Engineering Tripos Part IB, 2P8: Information Engineering, 2019-20

Lecturers

Professor R Cipolla, Dr R Turner, Dr G Vinnicombe [1]

Timing and Structure

Easter Term: Weeks 1-4 - 14 lectures + 2 examples classes, 4 lectures/week

Aims

The aims of the course are to:

- Provide a unified view of information engineering showing how signal processing, computer vision, machine learning and control relate to one another.
- Use example applications drawn from autonomous driving to provide concrete examples of important concepts and subareas of information engineering including computer vision, machine learning and reinforcement learning
- Introduce computer vision including algorithms for 3D reconstruction, registration and object recognition
- Introduce basic concepts in inference, learning and optimisation including maximum-likelihood estimation, Bayes' rule and gradient descent.
- Introduce basic algorithms for planning and the general area of sequential decision making / reinforcement learning

Objectives

As specific objectives, by the end of the course students should be able to:

- Provide example applications of machine perception, machine learning, and autonomous decision making systems.
- Understand the mathematical basis for perspective projection and feature detection; neural networks and parameter estimation; basic planning and reinforcement learning.
- Implement methods to solve simple computer vision and machine learning problems including object detection and segmentation and sequential decision making.

Content

A: Introduction to Autonomous Driving (1L) (Guest lecturer from industry)

- The anatomy of a self-driving car with description of autonomous driving hardware (the car, sensors, interfaces and actuators)
- Motivate the need for machine perception (computer vision), learning and decision making systems
- Important sub-problems in the data processing pipeline: object detection, localisation and mapping, prediction, planning and action
- Interaction with an example of a self-driving car in Cambridge

B: Machine Perception: Introduction to Computer Vision (6L) (R. Cipolla)

- An introduction to computer vision: reconstruction, registration and recognition
- Perspective projection

- Convolution with gaussians and derivatives of gaussians to provide bandpass filters.
- Edge detection using directional filters.
- Scale-space and image pyramids for feature detection
- The SIFT feature descriptor for matching image features.
- Demonstration of state-of-the-art object detection, semantic segmentation and localisation systems
- Examples paper and class

C: Machine Learning: Introduction to Deep Learning (5L) (R. Turner)

- Training a simple classifier: logistic regression and gradient descent
- Neural networks: Multi-layer perceptrons and back propagation
- Neural networks: convolutional neural networks
- Anatomy and training of a convolutional neural network in Tensorflow or PyTorch - network architecture, loss function, weight initialization, batch size, learning rate, epochs.
- Examples paper and class

D: Autonomous Decision Making: Introduction to Planning and Reinforcement Learning (4L) (G. Vinnicombe and A. Kendall)

- Introduction to planning: Shortest path problems, value functions and dynamic programming
- Introduction to reinforcement learning: Q-learning, actor-critic methods, approximations using neural networks
- Application of these methods to a self-driving car
- Demonstration of perception and planning algorithms running on a self-driving car
- Identification of open problems
- Examples paper

Booklists

Please see the Booklist for Part IB Courses [2] for references for this module.

Examination Guidelines

Please refer to Form & conduct of the examinations [3].

Last modified: 05/08/2020 08:36

Links
[1] mailto:rc10001@cam.ac.uk,ret26@cam.ac.uk,gv103@cam.ac.uk
[2] https://www.vle.cam.ac.uk/mod/book/view.php?id=364081&chapterid=44271
[3] https://teaching.eng.cam.ac.uk/content/form-conduct-examinations