Unit 6 - Expert Systems - Artificial Neural Networks

Artificial Neural networks

We have discussed the way in which an Artificial Neural Network (ANN) follows the general pattern of applying the ideas of Expert Systems (ES) to real situations and have evolved the following general model.

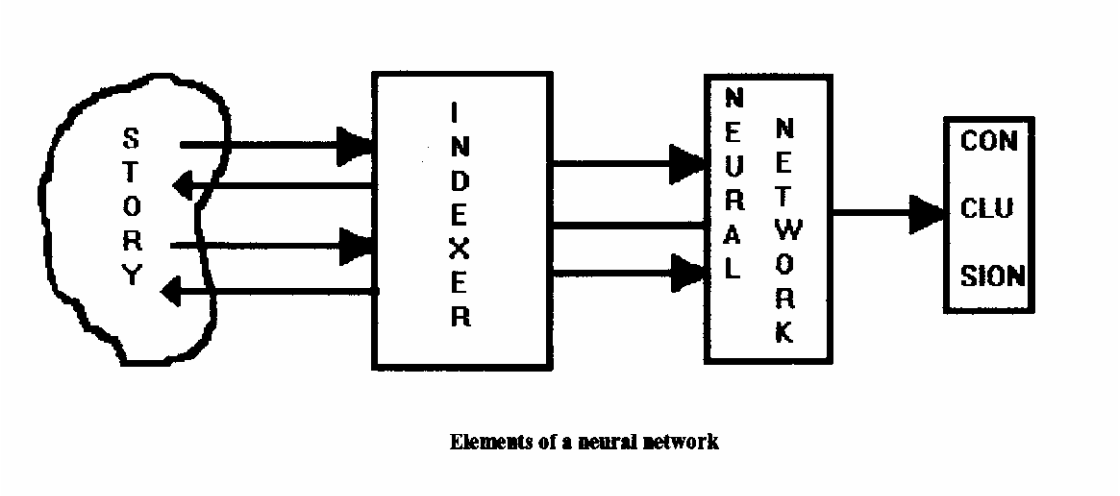

Fig 6.1 The elements of an Artificial Neural Network © J Hobson 1999

We now need to relate this to the specific ways in which an Artificial Neural Network actually behaves. Firstly we shall look at the functioning of the neurone in the human brain and then its electrical equivalent. Then we will look at the ways in which neurones can be joined together. Having done this we will consider how this is made into a device for analysing data.

The Human Side - The Neurone

Artificial Neural Networks are programmes, which try to mimic the way in which the human brain behaves. The basic building block of the brain activity is the neurone. It is estimated that there are between 50 billion and 150 billion neurones in the human brain and that each one is connected to 100 000 other neurones.
Collections of neurones are gathered together in networks, which are interconnected. These networks perform specialised activities.

