Mind the Gap

Computational Neuroscience

Nick Taylor

Department of Computing and Electrical Engineering

Heriot-Watt University

Edinburgh

Topics

- Computational neuroscience
- Neurophysiology
- Connectionism
- Common problems
- Points of divergence
- Bridging examples
- Concluding remarks

Computational Neuroscience

Computational neuroscience is characterised by its focus on understanding the nervous system as a computational device rather than by a particular experimental technique.

Experimentation and Modelling

- Neuronal Networks
- Sensory Systems
- Motor Systems
- Cerebral Cortex

Two Disciplines

- Neurophysiology
  - – Province of Biological Neuronal Network (BNN) Experimenters
- Connectionism
  - – Province of Artificial Neural Network (ANN) Modellers

Differing Perspectives

- BNN Experimenters' agenda
  - Understanding
    - Neurogenesis; Neurotransmitters; Plasticity
  - Pathology
    - Neuronal dysfunction; Diagnosis; Treatments
- ANN Modellers' agenda
  - Performance
    - Training/execution speeds; Reliability; Flexibility
  - Applicability
    - Architectures; Complexity; Fault tolerance

Topics

- Computational neuroscience
- Neurophysiology
- Connectionism
- Common problems
- Points of divergence
- Bridging examples
- Concluding remarks

Neurophysiology

- Background
- Axons, synapses & neurons
- Learning & synaptic plasticity
- Problems
- Summary

Background

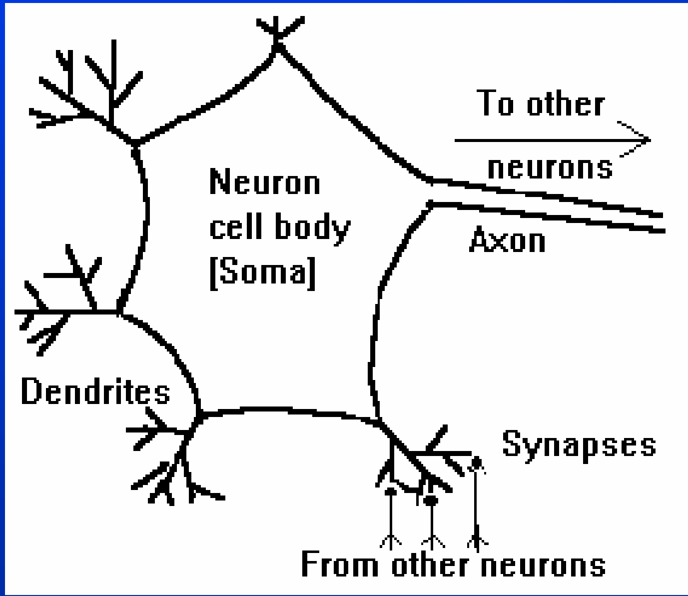

- The human brain contains about 1 billion neurons
- Each neuron is connected to thousands of others
- Neurons can be either excitatory or inhibitory

- Neurons perform very simple computations
- The computational power of the brain is derived from the complexity of the connections

Axons, Synapses and Neurons

- The primary mechanism for information transmission in the nervous system is the axon
- An axon relays all-or-nothing (binary) impulses
- Signal strength is determined from the frequency of the impulses
- An axon signal eventually arrives at a synapse
- A synapse may either attenuate or amplify the signal whilst transmitting it to a neuron
- A neuron accumulates the modified signals and produces an impulse on its own axon if the total synaptic input strength is sufficient

