1. Course Overview

EECE454 Introduction to Machine Learning Systems

2023 Fall, Jaeho Lee

What is Machine Learning?

Rough Idea - Finding Patterns

- Given some examples, human can find a pattern.

  ⇒ **Machine Learning.** Machine finds & use this pattern.

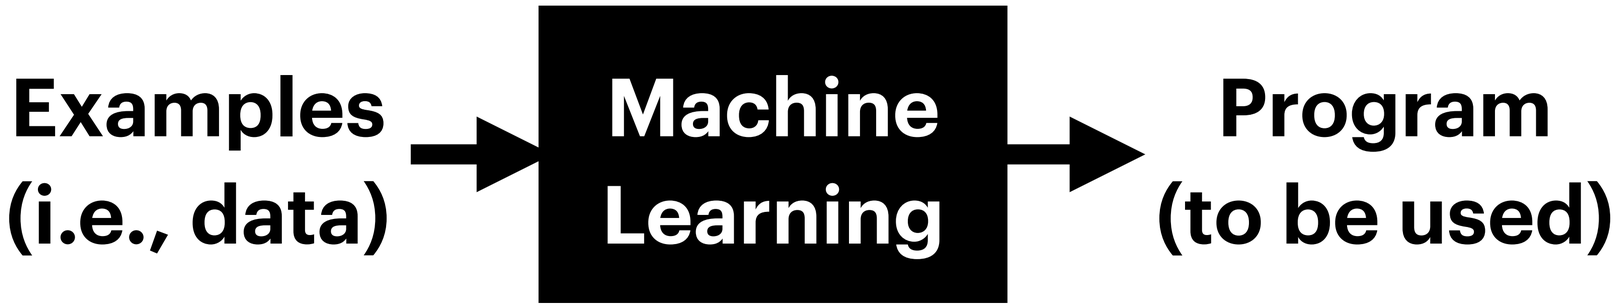

Example. Create a program that, given an image of a dog, returns the name of the dog specie.

Example. Create a program that, given a Netflix user and a movie, returns the expected user rating.

Example. Create a program that, given a text input, generate a “human-like” response.

