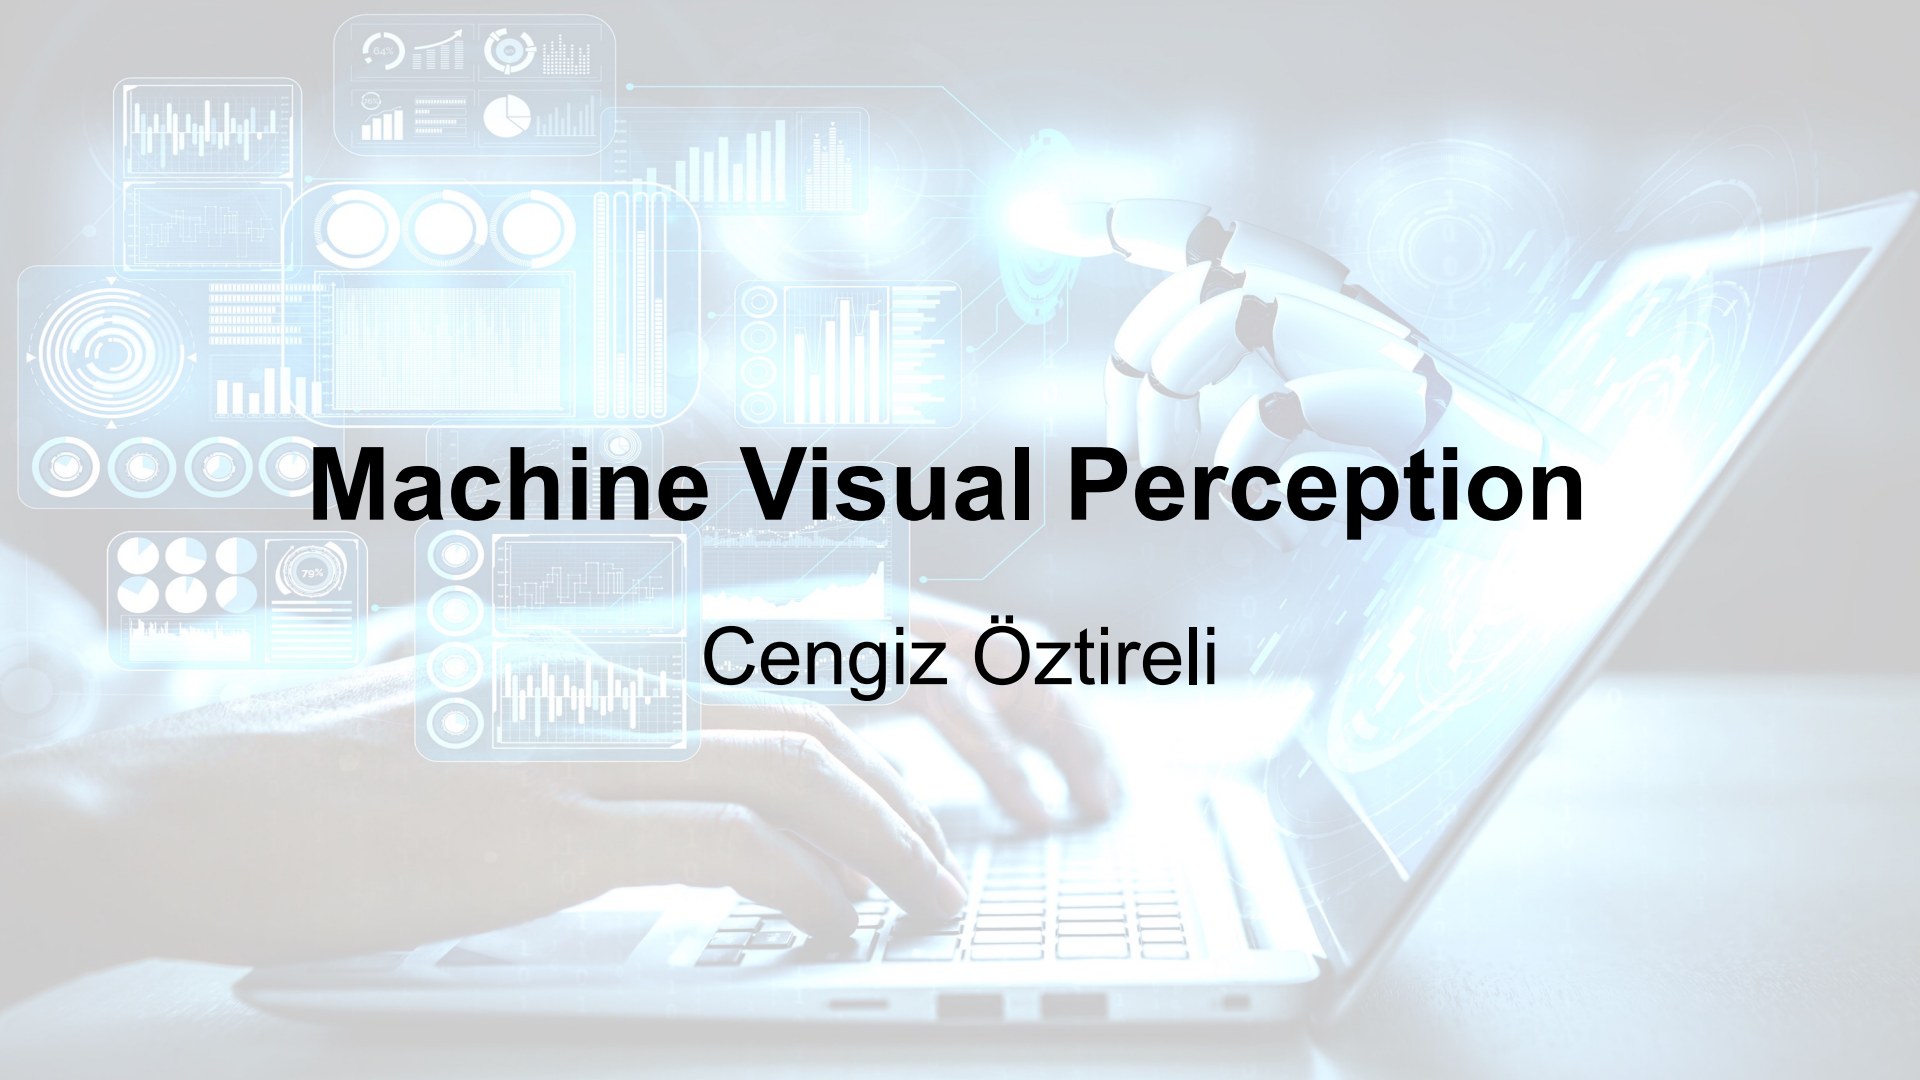

Machine Visual Perception

Cengiz Öztireli


Self
Driving


9:15 AM
84%
Smart Home

Course format

- 12 + 4 hours
- One practical exercise (20%)
- One course project (80%)
  - We will release a list of suggested projects
  - You may propose others, to be confirmed by us
  - In groups of 2-3, we will assign
