An introduction to Neural Networks

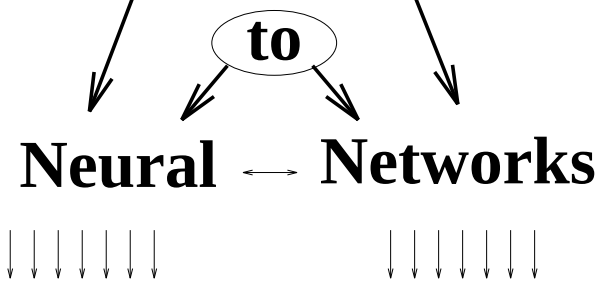

Ben J. A. Kröse

P. Patrick van der Smagt

Fifth edition

January 1993

©1993 The University of Amsterdam. Permission is granted to distribute *single* copies of this book for non-commercial use, as long as it is distributed as a whole in its original form, and the names of the authors and the University of Amsterdam are mentioned. Permission is also granted to use this book for non-commercial courses, provided the authors are notified of this beforehand.

The authors can be reached at

University of Amsterdam
Faculty of Mathematics & Computer Science
Kruislaan 403, NL–1098 SJ Amsterdam
THE NETHERLANDS
Phone: +31 20 525 7463
Fax: +31 20 525 7490
email: smagt@fwi.uva.nl, krose@fwi.uva.nl

Contents

Preface **9**

I FUNDAMENTALS **11**

1 Introduction **13**

2 Fundamentals **15**
2.1 A framework for distributed representation 15
2.1.1 Processing units 16
2.1.2 Connections between units 16
2.1.3 Activation and output rules 17
2.2 Network topologies 17
2.3 Training of artificial neural networks 18
2.3.1 Paradigms of learning 18
2.3.2 Modifying patterns of connectivity 18
2.4 Notation and terminology 19
2.4.1 Notation 19
2.4.2 Terminology 20

II THEORY **23**

3 Adaline and Perceptron **25**
3.1 The adaptive linear element (Adaline) 25
3.2 The Perceptron 26
3.3 Exclusive-or problem 27
3.4 Multi-layer perceptrons can do everything 28
3.5 Perceptron learning rule and convergence theorem 30
3.6 The delta rule 31

4 Back-Propagation **33**
4.1 Multi-layer feed-forward networks 33
4.2 The generalised delta rule 34
4.3 Working with back-propagation 36
4.4 Other activation functions 37
4.5 Deficiencies of back-propagation 38
4.6 Advanced algorithms 39

4.7 Applications 42

5 Self-Organising Networks **45**
5.1 Competitive learning 46
5.1.1 Clustering 46
5.1.2 Vector quantisation 49
5.1.3 Using vector quantisation 49
5.2 Kohonen network 52
5.3 Principal component networks 55
5.3.1 Introduction 55
5.3.2 Normalised Hebbian rule 56
5.3.3 Principal component extractor 56
5.3.4 More eigenvectors 57
5.4 Adaptive resonance theory 58
5.4.1 Background: Adaptive resonance theory 58
5.4.2 ART1: The simplified neural network model 58
5.4.3 ART1: The original model 61
5.5 Reinforcement learning 63
5.5.1 Associative search 63
5.5.2 Adaptive critic 64
5.5.3 Example: The cartpole system 65

6 Recurrent Networks **69**
6.1 The Hopfield network 70
6.1.1 Description 70
6.1.2 Hopfield network as associative memory 71
6.1.3 Neurons with graded response 72
6.2 Boltzmann machines 73

III APPLICATIONS **77**

7 Robot Control **79**
7.1 End-effector positioning 80
7.1.1 Camera–robot coordination is function approximation 81
7.2 Robot arm dynamics 86
7.3 Mobile robots 88
7.3.1 Model based navigation 88
7.3.2 Sensor based control 90

8 Vision **93**
8.1 Introduction 93
8.2 Feed-forward types of networks 94
8.3 Self-organising networks for image compression 94
8.3.1 Back-propagation 95
8.3.2 Linear networks 95
8.3.3 Principal components as features 96

8.4 The cognitron and neocognitron . . . 97
8.4.1 Description of the cells . . . 97
8.4.2 Structure of the cognitron . . . 98
8.4.3 Simulation results . . . 99
8.5 Relaxation types of networks . . . 99
8.5.1 Depth from stereo . . . 99
8.5.2 Image restoration and image segmentation . . . 101
8.5.3 Silicon retina . . . 101

IV IMPLEMENTATIONS 105

9 General Purpose Hardware 109
9.1 The Connection Machine . . . 110
9.1.1 Architecture . . . 110
9.1.2 Applicability to neural networks . . . 111
9.2 Systolic arrays . . . 112

10 Dedicated Neuro-Hardware 115
10.1 General issues . . . 115
10.1.1 Connectivity constraints . . . 115
10.1.2 Analogue vs. digital . . . 116
10.1.3 Optics . . . 116
10.1.4 Learning vs. non-learning . . . 117
10.2 Implementation examples . . . 118
10.2.1 Carver Mead's silicon retina . . . 118
10.2.2 LEP's LNeuro chip . . . 120

Author Index 123

Subject Index 126

List of Figures

2.1 The basic components of an artificial neural network. 16
2.2 Various activation functions for a unit. 17

3.1 The Adaline. 26
3.2 The Perceptron. 27
3.3 Single layer perceptron with two inputs. 28
3.4 Geometric representation of input space. 28
3.5 Solution of the XOR problem. 29

4.1 A multi-layer network with l layers of units. 34
4.2 The descent in weight space. 37
4.3 The periodic function $f(x) = \sin(2x)\sin(x)$ approximated with sine activation functions. 38
4.4 The periodic function $f(x) = \sin(2x)\sin(x)$ approximated with sigmoid activation functions. 39
4.5 Slow decrease with conjugate gradient in non-quadratic systems. 41

5.1 A simple competitive learning network. 46
5.2 Example of clustering in 3D with normalised vectors. 47
5.3 Determining the winner in a competitive learning network. 48
5.4 Vector quantisation tracks input density. 50
5.5 A network combining a vector quantisation layer with a 1-layer feed-forward neural network. This network can be used to approximate functions from $\mathbb{R}^4$ to $\mathbb{R}^3$, the input space $\mathbb{R}^4$ is discretised in two disjoint subspaces. 51
5.6 Gaussian neuron distance function. 53
5.7 A topology-conserving map converging. 54
5.8 The mapping of a two-dimensional input space on a one-dimensional Kohonen network. 54
5.9 Mexican hat 54
5.10 Distribution of input samples. 55
5.11 The ART architecture. 59
5.12 The ART1 neural network. 59
5.13 An example ART run. 61
5.14 Architecture of a reinforcement learning scheme with critic element 64
5.15 The cartpole system. 66

6.1 The auto-associator network. 69

7.1 An exemplar robot manipulator. 79
7.2 Indirect learning system for robotics. 82
7.3 The system used for specialised learning. 83
7.4 A Kohonen network merging the output of two cameras. 85
7.5 The neural model proposed by Kawato *et al.* 86
7.6 The neural network used by Kawato *et al.* 87
7.7 The desired joint pattern for joints 1. Joints 2 and 3 have similar time patterns. 88
7.8 Schematic representation of the stored rooms, and the partial information which is available from a single sonar scan. 89
7.9 The structure of the network for the autonomous land vehicle. 90

8.1 Input image for the network. 96
8.2 Weights of the PCA network. 96
8.3 The basic structure of the cognitron. 98
8.4 Cognitron receptive regions. 99
8.5 Two learning iterations in the cognitron. 100
8.6 Feeding back activation values in the cognitron. 103

9.1 The Connection Machine system organisation. 111
9.2 Typical use of a systolic array. 112
9.3 The Warp system architecture. 113

10.1 Connections between M input and N output neurons. 115
10.2 Optical implementation of matrix multiplication. 117
10.3 The photo-receptor used by Mead. 119
10.4 The resistive layer (a) and, enlarged, a single node (b). 120
10.5 The LNeuro chip. 121

Preface

This manuscript attempts to provide the reader with an insight in artificial neural networks. The absence of any state-of-the-art textbook forced us into writing our own; however, in the meantime a number of worthwhile textbooks have been published which can be used for background and in-depth information. At times, this manuscript may prove to be too thorough or not thorough enough for a complete understanding of the material. Further reading material can be found in some excellent text books which have appeared in the meantime[1–7].

Some of the material in this book, especially parts III and IV, contains timely material and thus may heavily change throughout the ages. The choice of describing robotics and vision as neural network applications coincides with the neural network research programs at the department of Computer Systems of the University of Amsterdam.

Much of the material presented in chapter 5 has been written by Joris van Dam and Anuj Dev at the University of Amsterdam. Also, Anuj contributed to material in chapter 8. Furthermore, we express our gratitude to those people out there in Net-Land who gave us feedback on this manuscript, especially Michiel van der Korst and Nicolas Maudit who pointed out quite a few of our goof-ups. We owe them many *kwartjes* for their help.

Amsterdam, January 1993
Patrick van der Smagt
Ben Kröse

References

[1] J. Hertz, A. Krogh & R. G. Palmer, *Introduction to the Theory of Neural Computation*, Addison Wesley, 1991.

[2] H. Ritter, T. Martinetz & K. Schulten, *Neural Computation and Self-Organizing Maps*, Addison-Wesley, 1992, *Neuronale Netze*, Addison-Wesley, 1991.

[3] P. D. Wasserman, *Neural Computing: Theory and Practice*, Van Nostrand Reinhold, New York, 1989.

[4] J. A. Anderson & E. Rosenfeld, eds., *Neurocomputing: Foundations of Research*, The MIT Press, Cambridge, MA, 1988.

[5] AFCEA International Press, *DARPA Neural Network Study*, 1988.

[6] D. E. Rumelhart & J. L. McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (volume 1), The MIT Press, 1986.

[7] J. L. McClelland & D. E. Rumelhart, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (volume 2), The MIT Press, 1986.

Part I

FUNDAMENTALS

1 Introduction

The 1980's have shown a renewed interest for neural networks (also known as 'connectionist models' or 'parallel distributed processing').

A first wave of interest emerged after the introduction of simplified neurons by McCulloch and Pitts in 1943 [1]. These neurons were presented as models of biological neurons and as conceptual components for circuits that could perform computational tasks.

When Minsky and Papert published their book *Perceptrons* in 1969 [2] in which they showed the deficiencies of perceptron models, most neural network funding was redirected and researchers left the field. Only a few researchers continued their efforts, most notably Teuvo Kohonen, Stephen Grossberg, James Anderson, and Kunihiko Fukushima.

The interest re-emerged only after some important theoretical results were attained in the early eighties (most notably the discovery of error back-propagation), and new hardware developments increased the processing capacities. This renewed interest is reflected in the number of scientists, the amounts of funding, the number of large conferences, and the number of journals associated with neural networks. Nowadays most universities have a neural networks group, within their psychology, physics, computer science, or biology departments.

Artificial neural networks can be characterised most adequately as 'computational models' with particular properties such as the ability to adapt or learn, to generalise, or to cluster or organise data, and which operation is based on parallel processing. However, many of the above-mentioned properties can be attributed to existing (non-neural) models; the intriguing question is to which extent the neural approach proves to be better suited for certain applications than existing models. To date an equivocal answer to this question is not found.