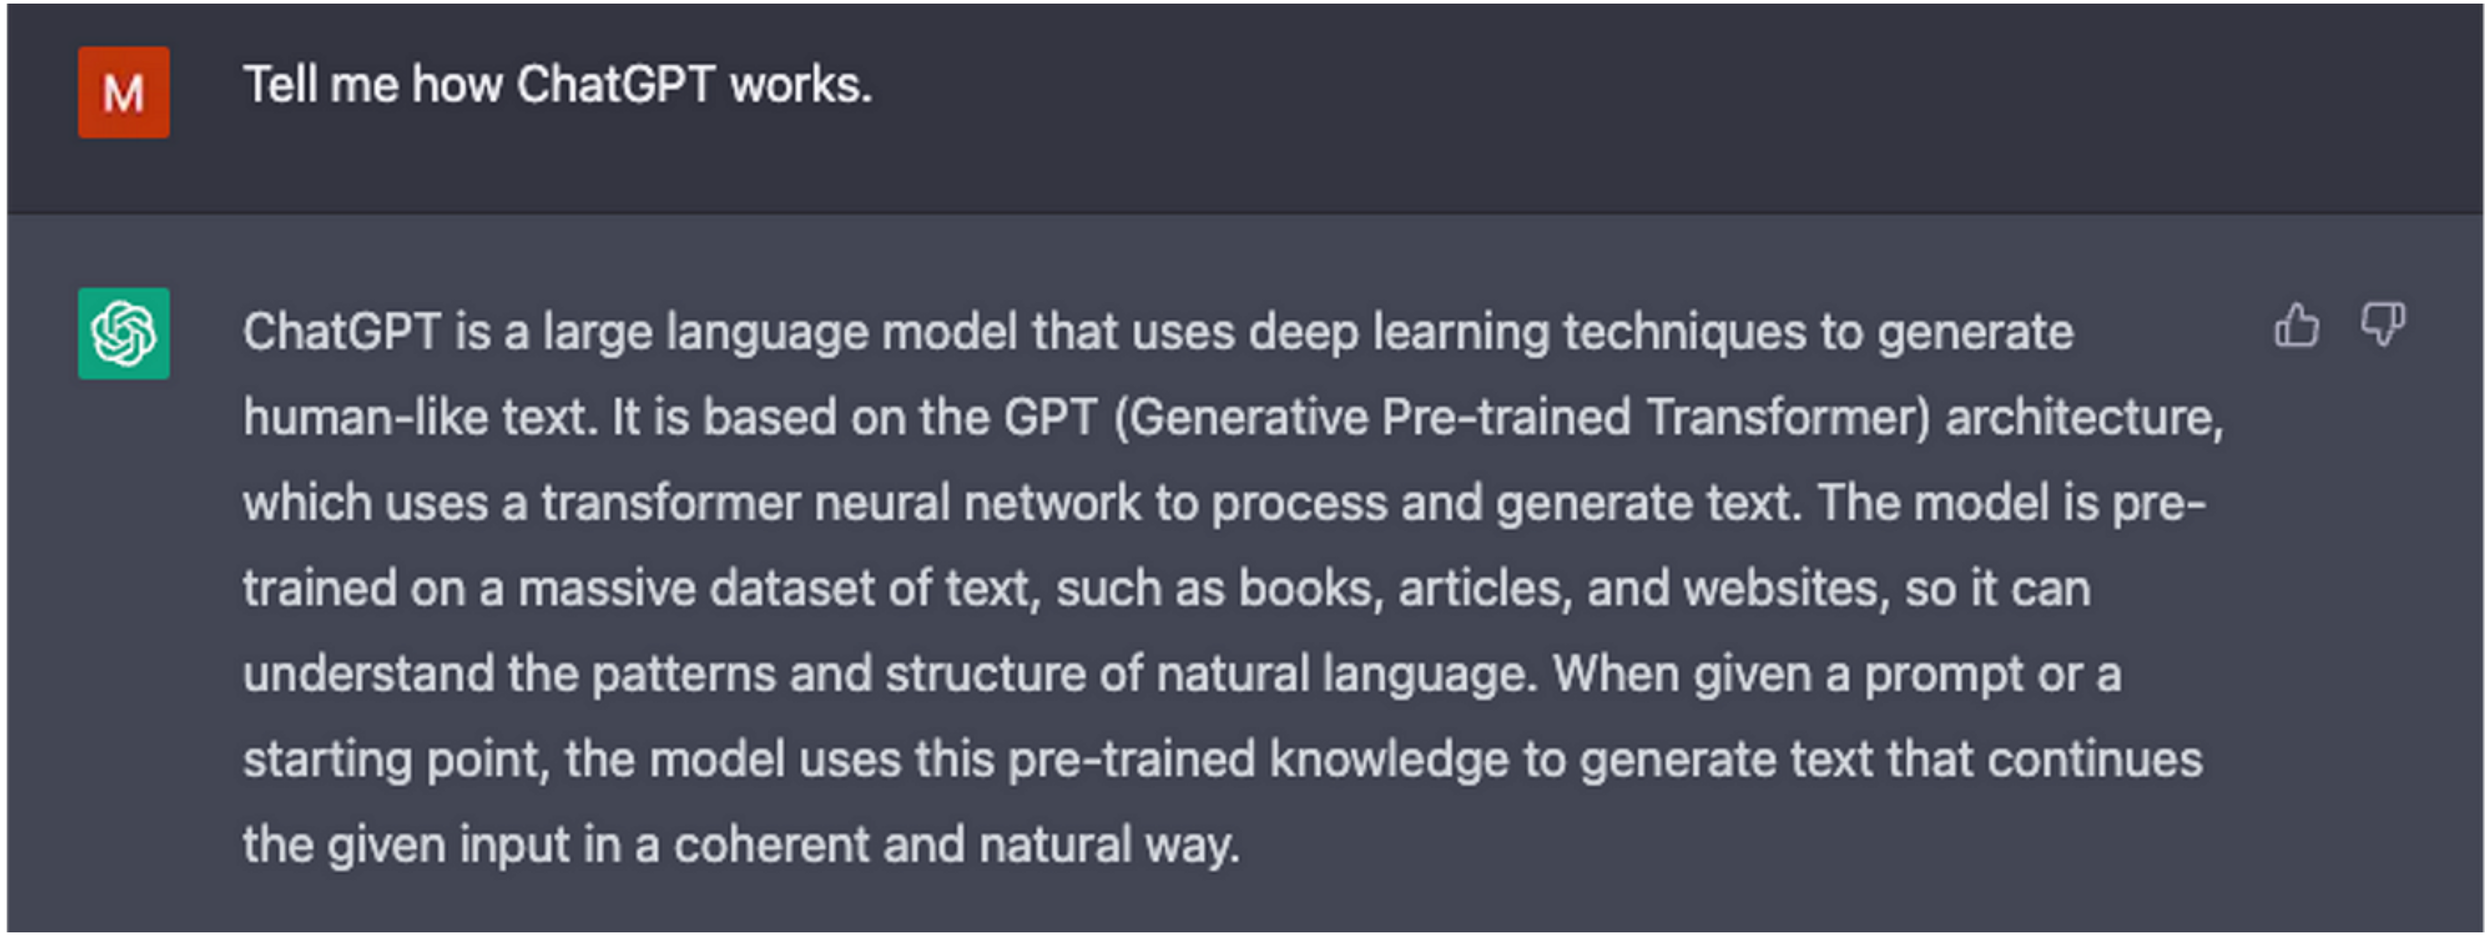

Why Machines?

Why machines?

- **Machines *use* the patterns, because...**
  - human attention is limited (self-driving cars)
  - humans are vulnerable (space mission)

- This step is called *Inference.*

Why machines?

- **Machines *find* the patterns, because...**
  - human are dumb (AlphaGo)
  - dataset is too big to handle (machine translation)
  - difficult to code human knowledge (dog classification)

- This step is called *Training*.
  - *Note*. Training data ≠ Test data (for inference)
    otherwise, we often call it “data mining”

Core Question

“How do we build a program,
that can discover & use patterns in the dataset?”

ML... How?

Lessons from an old perspective

- ML has its origin in ***Cybernetics*** (1947).
  - Coined by a control theorist Norbert Wiener.