  - More information will be on the webpage soon

Contributors

- Fangcheng Zhong (fz261)
- Param Hanji (pmh64)
- Chris Town (cpt23)

Machine Visual Perception

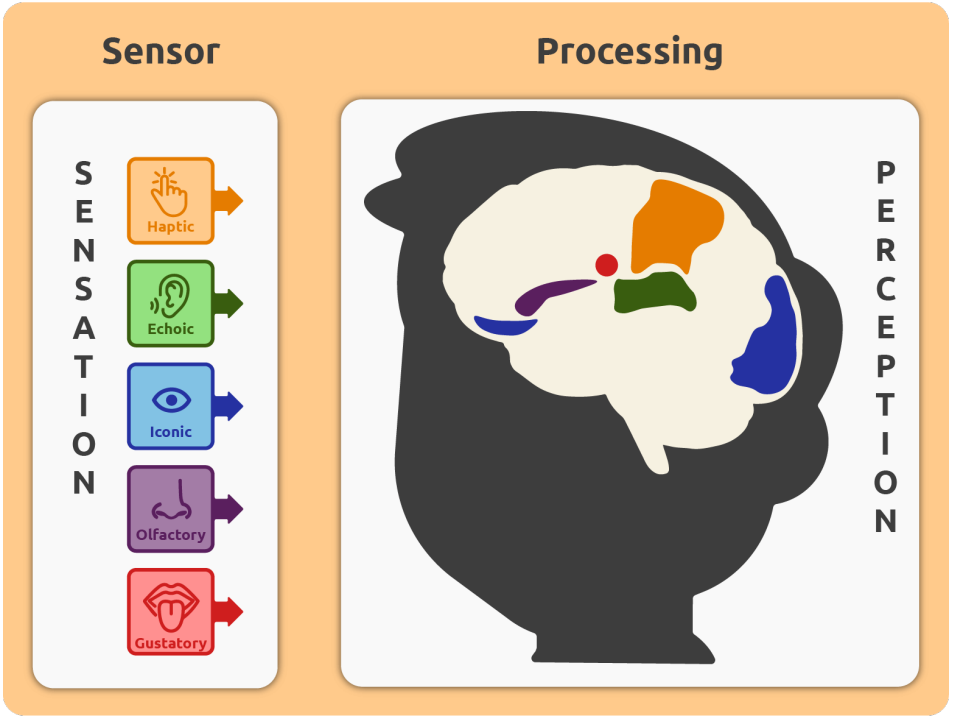

- Our goal is to enable a computer to "*see*" and "*understand*" the environment.
- Computer vision
- Machine/ deep learning