Often parallels with biological systems are described. However, there is still so little known (even at the lowest cell level) about biological systems, that the models we are using for our artificial neural systems seem to introduce an oversimplification of the 'biological' models.

In this course we give an introduction to artificial neural networks. The point of view we take is that of a computer scientist. We are not concerned with the psychological implication of the networks, and we will at most occasionally refer to biological neural models. We consider neural networks as a 'novel' computational scheme rather than anything else.

These lecture notes start with a chapter in which a number of fundamental properties are discussed. In chapter 3 a number of 'classical' approaches are described, as well as the discussion on their limitations which took place in the early sixties. Chapter 4 continues with the description of attempts to overcome these limitations and introduces the back-propagation learning algorithm. Self-organising networks are discussed in chapter 5. Chapter 6 discusses relaxation or recurrent networks. Chapters 7 and 8 focus on applications of neural networks in the fields of robotics and image processing respectively. Part IV discusses the hardware

aspects.

References

[1] W. S. McCulloch & W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bulletin of Mathematical Biophysics* 5 (1943), 115–133.

[2] M. Minsky & S. Papert, *Perceptrons: An Introduction to Computational Geometry*, The MIT Press, 1969.

2 Fundamentals

The artificial neural networks which we describe in this course are all variations on the parallel distributed processing (PDP) idea. The architecture of each network is based on very similar building blocks which perform the processing. In this chapter we first discuss these processing units and discuss different network topologies. Learning strategies—as a basis for an adaptive system—will be presented in the last section.

2.1 A framework for distributed representation

An artificial network consists of a pool of simple processing units which communicate by sending signals to each other over a large number of weighted connections.

A set of major aspects of a parallel distributed model can be distinguished (cf. Rumelhart and McClelland, 1986 [1,2]):

- a set of processing units ('neurons,' 'cells');
- a state of activation a_i for every unit, which also determines the output of the unit;
- connections between the units. Generally each connection is defined by a weight w_{ij} which determines the effect which the signal of unit j has on unit i;
- a propagation rule, which determines the effective input i_i of a unit from its external inputs;
- an activation function $\mathcal{F}_i$, which determines the new level of activation based on the effective input $i_i(t)$ and the current state $a_i(t)$ (i.e., the update);
- an external input or offset θ_i for each unit;
- a method for information gathering;
- an environment within which the system must operate, providing input signals and—if necessary—error signals.

Figure 2.1 illustrates these basics, some of which will be discussed in the next sections.

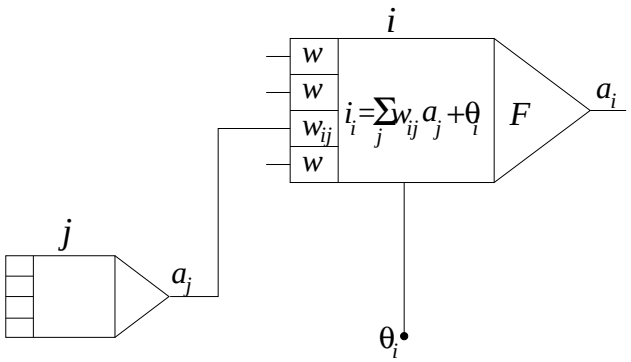

Figure 2.1: The basic components of an artificial neural network. The propagation rule used here is the 'standard' weighted summation. Note the definition of w_{ij}.

2.1.1 Processing units

Each unit performs a relatively simple job: receive input from neighbours or external sources and use this to compute an output signal which is propagated to other units. Apart from this processing, a second task is the adjustment of the weights. The system is inherently parallel in the sense that many units can carry out their computations at the same time.

Within neural systems it is useful to distinguish three types of units: *input* units which may receive data from outside the system, *output* units which send data out of the system and *hidden* units whose input and output signals remain within the system itself.

During operation, units can be updated either *synchronously* or *asynchronously*. With synchronous updating, all units update their activation simultaneously; with asynchronous updating, each unit has a (usually fixed) probability of updating its activation at a time t, and usually only one unit will be able to do this at a time. In some cases the latter model has some advantages.

2.1.2 Connections between units

In most cases we assume that each unit provides an additive contribution to the input of the unit with which it is connected. The total input to unit i is simply the weighted sum of the separate outputs from each of the connected units plus a *bias* or *offset* term θ_i:

$$i_i(t) = \sum_j w_{ij}(t)\, a_j(t) + \theta_i(t). \tag{2.1}$$

The contribution for positive w_{ij} is considered as an *excitation* and for negative w_{ij} as *inhibition*. In some cases more complex rules for combining inputs are used, in which a distinction is made between excitatory and inhibitory inputs. We call units with a propagation rule (2.1) *sigma units*.

A different propagation rule, introduced by Feldman and Ballard [3], is known as the propagation rule for the *sigma-pi unit*:

$$i_i(t) = \sum_j w_{ij}(t) \prod_k a_{j_k}(t) + \theta_i(t). \tag{2.2}$$

Often, the a_{j_k} are weighted before multiplication. Although these units are not frequently used, they have their value for gating of input, as well as implementation of lookup tables [4].

2.1.3 Activation and output rules

We also need a rule which gives the effect of the total input on the activation of the unit. We need a function $\mathcal{F}_i$ which takes the total input $i_i(t)$ and the current activation $a_i(t)$ and produces a new value of the activation of the unit i:

$$a_i(t+1) = \mathcal{F}_i(a_i(t), i_i(t)). \tag{2.3}$$

Often, the activation function is a nondecreasing function of the total input of the unit:

$$a_i(t+1) = \mathcal{F}_i(i_i(t)) = \mathcal{F}_i\left(\sum_j w_{ij}(t)\, a_j(t) + \theta_i(t)\right), \tag{2.4}$$

although activation functions are not restricted to nondecreasing functions. Generally, some sort of threshold function is used: a hard limiting threshold function (a sgn function), or a linear or semi-linear function, or a smoothly limiting threshold (see figure 2.2). For this smoothly limiting function often a sigmoid (S-shaped) function like

$$a_i = \mathcal{F}(i_i) = \frac{1}{1+e^{-i_i}} \tag{2.5}$$

is used. In some applications a hyperbolic tangent is used, yielding output values in the range $[-1, +1]$.

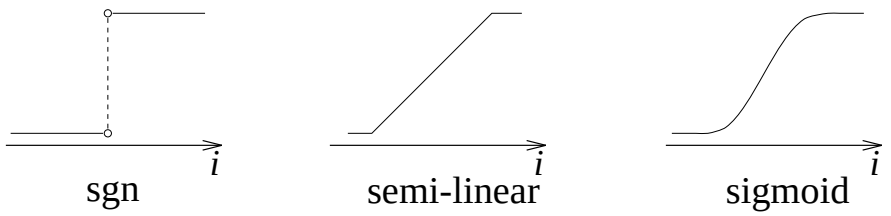

Figure 2.2: Various activation functions for a unit.

In some cases, the output of a unit can be a stochastic function of the total input of the unit. In that case the activation is not deterministically determined by the neuron input, but the neuron input determines the probability p that a neuron get a high activation value:

$$p(a_i \leftarrow 1) = \frac{1}{1+e^{-i_i/T}}, \tag{2.6}$$

in which T (cf. temperature) is a parameter which determines the slope of the probability function. This type of unit will be discussed more extensively in chapter 6.

In all networks we describe we consider the output of a neuron to be identical to its activation level.

2.2 Network topologies

In the previous section we discussed the properties of the basic processing unit in an artificial neural network. This section focuses on the pattern of connections between the units and the propagation of data.

As for this pattern of connections, the main distinction we can make is between:

- *Feed-forward* networks, where the data flow from input to output units is strictly feed-forward. The data processing can extend over multiple (layers of) units, but no feedback connections are present, that is, connections extending from outputs of units to inputs of units in the same layer or previous layers.

- *Recurrent* networks that do contain feedback connections. Contrary to feed-forward networks, the dynamical properties of the network are important. In some cases, the activation values of the units undergo a relaxation process such that the network will evolve to a stable state in which these activations do not change anymore. In other applications, the change of the activation values of the output neurons are significant, such that the dynamical behaviour constitutes the output of the network [5].

Classical examples of feed-forward networks are the Perceptron and Adaline, which will be discussed in the next chapter. Examples of recurrent networks have been presented by Anderson [6], Kohonen [7], and Hopfield [8] and will be discussed in chapter 6.

2.3 Training of artificial neural networks

A neural network has to be configured such that the application of a set of inputs produces (either 'direct' or via a relaxation process) the desired set of outputs. Various methods to set the strengths of the connections exist. One way is to set the weights explicitly, using *a priori* knowledge. Another way is to 'train' the neural network by feeding it teaching patterns and letting it change its weights according to some learning rule.

2.3.1 Paradigms of learning

We can categorise the learning situations in two distinct sorts. These are:

- *Supervised learning* or *Associative learning* in which the network is trained by providing it with input and matching output patterns. These input-output pairs can be provided by an external teacher, or by the system which contains the network (*self-supervised*).

- *Unsupervised learning* or *Self-organisation* in which an (output) unit is trained to respond to clusters of pattern within the input. In this paradigm the system is supposed to discover statistically salient features of the input population. Unlike the supervised learning paradigm, there is no *a priori* set of categories into which the patterns are to be classified; rather the system must develop its own representation of the input stimuli.

2.3.2 Modifying patterns of connectivity

Both learning paradigms discussed above result in an adjustment of the weights of the connections between units, according to some modification rule. Virtually all learning rules for models of this type can be considered as a variant of the *Hebbian* learning rule suggested by Hebb in his classic book *Organization of Behaviour* (1949) [9]. The basic idea is that if two units i and j are active simultaneously, their interconnection must be strengthened. If i receives input from j, the simplest version of Hebbian learning prescribes to modify the weight w_{ij} with

$$\Delta w_{ij} = \gamma a_i a_j, \tag{2.7}$$

where γ is a positive constant of proportionality representing the *learning rate*. Another common rule uses not the actual activation of unit i but the difference between the actual and desired activation for adjusting the weights:

$$\Delta w_{ij} = \gamma(d_i - a_i)a_j, \tag{2.8}$$

in which d_i is the desired activation provided by a teacher. This is often called the *Widrow-Hoff rule* or the *delta rule*, and will be discussed in the next chapter.

Many variants (often very exotic ones) have been published the last few years. In the next chapters some of these update rules will be discussed.

2.4 Notation and terminology

Throughout the years researchers from different disciplines have come up with a vast number of terms applicable in the field of neural networks. Our computer scientist point-of-view enables us to adhere to a subset of the terminology which is less biologically inspired, yet still conflicts arise. Our conventions are discussed below.

