Introduction to Neural Computation

Level 4/M

Neural Computation

Level 3

Website: http://www.cs.bham.ac.uk/~jxb/inc.html

Lecturer: Dr. John A. Bullinaria

Module Administration and Organisation

Neural Computation : Lecture 1 (part 1)

1. Aims and Learning Outcomes
2. Assessment
3. Lecture Plan
4. Recommended Books
5. Comments on Mathematical Requirements

Aims and Learning Outcomes

Aims

1. Introduce some of the fundamental techniques and principles of neural computation.
2. Investigate some common neural-based models and their applications.
3. Present neural network models in the larger context of state-of-the-art techniques of automated learning.

Learning Outcomes

1. Understand the relationship between real brains and simple artificial neural network models.
2. Describe and explain some of the principal architectures and learning algorithms of neural computation.
3. Explain the learning and generalization aspects of neural computation.
4. Demonstrate an understanding of the benefits and limitations of neural-based learning techniques in context of other state-of-the-art methods of automated learning.
5. Apply neural computation algorithms to specific technical and scientific problems. [L4 only]

Assessment

The assessment will depend on which module is being taken:

Neural Computation (Level 3)

100% 1.5 hour closed book examination in May/June

Resit (when allowed) - same as the normal assessment, but in August

Introduction to Neural Computation (Level 4/M)

80% 1.5 hour closed book examination in May/June

20% Continuous assessment by mini-project report with deadline in January. It will be given out in Week 6 and will involve a practical exercise based on the content of the first half of the module.

Resit (when allowed) - same as the normal assessment, but in August

The lecture content and examination will be the same for both modules.

Lecture Plan

Week	Lecture 1	Lecture 2
1	Introduction to Neural Networks and their History	Biological Neurons and Neural Networks, Artificial Neurons
2	Networks of Artificial Neurons, Single Layer Perceptrons	Learning and Generalization in Single Layer Perceptrons
3	Hebbian Learning, Gradient Descent Learning	The Generalized Delta Rule, Practical Considerations
4	Learning in Multi-Layer Perceptrons – Back-Propagation	Learning with Momentum, Conjugate Gradient Learning
5	Bias and Variance – Under-Fitting and Over-Fitting	Improving Generalization
6	Applications of Multi-Layer Perceptrons	Exercise Session
7	Recurrent Neural Networks	Radial Basis Function Networks: Introduction
8	Radial Basis Function Networks: Algorithms	Radial Basis Function Networks: Applications
9	Self Organizing Maps: Fundamentals	Self Organizing Maps: Properties and Applications
10	Learning Vector Quantization	Committee Machines
11	Model Selection and Evolutionary Optimization	Exercise Session

Main Recommended Books

Title	Author(s)	Publisher, Date	Comments
Neural Networks and Learning Machines	Simon Haykin	Pearson, 2009	Very comprehensive, but heavy in mathematics.
Neural Networks: A Comprehensive Foundation	Simon Haykin	Prentice Hall, 1999	Older edition of the above book, but often better.
Neural Networks for Pattern Recognition	Christopher Bishop	Clarendon Press, Oxford, 1995	This is the book I always use, but it doesn't cover the whole module.
The Essence of Neural Networks	Robrt Callan	Prentice Hall Europe, 1999	Concise introductory text.
An Introduction to Neural Networks	Kevin Gurney	UCL Press, 1997	Non-mathematical introduction.

Other Good Books

Title	Author(s)	Publisher, Date	Comments
Principles of Neurocomputing for Science and Engineering	F.M. Ham & I. Kostanic	McGraw Hill, 2001	Good advanced book, but rather mathematical.
Fundamentals of Neural Networks	Laurene Fausett	Prentice Hall, 1994	Good intermediate text.
Introduction to Neural Networks	R. Beale & T. Jackson	IOP Publishing, 1990	Former recommended book.
An Introduction to the Theory of Neural Computation	J. Hertz, A. Krogh & R.G. Palmer	Addison Wesley, 1991	Good all round book. Slightly mathematical.
Parallel Distributed Processing: Volumes 1 and 2	D.E. Rummelhart, J.L. McClelland, et al.	MIT Press, 1986	The original neural networks bible.
The Computational Brain	P.S. Churchland & T.J. Sejnowski	MIT Press, 1994	Good for computational neuroscience.

Comments on Mathematical Requirements

1. The easiest way to formulate and understand neural networks is in terms of mathematical concepts and equations.

2. Once you have the equations, it is fairly straightforward to convert them into C/C++/Java/MATLAB/etc. programs.

3. It is a prerequisite of this module to have passed A Level Mathematics or equivalent, and that level of mathematics will be assumed, though brief summaries of relevant mathematics will be provided as we go along.

4. It will be made clear in the lectures which aspects of the mathematics you may be expected to reproduce in the examinations, and which you should try to follow but will not be expected to reproduce.

5. Don't be surprised if you find it hard to become familiar with all the new notation – this is normal! However, it will be difficult to follow the lectures if you don't make an effort to understand the notation.

Introduction to Neural Networks and Their History

Neural Computation : Lecture 1 (part 2)

1. What is Neural Computation?
2. Why are Artificial Neural Networks Worth Studying?
3. Learning in Neural Networks
4. A Brief History of the Field
5. Artificial Neural Networks compared with Classical Symbolic AI
6. Some Current Artificial Neural Network Applications

What is Neural Computation ?

