


AITA : Biological Intelligence and Neural Networks

© John A. Bullinaria, 2003

1. How do Humans do Intelligent Things?
2. What are Neural Networks?
3. What are Artificial Neural Networks used for?
4. Introduction to Biological Neural Networks
5. Introduction to Artificial Neural Networks
6. Some Current Artificial Neural Networks Applications
7. Limitations of Artificial Neural Network Systems

How do Humans do Intelligent Things?

The human nervous system can be broken down into three stages that may be represented in block diagram form as:



Receptors collect information from the environment, and effectors generate interactions with the environment. The flow of information between them is represented by arrows – both forward and backward.

What we generally describe as “intelligence” is normally carried out in the central stage – the brain. The brain consists of a network of neurons, and the study of *neural networks* is now a major sub-field of AI.

What are Neural Networks ?

1. Neural Networks (NNs) are networks of neurons such as found in real (i.e. biological) brains.
2. Artificial neurons are crude approximations of the neurons found in 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 functioning brains. They may be physical devices, or simulated on conventional computers.
4. From a practical point of view, an 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 our ANNs are compared to real brains.

Why are Artificial Neural Networks worth studying?

1. They are extremely powerful computational devices (Turing equivalent, universal computers).
2. Massive parallelism makes them very efficient.
3. They can learn and generalize from training data – so there is no need for enormous feats of programming.
4. They are particularly fault tolerant – this is equivalent to the “graceful degradation” found in biological systems.
5. They are very noise tolerant – so they can cope with situations where normal symbolic systems would have difficulty.
6. In principle, they can do anything a symbolic/logic/rule based system can do, and 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 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.

Brains versus Computers : Some numbers

1. There are approximately 10 billion neurons in the human cortex, compared with 10 of thousands of processors in the most powerful parallel computers.
2. Each biological neuron is connected to several thousands of other neurons, similar to the connectivity in powerful parallel computers.
3. Lack of processing units can be compensated by speed. The typical operating speeds of biological neurons is measured in milliseconds (10^{-3} s), while a silicon chip can operate in nanoseconds (10^{-9} s).
4. The human brain is extremely energy efficient, using approximately 10^{-16} joules per operation per second, whereas the best computers today use around 10^{-6} joules per operation per second.
5. Brains have been evolving for tens of millions of years, computers have been evolving for tens of decades.

Learning in Neural Networks


There are many forms of neural networks. Most operate by passing neural ‘activations’ through a network of connected neurons. Individual neurons are very simple, but networks of them are very powerful.

They are able to form appropriate mappings between input and output representations. One of the most important 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 each input produces the correct output.

There are three broad types of learning: Supervised Learning (i.e. learning with a teacher), Reinforcement learning (i.e. learning with limited feedback), Unsupervised learning (i.e. learning with no help).

These ideas will become clearer by looking more closely at the relationship between biological and artificial neural networks.


Layout of a Biological Neural Network



Basic Components of Biological Neurons


1. The majority of *neurons* encode their activations or outputs as a series of brief electrical pulses (i.e. spikes or action potentials).
2. The neuron's *cell body (soma)* processes the incoming activations and converts them into output activations.
3. The neuron's *nucleus* contains the genetic material in the form of DNA. This is the same as in most types of cells, not just neurons.
4. *Dendrites* are fibres which emanate from the cell body and provide the receptive zones that receive activation from other neurons.
5. *Axons* are fibres acting as transmission lines that send activation to other neurons.
6. The junctions that allow signal transmission between the axons and dendrites are called *synapses*. The process of transmission is by diffusion of chemicals called *neurotransmitters* across the synaptic cleft.

Schematic Diagram of a Biological Neuron



The Basic Artificial Neuron


The basic artificial neuron is the following simplified model of a biological neuron:




1. A set of synapses (i.e. connections) brings in activations from other neurons.
2. The processing unit sums the inputs, and then applies a non-linear activation/squashing/transfer/threshold function $f(x)$.
3. An output line transmits the result to other neurons.

Artificial Neural Networks

The inputs into each neuron j are the outputs of each connected neuron i multiplied by the corresponding connection strength/weight w_{ij} . Any pattern of connectivity is allowed, but one usually takes a simplified *architecture* (i.e. layout) for the network, e.g. two or three layers of neurons with full connectivity between layers and no connections within layers. The activations of the first layer are the network inputs.



$$out_j = \text{Sigmoid}(\sum_i in_i w_{ij})$$



We usually start the network with random initial connection weights w_{ij} and use a *training algorithm* to update them iteratively so that the correct outputs are produced for each input pattern in a set of *training data*. In this way the networks *learn* how to perform appropriately.

ANNs compared with Classical Symbolic AI

The distinctions can put under three 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, the distinctions are becoming increasingly blurred. Hybrid systems are becoming increasingly common. Techniques exist for extracting symbolic rules from trained neural networks, and for embedding rules into neural systems.

Some Current Artificial Neural Network Applications

Brain modelling

Models of human development – help children with developmental problems

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

Neuropsychological models suggest – remedial actions for brain damaged patients

Real world applications

Financial modelling – predicting stocks, shares, currency exchange rates

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

Computer games – intelligent agents, backgammon, first person shooters

Control systems – autonomous adaptable robotics, microwave controllers

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

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

Noise reduction – function approximation, ECG noise reduction

Bioinformatics – protein secondary structure, DNA sequencing

Limitations of Artificial Neural Network Systems

ANNs are usually simulated in software and thus limited by the power of current computers. Even quite simple networks can take months to train.

Their long evolutionary history gives human brains a big advantage over ANNs – a lot of structure (e.g. modularity) and knowledge is innate, and does not need to be learned. Other factors (e.g. learning rates) have also been optimised over many generations. One can simulate evolution for our ANNs, but again we run into resource limitations.

Another problem is that choosing appropriate input and output representations for our ANNs can be far from straightforward. There remains considerable debate over what representations are used in real brains.

Often it makes sense to model processes at a higher level of abstraction. We know that our brains are neural networks and that they can manipulate symbols, but training an ANN is unlikely to be the best way to go about building an AI system that we have good reason to believe is most easily formulated in terms of manipulating symbols.

Overview and Reading

1. Human intelligence arises from the processes taking place in the neural networks that make up our brains. It is a natural AI strategy to study ANNs.
2. The two goals of ANN research are brain modelling and artificial system building. They both benefit from the transfer of ideas and techniques.
3. We had an introduction to the basic structures of biological and artificial neural networks, and saw how they process information and learn.
4. We ended with a look at some applications and limitations of ANNs.

Reading

1. Nilsson: Chapter 3
2. Negnevitsky: Chapter 6
3. Callan: Chapter 15
4. Rich & Knight: Chapter 18
5. Russell & Norvig: Section 20.5