

A futuristic scene with a robotic hand pointing at a laptop screen, overlaid with various data charts and graphs. The background is a soft blue gradient. The robotic hand is white and blue, pointing towards the right. The laptop screen is open, showing a glowing blue interface. Overlaid on the scene are numerous data visualization elements: bar charts, line graphs, pie charts, and circular progress indicators. Some of these elements are semi-transparent and appear to be floating in the air. The overall aesthetic is clean, modern, and high-tech.

Machine Visual Perception

Prof Cengiz Öztireli



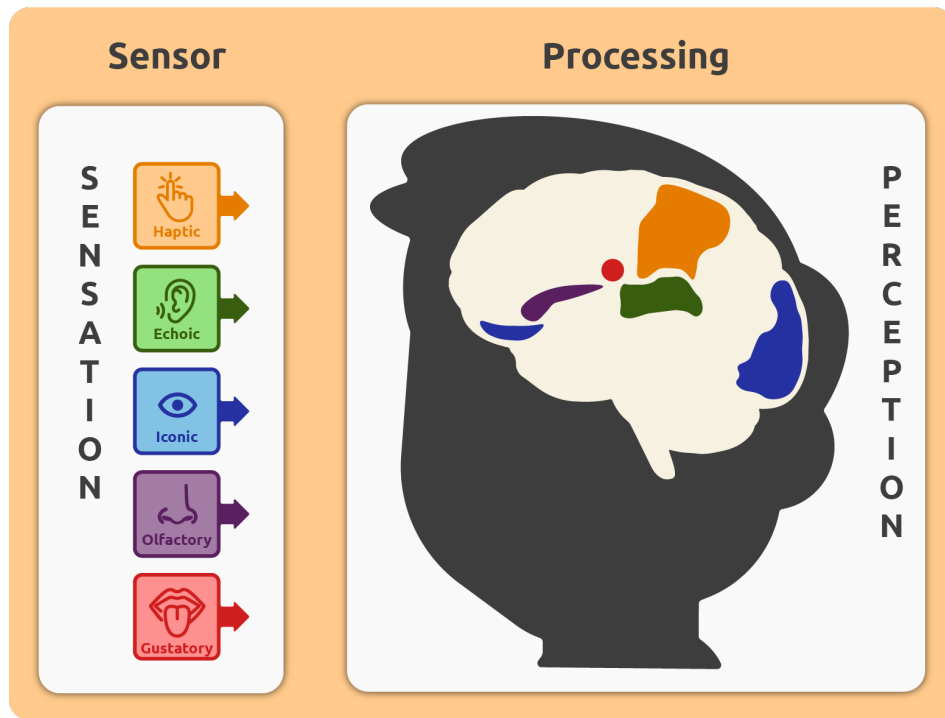
Course format

- 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 three
 - More information will follow on Moodle

