Biologically Inspired Computing: The Neural Network

John Bryden

School of Computing, University of Leeds, Leeds LS2 9JT, England
scs2jab@comp.leeds.ac.uk

1 Introduction

Artificial intelligence has been the inspiration and goal of computing since the discipline was first conceived by Alan Turing. Our understanding of the brain has increased in parallel with the development of computers capable of modelling its functions. While the human brain is vastly complex, too much so for the computation abilities of modern super computers, interesting results have been found while modelling the nervous system of smaller creatures such as the salamander [3].

Attempts to model biological neural networks have created the discipline of artificial neural networks which, through mathematical analysis and other methods, has grown out of its bio-inspired roots.

This paper discusses how neurobiology has inpired computer science by first looking at the biology of the neuron (section 2). A discussion of how biological neurons network (section 3) with each other will lead onto a discussion of the modelling of biological neural networks (section 4). The modelling of biological neural networks was the inspiration for the development of artificial neural networks (section 5).

2 The Neuron

The neuron is not only the basic building block of intelligence but it also forms a message passing and control system within multi-cellular organisms. The basic structure of a neuron can be seen as a body, called the *soma*, with one or many processes, *dendrites* or *axons*, branching out from it. The processes carry messages to and from the soma and terminate in *synapses*.

The neuron is only one cell, much like any other in the body. It has DNA code and is generated in much the same way most cells are generated; having said this neuron form and function is really very similar in all organisms even if the DNA has some differences. One major difference between neurons and most other cells in the human body is that neurons don't regenerate.

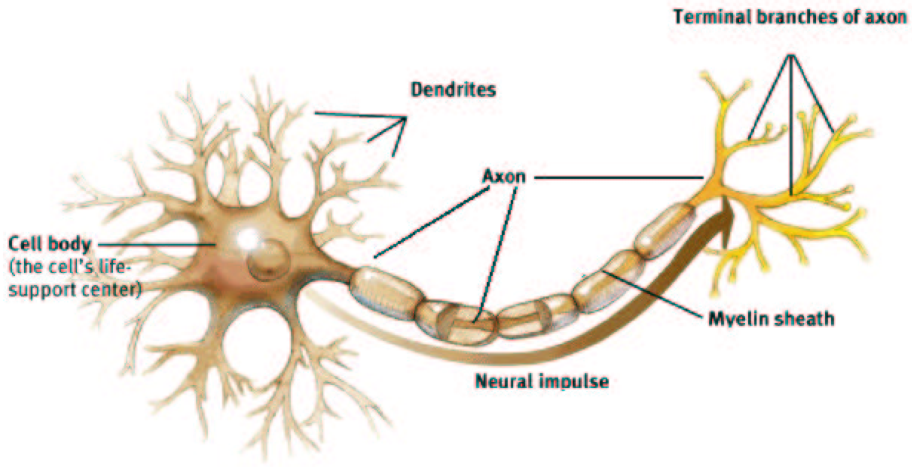

Fig. 1. A typical biological neuron [12]

2.1 Axons

Axons are the telegraph wires of the neural world. Most axons are unidirectional, conveying impulses from the soma to other neurons, muscles, etc. Axons in invertebrates can be bidirectional, impulses pass in both directions.

Impulses pass through axons using a domino effect. Axons have individual nodes which exist at intervals along the axon. When the electrical potential of a node reaches a certain threshold it will go to its action potential (or 'fire'), this is enough to trigger any nearby nodes. Signals only pass in one direction because once a node has fired it won't fire again for a certain amount of time.

The cell membrane is of crucial importance to the way the axons and synapses work. It contains a protein that pumps sodium out of the cell and pumps potassium in. The electrical potential of the cell is normally at -70 millivolts. If it gets to -50 millivolts ion channels are opened that let sodium flood into the cell. This creates a positive charge in that area of the cell (enough to trigger nearby areas) for a while until a second channel lets potassium flood out and the area will return to its normal potential. It will not fire again until the sodium / potassium ratio has been returned to normal by the pump.

Mielin shielding stops the pumping action on the axon at all points except at narrow gaps, this creates the nodes (called Nodes of Ranvier) along the process and hence the domino effect.

2.2 Dendrites

Dendrites are really quite similar to axons. In vertebrates they have the function of receiving signals from other axons or sensory organs. They are generally unidirectional, transmitting to the soma, but they can also transmit signals as well.

