

1 Artificial Neural Networks and their Biological Motivation

1.1 What is an Artificial Neural Network (ANN)?

Following the **Frequently Asked Questions** from the Usenet newsgroup `comp.ai.neural-nets` [5] we say that:

There is no universally accepted definition of an NN. But perhaps most people in the field would agree that an NN is a network of many simple processors (“units”), each possibly having a small amount of local memory. The units are connected by communication channels (“connections”) which usually carry numeric (as opposed to symbolic) data, encoded by any of various means. The units operate only on their local data and on the inputs they receive via the connections. The restriction to local operations is often relaxed during training.

Some NNs are models of biological neural networks and some are not, but historically, much of the inspiration for the field of NNs came from the desire to produce artificial systems capable of sophisticated, perhaps “intelligent”, computations similar to those that the human brain routinely performs, and thereby possibly to enhance our understanding of the human brain.

Most NNs have some sort of “training” rule whereby the weights of connections are adjusted on the basis of data. In other words, NNs “learn” from examples (as children learn to recognize dogs from examples of dogs) and exhibit some capability for generalization beyond the training data.

NNs normally have great potential for parallelism, since the computations of the components are largely independent of each other. Some people regard massive parallelism and high connectivity to be defining characteristics of NNs, but such requirements rule out various simple models, such as simple linear regression (a minimal feedforward net with only two units plus bias), which are usefully regarded as special cases of NNs.

According to Haykin, *Neural Networks: A Comprehensive Foundation* [1]:

A neural network is a massively parallel distributed processor that has a natural propensity for storing experimental knowledge and making it available for use. It resembles the brain in two respects:

1. Knowledge is acquired by the network through a learning process.
2. Interneuron connection strengths known as synaptic weights are used to store the knowledge.

We can also say that:

Neural networks are parameterised computational nonlinear algorithms for (numerical) data/signal/image processing. These algorithms are either implemented on a general-purpose computer or are built into a dedicated hardware.

1.2 Basic characteristics of biological neurons

(Adapted from Haykin [1])

- Biological neurons, the basic building blocks of the brain, are slower than silicon logic gates. The neurons operate in milliseconds which is about six orders of magnitude slower than the silicon gates operating in the nanosecond range.
- The brain makes up for the slow rate of operation with two factors:
 - a huge number of nerve cells (neurons) and interconnections between them. The number of neurons is estimated to be in the range of 10^{10} with $60 \cdot 10^{12}$ synapses (interconnections).
 - A function of a biological neuron seems to be much more complex than that of a logic gate.
- The brain is very energy efficient. It consumes only about 10^{-16} joules per operation per second, comparing with 10^{-6} J/oper·sec for a digital computer.

The brain is a highly complex, non-linear, parallel information processing system. It performs tasks like pattern recognition, perception, motor control, many times faster than the fastest digital computers.

- Consider an efficiency of the visual system which provides a representation of the environment which enables us to interact with the environment. For example, a complex task of perceptual recognition, e.g. recognition of a familiar face embedded in an unfamiliar scene can be accomplished in 100-200 ms, whereas tasks of much lesser complexity can take hours if not days on conventional computers.
- As another example consider an efficiency of the sonar system of a bat. Sonar is an active echo-location system. A bat sonar provides information about the distance from a target, its relative velocity and size, the size of various features of the target, and its azimuth and elevation.

The complex neural computations needed to extract all this information from the target echo occur within a brain which has the size of a plum.

The precision and success rate of the target location is rather impossible to match by radar or sonar engineers.

1.3 A (naïve) structure of biological neurons

A biological neuron, or a nerve cell, consists of (Figure 1.2 from Haykin [1]):




Figure 1–1: The pyramidal cell— a “prototype” of an artificial neuron.
synapses, dendrites, the cell body (or hillock), the axon.

Simplified functions of these very complex in their nature “building blocks” are as follow:

- The synapses are elementary signal processing devices.
 - A synapse is a biochemical device which converts a pre-synaptic electrical signal into a chemical signal and then back into a post-synaptic electrical signal.
 - The input pulse train has its amplitude modified by parameters stored in the synapse. The nature of this modification depends on the type of the synapse, which can be either inhibitory or excitatory.
- The postsynaptic signals are aggregated and transferred along the dendrites to the nerve cell body.
- The cell body generates the output neuronal signal, a spike, which is transferred along the axon to the synaptic terminals of other neurons.

The frequency of firing of a neuron is proportional to the total synaptic activities and is controlled by the synaptic parameters (weights).

