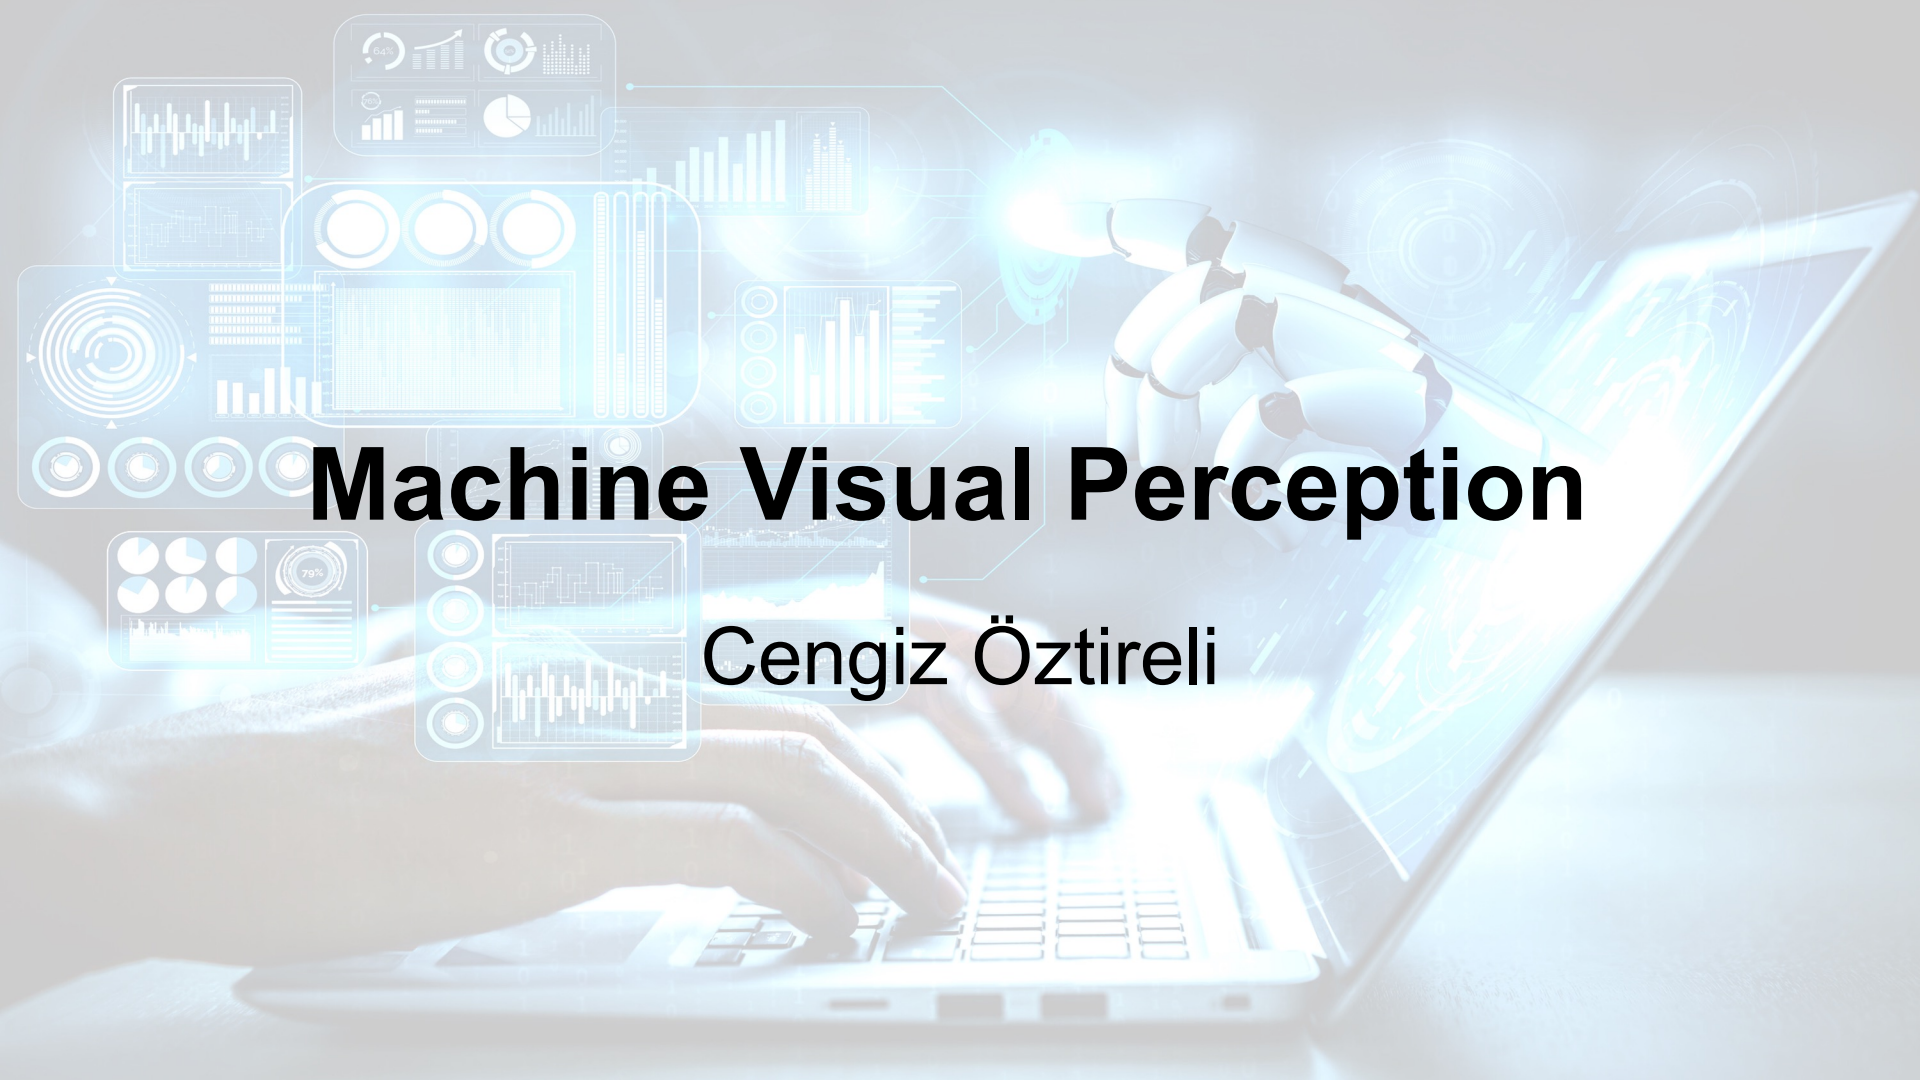

Machine Visual Perception

Cengiz Öztireli