2.3 Synapses

Synapses are the points where neurons communicate with each other. There are two main types of synapse, chemical and electrical.

Chemical Synapses Chemical synapses commonly link axons to dendrites. They are asymmetric and there is a small gap of approximately 200- to 300-Åbetween them [1] (p15). The presynaptic axon releases chemicals into the gap which will effect the behaviour of molecules on the membrane of the postsynaptic dendrite. There are many different chemicals, but the behaviour is generally that of either exciting the postsynaptic cell or inhibiting it. For example:

Acetylcholine - when released by the presynaptic cell Acetylcholine opens ion channels in the postsynaptic cell which allow 'free passage of both potassium and sodium across the membrane' [2] (p198). This produces a small positive pulse which increases the potential of the postsynaptic dendrite.

Gamma-aminobutyric acid (GABA) - GABA opens ion channels in the postsynaptic cell which allow passage of chloride ions. This reduces the potential of the postsynaptic dendrite. Chemical synapses have a time delay of about one millisecond which is quite slow. The main power of the brain is in the massive parallelisation of the neurons.

Electrical Synapses Electrical synapses, in contrast to chemical synapses, are symmetric. The gap between the cells is much smaller than the gap in chemical synapses. Ions can pass through channels directly from one cell to the other. This makes for much faster intercellular signalling.

2.4 Soma

Positive signals pass from the dendrites (and axons) to the soma where they increase the neuron's potential. If enough signals reach the soma to excite it (reach its action potential), then it will fire the whole neuron. The rate of signals arriving at the soma is important as the potential is constantly declining. The Soma also has many cell support functions.

2.5 Neuron Behaviour

In the absence of external stimulation, neurons commonly show three different types of behaviour: 'silent, beating and bursting' behaviour [1] (p59). Silent neurons are self-explanatory. Beating and bursting neurons change rate according to their inputs.

3 Biological Neural Networks

There are 100 billion neurons [2] (p190) in the human brain; each neuron is commonly connected to thousands of other neurons. Neuroscientists have had to look at much smaller organisms, where the connection graphs are much simpler, to understand how the networks actually work.

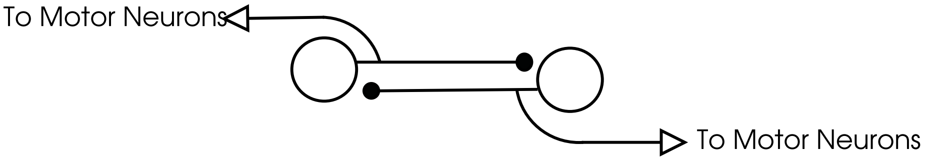

Fig. 2. Inhibitory Motor Neurons

Studies have been done on simple organisms such as the Lamprey and the Clione to show how the locomotion behaviour is produced. The two neurons in question (fig. 2) inhibit each other. They are special because they display the property of *postinhibitory rebound* [1] (p454). One neuron fires, inhibiting the other. As the inhibited cell returns to normal it will be more likely to fire action potentials. This process has inspired many artificial neural networks which model similar patterns.

4 Modelling Biological Neural Networks

The modelling of biological neural networks has been very important in the development of artificial neural networks. In 1943 Warren McCulloch and Walter Pitts [4] made an attempt to understand how the brain works. The neuroscience of the day was aware that the brain was made up of a very complex arrangement of neurons. They came up with a simple version of the neuron in attempt to understand the brain (fig. 3).

The McCulloch and Pitts (MCP) neuron has i inputs, each input (x_n) is weighted individually by multiplying by a weighting factor (w_n). The sum of the weighted inputs (h) is compared with a threshold (t), if it is greater than the threshold the neuron will have an output of 1, otherwise it will have an output of 0.

It is possible to see how this neuron models the behaviour of the biological neuron. The inputs are much like dendrites and can be connected to other neurons or inputs to the network. The output is much like an axon, it can be connected to other neurons or to an output of the network. Basically, if the sum of the inputs to the neuron is great enough then it will fire - much like its biological equivalent.