2.4.1 Notation

We use the following notation in our formulae. Note that not all symbols are meaningful for all networks, and that in some cases subscripts or superscripts may be left out (e.g., p is often not necessary) or added (e.g., vectors can, contrariwise to the notation below, have indices) where necessary. Vectors are indicated with a bold non-slanted font:

i $(, j, k, \ldots)$ the unit i $(, j, k, \ldots)$;

$\mathbf{x}^p$ the pth input pattern vector;

x_i^p the ith element of the pth input pattern vector;

$\mathbf{i}^p$ the input to a set of neurons when input pattern vector p is clamped (i.e., presented to the network); often: the input of the *network* by clamping input pattern vector p;

i_i^p the net input of neuron i due to input pattern vector p;

$\mathbf{o}^p$ the output of the *network* when input pattern vector p was input to the network;

o_i^p the ith element of the output of the network when input pattern vector p was input to the network;

$\mathbf{d}^p$ the desired output of the *network* when input pattern vector p was input to the network;

d_i^p the ith element of the desired output of the network when input pattern vector p was input to the network;

$\mathbf{a}^p$ the activation values of the *network* when input pattern vector p was input to the network;

a_i^p the activation values of element i of the network when input pattern vector p was input to the network;

W the matrix of connection weights;

$\mathbf{w}_i$ the weights of the connections which feed into unit i;

w_{ij} the weight of the connection from unit j to unit i;

$\boldsymbol{\mathcal{F}}$ the set of activation functions;

$\mathcal{F}_i$ the activation function associated with unit i;

γ_{ij} the learning rate associated with weight w_{ij};

θ the biases to the units;

θ_i the bias input to unit i;

U_i the threshold of unit i in $\mathcal{F}_i$;

E^p the error in the output of the network when input pattern vector p is input;

$\mathcal{E}$ the energy of the network.

2.4.2 Terminology

Output vs. activation of a unit. Since there is no need to do otherwise, we consider the output and the activation value of a unit to be one and the same thing. That is, the output of each neuron equals activation value. Note that the symbol o_i is usually reserved for the output of the *network* (i.e., an output unit i) and not of any unit i.

Bias, offset, threshold. These terms all refer to a constant (i.e., independent of the network input) term which is input to a unit. They may be used interchangeably, although the latter two terms are often envisaged as a property of the activation function. Furthermore, this external input is usually implemented (and can be written) as a weight from a unit with activation value 1.

Number of layers. In a feed-forward network, the inputs perform no computation and their layer is therefore not counted. Thus a network with one input layer, one hidden layer, and one output layer is referred to as a network with *two* layers. This convention is widely though not yet universally used.

Learn vs. store. A distinction can be made between *learning* and *storage* methods for information gathering. The former then would refer to methods which iteratively minimise some error or penalty function, whereas the latter refers to methods which analytically determine the parameters (weights) of the network. The reason for us to make this distinction is the following. Learning algorithms can be seen as a process in which the state of the network (viz. weights) relaxes to a certain optimum. These algorithms are mostly associated with networks in which the activation values are fed forward only. On the other hand, in recurrent networks the state of the network is described by the *activation values*, which also relax to some optimum. The weights of such networks are usually *a priori* determined.

References

[1] D. E. Rumelhart & J. L. McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (volume 1), The MIT Press, 1986.

[2] J. L. McClelland & D. E. Rumelhart, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (volume 2), The MIT Press, 1986.

[3] J. A. Feldman & D. H. Ballard, "Connectionist models and their properties," *Cognitive Science* 6 (1982), 205–254.

[4] B. W. Mel, *Connectionist Robot Motion Planning*, Academic Press, San Diego, CA, 1990.

[5] B. A. Pearlmutter, "Dynamic recurrent neural networks," School of Computer Science, Carnegie Mellon University, CMU–CS–90–196, Pittsburgh, PA 15213, December 1990.

[6] J. A. Anderson, "Neural models with cognitive implications," in *Basic Processes in Reading Perception and Comprehension Models*, D. LaBerge & S. J. Samuels, eds., Erlbaum, Hillsdale, NJ, 1977, 27–90.

[7] T. Kohonen, *Associative Memory: A System-Theoretical Approach*, Springer-Verlag, 1977.

[8] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the National Academy of Sciences* 79 (1982), 2554–2558.

[9] D. O. Hebb, *The Organization of Behaviour*, Wiley, New York, 1949.

Part II

THEORY

3 Adaline and Perceptron

This chapter describes some classical approaches to the neural computing and learning problem. We are not aiming at a complete survey of all related works but rather at setting a perspective of the main lines of earlier research on these topics.

In the first part of this chapter we discuss the representational power of the models, while in the second we will discuss their learning algorithms. It is important to distinguish between representation and learning. *Representation* refers to the ability of the neural system to simulate (i.e., *represent*) a specified function. On the other hand, *learning* requires the existence of a systematic procedure for adjusting the network weights to produce that function.

The two early models which will be described have the following main characteristics:

1. In the *Adaline*, a linear weighting is formed adaptively (by a supervised training method) which produces the desired response signal to different input patterns; the result is optimal in the sense of least squares.

2. In the *Perceptron*, there is an additional preprocessing network for the extraction of signal combinations (features) from the input pattern. The output circuits are made to discriminate between different input patterns, each one responding to a different class of patterns.

3.1 The adaptive linear element (Adaline)

The simplest 'learning' systems sum up their input signals linearly, whereby each signal is provided with a suitable *weighting parameter*. The 'adaptive linear element,' also named *Adaline*[1], suggested by Widrow and Hoff [1] is a classical example. In a simple physical implementation (fig. 3.1) this device comprises a set of controllable resistors connected to a circuit which can sum up currents caused by the input voltage signals. Usually the central block, the summer, is also followed by a quantiser which outputs either $+1$ of -1, depending on the polarity of the sum.

Although the adaptive process is here exemplified in a case when there is only one output, it may be clear that a system with many parallel outputs is directly implementable by multiple units of the above kind.

If the input conductances are denoted by w_i, $i = 0, 1, \ldots, n$, and the input and output

[1] ADALINE first stood for ADAptive LInear NEuron, but when neurons became less and less popular this acronym was changed to ADAptive LINear Element.

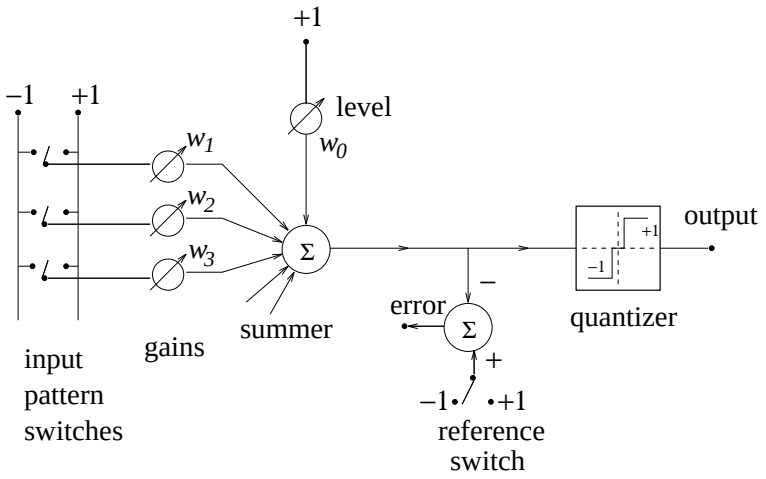

Figure 3.1: The Adaline.

signals by x_i and a, respectively, then the output of the central block is defined to be

$$a = \sum_{i=1}^{n} w_i x_i + \theta, \tag{3.1}$$

where $\theta = w_0$. The purpose of this device is to yield a given value a^p at its output when the set of values x_i^p, $i = 1, 2, \ldots, n$, is applied at the inputs. The problem is to determine the coefficients w_i, $i = 0, 1, \ldots, n$, in such a way that the input-output response is correct for a large number of arbitrarily chosen signal sets. If an exact mapping is not possible, the average error must be minimised, for instance, in the sense of least squares. An adaptive operation means that there exists a mechanism by which the w_i can be adjusted, usually iteratively, to attain the correct values. For the *Adaline*, Widrow introduced the delta rule to adjust the weights. This rule will be discussed in section 3.6.

3.2 The Perceptron

The Perceptron was proposed by Rosenblatt (1959) [2]. Perceptrons are feed-forward networks consisting of binary units. In their simplest form they consist of an N-element input layer ('retina') which feeds into a layer of M 'association,' 'mask,' or 'predicate' units ϕ_i, and a single output unit. The goal of the operation of the perceptron is to learn a given transformation $T : \{-1, 1\}^N \rightarrow \{-1, 1\}$ using learning samples with input $\mathbf{x}$ and corresponding output $o = T(\mathbf{x})$. In the original definition, the activity of the predicate units can be any function ϕ_i of the input layer $\mathbf{x}$ but the learning procedure only adjusts the connections to the output unit. The reason for this is that no recipe had been found to adjust the connections between $\mathbf{x}$ and ϕ_i. Depending on the functions ϕ_i, perceptrons can be grouped into different *families*. In [3] a number of these families are described and properties of these families have been described. The output unit of a perceptron is a linear threshold element which takes the output o according to the rule:

$$o = \mathcal{F}\left(\sum_{i=1}^{n} w_i \phi_i + \theta\right), \tag{3.2}$$

x_0	x_1	o
-1	-1	-1
-1	1	1
1	-1	1
1	1	-1

Table 3.1: Exclusive-or truth table.

in which $\mathcal{F}$ is the Heaviside or sgn function. In this particular case,

$$\mathcal{F}(i) = \begin{cases} 1 & \text{if } i > 0 \\ -1 & \text{otherwise.} \end{cases} \tag{3.3}$$

Rosenblatt (1959) [2] proved a remarkable theorem about perceptron learning (explained below) and in the early 60s perceptrons created a great deal of interest and optimism. The initial euphoria was replaced by disillusion after the publication of Minsky and Papert's *Perceptrons* in 1969 [3]. In this book they analysed the perceptron thoroughly and proved that there are severe restrictions on what perceptrons can represent.

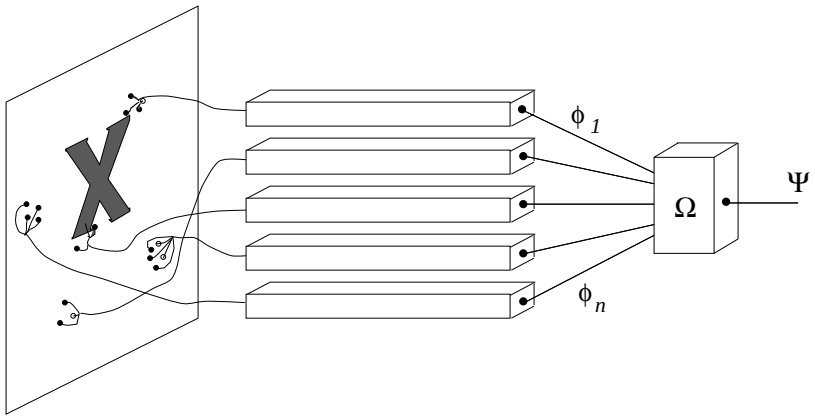

Figure 3.2: The Perceptron.

3.3 Exclusive-or problem

One of Minsky and Papert's most discouraging results shows that a single layer perceptron can not represent a simple exclusive-or function. Table 3.1 shows the desired relationships between inputs and output units for this function.

In a simple network with two inputs and one output, as depicted in figure 3.3, the net input is equal to:

$$i = w_1 x_1 + w_2 x_2 + \theta. \tag{3.4}$$

According to eq. (3.2), the output of the perceptron is zero when i is negative and equal to one when i is positive. In figure 3.4 a geometrical representation of the input domain is given.