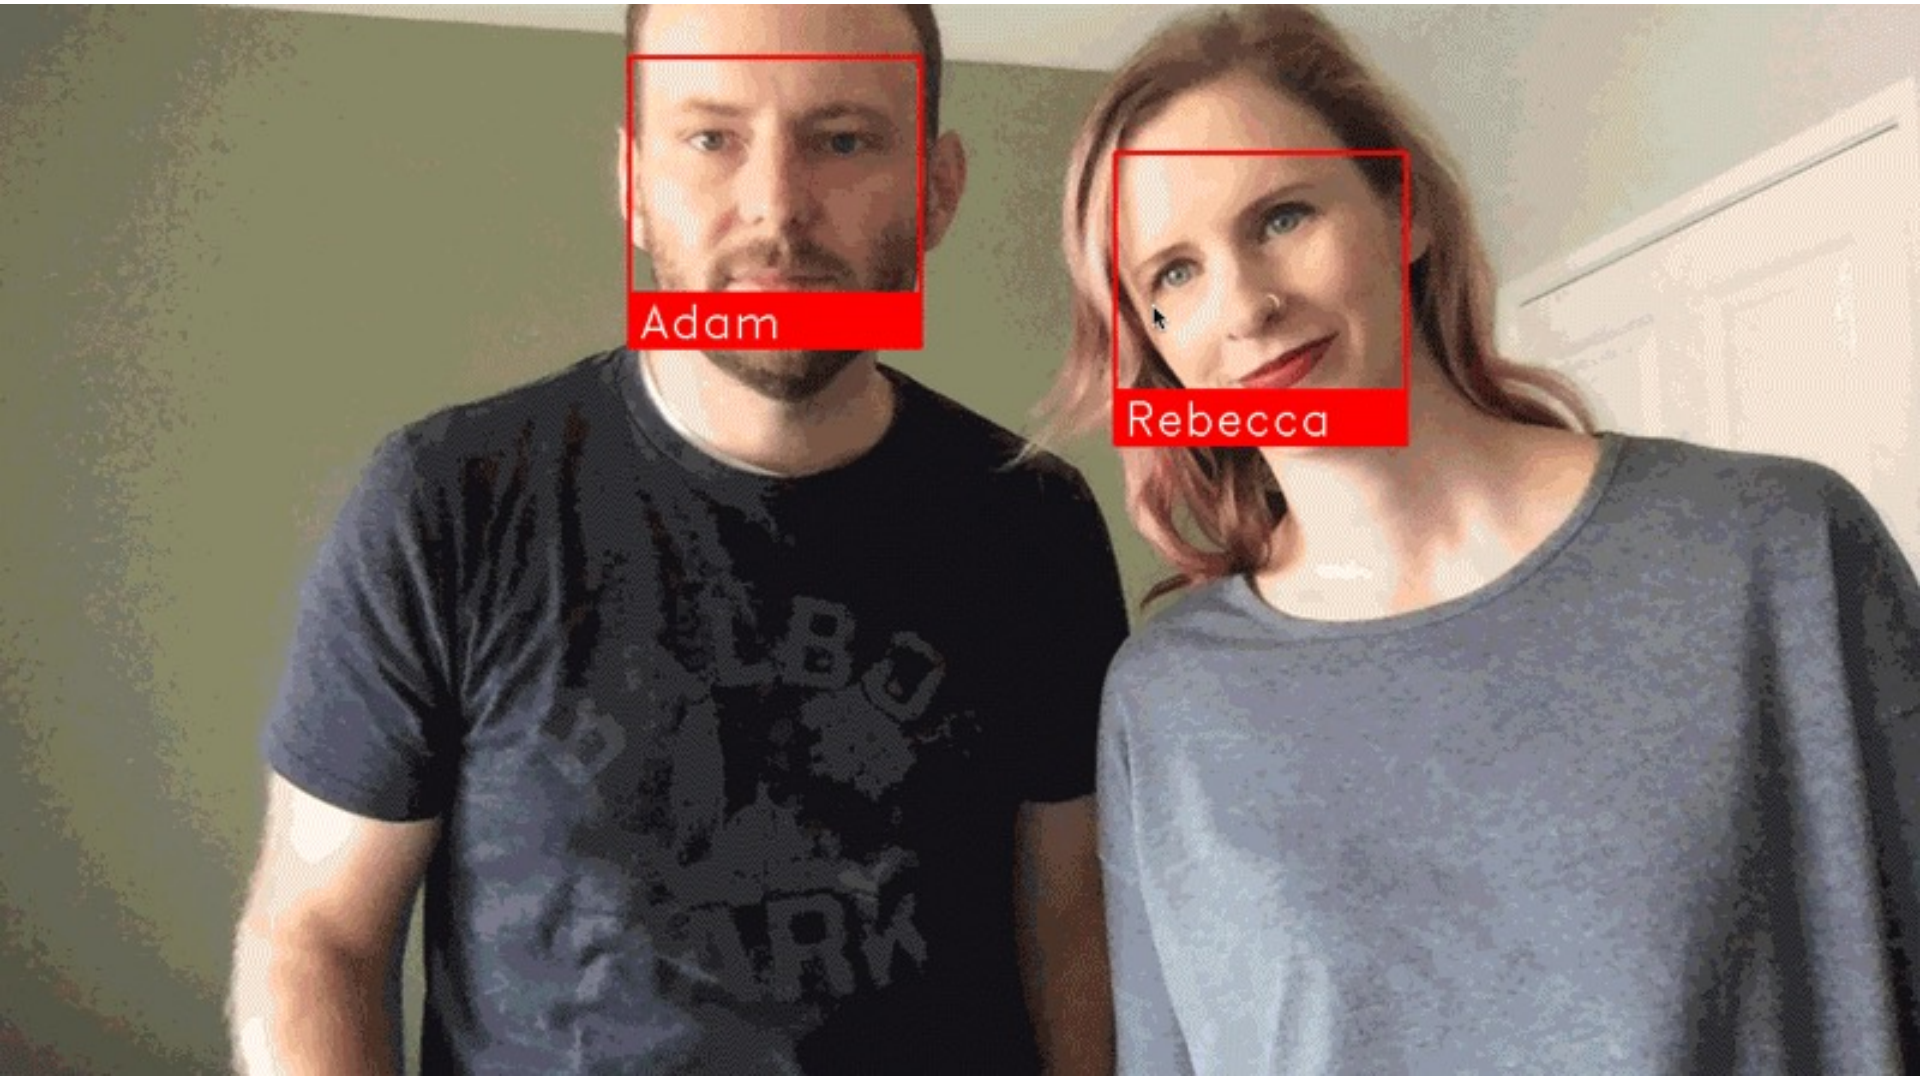
Contributors

- Dr Chris Town (cpt23)
- Dr Fangcheng Zhong (fz261)
- Dr Param Hanji (pmh64)
- Dr Jing Yang (jy496)
- Walter Wu (tw554)

