

Engineering Tripos Part IIB, 4F10: Deep Learning & Structured Data, 2017-18

Leader

[Prof M Gales](#) [1]

Lecturer

Prof M Gales, Dr J M Hernandez-Lobato

Timing and Structure

Michaelmas term. 14 lectures + 2 examples classes. Assessment: 100% exam

Prerequisites

Part IIA Modules 3F1, 3F3 and 3F8 advisable

Aims

The aims of the course are to:

- This module aims to teach the basic concepts of deep learning and forms of structure that can be used for generative and discriminative models. In addition the use of models for classifying structured data, such as speech and language, will be discussed

Objectives

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

- Understand the basic principles of pattern classification and deep learning;
- Understand generative and discriminative models for structured data;
- Understand the application of deep-learning to structured data;
- Be able to apply pattern processing techniques to practical applications.

Content

Introduction (1L)

Links with 3F8 and 4F13. General machine learning, examples of structured data, DNA, vision, speech and language processing.

Graphical Models and Conditional Independence (1L)

Graphical models and Bayesian networks. Simple inference examples. Dynamic Bayesian networks. Hidden Markov Models reminder.

Restricted Boltzman Machines (1L)

Structure of restricted Boltzman machines, contrastive divergence.

Deep Learning (2L)

Generative and discriminative deep models. Forms of network and activation functions. Convolutional neural networks, mixture-density neural networks. Optimisation approaches (first/second order methods, adaptive learning rates) and initialisation.

Deep Learning for Sequences (1L)

Recurrent neural networks, and long-short-term memory models. Variants of RNN including bidirectional RNNs. Use in generative and discriminative models.

Alternate Deep Network Architectures (1L)

Auto-encoders and variational extension, student-teacher training (possibly other examples: adversarial networks, siamese networks).

Discriminative Sequence Models (1L)

Conditional random fields and log-linear models (discuss maximum entropy models).

Support Vector Machines (2L)

Maximum margin classifiers, handling non-separable data, training SVMs, non-linear SVMs, kernel functions. Links with other kernel methods Gaussian Processes, Relevance Vector Machines. Multi-class SVMs and structured SVMs.

Kernels over Structured Data (1L)

Tree kernels, graph kernels, Fisher kernels. Relationship to RNNs.

Traditional and Bayesian Non-Parametric Techniques(2L)

- Classification and regression trees, parzen windows, K-nearest neighbours, nearest neighbour rule. Ensemble methods: random forests, bagging, boosting and model combination.

Application: Speech Processing (1L)

Example application of deep-learning to speech processing.

Booklists

Please see the [Booklist for Group F Courses](#) [2] for references for this module.

Examination Guidelines

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

UK-SPEC

This syllabus contributes to the following areas of the [UK-SPEC](#) [4] standard:

[Toggle display of UK-SPEC areas.](#)

GT1

Develop transferable skills that will be of value in a wide range of situations. These are exemplified by the Qualifications and Curriculum Authority Higher Level Key Skills and include problem solving, communication, and working with others, as well as the effective use of general IT facilities and information retrieval skills. They also include planning self-learning and improving performance, as the foundation for lifelong learning/CPD.

IA1

Apply appropriate quantitative science and engineering tools to the analysis of problems.

IA2

Demonstrate creative and innovative ability in the synthesis of solutions and in formulating designs.

KU1

Demonstrate knowledge and understanding of essential facts, concepts, theories and principles of their engineering discipline, and its underpinning science and mathematics.

KU2

Have an appreciation of the wider multidisciplinary engineering context and its underlying principles.

E1

Ability to use fundamental knowledge to investigate new and emerging technologies.

E2

Ability to extract data pertinent to an unfamiliar problem, and apply its solution using computer based engineering tools when appropriate.

E3

Ability to apply mathematical and computer based models for solving problems in engineering, and the ability to assess the limitations of particular cases.

P1

A thorough understanding of current practice and its limitations and some appreciation of likely new developments.

P3

Understanding of contexts in which engineering knowledge can be applied (e.g. operations and management, technology, development, etc).

P8

Ability to apply engineering techniques taking account of a range of commercial and industrial constraints.

US1

A comprehensive understanding of the scientific principles of own specialisation and related disciplines.

US2

A comprehensive knowledge and understanding of mathematical and computer models relevant to the engineering discipline, and an appreciation of their limitations.

Last modified: 25/09/2017 11:27

Source URL (modified on 25-09-17): <https://teaching17-18.eng.cam.ac.uk/content/engineering-tripos-part-iib-4f10-deep-learning-structured-data-2017-18>

Links

[1] <mailto:mjfg100@cam.ac.uk>

[2] <https://www.vle.cam.ac.uk/mod/book/view.php?id=364101&chapterid=55921>

[3] <https://teaching17-18.eng.cam.ac.uk/content/form-conduct-examinations>

[4] <https://teaching17-18.eng.cam.ac.uk/content/uk-spec>