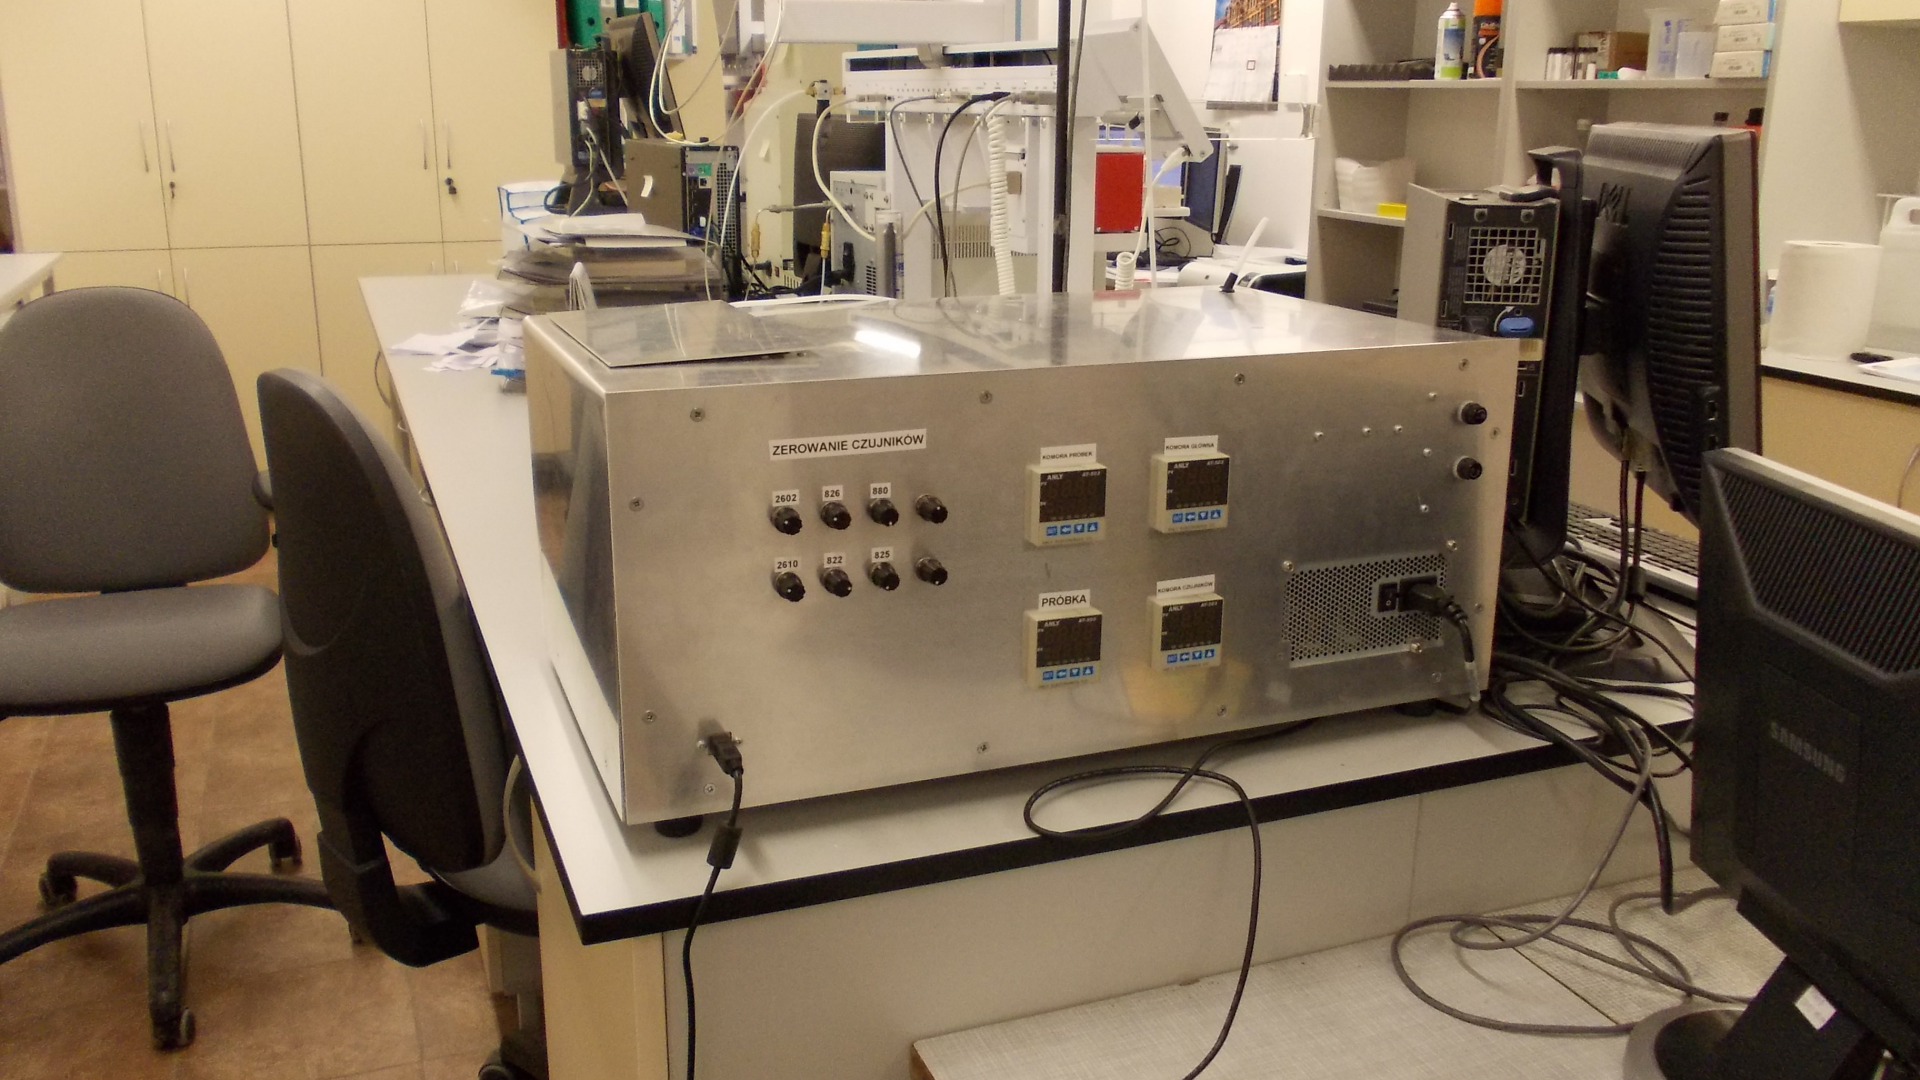
ZEROWANIE CZUJNIKÓW
2602
826
880
2610
822
825
PRÓBKA
SAMSUNG