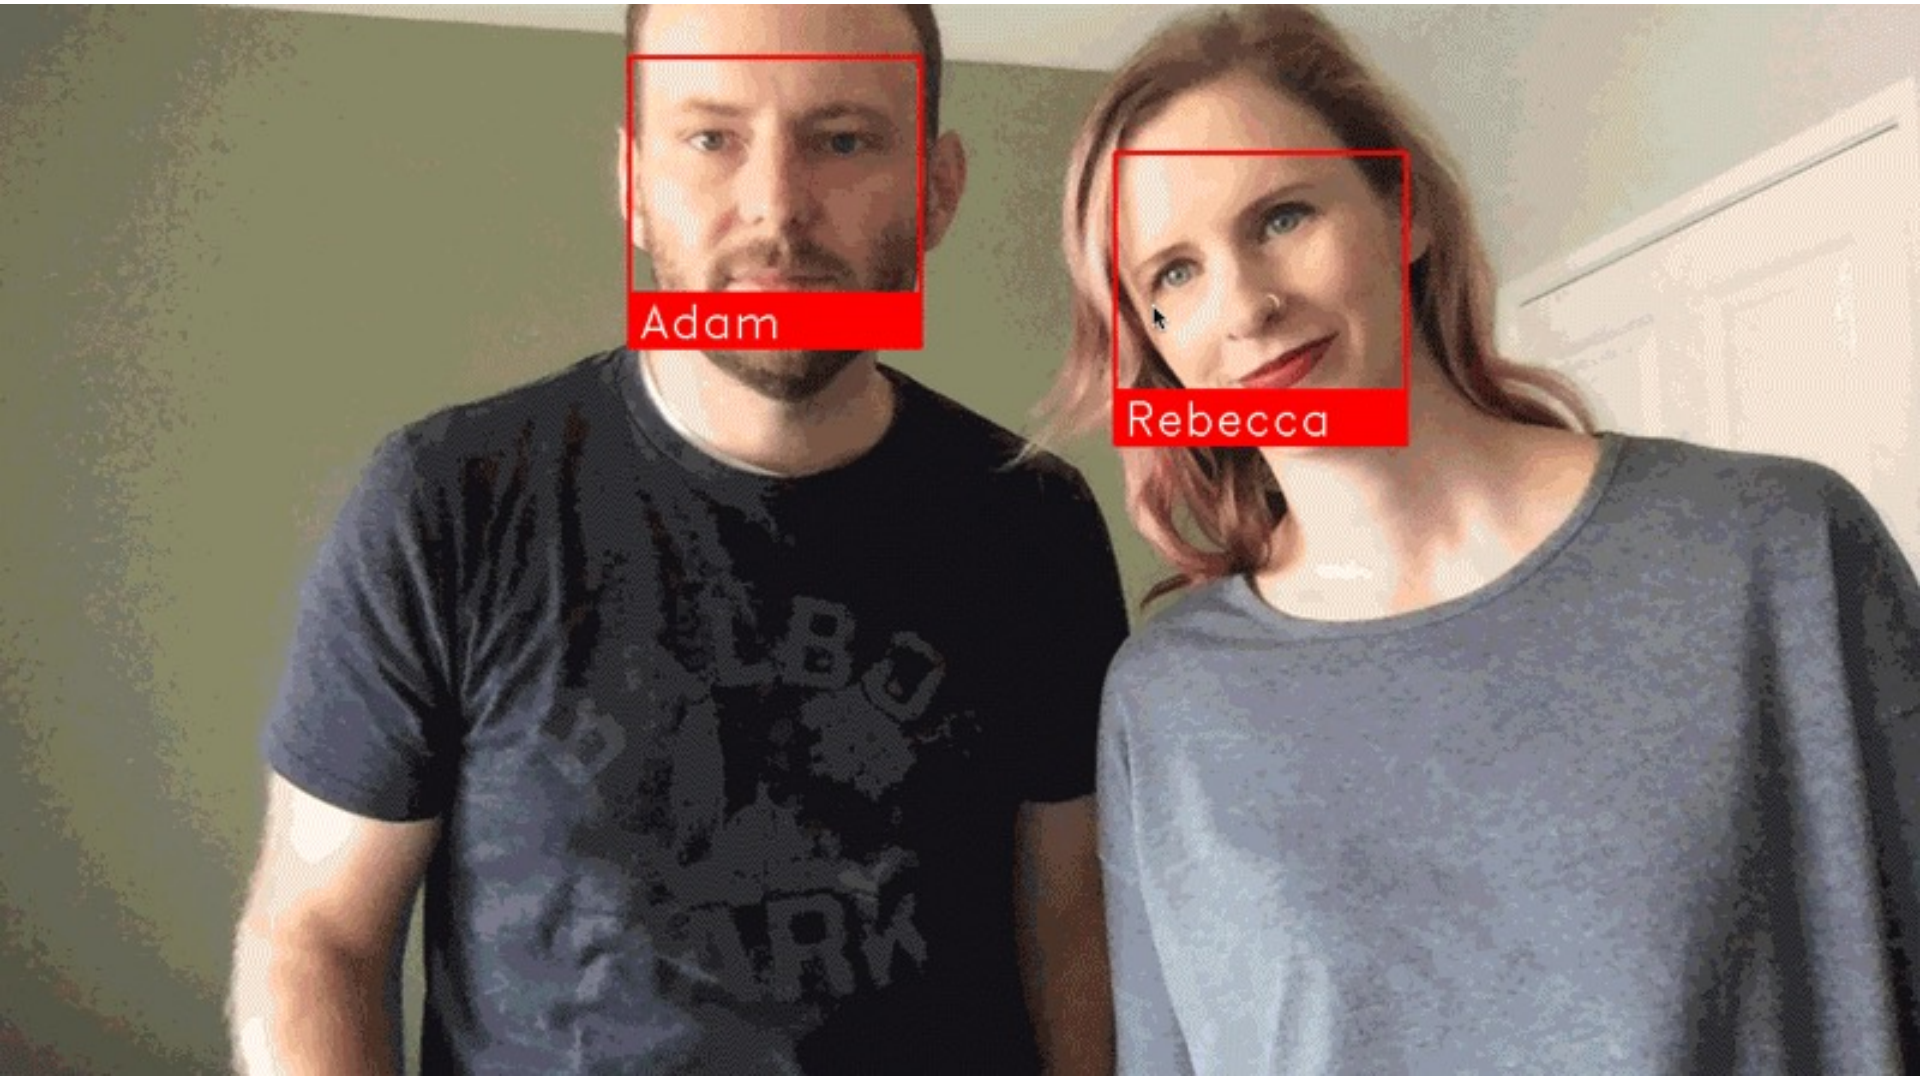
Adam
Rebecca

Assisted Diagnosis

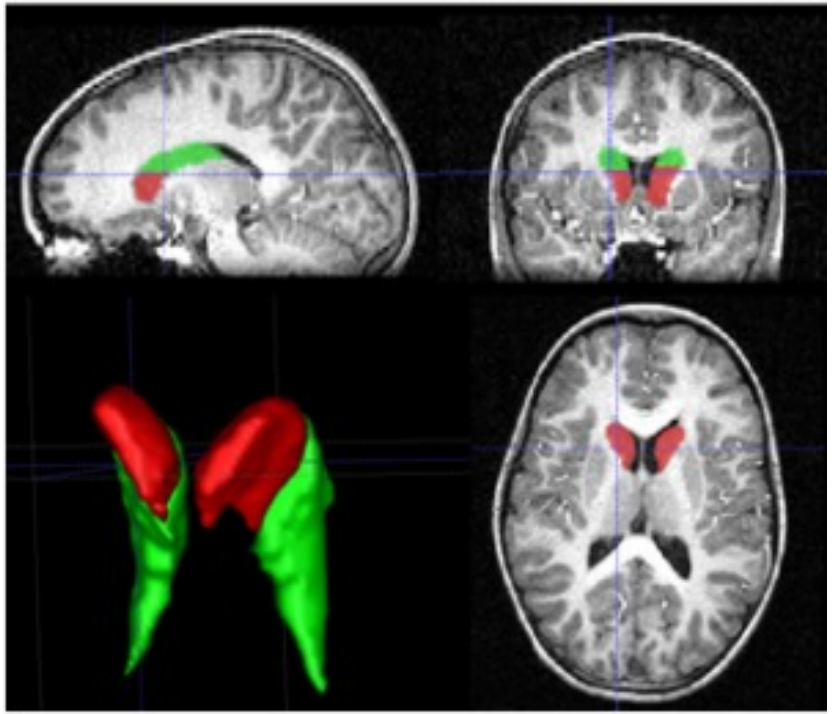

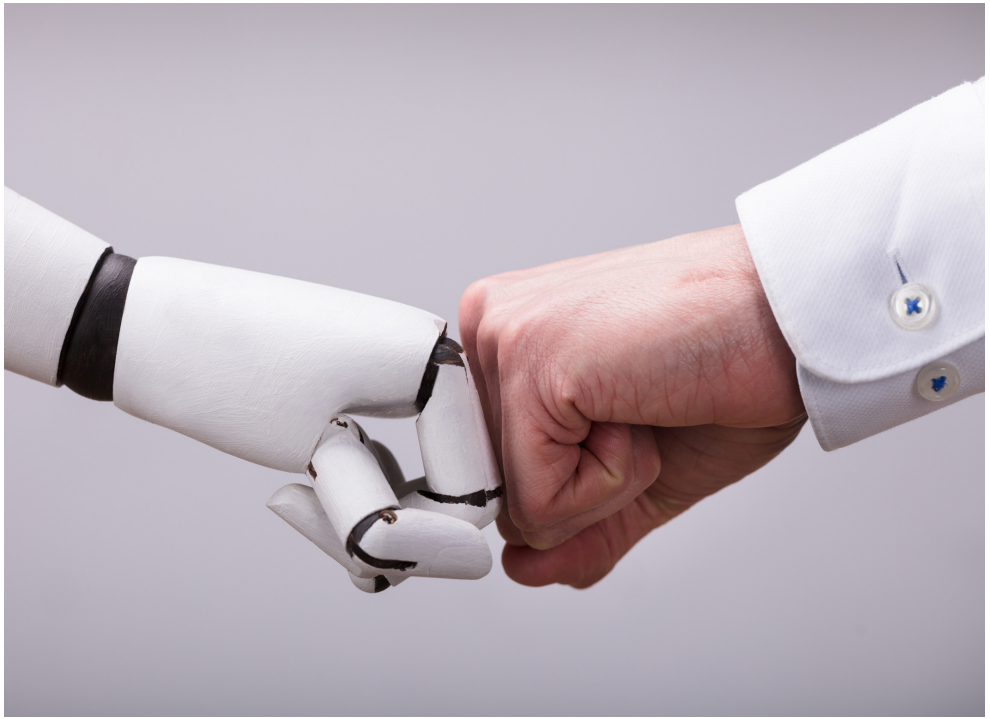


