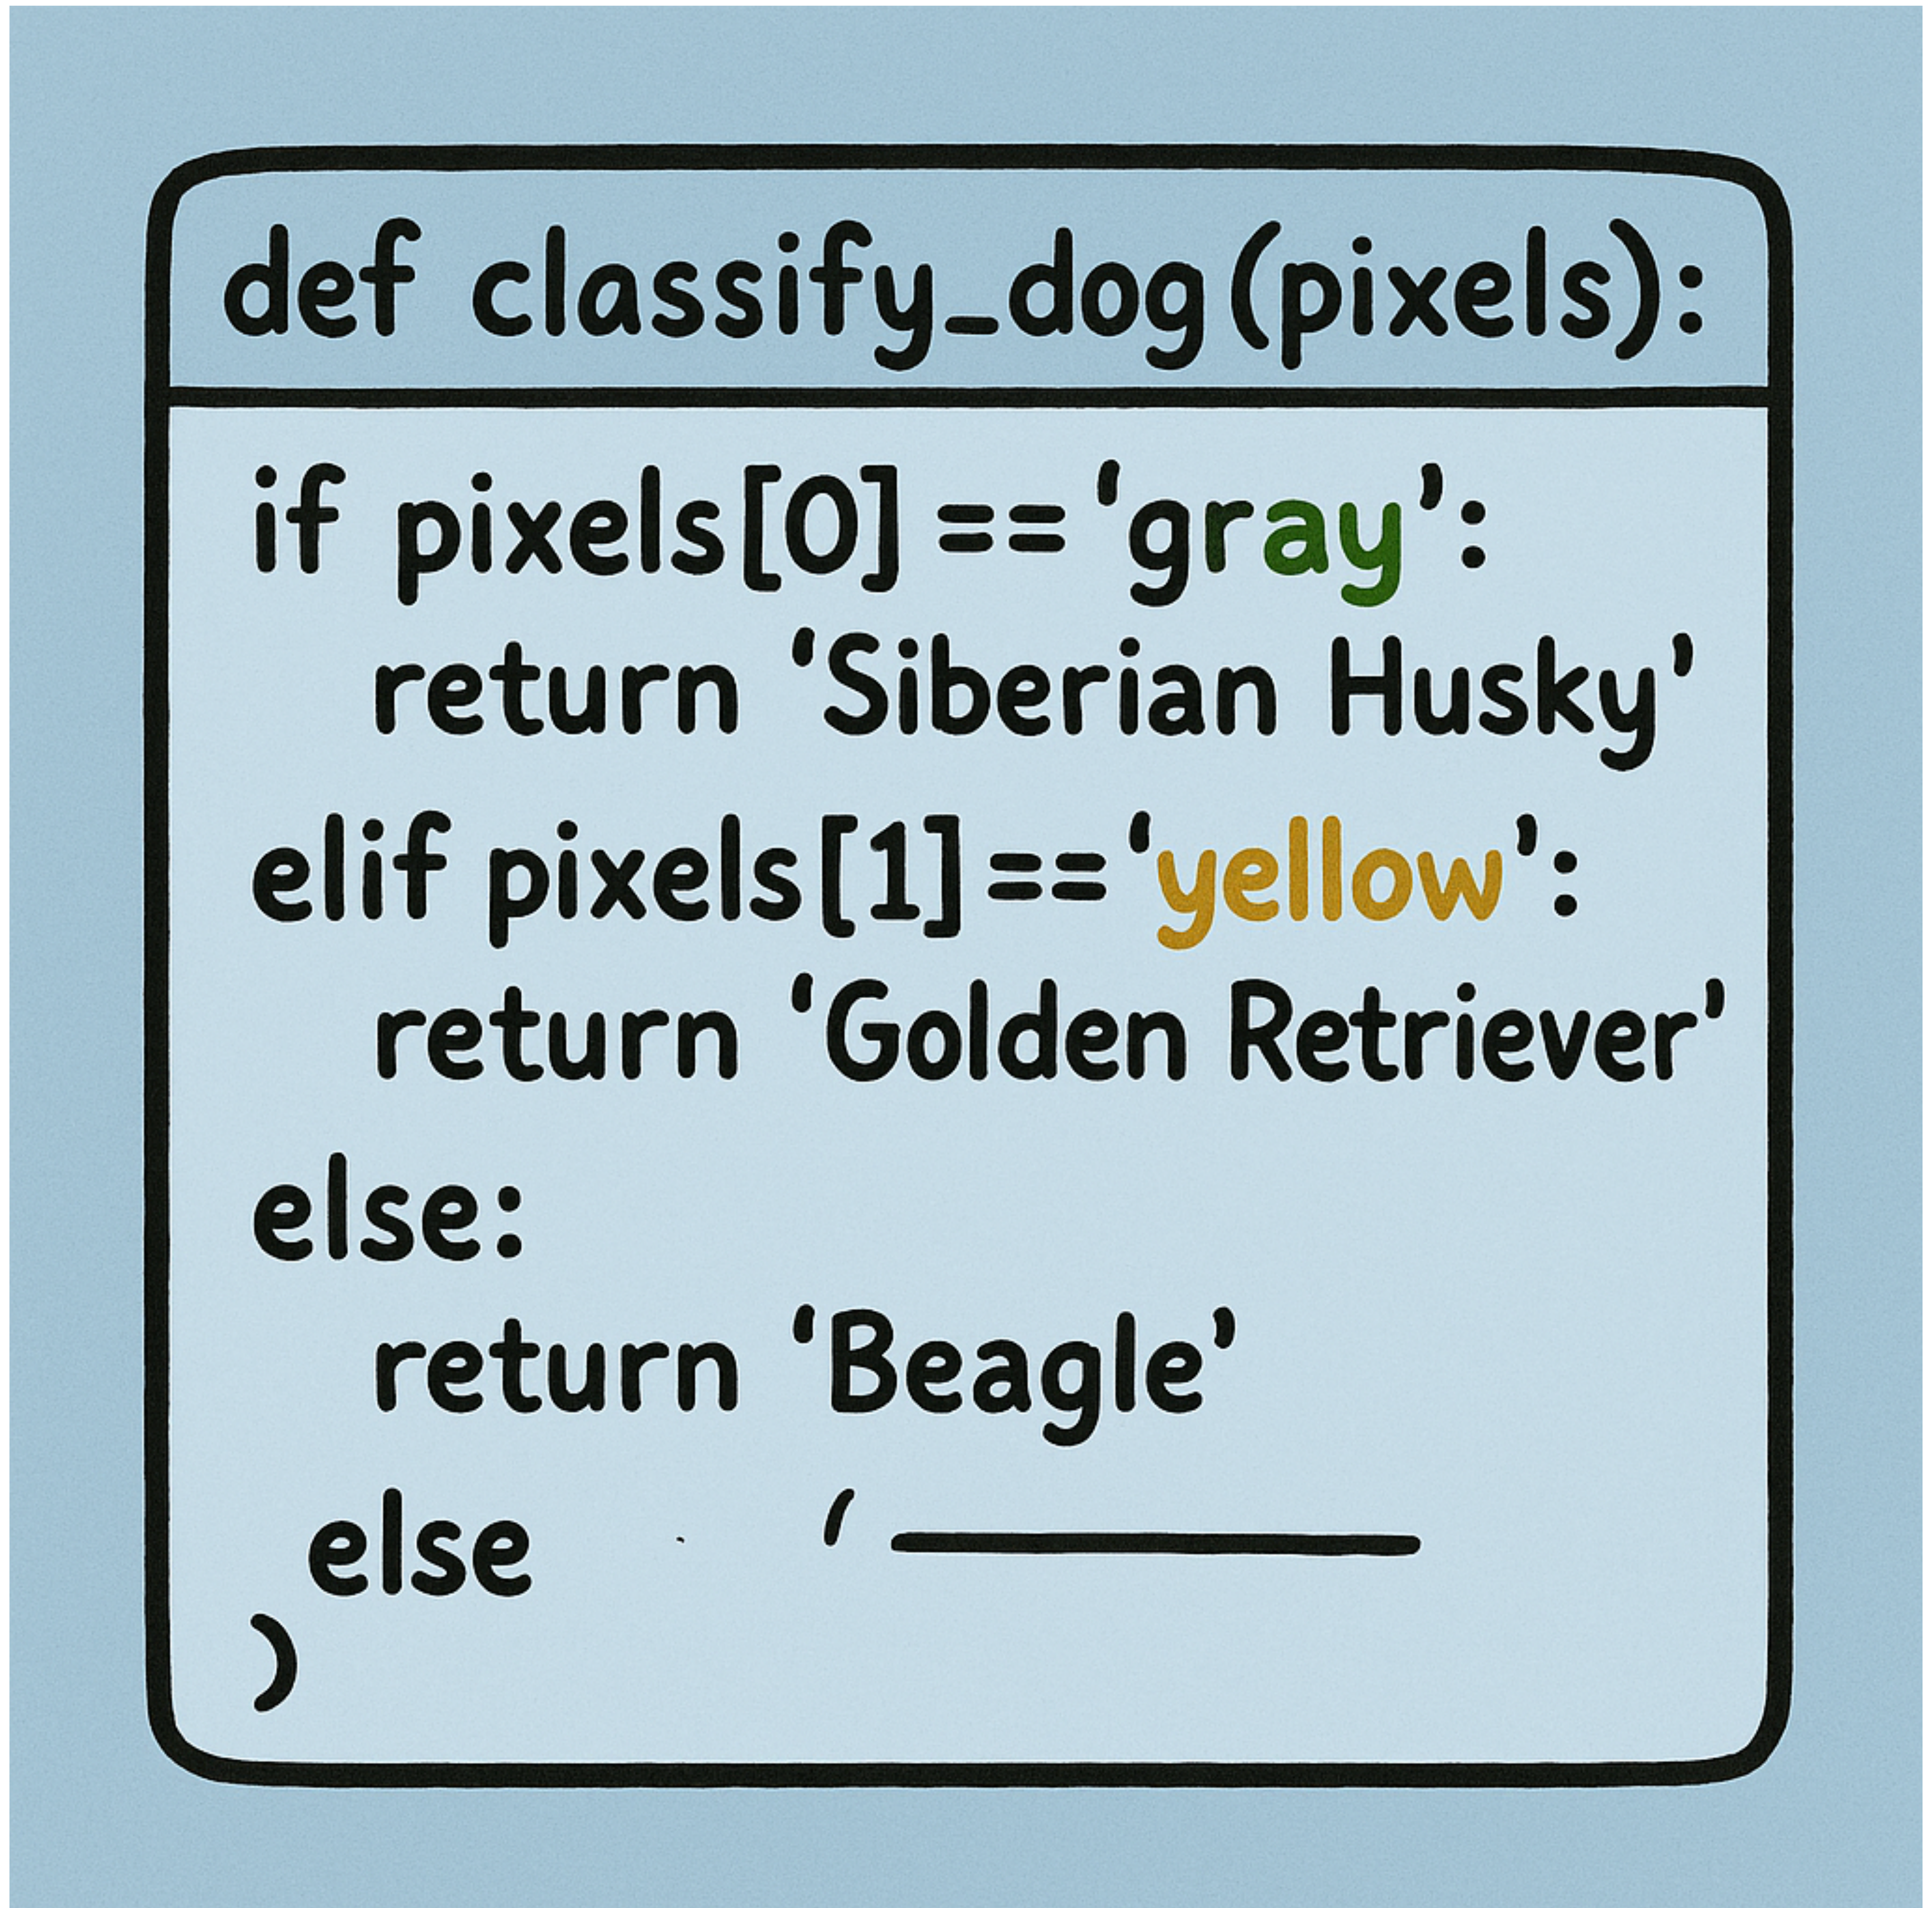

Why machine learning?