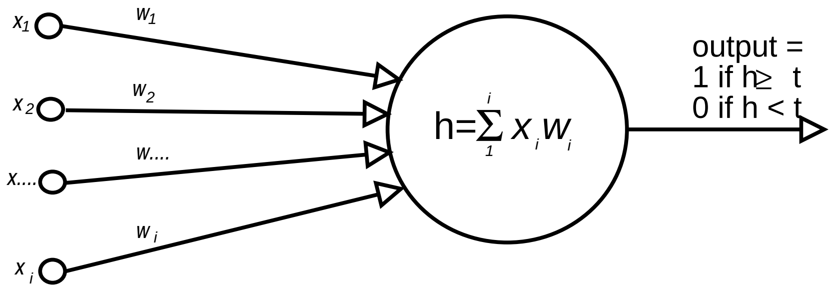

Fig. 3. The MCP (McCulloch Pitts) Neuron

The MCP neuron will fire repeatedly if its inputs are over its threshold; this is not quite as accurate as it could be. A better model would cover the concept of a neuron firing more repeatedly if its inputs are greater: more stimulation of a biological neuron would mean that it reaches its action potential in less time and therefore beats at a faster rate. Of course low stimulation won't fire a neuron at all and increases at the high stimulation end of the spectrum won't increase the rate much because a neuron takes some time to reset the sodium/potassium ratio after each pulse. The best model for this is a sigmoidal function [5].

$$out = \frac{1}{1+e^{-h}}$$

It is called a sigmoidal function because it is in an S shape (fig. 4). The gradient is close to one near the threshold (the threshold is zero in fig. 4) and is closer to 0 nearer the extremes. This model is the one that is most widely used in artificial neural networks today. The main problem researchers face is in the understanding of the behaviour of networks of the artificial neurons.

5 Artificial Neural Networks

5.1 Perceptrons

The early artificial neural networks (ANNs) were made up of MCP neurons. They were named *perceptrons*. A perceptron comprises of a number of MCP

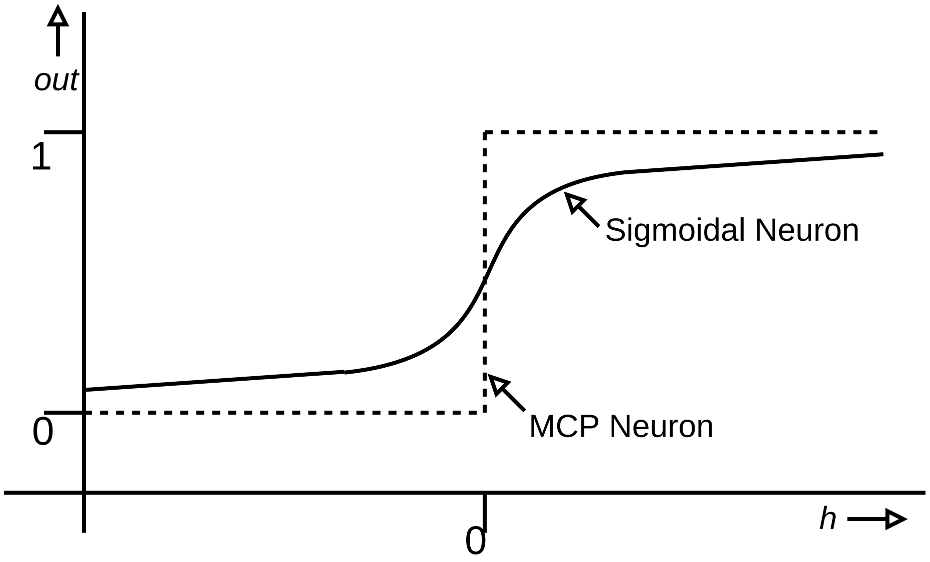

Fig. 4. The Sigmoid function compared with the MCP neuron's function

neurons all connected to a number of inputs. Each MCP neuron generates an individual output. There is no interconnection between the neurons in the simplest perceptrons.

It was proved by Rosenblatt [6] that if a perceptron can represent a function, then it can be taught to do this by the adjustment of the perceptron weights according to systematic rules. This generated great excitement in the perceptron community, however the hopes were dashed when Minsky [7] proved that perceptrons where actually very limited in what they could represent.

Take a single MCP neuron with two inputs x and y weighted with w_1 and w_2 respectively. The neuron will perform the calculation:

$$xw_1 + yw_2 = h$$

Since h is a constant, the equation draws a line in two-dimensional space. For all points one side of the line, the neuron will fire. For all points the other side, the neuron won't fire. Minsky considered getting an MCP neuron to model the XOR function. By plotting the four possible input states for x and y (0,0; 0,1; 1,0; 1,1) on a surface it is not possible to draw a threshold line so that the two valid inputs (0,1; 1,0) are one side of the line and the two invalid inputs (0,0; 1,1) are the other side of the line (fig. 5).