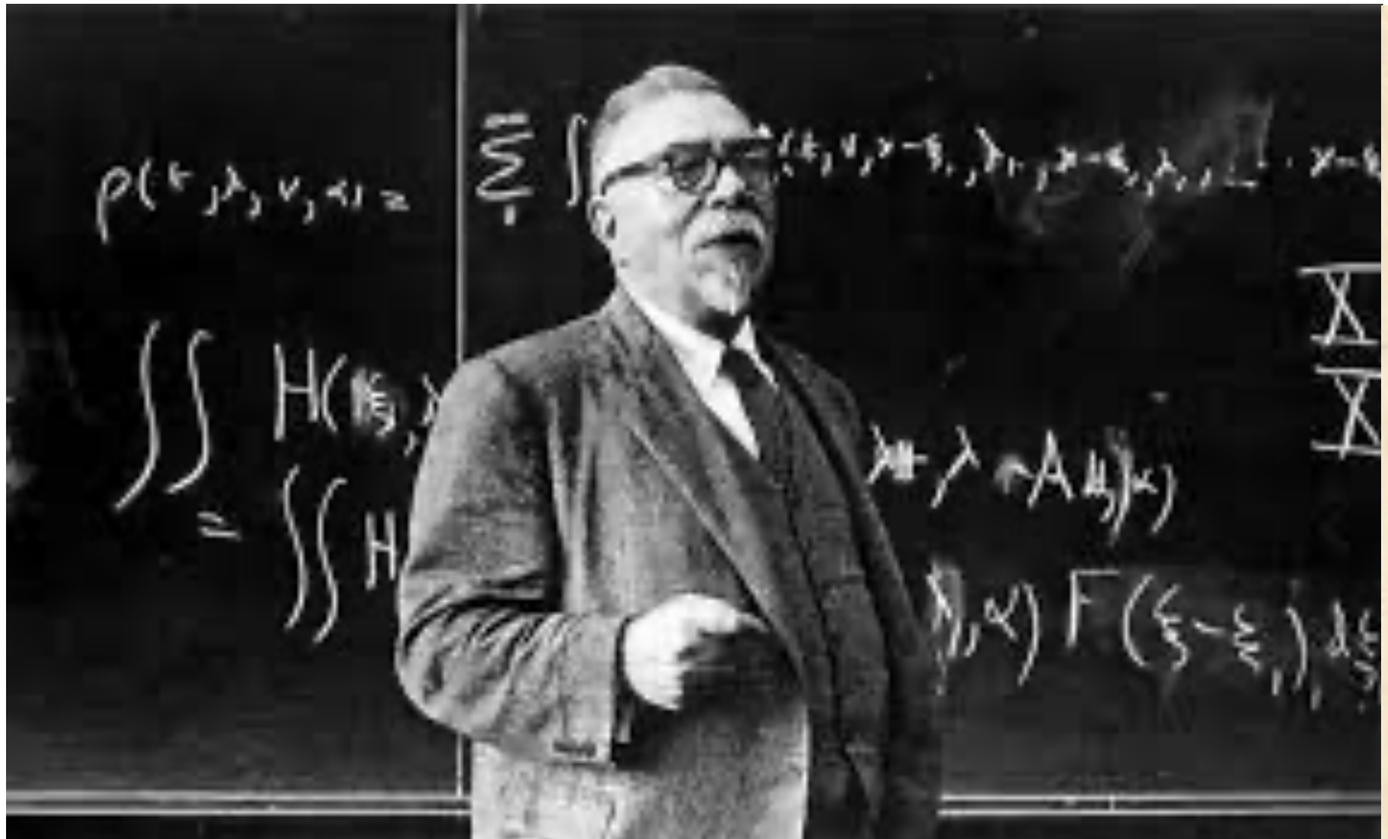

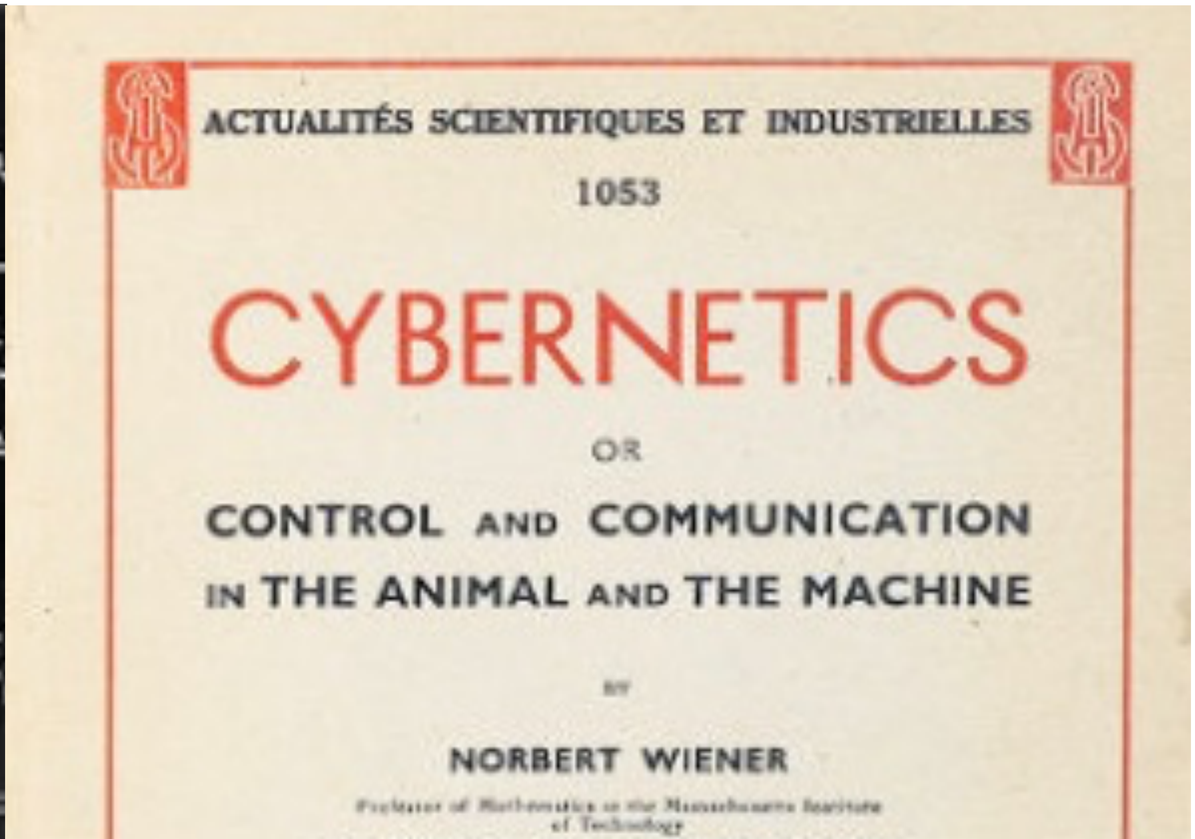

Cybernetics

... thus called “κυβερνήτης” (steering)

- Intelligence = “circular causal process (via feedback loop)”
  - proposed a holistic study of communication, control, and feedback mechanisms.