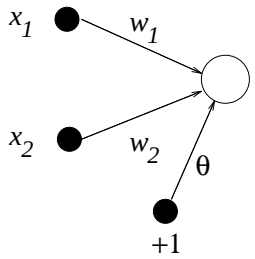

Figure 3.3: Single layer perceptron with two inputs.

For a constant θ, the output of the perceptron is equal to one on one side of the dividing line which is defined by:

$$w_1 x_1 + w_2 x_2 = -\theta \tag{3.5}$$

And equal to zero on the other side of this line.

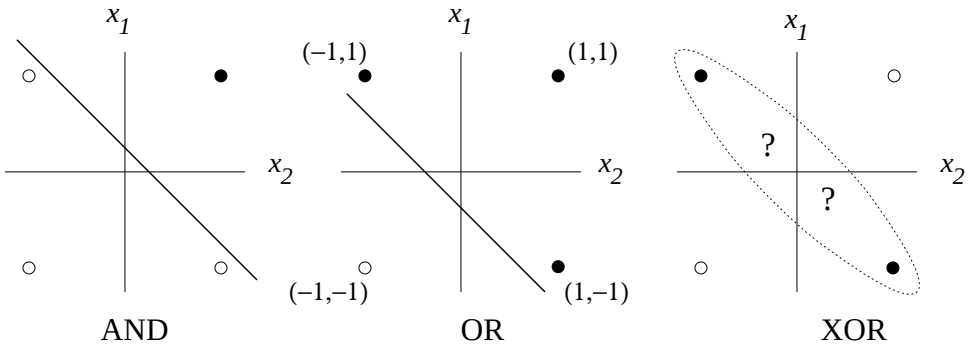

Figure 3.4: Geometric representation of input space.

To see that such a solution cannot be found, note figure 3.4. The input space consists of four points, and the two solid circles at $(1,-1)$ and $(-1,1)$ cannot be separated by a straight line from the two open circles at $(-1,-1)$ and $(1,1)$. The natural question to ask is, How can this problem be overcome? Minsky and Papert prove in their book that for binary inputs, any transformation can be carried out by adding a layer of predicates which are connected to all inputs. The proof is given in the next section.

For the specific XOR problem we geometrically show that by introducing *hidden units*, thereby extending the network to a *multi-layer perceptron*, the problem can be solved. Fig. 3.5a demonstrates that the four input points are now embedded in a three-dimensional space defined by the two inputs plus the single hidden unit. These four points are now easily separated by a linear manifold (plane) into two groups, as desired. This simple example demonstrates that adding hidden units increases the class of problems that are soluble by feed-forward, perceptron-like networks. However, by this generalisation of the basic architecture we have also incurred a serious loss; there is no longer a learning rule.

3.4 Multi-layer perceptrons can do everything

In the previous section we showed that by adding an extra hidden unit, the XOR problem can be solved. For binary units, one can prove that this architecture is able to perform any

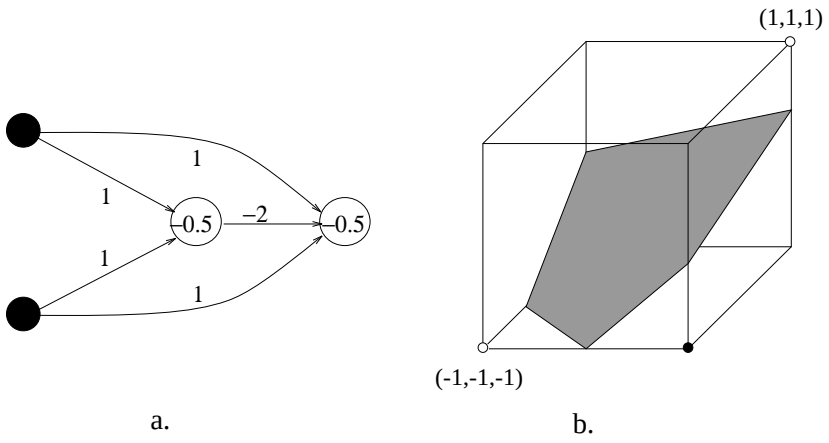

Figure 3.5: Solution of the XOR problem.
a) The perceptron of fig. 3.3 with an extra hidden unit. With the indicated values of the weights w_{ij} (next to the connecting lines) and the thresholds θ_i (in the circles) this perceptron solves the XOR problem. b) This is accomplished by mapping the four points of figure 3.4 onto the four points indicated here; clearly, separation (by a linear manifold) into the required groups is now possible.

transformation given the correct connections and weights. The most primitive is the next one. For a given transformation T, we can divide the set of all possible input vectors into two classes:

$$X^+ = \{\, \mathbf{x} \,|\, T(\mathbf{x}) = 1 \,\} \quad \text{and} \quad X^- = \{\, \mathbf{x} \,|\, T(\mathbf{x}) = -1 \,\}. \tag{3.6}$$

Since there are N input units, the total number of possible input vectors $\mathbf{x}$ is 2^N. For every $\mathbf{x}^p \in X^+$ a hidden unit i can be reserved of which the activation a_i is 1 if and only if the specific pattern p is present at the input: we can choose its weights w_{ij} equal to the specific pattern $\mathbf{x}^p$ and the bias θ_i equal to $1 - N$ such that

$$a_i^p = \text{sgn}\left(\sum_j w_{ij} x_j^p - N + \tfrac{1}{2} \right) \tag{3.7}$$

is equal to 1 for $\mathbf{x}^p = \mathbf{w}_i$ only. Similarly, the weights to the output neuron can be chosen such that the output is one as soon as one of the M predicate neurons is one:

$$o^p = \text{sgn}\left(\sum_{i=1}^{M} a_i + M - \tfrac{1}{2} \right). \tag{3.8}$$

This perceptron will give $o = 1$ only if $\mathbf{x} \in X^+$: it performs the desired mapping. The problem is the large number of predicate units, which is equal to the number of patterns in $X+$, which is maximally 2^N. Of course we can do the same trick for X^-, and we will always take the minimal number of mask units, which is maximally 2^{N-1}. A more elegant proof is given in [3], but the point is that for complex transformations the number of required units in the hidden layer is exponential in N.

3.5 Perceptron learning rule and convergence theorem

In the original description of the (multi-layer) perceptron, only the connections to the output unit are modified, since it was not clear which principle should be used according to which the connections from the input to the hidden layer should be modified. The perceptron learning rule is very simple and can be stated as follows:

1. Start with random weights for the connections;
2. Select an input vector $\mathbf{x}$ from the set of training samples;
3. If $o \neq T(\mathbf{x})$ (the perceptron gives an incorrect response), modify all connections w_i according to: $\Delta w_i = T(\mathbf{x})\phi_i$;
4. Go back to 2.

Note that the procedure is very similar to the Hebb rule; the only difference is that, when the network responds correctly, no connection weights are modified. Besides modifying the weights, we must also modify the threshold θ. This θ is considered as a connection w_0 between the output neuron and a 'dummy' predicate unit which is always on: $\phi_0 = 1$. Given the perceptron learning rule as stated above, this threshold is modified according to:

$$\Delta\theta = \begin{cases} 0 & \text{if the perceptron responds correctly;} \\ T(\mathbf{x}) & \text{otherwise.} \end{cases} \tag{3.9}$$

The most interesting of this rule is the existence of the associated perceptron convergence theorem, which states the following:

Theorem 1 *If there exists a set of connection weights $\mathbf{w}^*$ which is able to perform the transformation T, the perceptron learning rule will converge to some solution (which may or may not be the same as $\mathbf{w}^*$) in a finite number of steps for any initial choice of the weights.*

Proof *Given the fact that the length of the vector $\mathbf{w}^*$ does not play a role (because of the* sgn *operation), we take $\|\mathbf{w}^*\| = 1$. Because $\mathbf{w}^*$ is a correct solution, the value $|\mathbf{w}^* \cdot \boldsymbol{\phi}(\mathbf{x})|$, where $\cdot$ denotes dot or inner product, will be greater than 0 or: there exists a $\delta > 0$ such that $|\mathbf{w}^* \cdot \boldsymbol{\phi}(\mathbf{x})| > \delta$ for all inputs $\mathbf{x}$. Now define $\cos\alpha \equiv \mathbf{w} \cdot \mathbf{w}^*/\|\mathbf{w}\|$. When according to the perceptron learning rule, connection weights are modified at a given input $\mathbf{x}$, we know that $\Delta\mathbf{w} = T(\mathbf{x})\boldsymbol{\phi}(\mathbf{x})$, and the weight after modification is $\mathbf{w}' = \mathbf{w} + \Delta\mathbf{w}$. From this it follows that:*

$$\begin{aligned} \mathbf{w}' \cdot \mathbf{w}^* &= \mathbf{w} \cdot \mathbf{w}^* + T(\mathbf{x}) \cdot \mathbf{w}^* \cdot \boldsymbol{\phi}(\mathbf{x}) \\ &= \mathbf{w} \cdot \mathbf{w}^* + \operatorname{sgn}\left(\mathbf{w}^* \cdot \boldsymbol{\phi}(\mathbf{x})\right) \mathbf{w}^* \cdot \boldsymbol{\phi}(\mathbf{x}) \\ &> \mathbf{w} \cdot \mathbf{w}^* + \delta \end{aligned}$$

$$\begin{aligned} \|\mathbf{w}'\|^2 &= \|\mathbf{w} + T(\mathbf{x})\,\boldsymbol{\phi}(\mathbf{x})\|^2 \\ &= \mathbf{w}^2 + 2T(\mathbf{x})\,\mathbf{w} \cdot \boldsymbol{\phi}(\mathbf{x}) + \boldsymbol{\phi}^2(\mathbf{x}) \\ &< \mathbf{w}^2 + \boldsymbol{\phi}^2(\mathbf{x}) \quad \textit{(because } T(\mathbf{x}) = -\operatorname{sgn}[\mathbf{w} \cdot \boldsymbol{\phi}(\mathbf{x})] \textit{ !!)} \\ &= \mathbf{w}^2 + M. \end{aligned}$$

After t modifications we have:

$$\begin{aligned} \mathbf{w}(t) \cdot \mathbf{w}^* &> \mathbf{w} \cdot \mathbf{w}^* + t\delta \\ \|\mathbf{w}(t)\|^2 &< \mathbf{w}^2 + tM \end{aligned}$$

such that

$$\begin{aligned}\cos\alpha(t) &= \frac{\mathbf{w}^* \cdot \mathbf{w}(t)}{\|\mathbf{w}(t)\|} \\ &> \frac{\mathbf{w}^* \cdot \mathbf{w} + t\delta}{\sqrt{\mathbf{w}^2 + tM}}.\end{aligned}$$

From this follows that $\lim_{t\to\infty} \cos\alpha(t) = \lim_{t\to\infty} \frac{\delta}{\sqrt{M}}\sqrt{t} = \infty$, *while by definition* $\cos\alpha \leq 1$ *!*