- The pyramidal cell can receive 10^4 synaptic inputs and it can fan-out the output signal to thousands of target cells — the connectivity difficult to achieve in the artificial neural networks.

A biochemical synapse (see Figure 6.14 from G.M. Shepherd, Neurobiology, 1994)

1.4 Brain plasticity

- At the early stage of the human brain development (the first two years from birth) about 1 million synapses (hard-wired connections) are formed per second.
- Synapses are then modified through the learning process (plasticity of a neuron).
- In an adult brain plasticity may be accounted for by the above two mechanisms: creation of new synaptic connections between neurons, and modification of existing synapses.

1.5 What can you do with an NN and what not?

This text is adapted from NN_FAQ [5]:

In principle, NNs can compute any computable function, i.e., they can do everything a normal digital computer can do.

In practice, NNs are especially useful for classification and function approximation/mapping problems which are tolerant of some imprecision, which have lots of training data available, but to which hard and fast rules (such as those that might be used in an expert system) cannot easily be applied. Almost any mapping between vector spaces can be approximated to arbitrary precision by feedforward NNs (which are the type most often used in practical applications) if you have enough data and enough computing resources.

To be somewhat more precise, feedforward networks with a single hidden layer, under certain practically-satisfiable assumptions are statistically consistent estimators of, among others, arbitrary measurable, square-integrable regression functions, and binary classification devices.

NNs are, at least today, difficult to apply successfully to problems that concern manipulation of symbols and memory. And there are no methods for training NNs that can magically create information that is not contained in the training data.

1.6 Who is concerned with NNs? From NN_FAQ

- Computer scientists want to find out about the properties of non-symbolic information processing with neural nets and about learning systems in general.
- Statisticians use neural nets as flexible, nonlinear regression and classification models.
- Engineers of many kinds exploit the capabilities of neural networks in many areas, such as signal processing and automatic control.
- Cognitive scientists view neural networks as a possible apparatus to describe models of thinking and consciousness (High-level brain function).
- Neuro-physiologists use neural networks to describe and explore medium-level brain function (e.g. memory, sensory system, motorics).
- Physicists use neural networks to model phenomena in statistical mechanics and for a lot of other tasks.
- Biologists use Neural Networks to interpret nucleotide sequences.
- Philosophers and some other people may also be interested in Neural Networks for various reasons.

1.7 Taxonomy of neural networks

From the point of view of their **active** or **decoding** phase, artificial neural networks can be classified into **feedforward** (static) and **feedback** (dynamic, recurrent) systems.

From the point of view of their **learning** or **encoding** phase, artificial neural networks can be classified into **supervised** and **unsupervised** systems.

Feedforward supervised networks

This networks are typically used for function approximation tasks. Specific examples include:

- Linear recursive least-mean-square (LMS) networks
- Backpropagation networks
- Radial Basis networks

Feedforward unsupervised networks

This networks are used to extract important properties of the input data and to map input data into a “representation” domain. Two basic groups of methods belong to this category

- Hebbian networks performing the **Principal Component Analysis** of the input data, also known as the Karhunen-Loeve Transform.
- Competitive networks used to performed **Learning Vector Quantization**, or tessellation of the input data set.
Self-Organizing Kohonen Feature Maps also belong to this group.

Feedback networks

These networks are used to learn or process the temporal features of the input data and their internal state evolves with time. Specific examples include:

- Recurrent Backpropagation networks
- Associative Memories
- Adaptive Resonance networks

References

- [1] Simon Haykin. *Neural Networks – a Comprehensive Foundation*. Prentice Hall, New Jersey, 2nd edition, 1999. ISBN 0-13-273350-1.
- [2] H. Demuth and M. Beale. *Neural Network Toolbox. For use with MATLAB. User's Guide*. The MathWorks Inc., 2002. \$MATLAB/help/pdf_doc/nnet.pdf.
- [3] Martin T. Hagan, H Demuth, and M. Beale. *Neural Network Design*. PWS Publishing, 1996.
- [4] Hertz, Krogh, and Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, 1991. ISBN 0-201-51560-1.
- [5] W.S. Sarle, editor. *Neural Network FAQ*. Newsgroup: comp.ai.neural-nets, 2002. URL: <ftp://ftp.sas.com/pub/neural/FAQ.html>.
- [6] Mohamad H. Hassoun. *Fundamentals of Artificial Neural Networks*. The MIT Press, 1995. ISBN 0-262-08239-X.