Output
1:15
0:00 / 6:41


1:19 / 2:59

Machine Learning

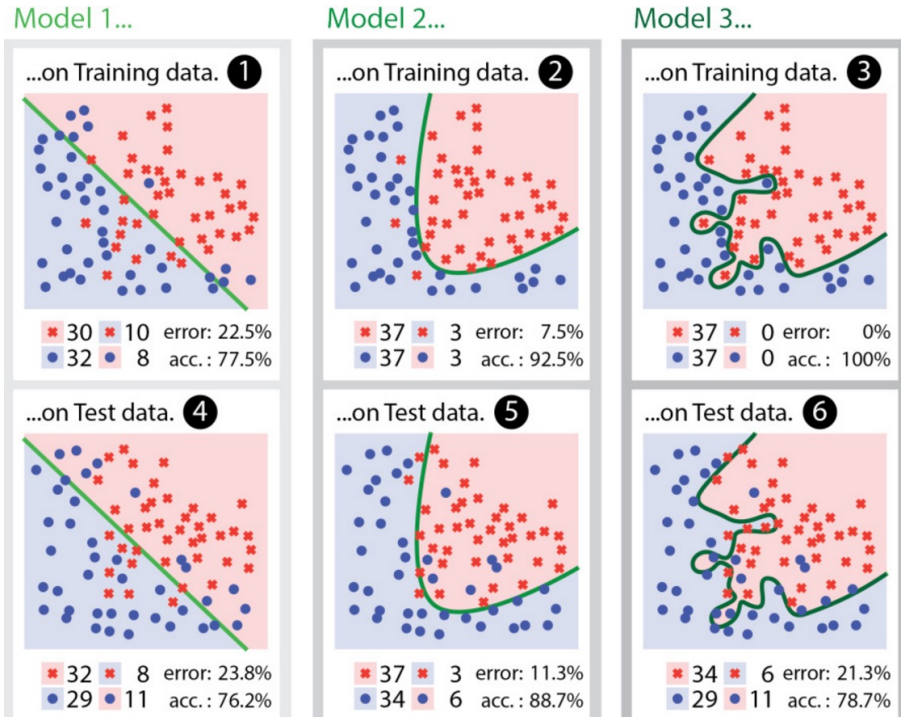

"I've narrowed it to two hypotheses: it grew or we shrunk."