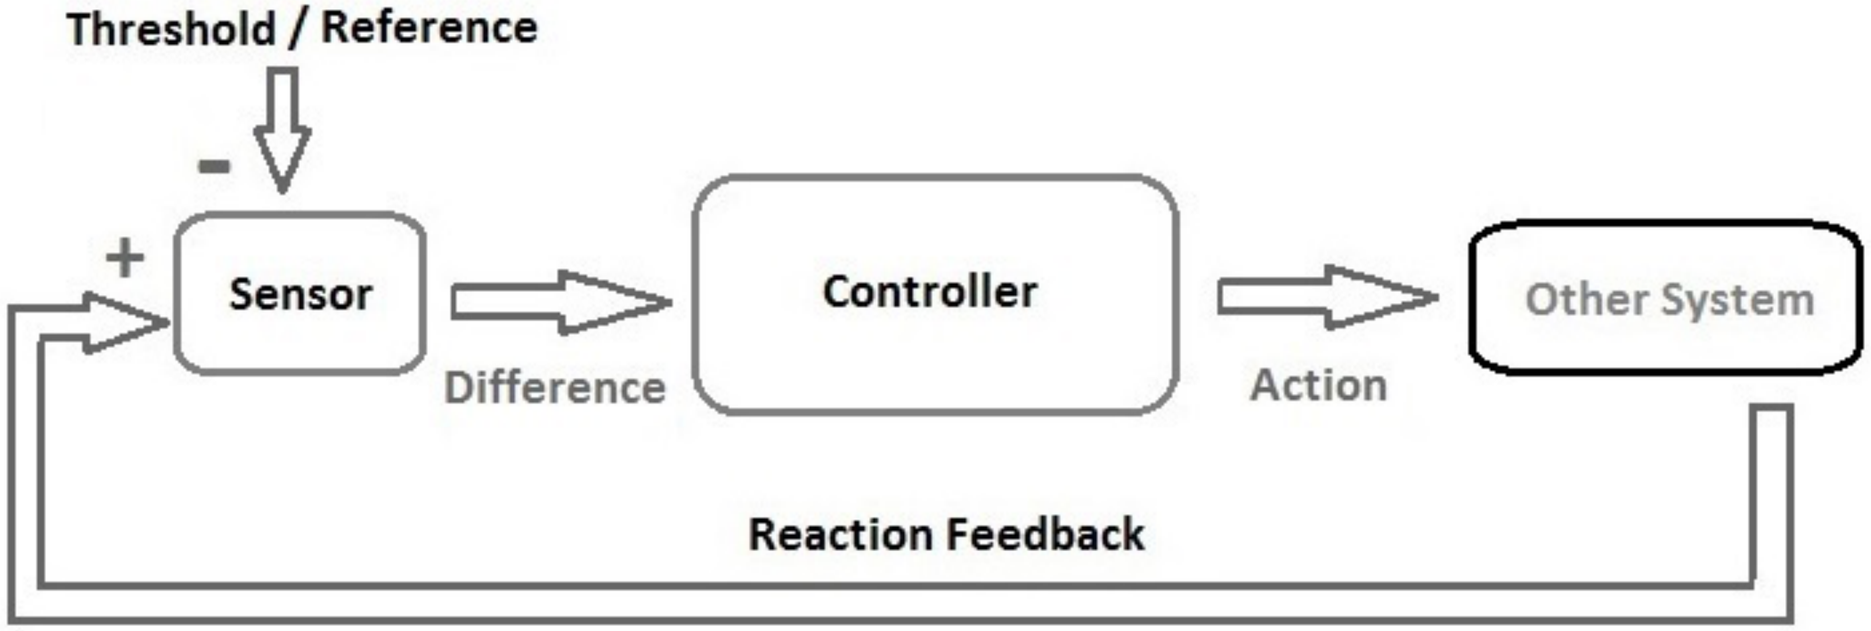

A Rough Picture

- The early cybernetics had all core concepts.
  - We have a **program** (model) **with** **changeable internal state** (model parameters).
  - We **find the right internal state** (optimize) through
    - testing the current program on training data
    - getting the feedback
    - modifying the state accordingly
      (can sometimes be done in one-shot)

This course

- Introduces many standard problems.
  - e.g., Classification, Regression, ...
- Describes classic models
  - e.g., linear, neural networks, ...
- Describes how to train (or optimize) them
  - e.g., gradient descent, closed-form solution ...

Administrivia

Team

- **Instructor.** Jaeho Lee 이재호
  jaeho.lee@postech.ac.kr
  #407, Engineering Building 2
  - ***What to ask.*** Coursework-related, anything else.
- **TA.** Minkyu Kim 김민규
  minkyu.kim@postech.ac.kr
  #404, Engineering Building 2
  - ***What to ask.*** Grading, HW, Attendance (cc Jaeho)

Time & Venue

- **Class.** #106, LG Hall.
  09:30AM — 10:45AM., Mon/Wed
- **Office Hrs.** #407, Engineering Building 2
  10:50AM—11:50AM, Wed
- **Web.** jaeho-lee.github.io ← for course materials
  PLMS ← for hw submissions

Grading

- Attendance (10%)
- Assignments (20% = 4 x 5%, subject to change)
- Mid-Term (30%)
- Final Project (40%)

- **Graduates.** Separately graded.
- **QE sit-ins.** Judged based on undergrads

Assignments

- **Attendance.** Random samples.
  ~5 people per day. (3 strikes = out)

- **Language.** Python + PyTorch
  - *Required.* Use PyTorch for matrix multiplication!
  - *Note.* Can use other languages (e.g., JAX) for the final project.
- **PyTorch Tutorial**
  https://pytorch.org/tutorials/

Prerequisites

- Not required, but I assume you know:
  - Calculus
  - Programming
  - Basic Linear Algebra
  - Basic Probability & Statistics
  - Basic Signals & Systems
- **Note.** Will try not to be presumptuous ;)

Textbook

- **Main.** Lecture notes*, based on
  "Mathematics for Machine Learning"
  Deisenroth, Faisal, and Ong
  free pdf: https://mml-book.github.io
- **More Theory.** "Patterns, Predictions, and Actions"
  Hardt and Recht
  free pdf: https://mlstory.org/
- **Coding.** Dive into Deep Learning (https://d2l.ai/)

* based on cool lecture notes by Prof. Tae-Hyun Oh / David Hsu

Honor Codes

Cheating means F

(No mercy)

- Sharing solutions → Not okay.
- Copying solutions → Definitely not.
- Discussion → you have me and TA.

Cheers