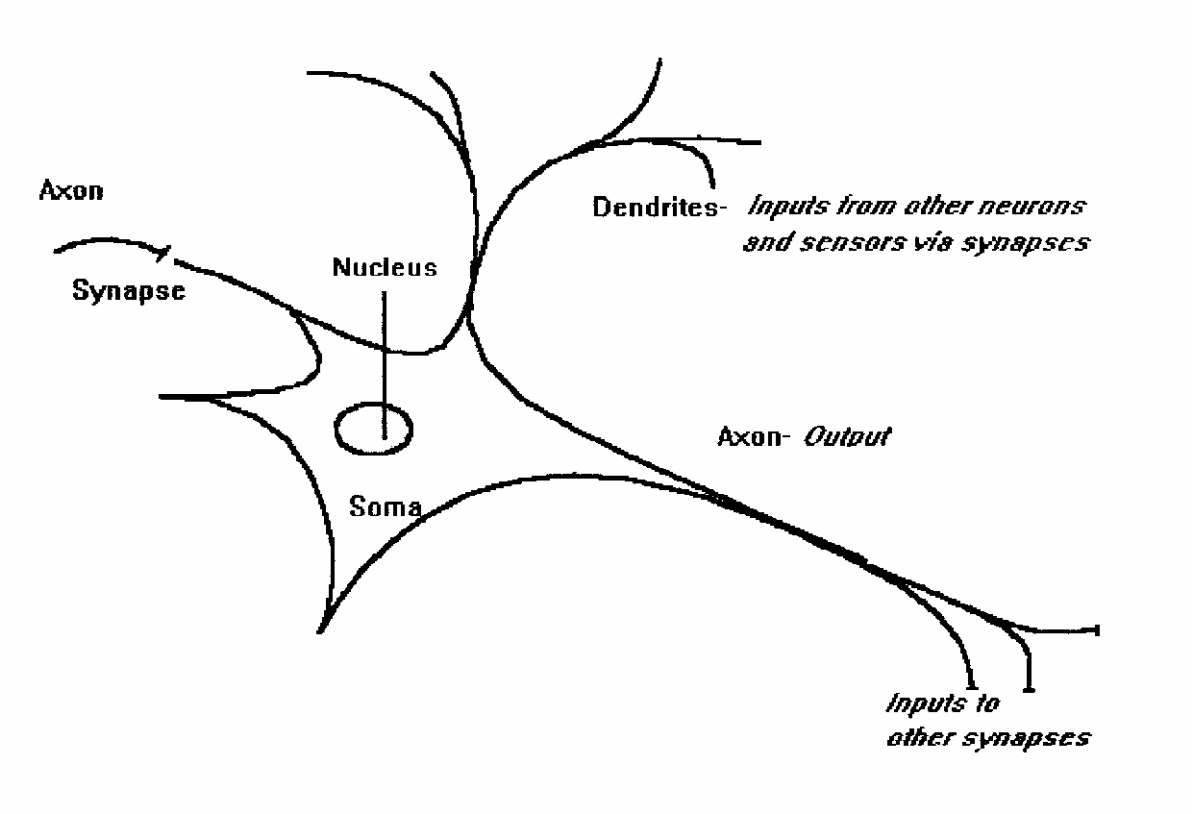

Fig 6.2 Schematic of a Neurone © J Hobson 1999

The neurone itself consists of a nucleus surrounded by a soma. The soma extends outwards into axons that terminate in dendrites. At the end of each dendrite there is a synapse that is the mechanism connecting one neurone with another.

The synapses can pass messages from one neurone to the next. These messages can be passed unchanged, enhanced or inhibited. The process of learning is one of training the synapse to do one of these three things.

A neurone will pass a message on if the level of messages coming in reaches a certain threshold level if this happens then the neurone is said 'to fire'. If the level of messages coming in does not reach the threshold then the message is not passed on and the neurone is said 'not to fire'.

The Artificial Neurone

The artificial neurone is the computing equivalent of the human neurone and is designed to perform in the same sort of way.

The basic mechanism of the artificial neurone is the same as in the human neurone and it is designed to pass messages on if the sum of the inputs is greater than a certain threshold level.

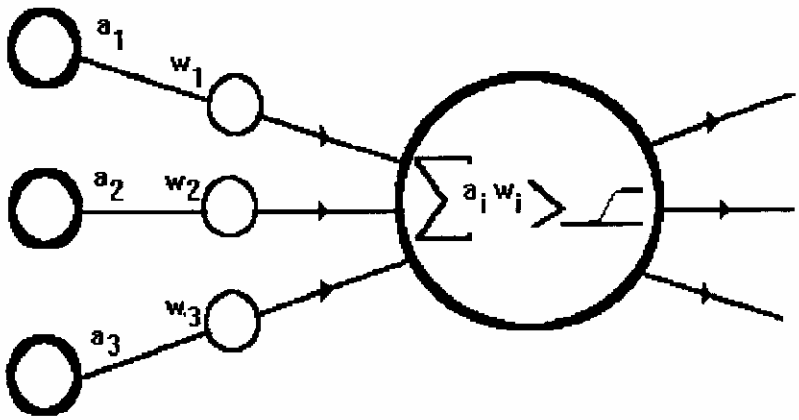

Fig 6.3 An Artificial Neurone © J Hobson 1999

Here if the sum of the inputs, $\Sigma a_i w_i$ is greater than a certain threshold the neurone fires.

