2 Introduction to Neural Networks

(Artificial) neural networks are information processing systems, whose structure and operation principles are inspired by the nervous system and the brain of animals and humans. They consist of a large number of fairly simple units, the so-called *neurons*, which are working in parallel. These neurons communicate by sending information in the form of activation signals, along directed connections, to each other.

A commonly used synonym for "neural network" is the term "connectionist model." The research area that is devoted to the study of connectionist models is called "connectionism." Furthermore, the expression "parallel distributed processing" can often be found in relation to (artificial) neural networks.

2.1 Motivation

(Artificial) neural networks are studied for various reasons: in (neuro-)biology and (neuro-)physiology, but also in psychology, one is mainly interested in their similarity to biological nervous systems. In these areas (artificial) neural networks are used as computational models with which one tries to simulate and thus to understand the mechanisms of nerve and brain functions. Especially in computer science, but also in other engineering sciences, one tries to mimic certain cognitive powers of humans (especially learning ability) using functional elements of the nervous system and the brain. In physics, certain mathematical models that are analogous to (artificial) neural networks are employed to describe specific physical phenomena. An example are models of magnetism, for instance, the Ising model.

As can already be seen from this brief list, the study of (artificial) neural networks is a highly interdisciplinary research area. However, in this book we widely neglect the use of (artificial) neural networks in physics (even though we draw on examples from physics to explain certain network models) and consider their biological basis only very briefly (see the next section). Rather we focus on the mathematical and

R. Kruse et al., *Computational Intelligence*, Texts in Computer Science,
DOI 10.1007/978-1-4471-7296-3_2

engineering aspects, particularly the use of (artificial) neural networks in the area of computer science that is commonly called "artificial intelligence."

While the reasons why biologists study (artificial) neural networks are fairly obvious, we may have to justify why neural networks are (or should be) studied in artificial intelligence. The reason is that the paradigm of classical artificial intelligence (sometimes called, in a somewhat pejorative manner, GOFAI — "good old-fashioned artificial intelligence") is based on a very strong hypothesis about how machines can be made to behave "intelligently." This hypothesis says that the essential requirement for intelligent behavior is the ability to manipulate symbols and symbol structures that are represented by physical structures. Here *symbol* means a token that refers to an object or a situation. This relation is interpreted in an operational manner: the system can perceive and/or manipulate the object referred to. This hypothesis was first formulated explicitly by Newell and Simon (1976):

> **Physical Symbol System Hypothesis**: A physical-symbol system has the necessary and sufficient means for general intelligent action.

As a matter of fact, classical artificial intelligence concentrated, based on this hypothesis, on symbolic forms of representing knowledge and in particular on propositional and predicate logic. (Artificial) neural networks, on the other hand, are no physical symbol systems, since they do not process *symbols*, but rather much more elementary *signals*, which, taken individually, rarely have a (clear) meaning. As a consequence, (artificial) neural networks are often called "sub-symbolic." However, if the ability to process symbols is necessary to produce intelligent behavior, then it is unnecessary to study (artificial) neural networks in artificial intelligence.

There is no doubt that classical artificial intelligence has achieved remarkable successes: nowadays computers can automatically solve many types of puzzles and brain-twisters and can play games like chess and Reversi (also known as Othello) on an extremely high level. However, when it comes to mimicking perception (seeing, hearing, etc.), computers usually perform fairly poorly compared to humans—at least if symbolic representations are relied upon: here computers are often too slow, too inflexible, and too little tolerant to noise and faults. We may conjecture that the problem is that in order to recognize patterns—a core task of perception—symbolic representations are not very well suited, because there are no adequate symbols on this level of processing. Rather "raw" (measurement) data needs to be structured and summarized before symbolic methods can effectively be applied. Hence it appears to be reasonable to examine the mechanisms of sub-symbolic information processing in natural intelligent systems—that is, animals and humans—in more detail and possibly to exploit these mechanisms to mimic intelligent behavior.

Additional arguments why studying (artificial) neural networks may be beneficial arise from the following observations:

- Expert systems that use symbolic representations usually become slower with a larger knowledge base, because larger sets of rules need to be traversed. Human experts, however, usually become faster. Maybe a non-symbolic representation (as it is used in natural neural networks) is more efficient.

- Despite the fairly long switching time of natural neurons (in the order of several milliseconds) essential cognitive tasks (like recognizing an object) are solved in a fraction of a second. If neural processing were sequential, only about 100 switching operations could be performed ("100-step rule"). Hence high parallelization must be present, which is easy to achieve with neural networks, but much more difficult to implement with other approaches.
- There is a large number of successful applications of (artificial) neural networks in industry, commerce, and finance.

2.2 Biological Background

As already mentioned, (artificial) neural networks are inspired by the structure and the operation principles of the nervous system and particularly the brain of animals and humans. In fact, the neural network models that we study in this book are not very close to their biological original, since they are too simplified to model the characteristics of natural neural networks correctly. Nevertheless we briefly consider natural neural networks here, because they formed the starting point for investigating artificial neural networks. The description follows Anderson (1995).

The nervous system of animals consists of the brain (in so-called "lower" life forms often only referred to as the "central nervous system"), the different sensory systems, which collect information from the different body parts (visual, auditory, olfactory, gustatory, thermal, tactile, etc., information), and the motor system, which controls movements. The greater part of information processing happens in the brain/central nervous system, although the amount of preprocessing outside the brain can be considerable, for example, in the retina of the eye.

W.r.t. processing information, the neurons are the most important components of the nervous system.[1] According to common estimates, there are about 100 billion (10^{11}) neurons in the human brain, of which a fairly large part is active in parallel. Neurons process information mainly by interacting with each other.