The conclusion is that there must be an upper limit $t_{\max}$ *for* t. *The system modifies its connections only a limited number of times. In other words: after maximally* $t_{\max}$ *modifications of the weights the perceptron is correctly performing the mapping.* $t_{\max}$ *will be reached when* $\cos\alpha = 1$. *If we start with connections* $\mathbf{w} = \mathbf{0}$,

$$t_{\max} = \frac{M}{\delta^2}. \tag{3.10}$$

□

3.6 The delta rule

An important generalisation of the perceptron training algorithm, presented by Widrow and Hoff as the 'least mean square' (LMS) learning procedure, extends this technique to continuous inputs and outputs. The LMS procedure, also known as the *delta rule*, has been applied most often with purely linear output units. For such a perceptron with a single output unit, the output is simply given by

$$a = \sum_j w_j x_j + \theta. \tag{3.11}$$

The error function, as indicated by the name least mean square, is the summed square error. That is, the total error, E, is defined to be

$$E = \sum_p E^p = \tfrac{1}{2}\sum_p (d^p - a^p)^2, \tag{3.12}$$

where the index p ranges over the set of input patterns and E^p represents the error on pattern p. The variable d^p is the desired (or *target*) output when pattern p is presented, and a^p is the actual output for this pattern. The LMS procedure finds the values of all the weights that minimise the error function by a method called *gradient descent*. The idea is to make a change in the weight proportional to the negative of the derivative of the error as measured on the current pattern with respect to each weight:

$$\Delta_p w_j = -\gamma \frac{\partial E^p}{\partial w_j} \tag{3.13}$$

where γ is a constant of proportionality. The derivative is

$$\frac{\partial E^p}{\partial w_j} = \frac{\partial E^p}{\partial a^p}\frac{\partial a^p}{\partial w_j}. \tag{3.14}$$

Because of the linear units (eq. (3.11)),

$$\frac{\partial a^p}{\partial w_j} = x_j \tag{3.15}$$

and

$$\frac{\partial E^p}{\partial a^p} = -(d^p - a^p) \tag{3.16}$$

such that

$$\Delta_p w_j = \gamma \delta^p x_j \tag{3.17}$$

where $\delta^p = d^p - a^p$ is the difference between the target output and the actual output for pattern p.

The delta rule modifies weight appropriately for target and actual outputs of either polarity and for both continuous and binary input and output units. These characteristics have opened up a wealth of new applications.

References

[1] B. Widrow & M. E. Hoff, "Adaptive switching circuits," in *1960 IRE WESCON Convention Record*, New York, 1960, 96–104.

[2] F. Rosenblatt, *Principles of Neurodynamics*, Spartan Books, New York, 1959.

[3] M. Minsky & S. Papert, *Perceptrons: An Introduction to Computational Geometry*, The MIT Press, 1969.

4 Back-Propagation

As we have seen in the previous chapter, a single-layer network has severe restrictions: the class of tasks that can be accomplished is very limited. In this chapter we will focus on feed-forward networks with layers of processing units.

Minsky and Papert [1] showed in 1969 that a two layer feed-forward network can overcome many restrictions, but did not present a solution to the problem of how to adjust the weights from input to hidden units. An answer to this question was presented by Rumelhart, Hinton and Williams in 1986 [2], and similar solutions appeared to have been published earlier (Parker (1985) [3]; Le Cun (1985) [4]; Werbos (1974) [5]).

The central idea behind this solution is that the errors for the units of the hidden layer are determined by back-propagating the errors of the units of the output layer. For this reason the method is often called the *back-propagation learning rule*. Back-propagation can also be considered as a generalisation of the delta rule for non-linear activation functions[1] and multi-layer networks.

4.1 Multi-layer feed-forward networks

A feed-forward network has a layered structure. Each layer consists of units which receive their input from units from a layer directly below and send their output to units in a layer directly above the unit. There are no connections within a layer. The N_i inputs are fed into the first layer of $N_{h,1}$ *hidden* units. The input units are merely 'fan-out' units; no processing takes place in these units. The activation of a hidden unit is a function $\mathcal{F}_i$ of the weighted inputs plus a bias, as given in in eq. (2.4). The output of the hidden units is distributed over the next layer of $N_{h,2}$ hidden units, until the last layer of hidden units, of which the outputs are fed into a layer of of N_o *output* units (see figure 4.1).

Although back-propagation can be applied to networks with any number of layers, as with networks with binary units (section 3.4) it has been shown [6–9] that only one layer of hidden units suffices to approximate any function with finitely many discontinuities to arbitrary precision, provided the activation functions of the hidden units are non-linear. In most applications a feed-forward network with a single layer of hidden units is used with a sigmoid activation function for the units.

[1] Of course, when linear activation functions are used, a multi-layer network is not more powerful than a single-layer network.

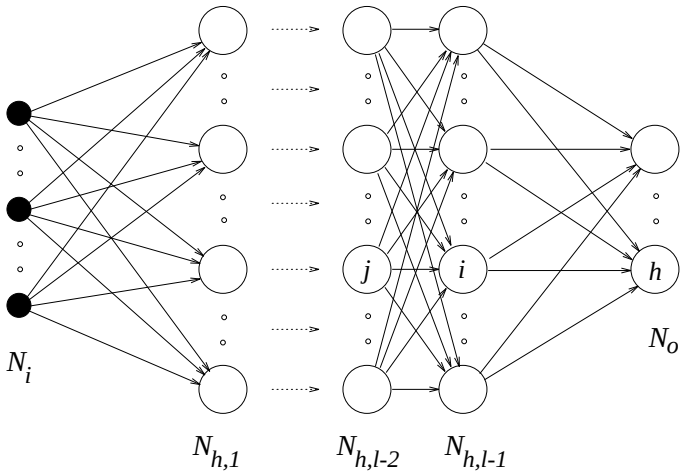

Figure 4.1: A multi-layer network with l layers of units.

4.2 The generalised delta rule

Since we are now using units with nonlinear activation functions, we have to generalise the delta rule which was presented in chapter 3 for linear functions to the set of non-linear activation functions. The activation is a differentiable function of the total input, given by

$$a_i^p = \mathcal{F}(i_i^p), \tag{4.1}$$

in which

$$i_i^p = \sum_j w_{ij} a_j^p + \theta_i. \tag{4.2}$$

To get the correct generalisation of the delta rule as presented in the previous chapter, we must set

$$\Delta_p w_{ij} = -\gamma \frac{\partial E^p}{\partial w_{ij}}. \tag{4.3}$$

The error measure E^p is defined as the total quadratic error for pattern p at the output units:

$$E^p = \tfrac{1}{2} \sum_{i=1}^{N_o} (d_i^p - a_i^p)^2, \tag{4.4}$$

while $E = \sum_p E^p$ is the overall measure of the error. We can write

$$\frac{\partial E^p}{\partial w_{ij}} = \frac{\partial E^p}{\partial i_i^p} \frac{\partial i_i^p}{\partial w_{ij}}. \tag{4.5}$$

By equation (4.2) we see that the second factor is

$$\frac{\partial i_i^p}{\partial w_{ij}} = a_j^p. \tag{4.6}$$

When we define

$$\delta_i^p = -\frac{\partial E^p}{\partial i_i^p}, \tag{4.7}$$

we will get an update rule which is equivalent to the delta rule as described in the previous chapter, resulting in a gradient descent on the error surface if we make the weight changes according to:

$$\Delta_p w_{ij} = \gamma \delta_i^p a_j^p. \tag{4.8}$$

The trick is to figure out what δ_i^p should be for each unit i in the network. The interesting result, which we now derive, is that there is a simple recursive computation of these δ's which can be implemented by propagating error signals backward through the network.

To compute δ_i^p we apply the chain rule to write this partial derivative as the product of two factors, one factor reflecting the change in error as a function of the output of the unit and one reflecting the change in the output as a function of changes in the input. Thus, we have

$$\delta_i^p = -\frac{\partial E^p}{\partial i_i^p} = -\frac{\partial E^p}{\partial a_i^p}\frac{\partial a_i^p}{\partial i_i^p}. \tag{4.9}$$

Let us compute the second factor. By equation (4.1) we see that

$$\frac{\partial a_i^p}{\partial i_i^p} = \mathcal{F}'(i_i^p), \tag{4.10}$$

which is simply the derivative of the squashing function $\mathcal{F}$ for the ith unit, evaluated at the net input i_i^p to that unit. To compute the first factor of equation (4.9), we consider two cases. First, assume that unit i is an output unit of the network. In this case, it follows from the definition of E^p that

$$\frac{\partial E^p}{\partial a_i^p} = -(d_i^p - a_i^p), \tag{4.11}$$

which is the same result as we obtained with the standard delta rule. Substituting this and equation (4.10) in equation (4.9), we get

$$\delta_i^p = (d_i^p - a_i^p)\, \mathcal{F}_i'(i_i^p) \tag{4.12}$$

for any output unit i. Secondly, if i is not an output unit, we do not readily know the contribution of the unit to the output error of the network. However, the error measure can be written as a function of the net inputs from hidden to output layer; $E^p = E^p(i_1^p, i_2^p, \ldots, i_i^p, \ldots)$ and we use the chain rule to write

$$\frac{\partial E^p}{\partial a_i^p} = \sum_{h=1}^{N_o} \frac{\partial E^p}{\partial i_h^p}\frac{\partial i_h^p}{\partial a_i^p} = \sum_{h=1}^{N_o} \frac{\partial E^p}{\partial i_h^p}\frac{\partial}{\partial a_i^p}\sum_{k=1}^{N_h} w_{hk} a_k^p = \sum_{h=1}^{N_o} \frac{\partial E^p}{\partial i_h^p} w_{hi} = -\sum_{h=1}^{N_o} \delta_h^p w_{hi}. \tag{4.13}$$

Substituting this in equation (4.9) yields

$$\delta_i^p = \mathcal{F}'(i_i^p) \sum_{h=1}^{N_o} \delta_h^p w_{hi}. \tag{4.14}$$

Equations (4.12) and (4.14) give a recursive procedure for computing the δ's for all units in the network, which are then used to compute the weight changes according to equation (4.8). This procedure constitutes the generalised delta rule for a feed-forward network of non-linear units.

4.3 Working with back-propagation

The application of the generalised delta rule thus involves two phases: During the first phase the input is presented and propagated forward through the network to compute the output values a_i^p for each output unit. This output is compared with the desired value, resulting in an error signal δ_i^p for each output unit. The second phase involves a backward pass through the network during which the error signal is passed to each unit in the network and appropriate weight changes are calculated.