Model of a Neuron

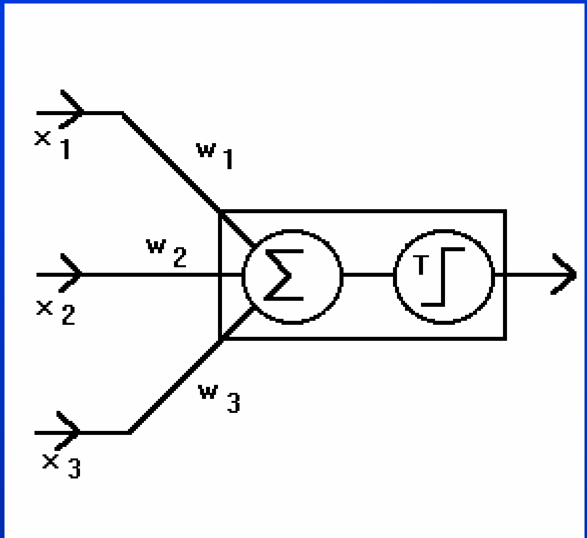

- McCulloch and Pitts model of a neuron (1943)
- Summation of weighted inputs
- Threshold, *T*, determines whether the neuron fires or not

- Firing rule:

$$\sum_i x_i w_i > T \quad \text{then} \quad \text{fire}$$

$$\leq T \quad \text{then} \quad \text{don't} \quad \text{fire}$$

Assemblies of Neurons

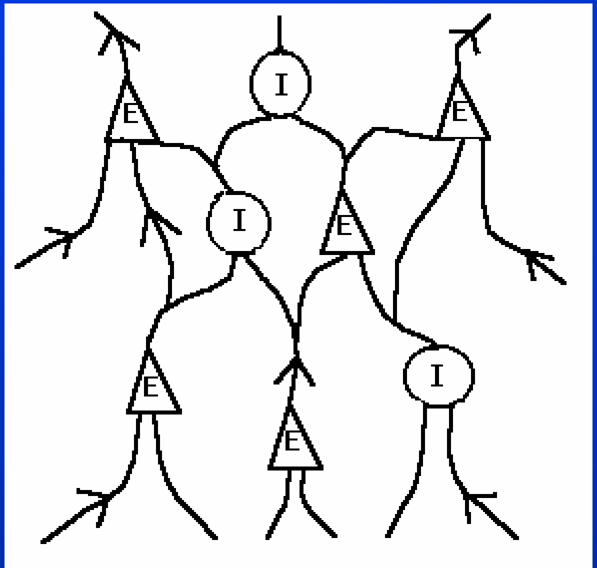

- Modifications to neuron assemblies can only be achieved by adjusting the attenuation or amplification which is applied at the synapses

- Hebb Rule (1949) [*after James (1890!)*]
  - If a particular input is always active when a neuron fires then the weight on that input should be increased
- Learning is achieved through synaptic plasticity

Learning & Synaptic Plasticity I

- Long-Term Potentiation (LTP)
  - Hebbian increases in synaptic efficacy (amplifications) have been recorded on
    - Active excitatory afferents to depolarised (firing) neurons
- Long-Term Depression (LTD)
  - Decreases in synaptic efficacy (attenuations) have been recorded on
    - Inactive excitatory afferents to depolarised (firing) neurons
    - Active excitatory afferents to hyperpolarised (non-firing) neurons
    - Active inhibitory afferents to depolarised (firing) neurons

Learning & Synaptic Plasticity II

- Nitric Oxide
  - Post-synaptic messenger discovered in 1990
  - Released by depolarised (firing) neurons
  - Can affect all active afferents in a local volume
- Consequences
  - NO makes it possible for one or more firing neurons to increase the synaptic efficacy of nearby neurons even if those nearby neurons aren’t firing
  - NO can boot-strap synaptic efficacies which have dropped beyond redemption back to viability

Problems

- Hebbian learning paradigm inadequate
- Scant information on plasticity of inhibitory synapses
- Little known about the implications of the NO discovery for more global forms of plasticity
- Frequency-based models and analyses practically non-existent
- Behaviour of populations of neurons very complex and difficult to investigate

Neurophysiology Summary

- Much is already known
  - Enough to build models
- Neurophysiological correlates for many computational requirements have been found
  - LTP, LTD, NO
- Much is still unknown
  - Enough to severely restrict the models
- NO research is still in its infancy
  - Wider implications yet to be investigated

Topics

- Computational neuroscience
- Neurophysiology
- Connectionism
- Common problems
- Points of divergence
- Bridging examples
- Concluding remarks