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 on the webpage

Contributors

- Yifan Liu (https://irfanicmll.github.io/)
- Chris Town (https://www.cl.cam.ac.uk/~cpt23/)
- Walter (Tianhao) Wu and Chengliang Zhou

Machine Visual Perception

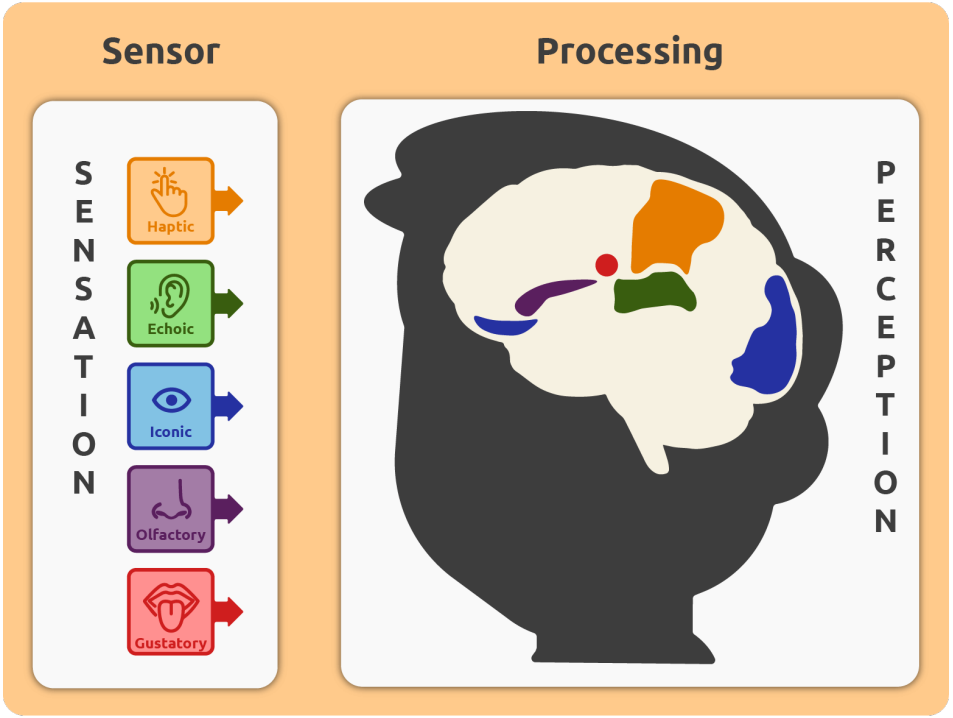

- Our goal is to enable a computer to "*see*" and "*understand*" the environment:
- Computer vision
- Machine/ deep learning