Weight adjustments with sigmoid activation function. The results from the previous section can be summarised in three equations:

- The weight of a connection is adjusted by an amount proportional to the product of an error signal δ, on the unit i receiving the input and the output of the unit j sending this signal along the connection:

$$\Delta_p w_{ij} = \gamma \delta_i^p a_j^p. \tag{4.15}$$

- If the unit is an output unit, the error signal is given by

$$\delta_i^p = (d_i^p - a_i^p)\, \mathcal{F}'(i_i^p). \tag{4.16}$$

  Take as the activation function $\mathcal{F}$ the 'sigmoid' function as defined in chapter 2:

$$a_i^p = \mathcal{F}(i_i^p) = \frac{1}{1 + e^{-i_i^p}}. \tag{4.17}$$

  In this case the derivative is equal to

$$\begin{aligned} \mathcal{F}'(i_i^p) &= \frac{\partial}{\partial i_i^p} \frac{1}{1 + e^{-i_i^p}} \\ &= \frac{1}{(1 + e^{-i_i^p})^2} (-e^{-i_i^p}) \\ &= \frac{1}{(1 + e^{-i_i^p})} \frac{e^{-i_i^p}}{(1 + e^{-i_i^p})} \\ &= a_i^p (1 - a_i^p). \end{aligned} \tag{4.18}$$

  such that the error signal for an output unit can be written as:

$$\delta_i^p = (d_i^p - a_i^p)\, a_i^p (1 - a_i^p). \tag{4.19}$$

- The error signal for a hidden unit is determined recursively in terms of error signals of the units to which it directly connects and the weights of those connections. For the sigmoid activation function:

$$\delta_i^p = \mathcal{F}'(i_i^p) \sum_{h=1}^{N_o} \delta_h^p w_{hi} = a_i^p (1 - a_i^p) \sum_{h=1}^{N_o} \delta_h^p w_{hi}. \tag{4.20}$$

Learning rate and momentum. The learning procedure requires that the change in weight is proportional to $\partial E^p/\partial w$. True gradient descent requires that infinitesimal steps are taken. The constant of proportionality is the learning rate γ. For practical purposes we choose a learning rate that is as large as possible without leading to oscillation. One way to avoid oscillation at large γ, is to make the change in weight dependent of the past weight change by adding a *momentum* term:

$$\Delta w_{ij}(t+1) = \gamma \delta_i^p a_j^p + \alpha \Delta w_{ij}(t), \tag{4.21}$$

where t indexes the presentation number and α is a constant which determines the effect of the previous weight change.

The role of the momentum term is shown in figure 4.2. When no momentum term is used, it takes a long time before the minimum has been reached with a low learning rate, whereas for high learning rates the minimum is never reached because of the oscillations. When adding the momentum term, the minimum will be reached faster.

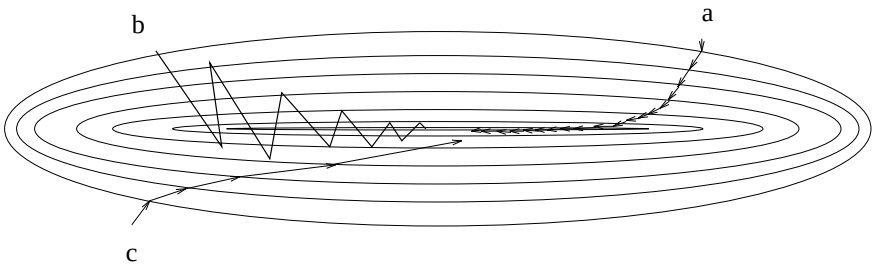

Figure 4.2: The descent in weight space. a) for small learning rate; b) for large learning rate: note the oscillations, and c) with large learning rate and momentum term added.

Learning per pattern. Although, theoretically, the back-propagation algorithm performs gradient descent on the total error only if the weights are adjusted after the full set of learning patterns has been presented, more often than not the learning rule is applied to each pattern separately, i.e., a pattern p is applied, E^p is calculated, and the weights are adapted ($p = 1, 2, \ldots, P$). There exists empirical indication that this results in faster convergence. Care has to be taken, however, with the order in which the patterns are taught. For example, when using the same sequence over and over again the network may become focused on the first few patterns. This problem can be overcome by using a permuted training method.

4.4 Other activation functions

Although sigmoid functions are quite often used as activation functions, other functions can be used as well. In some cases this leads to a formula which is known from traditional function approximation theories.

For example, from Fourier analysis it is known that any periodic function can be written as a infinite sum of sine and cosine terms (Fourier series):

$$f(x) = \sum_{n=0}^{\infty} (a_n \cos nx + b_n \sin nx). \tag{4.22}$$

We can rewrite this as a summation of sine terms

$$f(x) = a_0 + \sum_{n=1}^{\infty} c_n \sin(nx + \theta_n), \tag{4.23}$$

with $c_n = \sqrt{(a_n^2 + b_n^2)}$ and $\theta_n = \arctan(b/a)$. This can be seen as a feed-forward network with a single input unit for x; a single output unit for $f(x)$ and hidden units with an activation function $\mathcal{F} = \sin(i)$. The factor a_0 corresponds with the bias of the output unit, the factors c_n correspond with the weighs from hidden to output unit; the phase factor θ_n corresponds with the bias term of the hidden units and the factor n corresponds with the weights between the input and hidden layer.

The basic difference between the Fourier approach and the back-propagation approach is that the in the Fourier approach the 'weights' between the input and the hidden units (these are the factors n) are fixed integer numbers, whereas in the back-propagation approach these weights can take any value.

To illustrate the use of other activation functions we have trained a feed-forward network with one output unit, four hidden units, and one input with ten patterns drawn from the function $f(x) = \sin(2x)\sin(x)$. The result is depicted in Figure 4.3. The same function (albeit with other learning points) is learned with a network with eight (!) sigmoid hidden units (see figure 4.4).

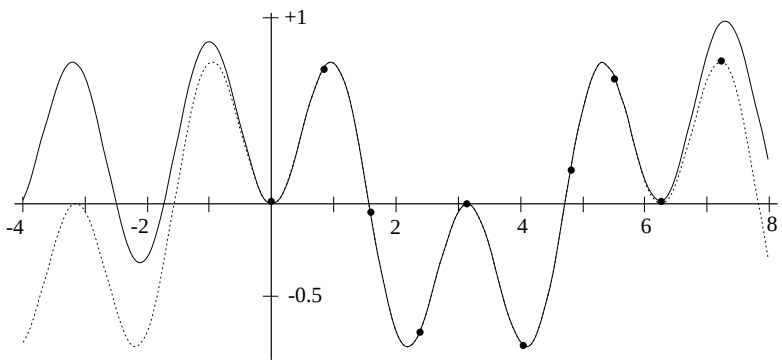

Figure 4.3: The periodic function $f(x) = \sin(2x)\sin(x)$ approximated with sine activation functions.
(Adapted from [10].)

4.5 Deficiencies of back-propagation

Despite the apparent success of the back-propagation learning algorithm, there are some aspects which make the algorithm not guaranteed to be universally useful. Most troublesome is the long training process. This can be a result of a non-optimum learning rate and momentum. A lot of advanced algorithms based on back-propagation learning have some optimised method to adapt this learning rate, as will be discussed in the next section. Outright training failures generally arise from two sources: network paralysis and local minima.

Network paralysis. As the network trains, the weights can be adjusted to very large values. The total input of a hidden unit or output unit can therefore reach very high (either

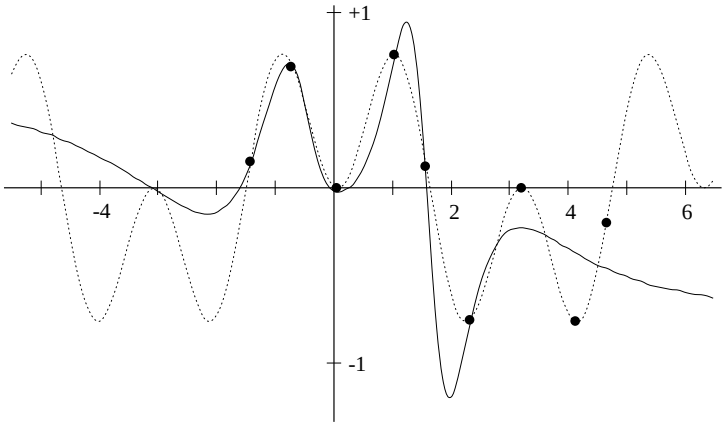

Figure 4.4: The periodic function $f(x) = \sin(2x)\sin(x)$ approximated with sigmoid activation functions.
(Adapted from [10].)

positive or negative) values, and because of the sigmoid activation function the unit will have an activation very close to zero or very close to one. As is clear from equations (4.19) and (4.20), the weight adjustments which are proportional to $a_j^p(1 - a_j^p)$ will be close to zero, and the training process can come to a virtual standstill.

Local minima. The error surface of a complex network is full of hills and valleys. Because of the gradient descent, the network can get trapped in a local minimum when there is a much deeper minimum nearby. Probabilistic methods can help to avoid this trap, but they tend to be slow. Another suggested possibility is to increase the number of hidden units. Although this will work because of the higher dimensionality of the error space, and the chance to get trapped is smaller, it appears that there is some upper limit of the number of hidden units which, when exceeded, again results in the system being trapped in local minima.

4.6 Advanced algorithms

Many researchers have devised improvements of and extensions to the basic back-propagation algorithm described above. It is too early for a full evaluation: some of these techniques may prove to be fundamental, others may simply fade away. A few methods are discussed in this section.

Maybe the most obvious improvement is to replace the rather primitive steepest descent method with a *direction set* minimisation method, e.g., conjugate gradient minimisation. Note that minimisation along a direction $\mathbf{u}$ brings the function f at a place where its gradient is perpendicular to $\mathbf{u}$ (otherwise minimisation along $\mathbf{u}$ is not complete). Instead of following the gradient at every step, a set of n directions is constructed which are all conjugate to each other such that minimisation along one of these directions $\mathbf{u}_j$ does not spoil the minimisation along one of the earlier directions $\mathbf{u}_i$, i.e., the directions are non-interfering. Thus one minimisation in the direction of $\mathbf{u}_i$ suffices, such that n minimisations in a system with n degrees of freedom bring this system to a minimum (provided the system is quadratic). This is different from gradient descent, which directly minimises in the direction of the steepest descent [11].

Suppose the function to be minimised is approximated by its Taylor series

$$\begin{aligned} f(\mathbf{x}) &= f(\mathbf{p}) + \sum_i \left.\frac{\partial f}{\partial x_i}\right|_{\mathbf{p}} x_i + \tfrac{1}{2}\sum_{i,j} \left.\frac{\partial^2 f}{\partial x_i \partial x_j}\right|_{\mathbf{p}} x_i x_j + \cdots \\ &\approx \tfrac{1}{2}\mathbf{x}^T\mathbf{A}\mathbf{x} - \mathbf{b}^T\mathbf{x} + c \end{aligned}$$

where T denotes transpose, and

$$c \equiv f(\mathbf{p}) \qquad \mathbf{b} \equiv -\nabla f|_{\mathbf{p}} \qquad [\mathbf{A}]_{ij} \equiv \left.\frac{\partial^2 f}{\partial x_i \partial x_j}\right|_{\mathbf{p}}. \tag{4.24}$$

$\mathbf{A}$ is a symmetric positive definite[2] $n \times n$ matrix, the *Hessian* of f at $\mathbf{p}$. The gradient of f is

$$\nabla f = \mathbf{A}\mathbf{x} - \mathbf{b}, \tag{4.26}$$

such that a change of $\mathbf{x}$ results in a change of the gradient as

$$\delta(\nabla f) = \mathbf{A}(\delta\mathbf{x}). \tag{4.27}$$

Now suppose f was minimised along a direction $\mathbf{u}_i$ to a point where the gradient $-\mathbf{g}_{i+1}$ of f is perpendicular to $\mathbf{u}_i$, i.e.,

$$\mathbf{u}_i^T\mathbf{g}_{i+1} = 0, \tag{4.28}$$

and a new direction $\mathbf{u}_{i+1}$ is sought. In order to make sure that moving along $\mathbf{u}_{i+1}$ does not spoil minimisation along $\mathbf{u}_i$ we require that the gradient of f remain perpendicular to $\mathbf{u}_i$, i.e.,

$$\mathbf{u}_i^T\mathbf{g}_{i+2} = 0; \tag{4.29}$$

otherwise we would once more have to minimise in a direction which has a component of $\mathbf{u}_i$. Combining (4.28) and (4.29), we get

$$0 = \mathbf{u}_i^T(\mathbf{g}_{i+1} - \mathbf{g}_{i+2}) = \mathbf{u}_i^T\delta(\nabla f) = \mathbf{u}_i^T\mathbf{A}\mathbf{u}_{i+1}. \tag{4.30}$$

When eq. (4.30) holds for two vectors $\mathbf{u}_i$ and $\mathbf{u}_{i+1}$ they are said to be *conjugate*.