Paradigms of Machine Learning

- Supervised learning
  - Given a training set of N examples

$$(x_1, y_1), (x_2, y_2), ..., (x_N, y_N)$$

  - Minimize the discrepancy between the predicted output and the given output (ground truth)

Paradigms of Machine Learning

- Supervised learning

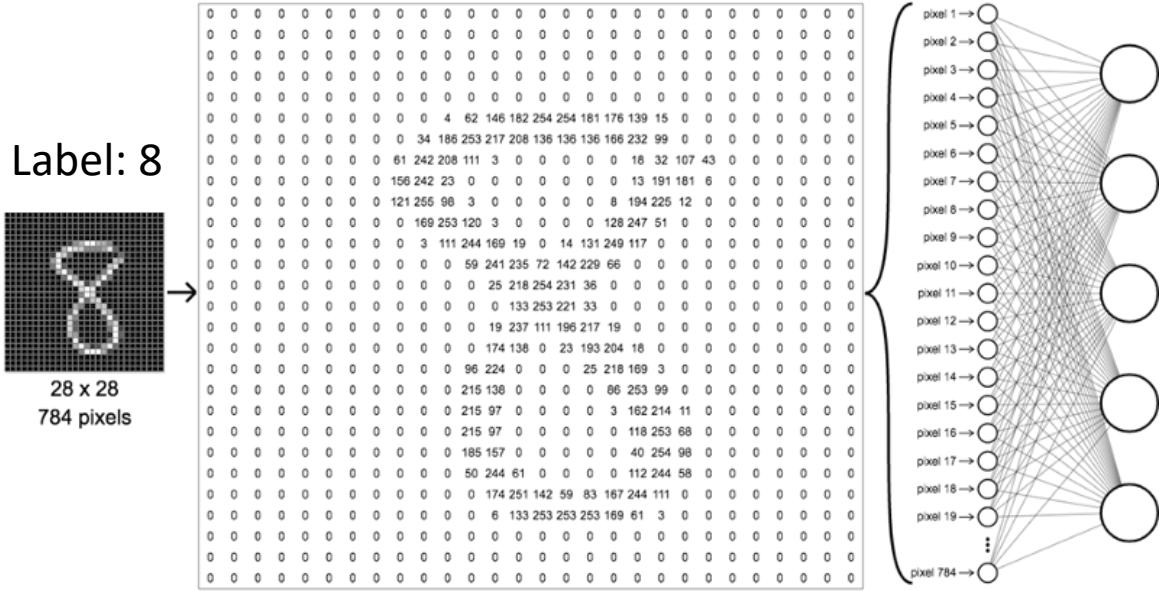

Given the pixels of a hand-written number, identify the number.

Paradigms of Machine Learning

- Un-supervised learning
  - Learning patterns without specific target output values
  - Example: clustering

