Neural Networks

Stephen Bridges

June 13, 2011

All species of the animal kingdom have brains.
These organs are all capable of doing tasks that computers find particularly complicated:

- pattern recognition
- knowledge storage
- fine control of extremities

Reactions are also fast, the brain operates at about 100Hz.
They consist of a large number of cells, called neurons, each making many connections. These connections can change over time.
The human brain: 10^{12} cells in it, 10^{15} connections between them.

What is a neuron?

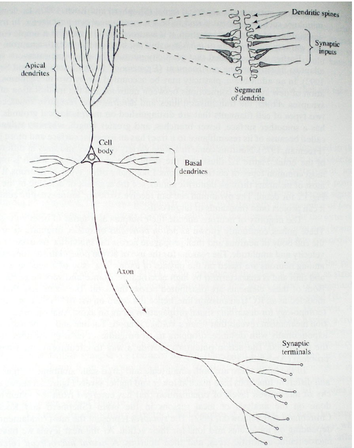

FIGURE 1.2 The pyramidal cell.

Has a body, the soma. A large number of inputs, called dendrites.
A single output, called an axon.
The output can be to a muscle or other neurons
There are no physical connections between neurons, but small gaps called synapses.

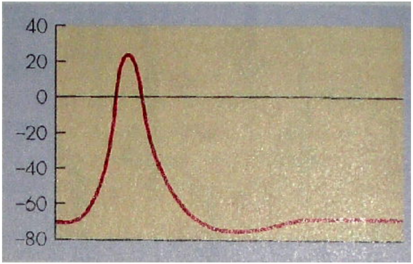

Neuron typically rests at a negative potential of -70mV. Depolarising it by a small amount causes no change.
Applying a greater positive charge causes a massive positive spike to 20mV. Spikes then propagate through the network.

Voltages depend on the flow of sodium (in) and potassium (out) ions through the cell, producing the complex waveform to the left.

Signals cross the synaptic gap via chemical carriers: neurotransmitters.

- Peptides (e.g. endorphins)
- Amino acids (e.g. epinephrine)
- Gases (e.g. nitric NO_2)
- many more

They are released by the signalling neuron, and receptors on the other open channels in the neuron membrane.
Many different types; even on the same neuron. All have subtly different net effects. Various poisons work by blocking a neurotransmitter, preventing brain working or muscles triggering.

McCulloch-Pitts model

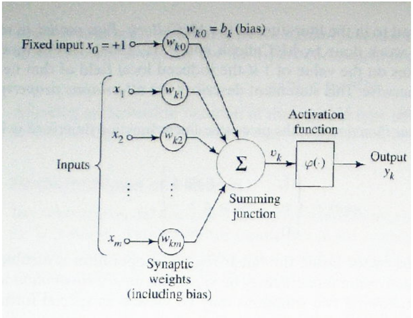

In 1943 McCulloch and Pitts devised a mathematical model of a neuron.

- Number of inputs
- Each multiplied by a weight
- Summed together
- Total is thresholded

$$v_k = \sum_{j=1}^{m} w_{kj} x_j + b_k$$

The final field is a bias weight.

Classification

Being able to draw boundaries in a 'graph' of all input signals and mark desired outputs for those boundaries.

If this can't be done in a single continuous line, the problem is non-linearly seperable (e.g. XOR vs OR).

Inputs and outputs have many dimensions. Effectively drawing a hyperplane through N-dimensional space.

Perceptrons

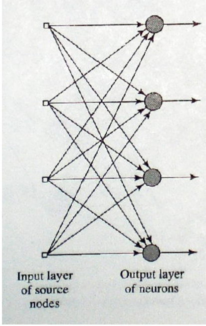

Feedforward network (no loops or surprises), single-layer, or multi-layer. Middle layers called "hidden" neurons. Used for

classification of input data - taking many inputs and producing a very small set of outputs. Hidden layers needed for more complex classifications (non-linearly separable).

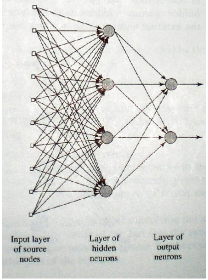

Other types

Networks with feedback

Neural nets don't just have to feedforward. If inputs are connected to outputs, corrupt patterns fed in can cause the learnt patterns to be recalled.

Self-Organising Maps

Self-organising maps are two dimension networks where adjacent neurons can also excite each other (like the visual cortex). This causes feedback within the network until a single neuron "wins".

Support Vector Machines

Much more mathmetically complex, able to in theory find ideal (for training data) solutions in classifying an error space.

Hebbian Learning

> *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 changes take place in one or both cells such that A's efficiency as one of the cells firing B, is increased.*

This is how the hippocampus appears to learn - the hippocampus is used in spatial awareness and the formation of long-term memory. There is no teacher.

Supervised learning

1. Feed input into network
2. Compare output to desired output
3. Adjust weights based on error

Simpler than it sounds, actually incredibly difficult - global minimisation problem.
Also must not overtrain the network, then it learns the training data, and not the problem you're trying to solve.

TD-Gammon

Neural network that plays backgammon.

Unlike chess, a smaller problem space, with a random element (dice). A MLP used with one hidden layer, 40-80 nodes wide. Input and ouput layers are board layouts.

Trained by playing itself. Temporal-difference learning used to measure neural network's moves in a game. The result of the game is used to measure the individual moves.

Results: played at grand-master level, and learnt new techniques now used by puny humans!

References

1. James Kalat, "Biological Psychology"
2. Simon Haykin, "Neural Networks"
3. Hugh Wilson, "Spikes, Decisions, and Actions"