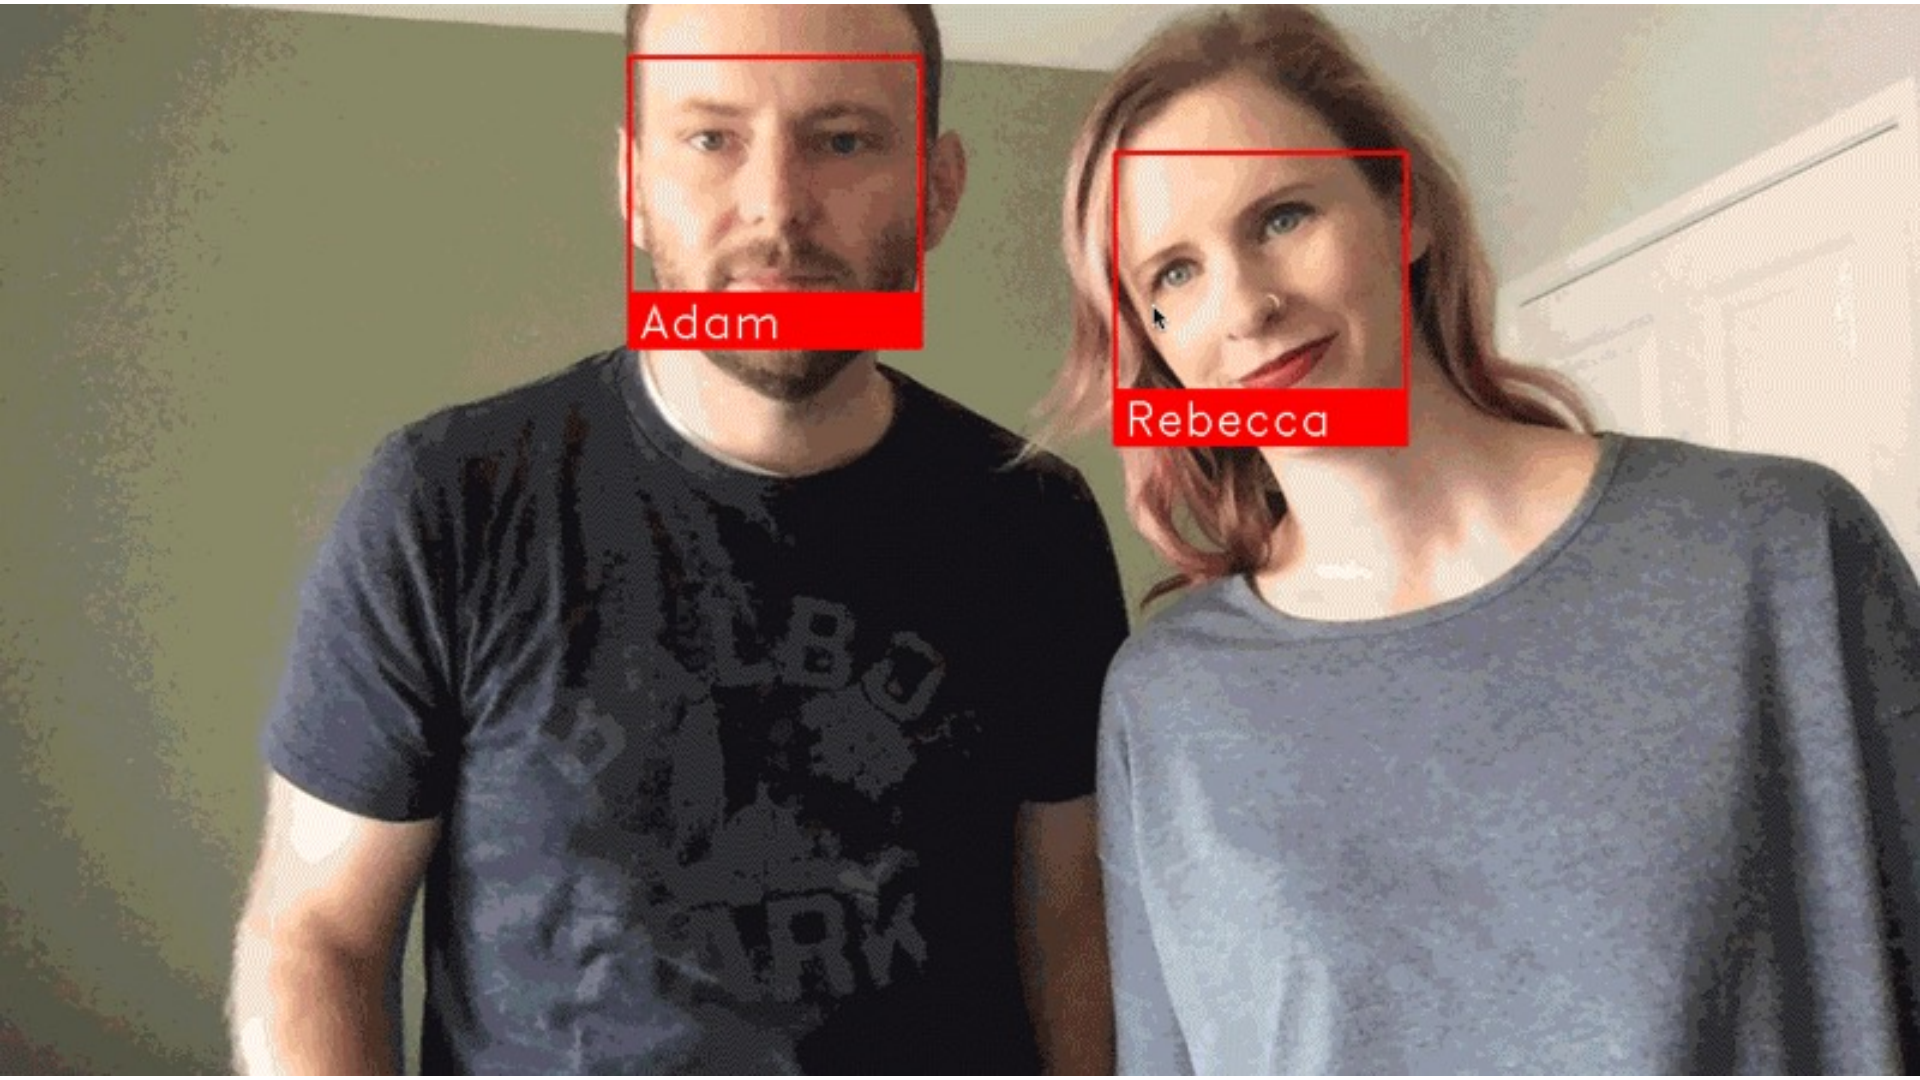
Adam
Rebecca

Assisted Diagnosis

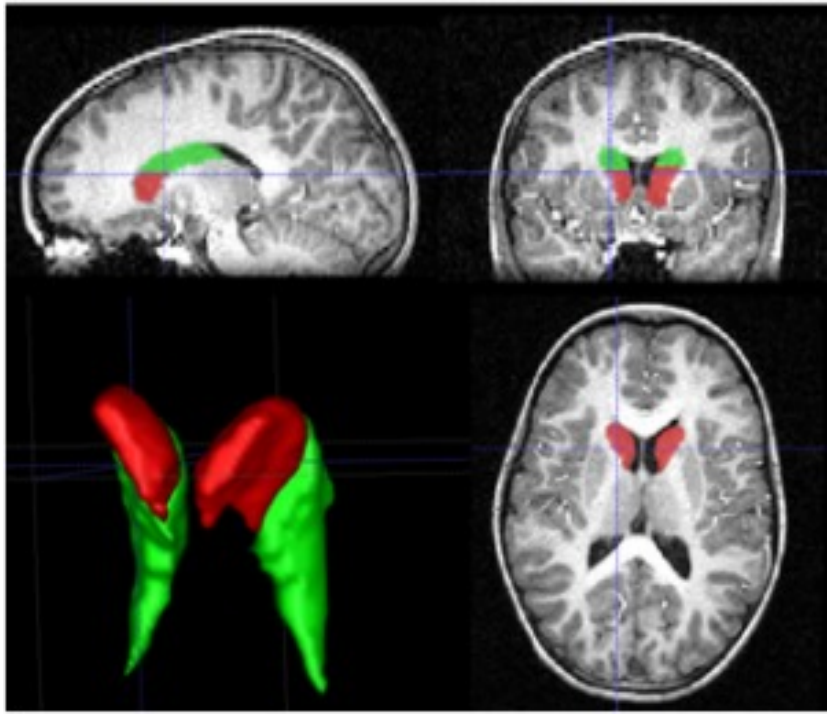

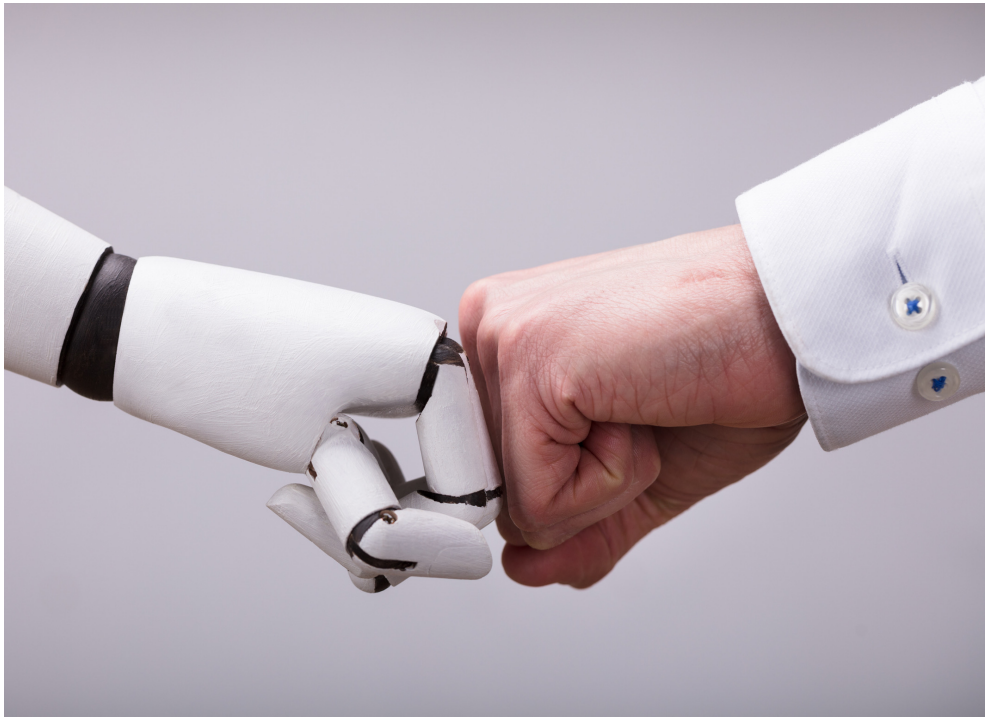