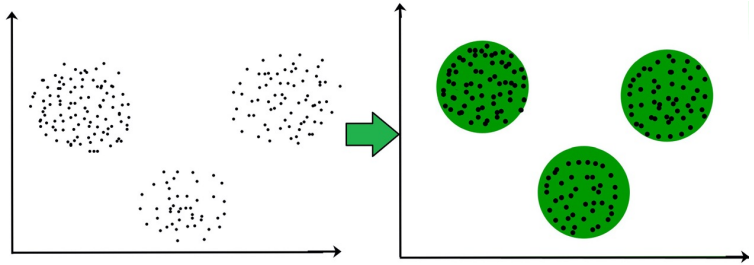

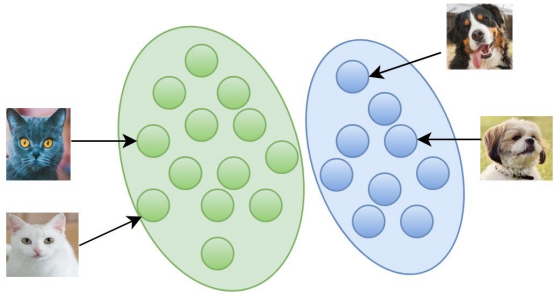

Paradigms of Machine Learning

- Un-supervised learning

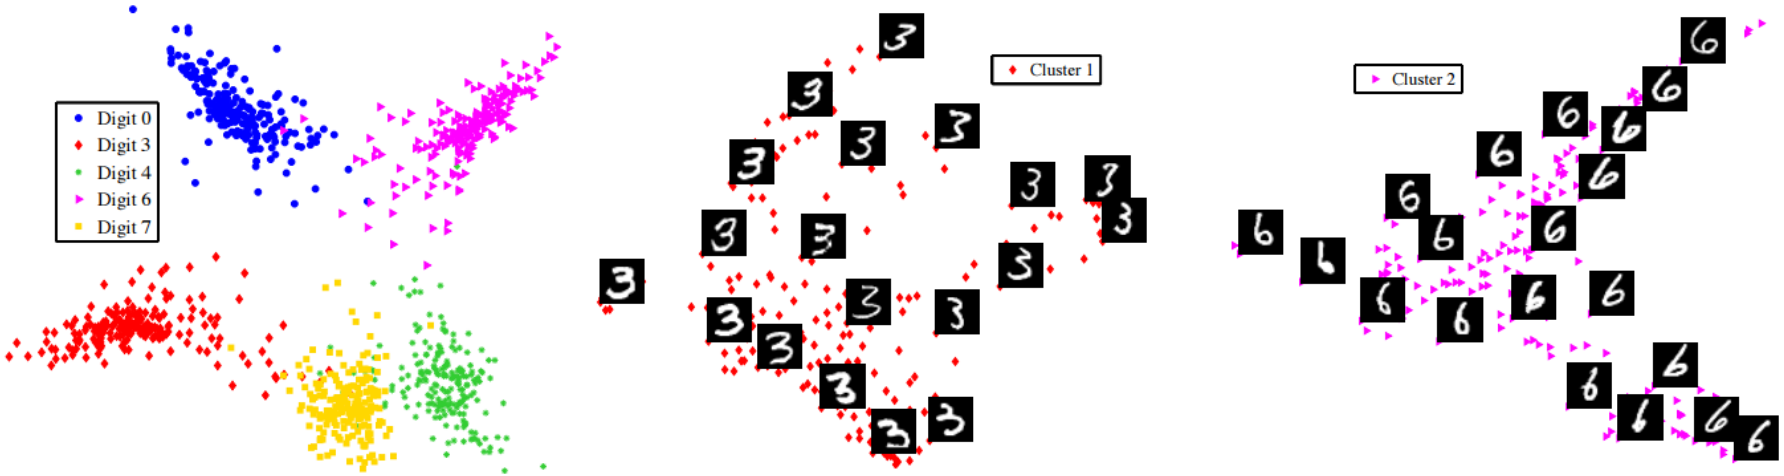

Paradigms of Machine Learning

- Un-supervised learning

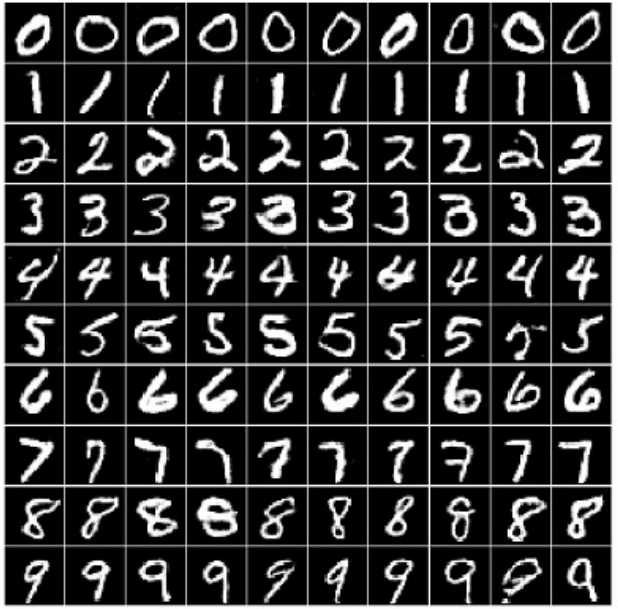

Paradigms of Machine Learning

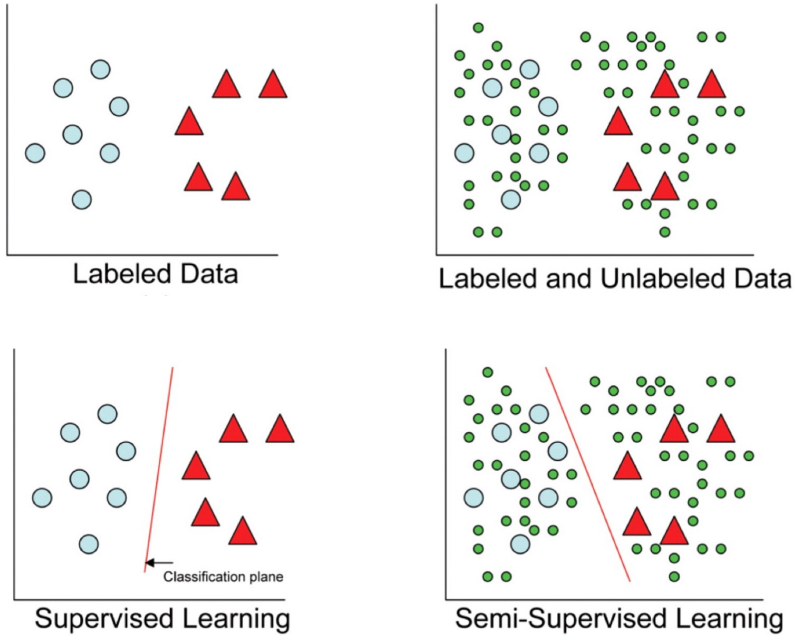

- Semi-supervised
  - Learning in the case of sparse labeled (supervised) data
  - Use accessible data to improve decision boundaries and better classify unlabeled data

Paradigms of Machine Learning

- Self-supervised learning

Learn data representations via auxiliary tasks.