Now, starting at some point $\mathbf{p}_0$, the first minimisation direction $\mathbf{u}_0$ is taken equal to $\mathbf{g}_0 = -\nabla f(\mathbf{p}_0)$, resulting in a new point $\mathbf{p}_1$. For $i \geq 0$, calculate the directions

$$\mathbf{u}_{i+1} = \mathbf{g}_{i+1} + \gamma_i\mathbf{u}_i, \tag{4.31}$$

where γ_i is chosen to make $\mathbf{u}_i^T\mathbf{A}\mathbf{u}_{i-1} = 0$ and the successive gradients perpendicular, i.e.,

$$\gamma_i = \frac{\mathbf{g}_{i+1}^T\mathbf{g}_{i+1}}{\mathbf{g}_i^T\mathbf{g}_i} \qquad \text{with} \qquad \mathbf{g}_k = -\nabla f|_{\mathbf{p}_k} \quad \text{for all } k \geq 0. \tag{4.32}$$

Next, calculate $\mathbf{p}_{i+2} = \mathbf{p}_{i+1} + \lambda_{i+1}\mathbf{u}_{i+1}$ where λ_{i+1} is chosen so as to minimise $f(\mathbf{p}_{i+2})$[3].

It can be shown that the $\mathbf{u}$'s thus constructed are all mutually conjugate (e.g., see [12]). The process described above is known as the *Fletcher-Reeves* method, but there are many variants which work more or less the same [13–15].

[2] A matrix $\mathbf{A}$ is called *positive definite* if $\forall \mathbf{y} \neq \mathbf{0}$,

$$\mathbf{y}^T\mathbf{A}\mathbf{y} > 0. \tag{4.25}$$

[3] This is not a trivial problem (see [11, ch. 10]). However, line minimisation methods exist with super-linear convergence (see footnote 4).

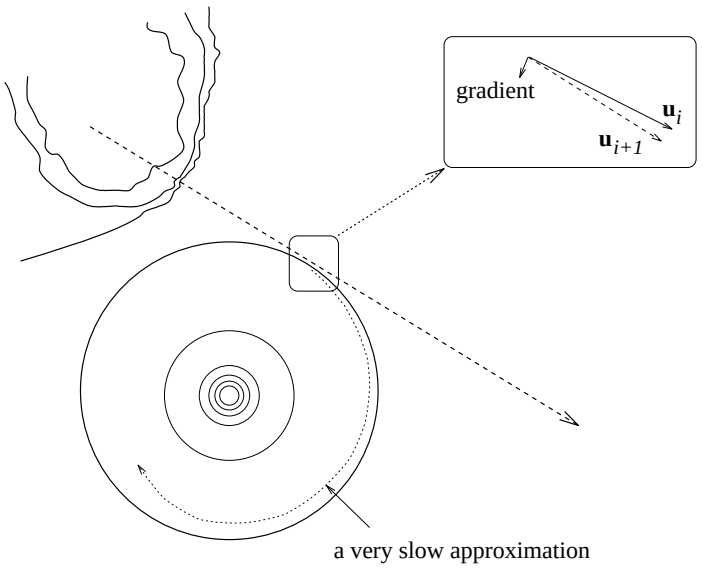

Figure 4.5: Slow decrease with conjugate gradient in non-quadratic systems. The hills on the left are very steep, resulting in a large search vector $\mathbf{u}_i$. When the quadratic portion is entered the new search direction is constructed from the previous direction and the gradient, resulting in a spiraling minimisation. This problem can be overcome by detecting such spiraling minimisations and restarting the algorithm with $\mathbf{u}_0 = -\nabla f$.

Although only n iterations are needed for a quadratic system with n degrees of freedom, due to the fact that we are not minimising quadratic systems, as well as a result of round-off errors, the n directions have to be followed several times (see figure 4.5). Powell [15] introduced some improvements to correct for behaviour in non-quadratic systems. The resulting cost is $O(n)$ which is significantly better than the linear convergence[4] of steepest descent.

Some improvements on back-propagation have been presented based on an independent adaptive learning rate parameter for each weight.

Van den Boomgaard and Smeulders [16] show that for a feed-forward network without hidden units an incremental procedure to find the optimal weight matrix W needs an adjustment of the weights with

$$\Delta W(t+1) = \gamma(t+1)\left(\mathbf{d}(t+1) - W(t)\,\mathbf{x}(t+1)\right)\mathbf{x}(t+1), \tag{4.33}$$

in which γ is not a constant but an variable $(N_i+1)\times(N_i+1)$ matrix which depends on the input vector. By using *a priori* knowledge about the input signal, the storage requirements for γ can be reduced.

Silva and Almeida [17] also show the advantages of an independent step size for each weight in the network. In their algorithm the learning rate is adapted after every learning

[4] A method is said to converge linearly if $E_{i+1} = cE_i$ with $c < 1$. Methods which converge with a higher power, i.e., $E_{i+1} = c(E_i)^m$ with $m > 1$ are called *super-linear*.

pattern:

$$\gamma_{ij}(t+1) = \begin{cases} u\gamma_{ij}(t) & \text{if } \frac{\partial E(t+1)}{\partial w_{ij}} \text{ and } \frac{\partial E(t)}{\partial w_{ij}} \text{ have the same signs;} \\ d\gamma_{ij}(t) & \text{if } \frac{\partial E(t+1)}{\partial w_{ij}} \text{ and } \frac{\partial E(t)}{\partial w_{ij}} \text{ have opposite signs.} \end{cases} \tag{4.34}$$

where u and d are positive constants with values slightly above and below unity, respectively. The idea is to decrease the learning rate in case of oscillations.

4.7 Applications

Back-propagation has been applied to a wide variety of research applications. Sejnowski and Rosenberg (1987) [18] produced a spectacular success with NETtalk, a system that converts printed English text into highly intelligible speech.

A feed-forward network with one layer of hidden units has been described by Gorman and Sejnowski (1988) [19] as a classification machine for sonar signals.

Another application of a multi-layer feed-forward network with a back-propagation training algorithm is to learn an unknown function between input and output signals from the presentation of examples. It is hoped that the network is able to generalise correctly, so that input values which are not presented as learning patterns will result in correct output values. An example is the work of Josin [20], who used a two-layer feed-forward network with back-propagation learning to perform the inverse kinematic transform which is needed by a robot arm controller (see chapter 7).

References

[1] M. Minsky & S. Papert, *Perceptrons: An Introduction to Computational Geometry*, The MIT Press, 1969.

[2] D. E. Rumelhart, G. E. Hinton & R. J. Williams, "Learning representations by back-propagating errors," *Nature* 323 (1986), 533–536.

[3] D. B. Parker, Massachusetts Institute of Technology, Center for Computational Research in Economics and Management Science, TR–47, Cambridge, MA, 1985.

[4] Y. Le Cun, "Une procedure d'apprentissage pour reseau a seuil assymetrique," *Proceedings of Cognitiva* 85 (1985), 599–604.

[5] P. J. Werbos, "Beyond regression: New tools for prediction and analysis in the behavioral sciences," Harvard University, Masters Thesis, 1974.

[6] K. Hornik, M. Stinchcombe & H. White, "Multilayer feedforward networks are universal approximators," *Neural Networks* 2 (1989), 359–366.

[7] K.-I. Funahashi, "On the approximate realization of continuous mappings by neural networks," *Neural Networks* 2 (1989), 193–192.

[8] G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Mathematics of Control, Signals, and Systems* 2 (1989), 303–314.

[9] E. J. Hartman, J. D. Keeler & J. M. Kowalski, "Layered neural networks with Gaussian hidden units as universal approximations," *Neural Computation* 2 (Summer 1990), 210–215.

[10] M. M. Dastani, *Functie-Benadering met Feed-Forward Netwerken*, Universiteit van Amsterdam, Faculteit Computer Systemen, June 1991, graduation thesis.

[11] W. H. Press, B. P. Flannery, S. A. Teukolsky & W. T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, Cambridge, 1986.

[12] J. Stoer & R. Bulirsch, *Introduction to Numerical Analysis*, Springer-Verlag, New York–Heidelberg–Berlin, 1980, *Einführung in die Numerische Mathematik*, Springer-Verlag, Heidelberg–Berlin, 1972, 1976.

[13] M. R. Hestenes & E. Stiefel, "Methods of conjugate gradients for solving linear systems," *Nat. Bur. Standards J. Res.* 49 (1952), 409–436.

[14] E. Polak, *Computational Methods in Optimization*, Academic Press, New York, 1971.

[15] M. J. D. Powell, "Restart procedures for the conjugate gradient method," *Mathematical Programming* 12 (April 1977), 241–254.

[16] R. van den Boomgaard & A. Smeulders, "Self learning image processing using a-priori knowledge of spatial patterns," in *Proceedings of the I.A.S. Conference*, T. Kanade, F. C. A. Groen & L. O. Hertzberger, eds., Amsterdam, December 1989, 305–314.

[17] F. M. Silva & L. B. Almeida, "Speeding up backpropagation," in *Advanced Neural Computers*, R. Eckmiller, ed., North-Holland, 1990, 151–160.

[18] T. J. Sejnowski & C. R. Rosenberg, "NETtalk: A parallel network that learns to read aloud," The John Hopkins University Electrical Engineering and Computer Science Department, JHU/EECS–86/01, 1986.

[19] R. P. Gorman & T. J. Sejnowski, "Analysis of hidden units in a layered network trained to classify sonar targets," *Neural Networks* 1 (1988), 75–89.

[20] G. Josin, "Neural-space generalization of a topological transformation," *Biological Cybernetics* 59 (1988), 283–290.

5 Self-Organising Networks

In the previous chapters we discussed a number of networks which were trained to perform a mapping $\mathcal{F} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ by presenting the network 'examples' $(\mathbf{x}^p, \mathbf{o}^p)$ with $\mathbf{o}^p = \mathcal{F}(\mathbf{x}^p)$ of this mapping. However, problems exist where such training data, consisting of input and desired output pairs are not available, but where the only information is provided by a set of input patterns $\mathbf{x}^p$. In these cases the relevant information has to be found within the (redundant) training samples $\mathbf{x}^p$.