A **neuron** is a cell that collects and transmits electrical activity. Neurons exist in many different shapes and sizes. Nevertheless, one can derive a "prototypical" neuron that resembles all kinds of neurons to some degree (although this is a fairly severe simplification). This prototype is shown schematically in Fig. 2.1. The **cell body** of the neuron, which contains the **nucleus**, is also called **soma**. It has a diameter of about 5–100 μm (micrometer, 1 μm = 10^{-6} m). From the cell body extend several short, heavily ramified branches that are called **dendrites**. In addition, it has a long extension called **axon**. The axon can be between a few millimeters and one

[1]The nervous system consists not only of neurons, not even for the largest part. Besides neurons there are various other cells, for instance, the so-called glia cells, which have a supporting function.

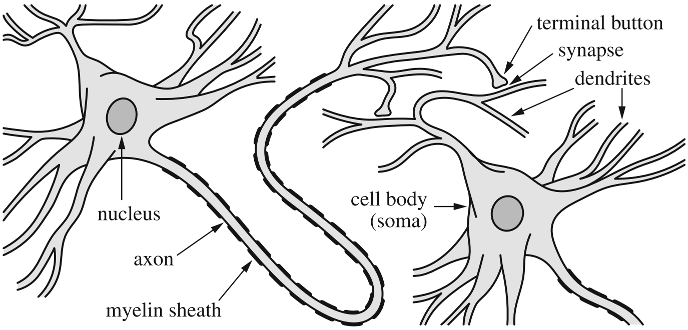

Fig. 2.1 Prototypical structure of biological neurons

meter long. Axon and dendrites differ in the structure and the properties of the **cell membrane**. In particular, the axon is often covered by a **myelin sheath**.

The axons are the fixed paths along which neurons communicate with each other. The axon of a neuron leads to the dendrites of other neurons. At its end the axon is heavily ramified and possesses at the ends of these branches **terminal buttons**. Each terminal button almost touches a dendrite or the cell body of another neuron. The gap between a terminal button and a dendrite is usually between 10 and 50 nm (nanometer; 1 nm = 10^{-9} m) wide. Such a place, at which an axon and a dendrite almost touch each other, is called **synapse**.

The most common form of communication between neurons is that a terminal button of the axon releases certain chemicals, the so-called **neurotransmitters**, which act on the membrane of the receiving dendrite and change its polarization (its electrical potential). Usually the inside of the cell membrane, which encloses the whole neuron, is about 70 mV (millivolts; 1 mV = 10^{-3} V) more negative than its outside, because the concentration of negative ions is greater on the inside, while the concentration of positive ions is greater on the outside. Depending on the type of the released neurotransmitter, the potential difference may be reduced or increased on the side of the dendrite. Synapses that reduce the potential difference are called **excitatory**, those that increase it are called **inhibitory**.

In an adult human all connections between neurons are completely established and no new connections are created (again this is a severe simplification). An average neuron possesses between 1000 and 10,000 connections to other neurons. The change of the electrical potential that is caused by a single synapse is fairly small, but the individual excitatory and inhibitory effects can accumulate (counting the excitatory influences as positive and the inhibitory ones as negative). If the excitatory net input is large enough, the potential difference in the cell body can be significantly reduced. If the reduction is large enough, the axon's base is depolarized. This depolarization is caused by positive sodium ions entering the cell. As a consequence, the inside of the cell becomes temporarily (for about one millisecond) more positive than its outside. Afterwards the potential difference is rebuilt by positive potassium ions leaving the cell. Finally, the original distribution of sodium and potassium ions is reestablished by special ion pumps in the cell membrane.

The sudden, temporary change of the electrical potential, which is called **action potential**, propagates along the axon. The propagation speed lies between 0.5 and 130 m/s, depending on the properties of the axon. In particular, it depends on how heavily the axon is covered with a myelin sheath (the more myelin, the faster the action potential is propagated). When this nerve impulse reaches the end of the axon, it causes neurotransmitters to be released at the terminal buttons, thus passing the signal on to the next cell, where the process is repeated.

In summary: changes of the electrical potential are accumulated at the cell body of a neuron and, if they reach a certain threshold, are propagated along the axon. This nerve impulse causes that neurotransmitters are released by the terminal buttons at the end of the axon, thus inducing a change of the electrical potential in the receiving neuron. Even though this description is heavily simplified, it captures the essentials of neural information processing on the level of individual neurons.

In the human nervous system information is encoded by continuously changing quantities, primarily two: the electrical potential of the neuron's membrane and the number of nerve impulses that a neuron transmits per second. The latter is also called the **firing rate** of the neuron. It is commonly assumed that the number of impulses is more important than their shape (in the sense of a change of the electrical potential), although competing theories of **neural coding** exist. A neuron can emit 100 or even more impulses per second. The higher the firing rate, the higher the influence a neuron has on connected neurons. However, in artificial neural networks this frequency coding of information is usually not emulated.

References

J.R. Anderson, *Cognitive Psychology and its Implications*, 4th edn. (Freeman, New York, NY, USA, 1995)

M.A. Boden (ed.), *The Philosophy of Artificial Intelligence* (Oxford University Press, Oxford, United Kingdom, 1990)

A. Newell, H.A. Simon. Computer Science as Empirical Enquiry: Symbols and Search. *Communications of the Association for Computing Machinery* 19. Association for Computing Machinery, New York, NY, USA 1976. Reprinted in [Boden 1990], 105–132

http://www.springer.com/978-1-4471-7294-9

Computational Intelligence
A Methodological Introduction
Kruse, R.; Borgelt, C.; Braune, C.; Mostaghim, S.; Steinbrecher, M.
2016, XIII, 564 p. 255 illus., Hardcover
ISBN: 978-1-4471-7294-9