Output
1:15
0:00 / 6:41


1:19 / 2:59

Machine Learning

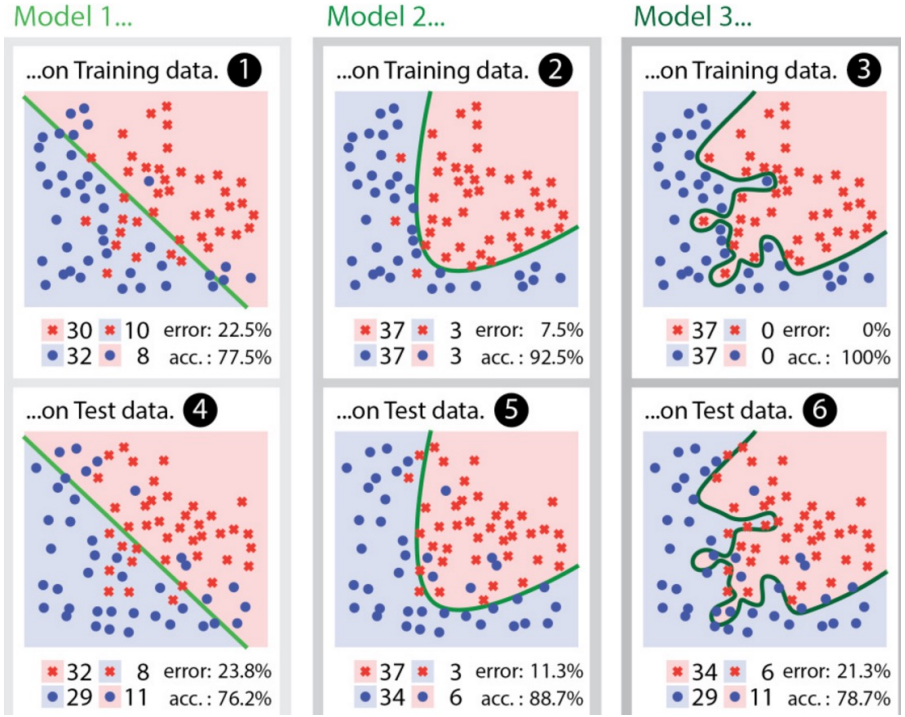

"I've narrowed it to two hypotheses: it grew or we shrunk."