- Summing up, ML is useful because it can **scale up intelligence**
  - Super-human intelligence
    - using massive datasets
  - Massive deployment
    - inexpensive, consistent and robust
  - Can utilize massive sensory inputs
    - e.g., multiple sensors

“Human” in Machine Learning

- One thing to remember is that we **DON’T** expect

machine intelligence = human intelligence

  - Different sensory inputs
    - e.g., still don’t have sensors as dense as human
  - Machine should be better than human

"Human" in Machine Learning

- Rather, human is:
  - **Proof-of-concept.** that some tasks are indeed solvable
    - e.g., self-driving is solvable with pure vision, not radar
  - **Source of data / supervision.** a black-box we want to approximate
    - e.g., chatbots generating human-like responses
  - **End-user.** Someone who machine wants to assist and make happy
    - e.g., robotic pet

In this course

In this course…

- Study the **basics of machine learning**
  - Basic framework
    - Tasks, dataset, and mathematical theories
  - Algorithms
    - Classic algorithms (e.g., linear models)
    - Deep learning
    - Hands-on experiences
  - Frontiers
    - Application on specific domains (vision, language, robotics)
    - Challenges to be solved

Week-by-week

- Let's see the webpage: https://jaeho-lee.github.io/docs/teaching/fall25ml/

Broader scope

- Machine intelligence can be broken down into three parts:
  - This course (ML). Reasoning + Little bit of everything (Me)
  - **Computer Vision.** Perception (Prof. Kwang In Kim)
  - **Robot Learning.** Action (Prof. Hyemin Ahn)

Administrivia

Instructor

- Jaho Lee 이재호
    - jaeho.lee@postech.ac.kr
- **Positions**
    - Assistant Professor @ POSTECH EE 22.03–Present
    - Research Scientist @ Google 23.09–25.08
    - Ph.D. @ University of Illinois Urbana-Champaign
- **Roles**
    - Lectures, Q&A

Teaching Assistant

- Minseok Kim 김민석
  - kms2914@postech.ac.kr
  - M.S. candidate in EE
- Taesun Yeom 염태선
  - tsyeom@postech.ac.kr
  - Ph.D. candidate in AI
- **Roles**
  - Attendance, Grading, Assignments

Attendance

- Please use the **electronic attendance** system.
    - Counted begins on September 11.

- You'll get an **F** if you miss **more than 3** classes
    - "I forgot to … " doesn't mean it's okay
    - "I made mistakes in electronic attendance" does not count
    - Your research group's business doesn't mean it is okay to miss

Location & Hours

- **Lectures**
  - LG Hall 105
  - Mondays & Wednesdays, 9:30AM — 11:00AM
- **Office Hours**
  - GoAround Coffee, RIST
  - Wednesdays, 5:00PM — 6:00 PM
- **Materials**
  - https://jaeho-lee.github.io Lecture notes
  - PLMS Assignments

Prerequisites

- **I'll assume that you know:**
  - Calculus
  - Basic linear algebra
  - Basic probability & random variables
  - Signals & Systems
  - Python programming

Textbook

- **Main**
  - Lecture slides

- **Supplementary**
  - Mathematics for Machine Learning https://mml-book.github.io/
  - Understanding Deep Learning https://udlbook.github.io/udlbook/
  - Patterns, Predictions, and Actions https://mlstory.org/
  - Dive into Deep Learning https://d2l.ai/

Grading

- **Attendance (10%)**
    - Use electronic attendance system
- **Assignments (30%)**
    - Planning to give you at least 3 homeworks
- **Mid-Term (30%)**
- **Final Project (30%)**
    - Deep learning project

* Graduate students are graded separately

* QE sit-ins will be evaluated against UGs

Final Project

- Local Kaggle Competition https://www.kaggle.com/competitions
  - Presentations in the final week

Competitions

Search competitions Filters

Featured ✕

Results Recently Launched

NeurIPS - Ariel Data Challenge 2025 $50,000
Derive exoplanet signals from Ariel's optical instruments
Featured · Code Competition · 51 Teams · 3 months to go

Google - The Gemma 3n Impact Challenge $150,000
Explore the newest Gemma model and build your best products for a better world
Featured · 0 Teams · A month to go

NeurIPS - Open Polymer Prediction 2025 $50,000
Predicting polymer properties with machine learning to accelerate sustainable materials research.
Featured · Code Competition · 794 Teams · 3 months to go

CMI - Detect Behavior with Sensor Data $50,000
Predicting Body Focused Repetitive Behaviors from a Wrist-Worn Device
Featured · Code Competition · 1386 Teams · 2 months to go

Honor codes

- You’ll get **F** if:
  - Sharing solutions
  - Copying solutions
  - “Collaborating” with your friends for homeworks
    - Use me & TA, instead
  - Using GPT for homework
  - Miss **more than 3** classes