Paradigms of Machine Learning

- Reinforcement Learning
  - Learning what actions to take in order to maximise some **reward**

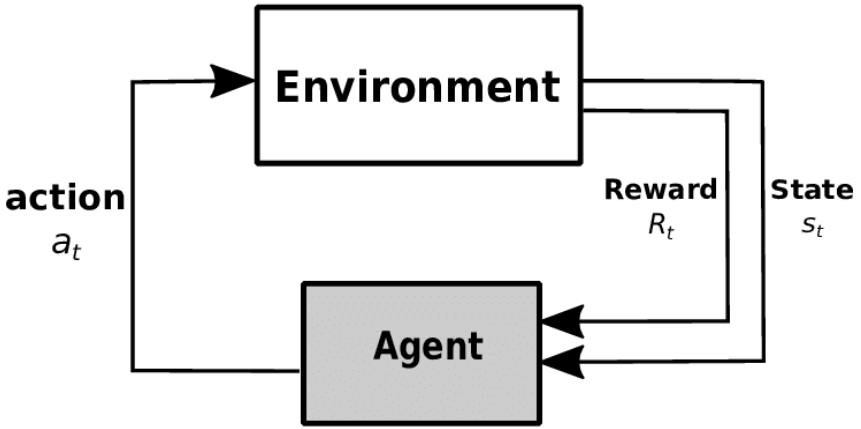

Paradigms of Machine Learning

- Reinforcement Learning
  - Learning what actions to take in order to maximise some **reward**

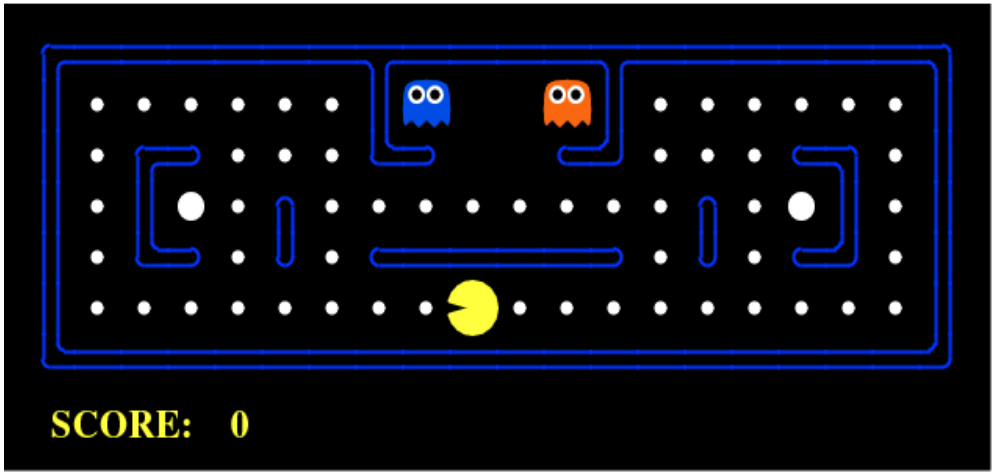

Input type

- What can a machine “see”?
  - Image: an array of pixels
  - Video: a list of images
  - Point clouds

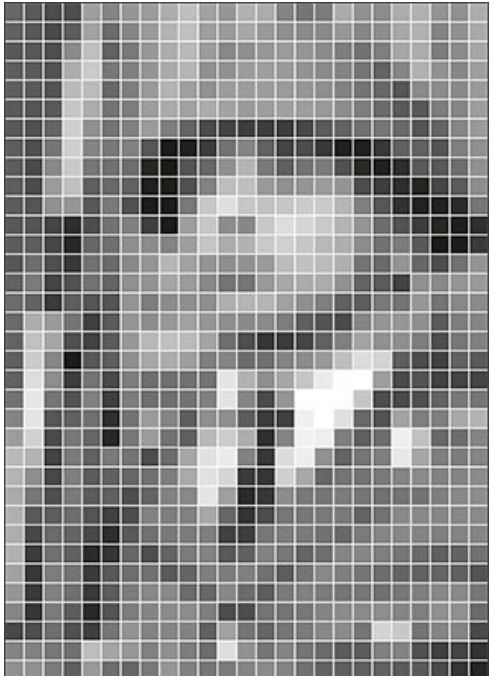

Input type

- What can a machine "see"?
  - Image: an array of pixels
  - Video: a list of images
  - Point clouds