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

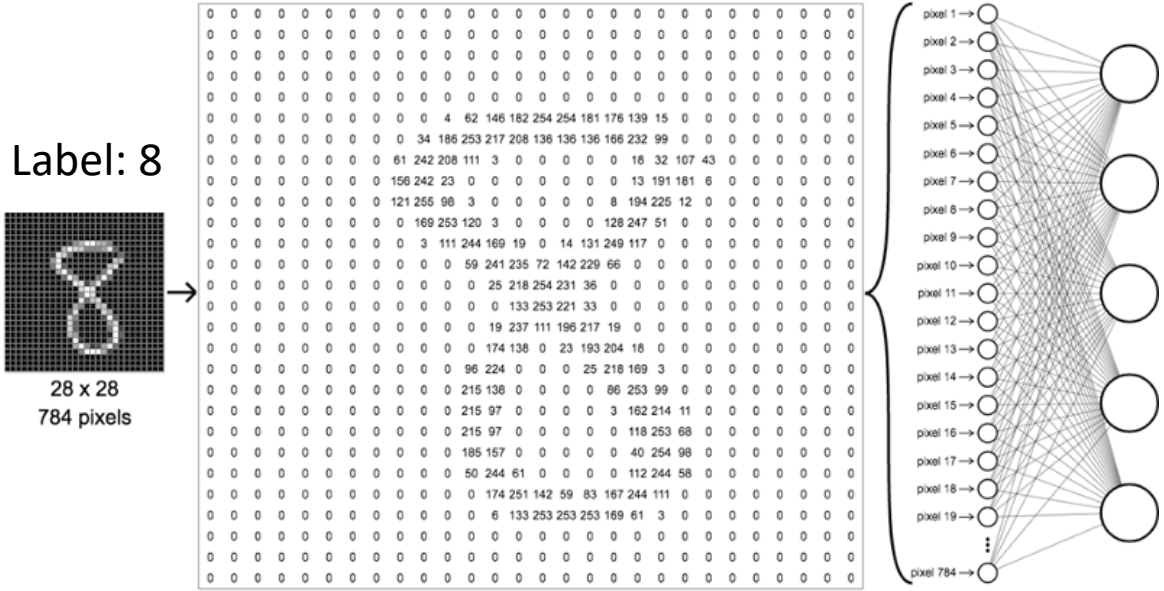

Given the pixels of a hand-written number, identify the number.