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







1:19 / 2:59



Paradigms of Machine Learning



- Supervised learning
 - Given a training set of N examples

$$(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)$$

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

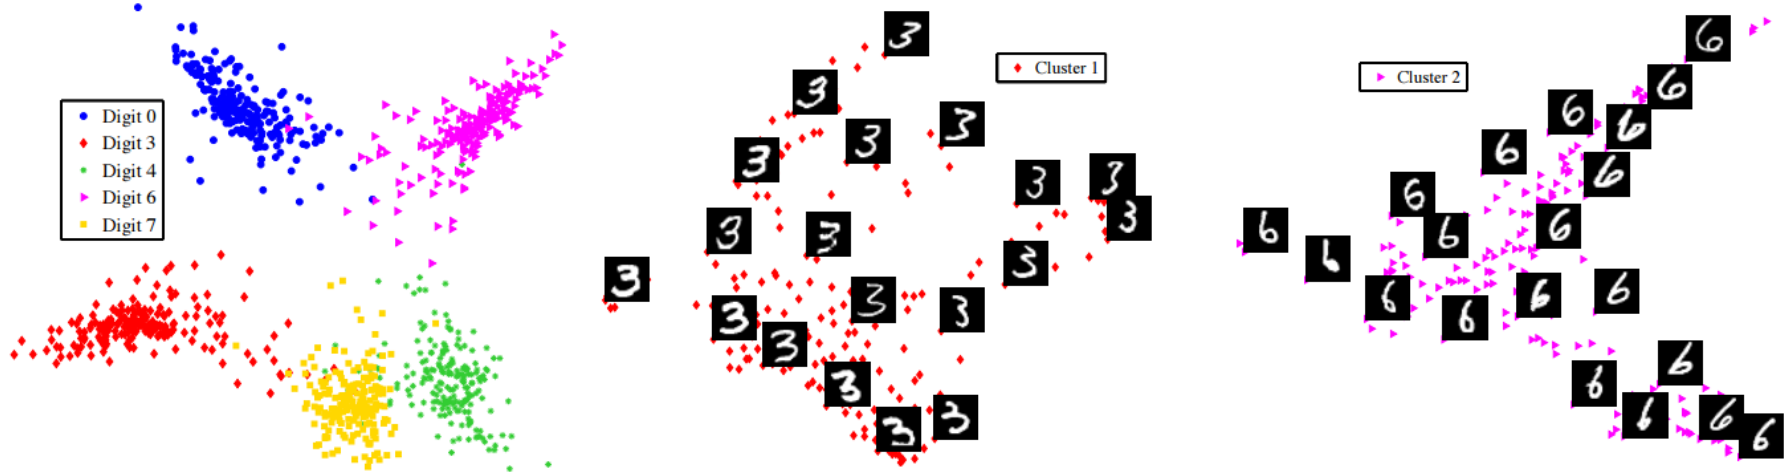
Paradigms of Machine Learning

- 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




Labeled Data



Labeled and Unlabeled Data



Supervised Learning



Semi-Supervised 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

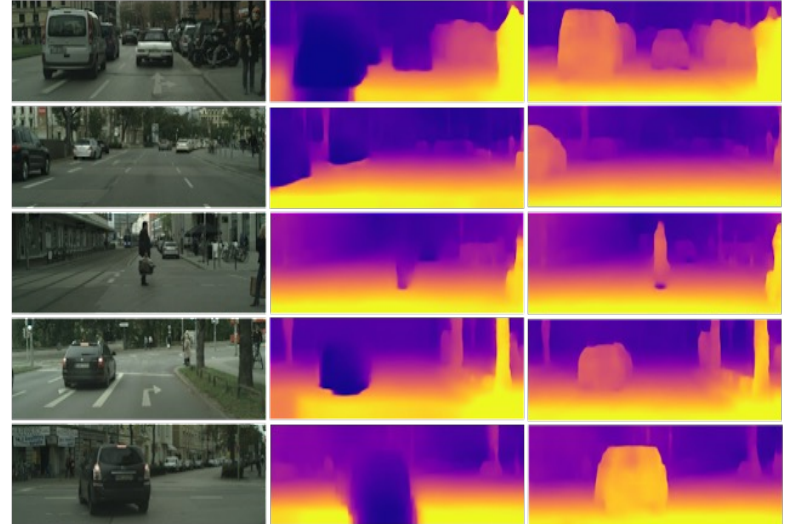


Classification



Regression

Regression



Classification



Bird



Flamingo



Cock




Cat



Egyptian cat



Persian cat



Practical Exercise

- Two parts: classification & regression
- Details will be on the course web page and Moodle
- Questions? Please ask the demonstrators and myself.
- Deadline: 10 November 2023, 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 2023, 12:00 PM