From this it can be shown how multi-neuron single layer perceptrons are incapable of separating points in n-dimensional space. This and other revelations dampened much of the enthusiasm about ANNs. The networks were too stable

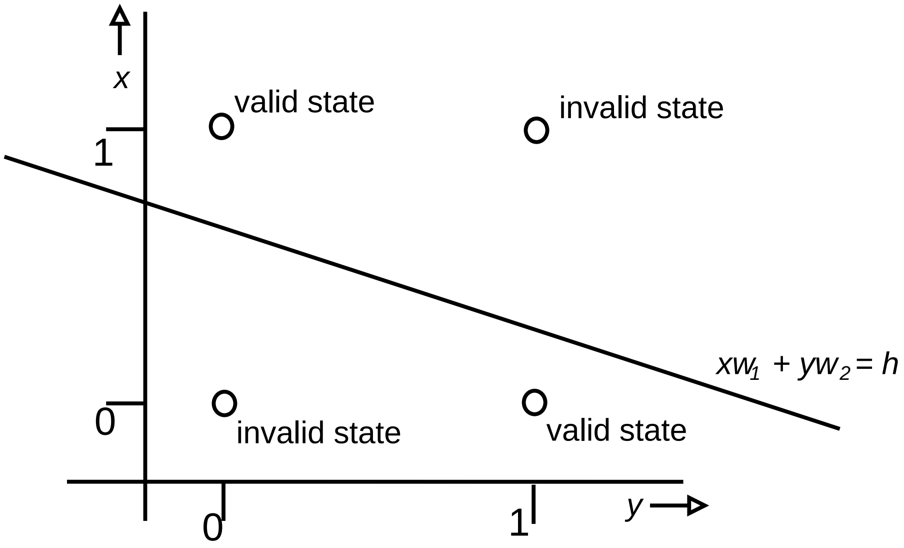

Fig. 5. Minsky's exclusive-or problem

and not capable of any kind of fuzzy logic. ANNs remained in the doldrums until 1982 when John Hopfield [8] published his article about Attractor Neural Nets.

5.2 Attractor Neural Nets

An attractor neural net (sometimes also called a Hopfield net) is a network where all the neurons are both inputs and outputs. Each neuron i is connected to each other neuron j such that the weight $w_{ij} = w_{ji}$. That is, each neuron pair has the same weight in both directions.

Taking a set of n MPC neurons in an attractor net, the net will have 2^n states (each neuron can either be at one or two). Hopfield defined a concept of e nergy e_{ij} for each pair of neurons x_i and x_j being:

$$e_{ij} = -w_{ij}x_ix_j$$

If w_{ij} is positive then clearly the two neurons will tend to make each other fire. The state of both neurons being on, being the most stable, will have the least energy.

By summing the energies of all the node pairs in the net, the energy of the whole network can be expressed. Activating a neuron (summing its inputs and testing against its threshold) will change the state and therefore the energy of the system. It was shown that activating the neuron will always reduce the energy of the system. From this it can be deduced that all attractor nets will find their way to a stable (least energy) state.

This means that by adjusting the weights using a suitable training algorithm patterns can be programmed into the network so that they are stable states. Given an initial state the network will recurse until it reaches one of its stable states. This would hopefully be one close to the initial state

While Hopfield's networks were less biologically inspired due to the fact that they are a departure from the biological model (with the neuron pairs having the same weights), they show the sort of fuzzy logic that is seen in biology. These abilities to eliminate noise, include multiple stable states and perform optimisations were revolutionary and started off a new wave of artificial neural network research. Hopfield nets have been shown to be able to solve problems such as the travelling salesman problem [10].

5.3 Recurrent Neural Nets

Hopfield nets are a subset of recurrent neural nets. There are three types of artifical (generally sigmoid) neuron in a recurrent net, input, output and hidden. The hidden neurons are all connected to each other, themselves and every input and output neuron, with every connection having an individual weight. Recurrent nets are more complicated than the other nets that were discussed earlier. That makes them harder to train. They are however closest to the biological model.