In the first computer models of the neurone the mathematical form was given by a (Heavyside) step function. This was shown, by Papert and Minsky, to be incapable of solving the XOR problem. However it is possible to overcome this deficiency by replacing the Heavyside step function by a sigmoid curve. The difference between a Heavyside step function and a sigmoid curve is shown in figures 6.4 and 6.5.

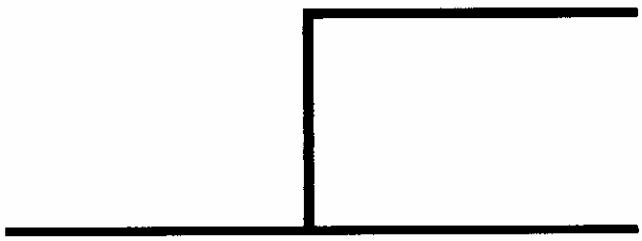

Fig 6.4 Heavyside Step Function © J Hobson 1999

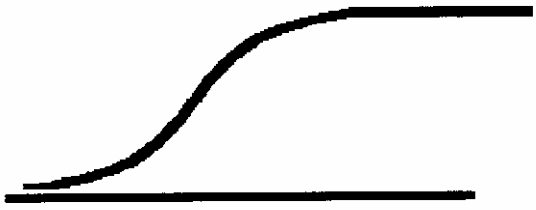

Fig 6.5 Sigmoid Curve © J Hobson 1999

Networks of Neurones

When neurones are connected together they form networks as shown in figure 6.6, where neurones are connected together to form a network. In this type of network the nodes on the left handside are the input nodes and those on the right are the output nodes. The two layers in the middle are called the hidden layers.

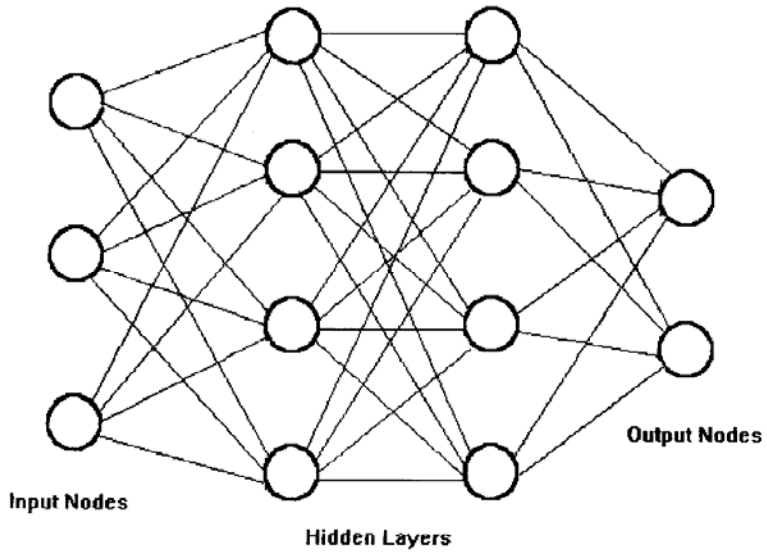

Fig 6.6 Artificial Neural Network with Two Hidden Layers © J Hobson 1999

I am lead to believe that it is theoretically unnecessary to have more than one hidden layer but this one hidden layer may need to be infinitely long (sic) (I have no reference for this).

Creating an Index

The first stage in using an artificial neural network is to create an index. The index is of the significant points from the story and will vary depending upon what you are trying to do. If you were trying to monitor credit cards for fraud then you might need the card number, the time of last use, the place of last use and the amount. Where as if you were trying to predict the outcome of football matches then you might use the names of the teams, who was playing at home, and the result as a home win, a draw or an away win.