Connectionism

- Background
- Architectures
- Applications
- Problems
- Summary

Background

- Artificial Neural Networks (ANNs) are inspired, but not constrained, by biological neuronal networks
- Two very commonly used architectures
  - The Hopfield Network
    - Single layer, total connectivity within layer, auto-associative
  - The Multi-Layer Perceptron
    - Multiple layer, total connectivity between adjacent layers, no connectivity within layers, hetero-associative

The Hopfield Network

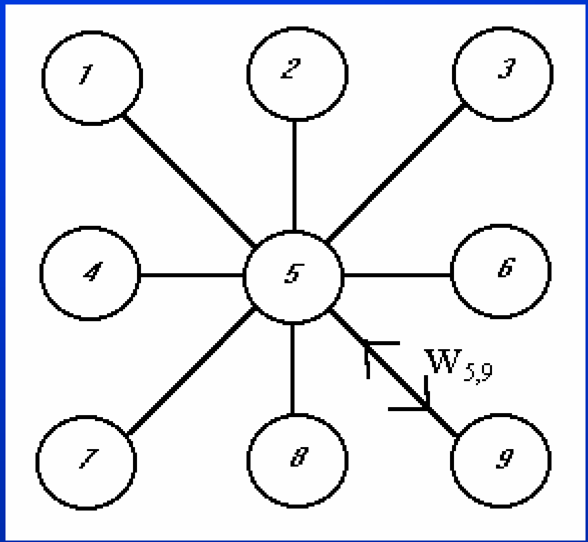

- Each node connected to every other node in the network
- Symmetric weights on connections ($w_{5,9} = w_{9,5}$)
- Node activations either -1 or +1

- Training performed in one pass: $w_{i,j} = \frac{1}{N} \Sigma\, p_i\, p_j$

- Execution performed iteratively: $s_i = \text{sign}\, \{\Sigma\, w_{i,j}\, s_j\}$

The Multi-Layer Perceptron

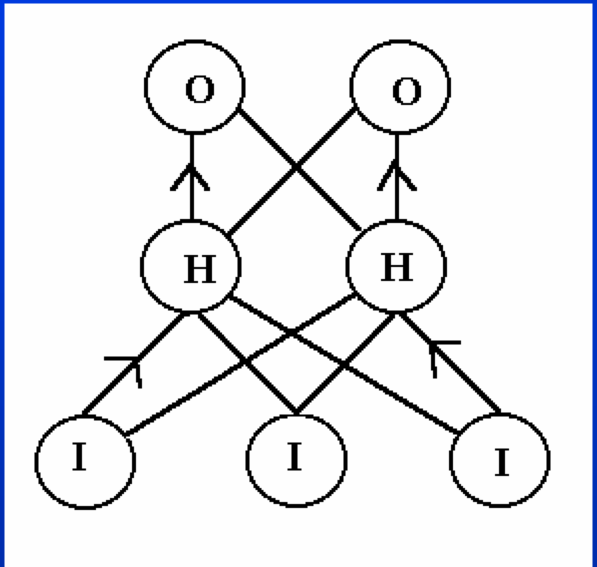

- Each node connected to every node in adjacent layers
- Connections feed forward from input nodes (I), through hidden nodes (H) to output nodes (O)

- Training performed iteratively: $\Delta w_{j,i} = \eta\, \delta_j\, s_i$

- Execution performed in one pass: $s_i = f(\Sigma\, w_{i,j}\, s_j)$

Hopfield Applications

- Content Addressable Memory
  - Partial patterns can be completed to reproduce previously learnt patterns in their entirety
    - Partially incorrect patterns are simply partial patterns
- Optimisation
  - Learnt patterns are simply attractors - minima of some energy function defined in terms of the $w_{i,j}$ and s_i variables
    - Using the objective function in an optimisation problem as the energy function, with suitably defined weights and activation equations, a Hopfield network can find minima of the objective function

MLP Applications

- Classification/Mapping
  - Kolmogorov's Mapping Neural Network Existence Theorem (Hecht-Nielsen)

    Any continuous function, $f:[0,1]^n \rightarrow \Re^m$, *can be implemented exactly by a three-layer MLP having* n *input units,* $(2n+1)$ *hidden units and* m *outputs.*