5.4 Training Neural Nets

The classic way of training perceptrons is by supervised learning. The network is given an input and the output is checked against the expected output. Using linear analysis, the neuron weights are checked to find which weight was most responsible for the error. That weight is then adjusted. It is unknown whether there are any cases of biological neural nets learning through this kind of supervision. It doesn't appear that it is a biologically inspired algorithm however.

Attractor neural nets are trained using a Hebbian algorithm. The basic prescription is that if two neurons in a pattern are seen to be the same, their weight is increased; if the two neurons are different their weight is decreased. This changing of weights means that the sought after patterns have lower energies (and we know attractor nets tend towards lower energies). An averaging of all the weights generated in this way for all of the patterns will generate the weights for the attractor net. Of course, as the number of stable patterns increases, the less effective the net becomes at distinguishing between them.

Donald Hebb noted in his book The Organization of Behaviour [9] (p62) that

> *When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.*

There is clearly an analogue between the biological behaviour noted by Hebb and the attractor net's prescribed training algorithm. Neurons (biological and artificial) that commonly fire each other reinforce each other. In this way, an attractor network can actually learn 'on the fly'. The more often an input pattern is repeated, the more stable it will become in the network.

There is a biologically inspired way of training recurrent neural nets that takes its methodology from another field of bio-inspired computing - genetic algorithms. The set of weights for a neural net can be seen as a genome. By introducing a fitness function that differentiates networks from each other according to how well they perform a task, the genomes of the more successful networks can be used as parents for future generations. The genomes are combined according to some algorithm, commonly taking half of one parent's and half of the other's genome. Mutations are introduced to make sure the genetic algorithms have more complete coverage of the solution space.

6 Conclusion

Many great strides have been made in artificial neural networks. They are currently used by many different disciplines for many different purposes. Some applications are more biologically related, others are more engineering related.

Neural networks have been used to model the brains of animals. Interesting results have been found about the behaviour and evolution of animals. An example of the sort of work done would be a project by Quinn and Noble [11] who used a small neural net and a genetic algorithm to model the behaviour of animals competing with each other for food.

The analogy drawn by Hopfield of neural network states having energy has proved useful for physicists and mathematicians in the field of thermodynamics. Engineers and computer scientists have moved away from biology and have applied artificial neural nets in real-world engineering situations. The abilities of ANNs to provide pattern matching, pattern classification, optimisation and associative learning have all proved useful.

The biological neural network is still the most advanced complex adaptive system known to man, and still holds many secrets which will inspire computer science in the future. One day, computers may even be thought of as being alive.

References

1. Levitan, I.B., Kaczmarek, L.K. (1997). The Neuron. OUP, Oxford.
2. Wasserman, P.D. (1989). Neural Computing, Van Nostrand Reinhold, New York.
3. Biologically Inspired Robotics Group (BIRG) at EPFL. [online], available at http://lslwww.epfl.ch/birg/salamander.shtml
4. McCulloch, W. S. and Pitts, W. H. (1943). A logical calculus of the ideas immanent in nervous activity. In the Bulletin of Mathematical Biophysics, 5:115-133. University of Chicago Press, Chicago.

5. Hopfield, J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. Proceedings of the National Academy of Sciences of the USA, 81:3088 - 3092. The National Academy of Sciences, Washington DC.
6. Rosenblatt, F. (1962). Principles of neurodynamics. Spartan Books, New York.
7. Misky, M. L., Paper S. (1969). Perceptrons. MIT Press, Cambridge MA
8. Hopfield, J.J. (1982). Neural Networks and Physical Systems with Emergent Collective Computational Properties. Proceedings of the National Academy of Sciences of the USA, 79: 2554 - 2558. The National Academy of Sciences, Washington DC.
9. Hebb, D. O. (1949). The organization of behavior : a neuropsychological theory. Wiley, New York.
10. Hopfield, J. and Tank, D. (1985). Neural computation of decisions in optimization problems. Biological Cybernetics, 52:141 - 152. Springer, Berlin.
11. Quinn, M., Noble, J. (2001). Modelling Animal Behaviour in Contests: Conventions for Resource Allocation. Presented at the European Conference on Artificial Life 2001. [online], available at http://www.comp.leeds.ac.uk/jasonn/COMP5320/ecalTalk.pdf
12. Spencer, S. Lecture on Behavioural Neuroscience. [online], available at http://www.arts.uwaterloo.ca/ sspencer/psych101/lecture2/sld002.htm