Training a Neural Network

In order to make a neural network work it is necessary to train it. To do this a set of training cases is needed. This is a set of cases where the significant details of the story have been extracted from the story and made into an index, the result of the story is also known and these results are also indexed (see the football match). This means that both the input nodes are known and the output nodes. These are fed into the neural network

The training itself is a process that teaches the network to produce the known results from the known inputs. This process is known as back propagation. It is equivalent to finding the values of w1, w2 and w3 in figure 6.3.

Training is not an instantaneous process and may take a considerable time. Initial values are created for w1, w2 and w3, these are refined until they produce the correct results within an acceptable level of error.

Using a Neural Network

When the network has been trained it is ready for use.

Using the network after it has been trained is a matter of using the index as a basis for classifying the relevant parts of the story and then setting the input nodes accordingly. The neural network is then run and as a result the output nodes are classified with degrees of certainty. So output node 1 could be given a certainty of 80%.

Quite what the degree of certainty means depends upon the programmer and it is probably better to treat the degrees of certainty as a means of ranking the outputs from likely to unlikely.

Uses of Neural Networks

Neural networks are generally of use where an underlying rule system does not exist or is surrounded by generalities. So you could say what happens in general but the outcome is very dependent upon the particular circumstances. This might be the case when a chemical process, such as a cement furnace is being controlled.

Rule based systems are also slow to respond and so when the speed of response is critical a neural network might be substituted. An example of this could be a system to guide an aircraft through SAM missile sites where their positions are fed automatically to the system and the aircraft is automatically guided.

Practical Uses of Neural Networks

Manpower (a contractor supplier agency) is using neural networks **[a3]** to match contracts to their contractors. It is used to supplement the SQL search engine which was only providing exact matches to search criteria.

The use of neural networks is being enlarged by the SAS Institute **[a3]** by searching for fraudulent claims in cv's . This technology could also be used to search for errors in witness statements.

If neural networks are seen as pattern recognition software that is statistically based pattern processing algorithms **[a4]**. This approach is used in fraud detection and is based on the assumption that most transactions are not fraudulent, so fraudulent transactions may not share the same characteristics. Pattern recognition algorithms may be able to detect these anomalies. Examples of their usage are by telecommunications companies to detect fraudulent phone applications and to detect stolen phones by looking at deviations from established usage patterns. This type of application can reduce fraud by 30 to 50%.

References

(these are here for reference and may not be still available on the internet)

	Author	Title	Publisher	Date	Page	ISBN
[b1]	Turban and Aronson	Decision Support Systems and Intelligent systems; Neural computing: The Basics			647	

Articles

	Author	Title	Publication	Date	Page	ISSN
[a1]	Cochrane P	Axon dendrite questions?	Daily Telegraph	18 Nov 99	Connected p10	
[a2]	Anon	Recognise the potential of neural networks	Computing	17 Aug 2000		
[a3]	Doyle E	Manpower uses AI to match contracts	Computer Weekly	31 Jan 02	12	
[a4]	Harvey F	A key role in detecting fraud patterns	FT	23 Jan 02	FT Telecoms III	
[a5]	Anon	Recognise the potential ofm neuralnetworks	Computing	17 Aug 2000	39	
[a6]	Anon	Learn in the workplace	Computing	2 Nov 2000	67	
[a7]	Gregory H	Tesco online snoop plan	Computer Weekly	10 Aug 2000	10	
[a8]	Cochrane P	Axon dendrite questions	Daily Telegraph	18 Nov 1999	10	

Websites

	Web site	Date	Comment
[w1]	http://www.ncaf.co.uk/index.htm	2000	
[w2]	http://www.neuralt.com/solution/solutionscomplete.htm	2000	
[w3]	http://www.emsl.pnl.gov:2080/proj/neuron/neural/	2000	
[w4]	http://www.download.com	2000	
[w5]	http://www.newscientist.com/news/news.jsp?id=ns22612	14 Nov 02	Copy on file