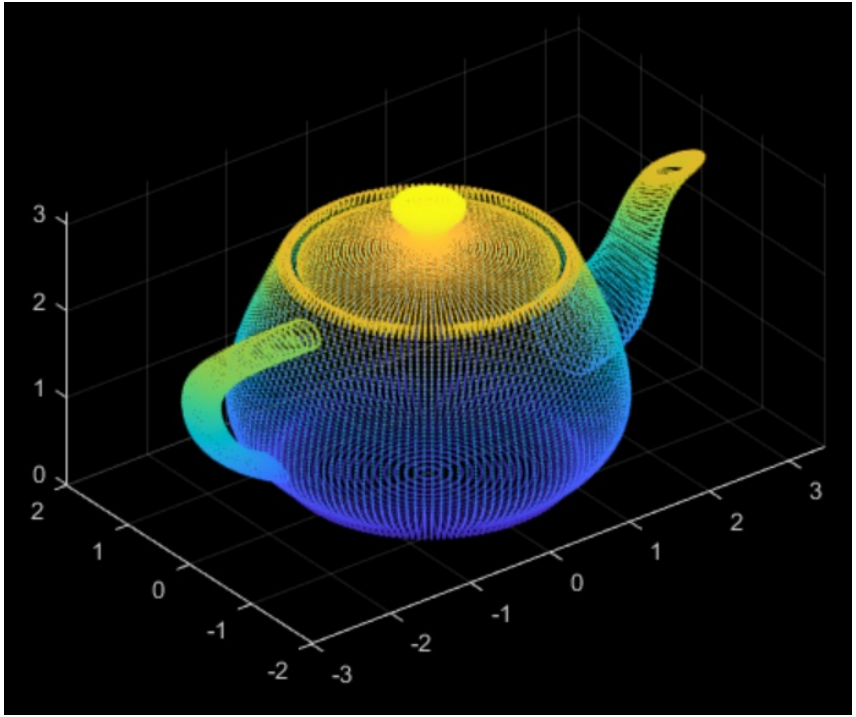

Output type

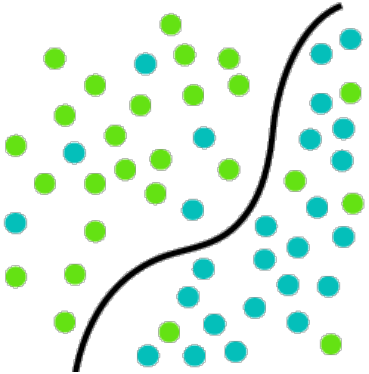

Classification

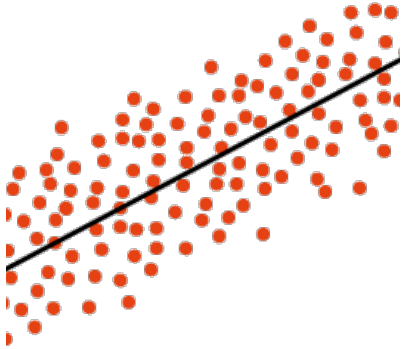

Regression

Regression

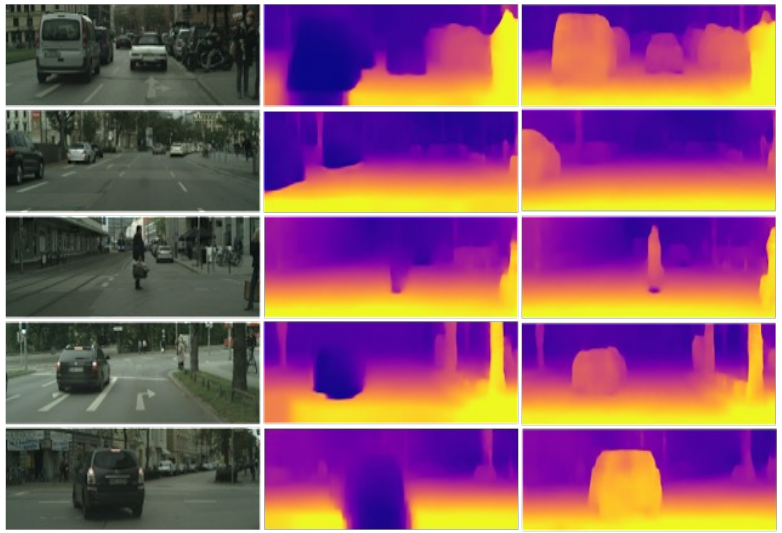

Classification

Bird

Cat

Flamingo

Cock

Egyptian cat

Persian cat

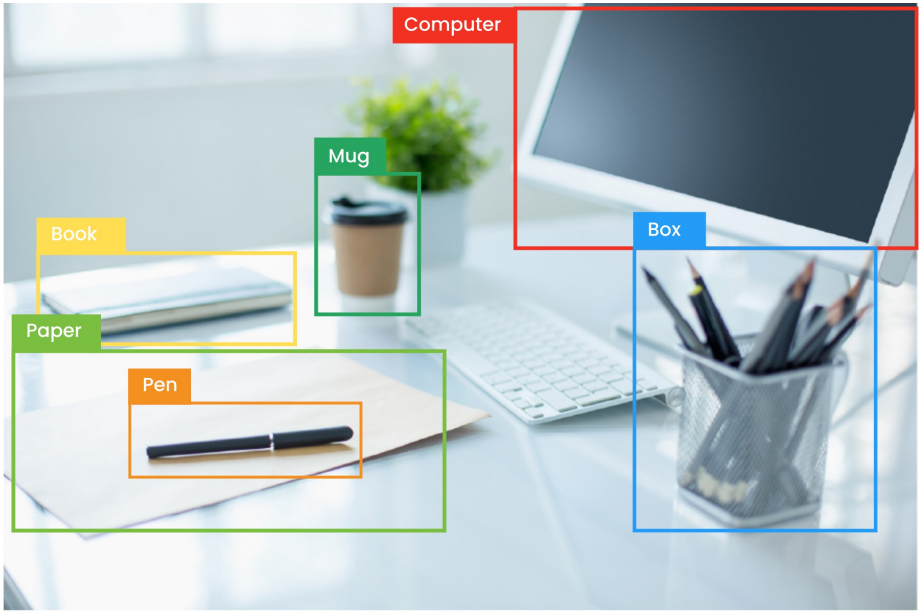

Practical Exercise

- Two parts: classification & regression
- Details will be on the course web page and Moodle
- Any questions? Please ask Param, Fangcheng, myself.
- Release: 19 October 2022, 12:00 PM
- Deadline: 11 November 2022, 12:00 PM

Course Project

- Main task:
  - implementing a paper or,
  - extending an existing implementation or,
  - applying an implementation to a new task and dataset.
- Project deliverables:
  - A project plan
  - Project presentation
  - Project report
  - Implementation
- Project deadline: 1 December 2022, 12:00 PM

Course Project

- Timeline:
  - This week on Friday we will release project proposals
  - By this week Friday
    - send names and CRSids of group members to aco41
    - only one email per group
  - By next Friday
    - send a short project proposal (template, also on Moodle)

- **All announcements are on Moodle!**