Paradigms of Machine Learning

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

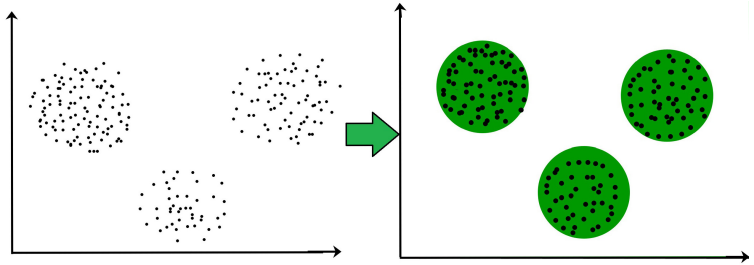

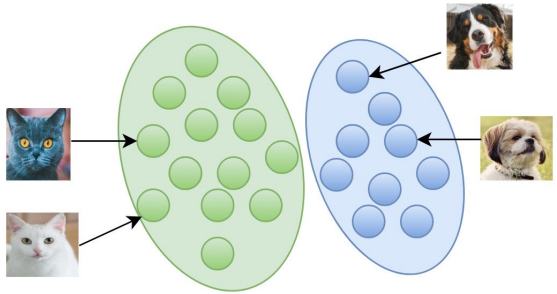

Paradigms of Machine Learning

- Un-supervised learning

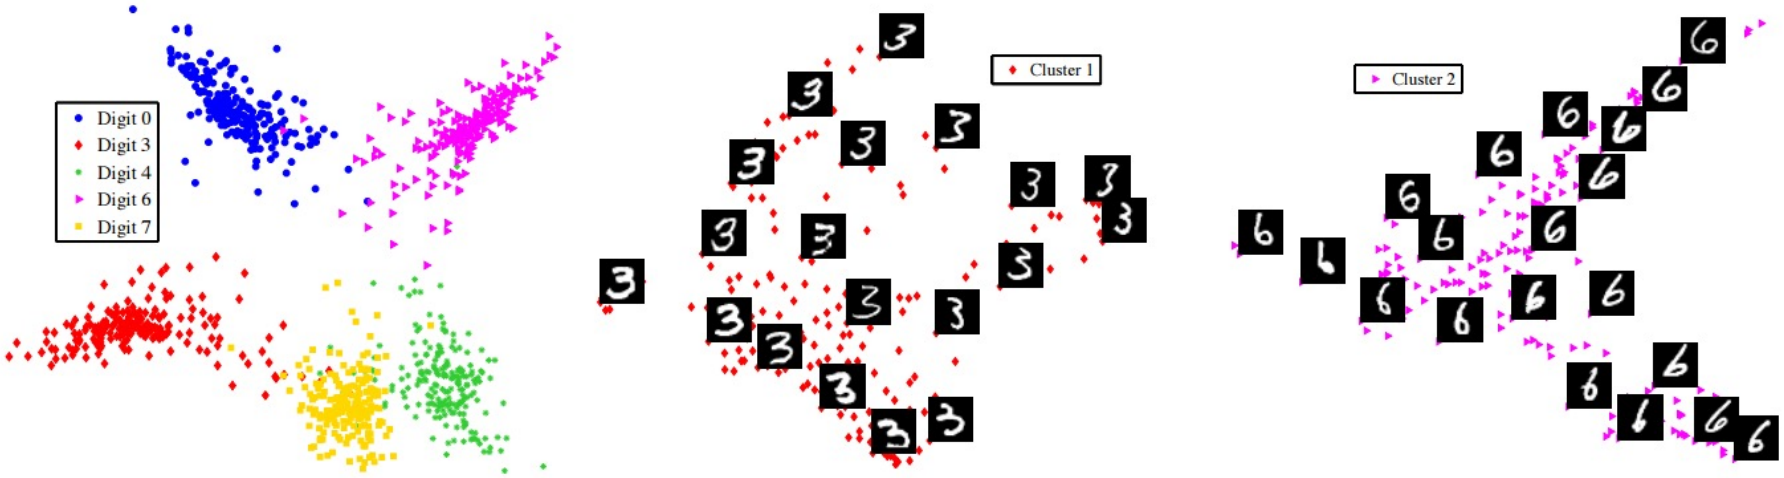

Paradigms of Machine Learning

- Un-supervised learning

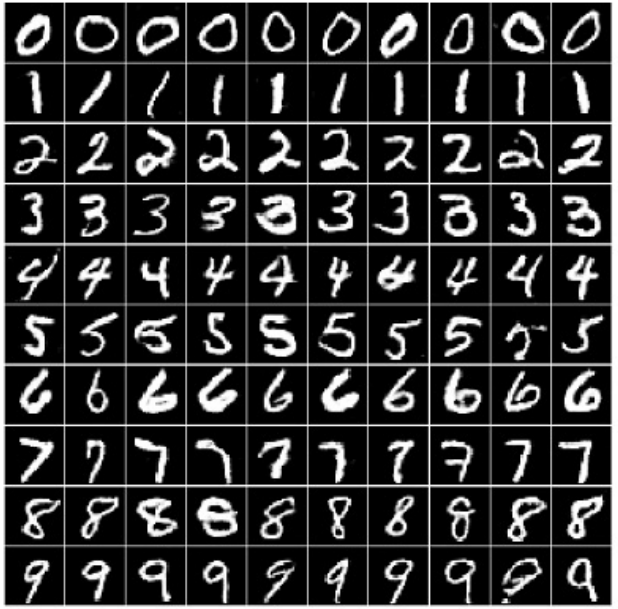

Paradigms of Machine Learning

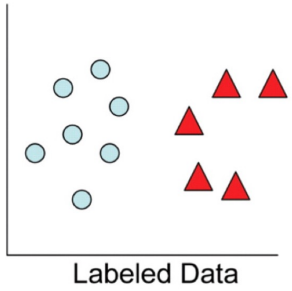

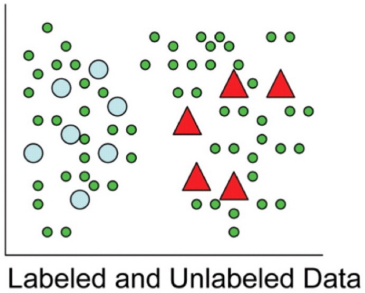

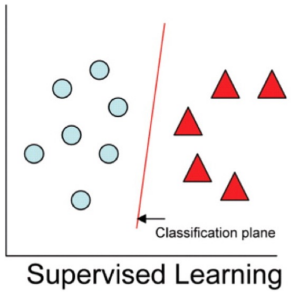

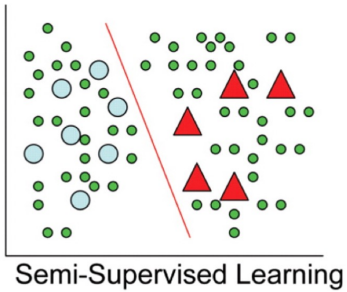

- Semi-supervised
  - Learning in the case of sparse labeled (supervised) data
  - Use accessible data to improve decision boundaries and better classify unlabeled data

