CS2411 - Subsymbolic Processing and Neural Networks

Lecturer: Jon Shapiro, room 2.34, extension 6253, e-mail jls.

Textbooks: The readings cover the core of the material and contain material which you are expected to know. I recommend that you use one of these two books.

- Computer Systems That Learn, Weiss and Kulikowski, 1991.
  This book is the closest to the course, but is at a more elementary level. It covers validation and testing very well.

or

- Adaptive Pattern Recognition and Neural Networks, Yoh-Han Pao (1989).
  This is very good book. It is better on neural networks; less good than the above on Bayesian approaches. It contains more mathematical detail and advanced material.

Handouts: Unfortunately, neither book covers all of the material from the course. For this material, handouts will be given.

Course Description

Lectures: 24 lectures.

Examples classes: There are five one-hour examples classes for this course. These occur on week B Mondays at 4 pm and week B Fridays at 3 pm in room LF17.

Labs: There are five two-hour lab sessions. These occur on week A Mondays at 3-5 pm in the Engineering lab. If there are too many students on the course to fit into one lab session, the other session will take place on week A Fridays at 3-5 in the Engineering lab. There are three lab exercises for this course Exercise 1 is on Bayesian classification; this lasts for two lab sessions. Exercise 2 is on multi-layer perceptron neural networks; it lasts for two lab sessions. Exercise 3 is on genetic algorithms; it lasts for a single session. The labs use the Matlab software on Sun workstations – a tutorial for this package will be handed out during lectures.

The labs will be marked in the lab sessions, that is you will show me or a demonstrator your results. In addition, send the following to the lab archive:

Exercise 1: evaluate.m — the Matlab code which computes the feature which is used to classify the images

Exercise 2: generalise.txt — a text file containing the generalisation performance for different numbers of training patterns used to answer the question, “How many training patterns are required to get a generalisation error of 10%”.

Exercise 3: fitness.m — the fitness function of the genetic algorithm as a matlab function.

Reading by Syllabus Topic

Listed below are the relevant sections of the recommended reading and the supplemental readings for various topics of the course. It is expected that you will need to read at least one piece of supportive material for each topic.

Introduction(Lectures 1): Weiss and Kulikowski (WK) Chapter 1; Pao Chapter 1.

Bayesian Decision and Classification(Lectures 2,3,4): Probability theory handout, WK Chapter 3; Pao Chapter 2.

Validation(Lecture 5): WK Chapter 2.

Generalisation(Lecture 6): none

Single-Layer Neural Networks (Lectures 7,8): Perceptrons handout; WK Chapter 4; Pao Chapter 5.

Multi-Layer Neural Networks (Lectures 9 – 11): Backpropagation handout; WK Chapter 4; Pao Chapter 6.

Learning from Non-numeric Data(Lectures 12 – 14: Information theory handout, WK Chapter 5; Pao Chapter 4 (and possibly 3).

Non-symbolic Search Techniques(Lectures 15, 16): .

Genetic Algorithms(Lectures 17, 18): Genetic Algorithms handout.

Unsupervised learning(Lectures 19 – 21): Unsupervised learning handout.

Schedule of topics

Below is a schedule of the topics of the course.

Week Starting:	Lecture:	Examples Class:	Lab:
Th Sept 24 (A)	(away) Introduction		
Th Oct 1 (B)	Bayesian approaches Bayes II	Probability Review	
Th Oct 8 (A)	Bayes III Validation		Bayes lab 1
Th Oct 15 (B)	Generalization Neural Network approaches	Bayes	
Th Oct 22 (A)	Perceptrons Multi-layer perceptrons		Bayes lab 2
Th Oct 29 (B)	Applications of MLPs Applications II	Neural Networks	
Mon Nov 9 (A)	Decision Trees ID3		Neural Networks lab 1
Mon Nov 16 (B)	ID3 Search Techniques	Decision Trees	
Mon Nov 23 (A)	Search Genetic Algorithms		Neural Networks Lab 2
Mon Nov 30 (B)	Genetic Algorithms Unsupervised Learning	Genetic Algorithms	
Mon Dec 7 (A)	Unsupervised Learning Unsupervised Learning		Genetic Algorithms lab
Mon Dec 14 (B)	Review I Review II		Lab Catchup

Other Sources

Here are some other sources of information about the topics in the course. This includes both simple and advanced treatments of the material; the simpler texts are listed first.

- Neural Computing, Beale and Jackson, 1990

  An elementary book on neural networks. Good treatment of Bayesian classification, neural networks, and some of the search aspects of the course at a simple level. Contains nothing on validation or non-numerical methods.

- Machine learning, Neural and Statistical Classification, D. Michie, D.J. Spiegelhalter and C. C. Taylor, 1994

  Not a textbook. A collection of articles describing current applications of and approaches to the techniques studied in the course.

- Techniques in Computational Learning C. J. Thornton, 1992.

  A book on methods of learning from data, including supervised learning and unsupervised learning in neural networks, and learning in decision trees, as well as learning methods not covered in this course. It contains nothing on generalisation, validation, or Bayesian methods, unfortunately.

- Modern Heuristic Techniques for Combinatorial Problems, C. Reeves (editor), 1993.

  An overview of search techniques including simulated annealing and genetic algorithms.

- Genetic Algorithms in Search, Optimization, and Machine Learning, D. Goldberg, 1989.

  The standard reference on Genetic Algorithms.

- Neural Networks, A comprehensive Foundation, Simon Haykin, 1994. Second edition, 1998.

  An excellent book on neural networks. The only problem with it is that it contains almost no mention of applications. Advanced and comprehensive. Second edition is very up to date.

- Neural Networks for Pattern Recognition, Christopher Bishop, 1996

  A possible future textbook for this course — this book contains an excellent treatment of Bayesian approaches, neural networks, and generalisation, as well as other aspects not treated here. Somewhat more advanced in its treatment, its study will give a deep understanding of the material.

- Pattern Classification and Scene Analysis, Duda and Hart, 1973

  The classic work on traditional statistical methods and Bayesian methods.

Here is the alternative reading by syllabus topic.

Bayesian Classification: Beale and Jackson, Chapter 2; Michie, Spiegelhalter and Taylor (MST) Chapters 2,3,4; Bishop Chapters 1 and 2.

Generalisation and Validation: MST Chapter 7; Bishop Chapters 1.5, 9.1, 9.2, and 9.8.

Supervised Learning Neural Networks: Beale and Jackson Chapters 3 and 4; MST Chapter 6; Thornton Chapters 9 – 14; Bishop Chapters 3 and 4.

Unsupervised Learning in Neural Networks: Beale and Jackson Chapter 5; Thornton Chapter 18.

Learning from Non-Numeric Data: MST Chapter 5; Chapter 5 and 6.

Subsymbolic Search Techniques and Genetic Algorithms: Reeves Chapters 1, 2, and 4; Goldberg Chapter 2.