1. ***Neural Computation*** is a general ***Machine Learning*** approach that involves processing information in a similar manner to the networks of neurons (i.e. ***Neural Networks***) found in human/animal brains.

2. ***Artificial Neurons*** are crude approximations of the neurons found in biological brains. They may be physical devices, or purely mathematical constructs.

3. ***Artificial Neural Networks*** (ANNs) are networks of Artificial Neurons, and hence constitute crude approximations to parts of real brains. They may be physical devices, or simulated on conventional computers.

4. From a practical point of view, any ANN is just a parallel computational system consisting of many simple processing elements connected together in a specific way in order to perform a particular task.

5. One should never lose sight of how crude the approximations are, and how over-simplified ANNs are compared to the neural networks in real brains.

Why are Artificial Neural Networks worth studying?

1. Even though individual artificial neurons are very simple, networks of them can be shown to be extremely powerful computational devices (Turing equivalent, universal computers).

2. Very simple ANNs can be set up to learn and generalize well – so they can perform difficult tasks without the need for enormous feats of programming.

3. Their massive parallelism can make them very efficient.

4. They are particularly fault tolerant – this is equivalent to the "graceful degradation" found in biological brains.

5. They are very noise tolerant – so they can cope with situations where normal symbolic (rule-based) systems would have difficulty.

6. In principle, they can do anything a symbolic/logic system can do, and a lot more. Though, in practice, getting them to do it can be rather difficult…

What are Artificial Neural Networks used for?

As with the field of AI in general, there are two basic goals for neural network research:

Brain modelling : The *scientific* goal of building models of how real brains work. This can potentially help us understand the nature of human intelligence, formulate better teaching strategies, or better remedial actions for brain damaged patients.

Artificial System Building : The *engineering* goal of building efficient systems for real world applications. This may make machines more powerful, relieve humans of tedious tasks, and may even improve upon human performance.

These should not be thought of as competing goals. We often use exactly the same neural networks and techniques for both. Frequently progress is made when the two approaches are allowed to feed into each other. There are fundamental differences though, e.g. the need for biological plausibility in brain modelling, and the need for computational efficiency in artificial system building.

Learning in Neural Networks

There are many different types of neural networks. Most operate by passing neural "activations" through a network of connected neurons.

One of the most useful and powerful features of neural networks is their ability to ***learn*** and ***generalize*** from a set of training data. They adapt the strengths/weights of the connections between neurons so that their final output activations are optimized.

There are three broad types of learning:

1. Supervised Learning (i.e. learning with a teacher)

2. Reinforcement learning (i.e. learning with limited feedback)

3. Unsupervised learning (i.e. learning with no help)

This module will study in some detail the most common learning algorithms for the most common types of neural network.

A Brief History of the Field

1943 McCulloch and Pitts proposed the McCulloch-Pitts neuron model

1949 Hebb published his book *The Organization of Behavior*, in which the Hebbian learning rule was proposed.

1958 Rosenblatt introduced the simple single layer networks now called Perceptrons.

1969 Minsky and Papert's book *Perceptrons* demonstrated the limitation of single layer perceptrons, and almost the whole field went into hibernation.

1982 Hopfield published a series of papers on Hopfield networks.

1982 Kohonen developed the Self-Organising Maps that now bear his name.

1986 The Back-Propagation learning algorithm for Multi-Layer Perceptrons was re-discovered and the whole field took off again.

1990s The sub-field of Radial Basis Function Networks was developed.

2000s The power of Neural Network Ensembles, Support Vector Machines, Bayesian Techniques, Simulated Evolution, and Deep Learning becomes apparent.

ANNs compared with Classical Symbolic AI

The distinctions can be categorized under three broad headings:

1. Level of Explanation
2. Processing Style
3. Representational Structure

These lead to a traditional set of dichotomies:

1. Sub-symbolic vs. Symbolic
2. Non-modular vs. Modular
3. Distributed representation vs. Localist representation
4. Bottom up vs. Top Down
5. Parallel processing vs. Sequential processing

In practice, however, the distinctions are becoming increasingly blurred.

Some Current Artificial Neural Network Applications

Brain modelling

Models of human development – to help children with developmental problems

Simulations of adult performance – aiding our understanding of how the brain works

Neuropsychological models – suggesting remedial actions for brain damaged patients

Artificial Life simulations – clarifying brain evolution, structure, growth, etc.

Real world applications

Pattern recognition – speech recognition, hand-writing recognition, sonar signals

Data analysis – data compression, data mining, PCA, GTM

Noise reduction – function approximation, ECG noise reduction

Control systems – autonomous adaptable robots, microwave controllers

Computer games – intelligent agents, backgammon, first person shooters

Financial modelling – predicting stocks, shares, currency exchange rates

Other time series prediction – climate, weather, airline marketing tactician

Overview and Reading

1. Artificial Neural Networks are extremely powerful (Turing equivalent) computational systems consisting of many simple processing elements connected together to perform tasks analogously to biological brains.
2. They are massively parallel, which makes them efficient, robust, fault tolerant and noise tolerant.
3. They can learn from training data and generalize to new situations.
4. They are useful for brain modelling and real world applications involving pattern recognition, function approximation, prediction, …

Reading

1. Haykin-1999: Sections 1.1, 1.8, 1.9
2. Gurney: Sections 1.1, 1.2, 1.3
3. Beale & Jackson: Sections 1.1, 1.2, 1.3, 1.4
4. Ham & Kostanic: Sections 1.1, 1.2