Paradigms of Machine Learning

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

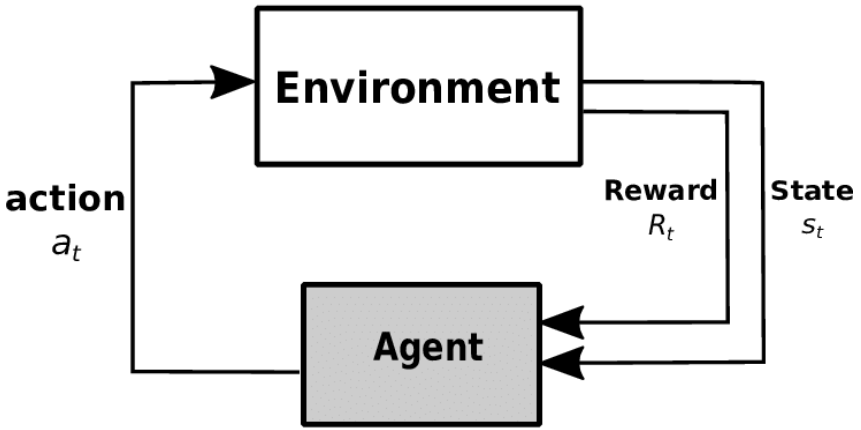

Paradigms of Machine Learning

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

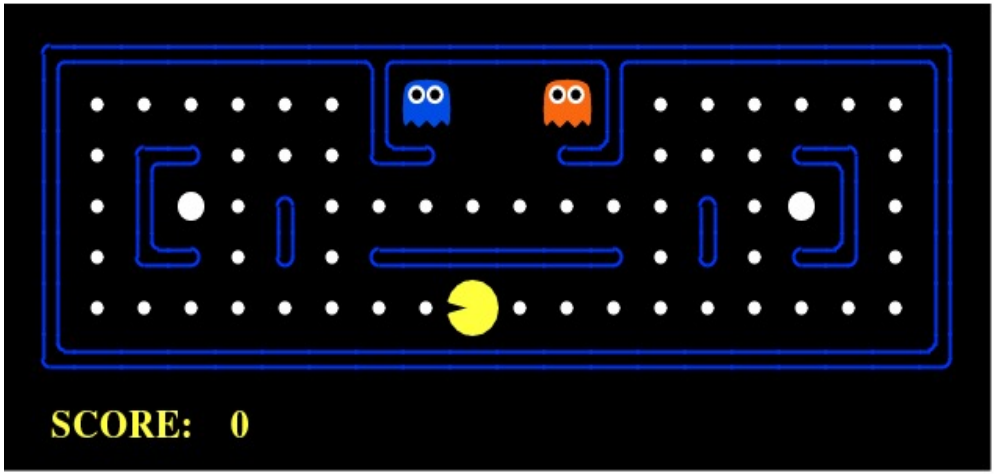

Input type

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

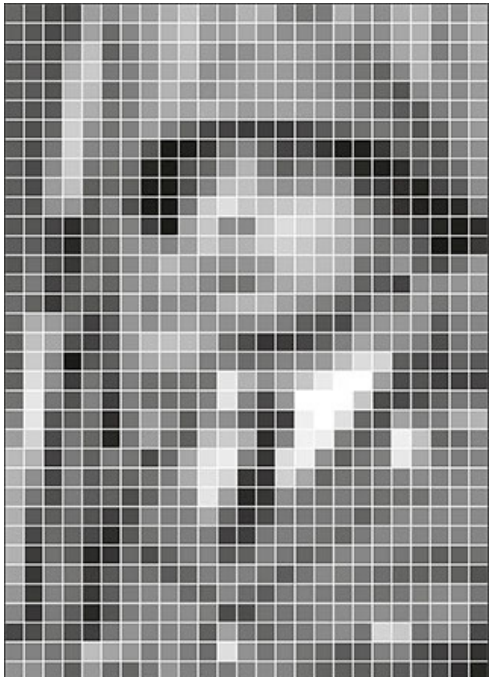

Input type

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

Input type

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

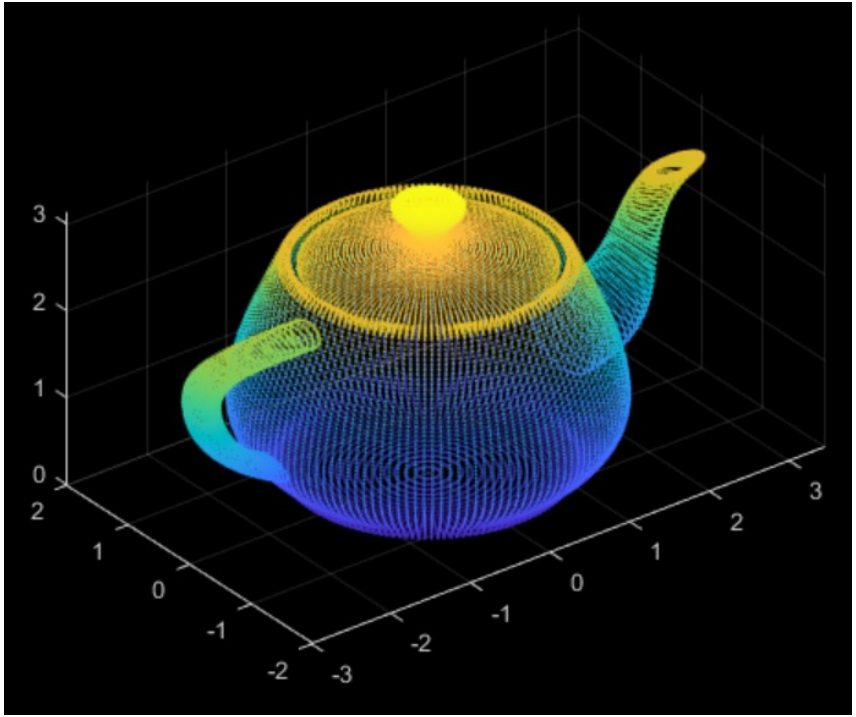

Output type

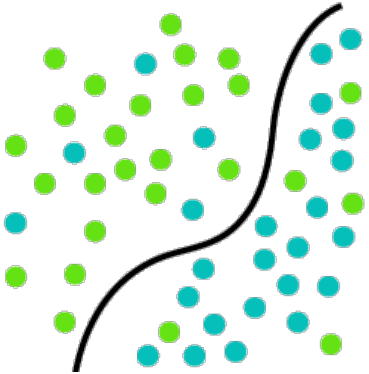

Classification

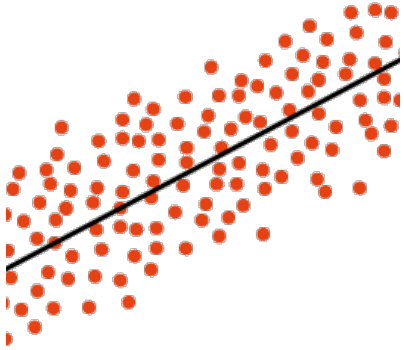

Regression

Regression

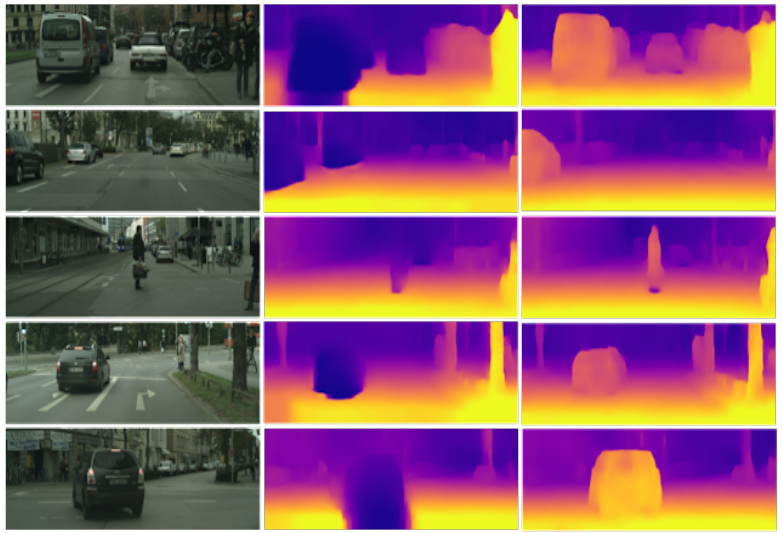

Classification

Bird

Cat

Flamingo

Cock

Egyptian cat

Persian cat