  - Applications are legion
    - Classification into categories by attribute values
    - Character recognition
    - Speech synthesis (NETtalk)
    - Vehicle navigation (ALVINN)

Problems

- Local minima
  - Hopfield: Linear combinations of learnt patterns or optimal solutions become attractors
  - MLP: Gradient descent training is the inverse of Hill-climbing search and is just as susceptible to local minima as the latter is to local maxima
- Limited storage capacity (Hopfield)
  - Less than $N/ln(N)$ patterns can be memorised safely
- Over-training (MLP)
  - Too many free variables ($w_{i,j}$) thwart generalisation

Connectionism Summary

- Neurologically inspired
  - Biological neurons and assemblies of neurons
- Broad applicability
  - Various architectures and training paradigms
- Readily implemented
  - Simple algorithms and data structures
- Reliability problems
  - Sub-optimality, capacity limitations, over-training, *Black Box* naivety

Topics

- Computational neuroscience
- Neurophysiology
- Connectionism
- Common problems
- Points of divergence
- Bridging examples
- Concluding remarks

Common Problems (with Hebb)

- Local learning
  - Hebbian rules for synaptic adjustment imply local learning only (NO)
- Weight decrementing
  - Hebbian rules provide no mechanism for weights to decrease (LTD)
- Weight resurrection
  - Hebbian rules provide no mechanism for increasing the weights to a neuron whose weights are so low that it never fires (NO + LTP)

Points of Divergence

- Neurophysiological facts violated in ANNs
  - Asymmetric connectivity (Hopfield)
  - Partial connectivity (Hopfield & others)
  - Homogeneity of neuron efferents (All ANNs)
  - Immutability of neuron type (All ANNs)
- Computational requirements lacking evidence
  - Increases in inhibitory weights
  - Availability of global information during learning
  - Incremental learning methods (ANN problem)

Bridging Examples

- Two parts of the brain have received extensive study by both the BNN and ANN disciplines
  - Visual Cortex
  - Hippocampal Formation
- Other aspects of the nervous system have been targetted by the modellers
  - Willshaw (1981)
    - Model for the innervation of skeletal muscle
    - Suggested experimental work to acquire more data
  - Associative nets in general

Examples -Visual Cortex

- Hubel & Wiesel (1968)
  - Identified individual neurons which responded to dark or light bars at specific orientations
- von der Marlsburg (1973)
  - Model of Hubel & Wiesel’s work
- Marr (1976-1982, died 1980)
  - Computer models based on neurophysiological and psychological findings
- Fukushima (1982)
  - Neocognitron ANN model based on the visual system

Examples - Hippocampus

- Marr (1970)
  - Theory of the hippocampus as a simple memory
  - Made predictions to guide neurophysiologists
- O'Keefe & Nadel (1978)
  - Model of the hippocampus for spatial mapping
  - Made refutable predictions about behaviour
- Willshaw & Buckingham (1990)
  - Further implications of Marr's theory
- Traub & Miles (1991)
  - Collective neuronal behaviour and synchronisation

Examples - Associative Nets

- Marr (1969)
  - Theory of the cerebellum as an associative net
- Tyrrell & Willshaw (1992)
  - Model of the cerebellum as an associative net
- Heriot-Watt Studies (1992 - present)
  - Modified Hopfield Networks
    - Asymmetric connectivity mitigates the negative effects of increasingly partial connectivity
    - Asymmetric connectivity improves tolerance to noise
    - Partial connectivity with immutable neuron type (10% inhibitory) and homogeneity of efferents improves recall

Topics

- Computational neuroscience
- Neurophysiology
- Connectionism
- Common problems
- Points of divergence
- Bridging examples
- Concluding remarks

Concluding Remarks

- BNN experimental results have to be correctly placed within a complex body of knowledge
- BNN experimenters can test theories and obtain pointers from ANN models
- ANN models lag woefully behind their BNN counterparts
- ANN modellers can obtain pointers from BNN experiments
- The more Computer Scientists, Mathematicians and Physicists learn about BNN the better