Some examples of such problems are:

- clustering: the input data may be grouped in 'clusters' and the data processing system has to find these inherent clusters in the input data. The output of the system gives the cluster label of the input pattern (discrete output);
- vector quantisation: this problem occurs when a continuous space has to be discretised. The input of the system is the n-dimensional vector $\mathbf{x}$, the output is a discrete representation of the input space. The system has to find optimal discretisation of the input space;
- dimensionality reduction: the input data are grouped in a subspace which has lower dimensionality than the dimensionality of the data. The system has to learn an optimal mapping, such that most of the variance in the input data is preserved in the output data;
- feature extraction: the system has to extract features from the input signal. This often means a dimensionality reduction as described above.

In this chapter we discuss a number of neuro-computational approaches for these kinds of problems. Training is done without the presence of an external teacher. The unsupervised weight adapting algorithms are usually based on some form of global competition between the neurons.

There are very many types of self-organising networks, applicable to a wide area of problems. One of the most basic schemes is *competitive learning* as proposed by Rumelhart and Zipser [1]. A very similar network but with different emergent properties is the topology-conserving map devised by Kohonen. Other self-organising networks are ART, proposed by Carpenter and Grossberg [2,3], and Fukushima's cognitron [4,5].

This chapter is concluded with a description of the reinforcement learning procedure. This form of learning is in fact a supervised learning method, but does not need a set of 'correct' or 'target' output values for each input pattern as was the case for the supervised networks in the previous chapters.

5.1 Competitive learning

5.1.1 Clustering

Competitive learning is a learning procedure that divides a set of input patterns in clusters that are inherent to the input data. A competitive learning network is provided only with input vectors $\mathbf{x}$ and thus implements an *unsupervised* learning procedure. We will show its equivalence to a class of 'traditional' clustering algorithms shortly. Another important use of these networks is vector quantisation, as discussed in section 5.1.2.

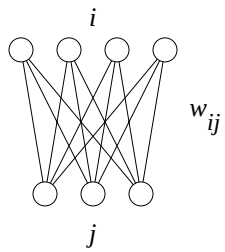

Figure 5.1: A simple competitive learning network. Each of the four outputs i is connected to all inputs j.

An example of a competitive learning network is shown in figure 5.1. All output units i are connected to all input units j with weights w_{ij}. The weight vectors are initialised to random values. For the time being, we assume that both input vectors $\mathbf{x}$ and weight vectors $\mathbf{w}_i$ are normalised to unit length. Each output unit i calculates its activation value a_i according to

$$a_i = \sum_j w_{ij} x_j = {\mathbf{w}_i}^T \mathbf{x}. \tag{5.1}$$

In a next pass, output neuron k is selected with maximum activation

$$\forall i \neq k: \qquad a_i \leq a_k. \tag{5.2}$$

Activations are reset such that $a_k = 1$ and $a_i = 0$, $i \neq k$. This is is the *competitive* aspect of the network, and we refer to the output layer as the *winner-take-all* layer. The winner-take-all layer is usually implemented in software by simply selecting the output neuron with highest activation value. This function can also be performed by a neural network known as MAXNET [6]. In MAXNET, all neurons i are connected to other units i' with inhibitory links and to itself with an excitatory link:

$$w_{i,i'} = \begin{cases} -\epsilon & \text{if } i \neq i' \\ +1 & \text{otherwise.} \end{cases} \tag{5.3}$$

It can be shown that this network converges to a situation where only the neuron with highest initial activation survives, whereas the activations of all other neurons converge to zero. From now on, we will simply assume a winner k is selected without being concerned which algorithm is used.

Once the winner k has been selected, the weights are updated according to:

$$\mathbf{w}_k(t+1) = \frac{\mathbf{w}_k(t) + \gamma(\mathbf{x}(t) - \mathbf{w}_k(t))}{\|\mathbf{w}_k(t) + \gamma(\mathbf{x}(t) - \mathbf{w}_k(t))\|} \tag{5.4}$$

where the divisor ensures that all weight vectors $\mathbf{w}$ are normalised. Note that only the weights of winner k are updated.

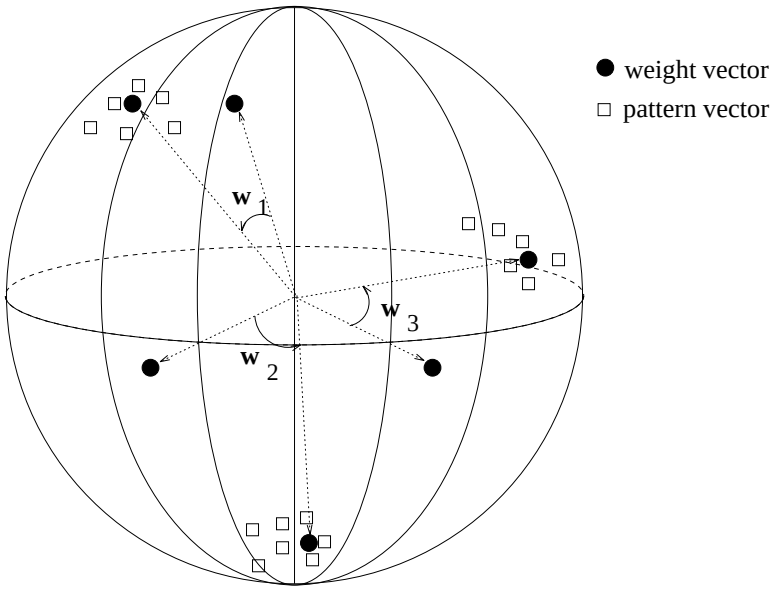

Figure 5.2: Example of clustering in 3D with normalised vectors, which all lie on the unity sphere. The three weight vectors are rotated towards the centres of gravity of the three different input clusters.

The weight update given in equation (5.4) effectively rotates the weight vector $\mathbf{w}_i$ towards the input vector $\mathbf{x}$. Each time an input $\mathbf{x}$ is presented, the weight vector closest to this input is selected and is subsequently rotated towards the input. Consequently, weight vectors are rotated towards those areas where many inputs appear: the clusters in the input. This procedure is visualised in figure 5.2.

Previously it was assumed that both inputs $\mathbf{x}$ and weight vectors $\mathbf{w}$ were normalised. In figure 5.3 it is shown how the algorithm would fail if unnormalised vectors were to be used. Naturally one would like to accommodate the algorithm for unnormalised input data. To this end, the winning neuron k is selected with its weight vector $\mathbf{w}_k$ *closest* to the input pattern $\mathbf{x}$, using the Euclidean distance measure:

$$k : \|\mathbf{w}_k - \mathbf{x}\| \leq \|\mathbf{w}_j - \mathbf{x}\| \quad \forall j. \tag{5.5}$$

It is easily checked that equation (5.5) reduces to (5.1)/(5.2) if all vectors are normalised. The Euclidean distance norm is therefore a more general case of equations (5.1)/(5.2). Instead of *rotating* the weight vector towards the input as performed by equation (5.4), the weight update must be changed to implement a *shift* towards the input:

$$\mathbf{w}_k(t+1) = \mathbf{w}_k(t) + \gamma(\mathbf{x}(t) - \mathbf{w}_k(t)). \tag{5.6}$$

Both the Euclidean distance norm and the use of normalised weight vectors are common in neural networks. In other clustering applications, especially in psychology, much more attention has been paid to the choice of an appropriate distance function on the input space

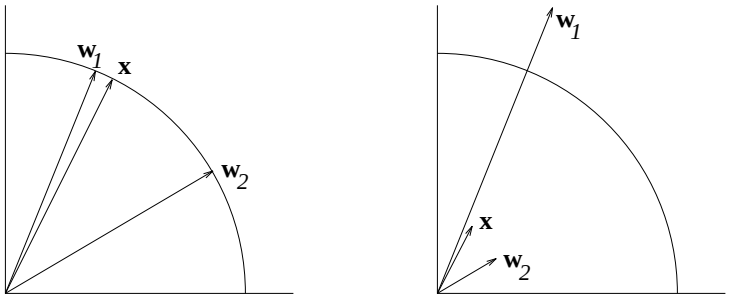

Figure 5.3: Determining the winner in a competitive learning network. a. Three normalised vectors. b. The three vectors having the same directions as in a., but with different lengths. In a., vectors $\mathbf{x}$ and $\mathbf{w}_1$ are nearest to each other, and their dot product $\mathbf{x}^T\mathbf{w}_1 = |\mathbf{x}||\mathbf{w}_1|\cos\alpha$ is larger than the dot product of $\mathbf{x}$ and $\mathbf{w}_2$. In b., however, the pattern and weight vectors are not normalised, and in this case $\mathbf{w}_2$ should be considered the 'winner' when $\mathbf{x}$ is applied. However, the dot product $\mathbf{x}^T\mathbf{w}_1$ is still larger than $\mathbf{x}^T\mathbf{w}_2$.

of patterns $\mathbf{x}$ in close harmony with the choice of appropriate input vectors $\mathbf{x}$. Different choices obviously lead to different clusterings of the input data and it is the psychologist's insight in the clustering process that motivates the choice of particular input vectors and distance measures. This is only to say, that competitive learning networks could also be fed with these other distance functions by changing the distance norm if need arises, though it is not common practice at this time.

Especially if the input vectors are drawn from a large or high-dimensional input space, it is not beyond imagination that a randomly initialised weight vector $\mathbf{w}_i$ will never be chosen as the winner and will thus never be moved and never be used. Therefore, it is customary to initialise weight vectors to a set of input patterns $\{\mathbf{x}\}$ drawn from the input set at random. Another more thorough approach that avoids these and other problems in competitive learning is called *leaky learning*. This is implemented by expanding the weight update given in equation (5.6) with

$$\mathbf{w}_l(t+1) = \mathbf{w}_l(t) + \gamma'(\mathbf{x}(t) - \mathbf{w}_l(t)) \qquad \forall l \neq k \tag{5.7}$$

with $\gamma' \ll \gamma$ the *leaky learning rate*. Recently a somewhat similar method emerges called *frequency sensitive competitive learning* [7]. In this algorithm, each neuron records the number of times it is selected winner. The more often it wins, the less sensitive it becomes to competition. Conversely, neurons that consistently fail to win increase their chances of being selected winner.

Earlier it was claimed, that a competitive network performs a clustering process on the input data. I.e., input patterns are divided in disjoint clusters such that similarities between input patterns in the same cluster are much bigger than similarities between inputs in different clusters. Similarity is measured by a distance function on the input vectors, as discussed before. A common criterion to measure the quality of a given clustering is the *square error criterion*, given by

$$E = \sum_p \|\mathbf{w}_k - \mathbf{x}^p\|^2 \tag{5.8}$$

where k is the winning neuron when input $\mathbf{x}^p$ is presented. The weights $\mathbf{w}$ are interpreted as *cluster centres*. It is not difficult to show that competitive learning indeed seeks to find a minimum for this square error by following the negative gradient of the error-function:

Theorem 2 *The error function for pattern* $\mathbf{x}^p$

$$E^p = \tfrac{1}{2}\sum_j (w_{kj} - x_j^p)^2, \tag{5.9}$$

where k is the winning unit, is minimised by the weight update rule in eq. (5.6).

Proof *As in eq. (3.13), we calculate the effect of a weight change on the error function. So we have that*

$$\Delta_p w_{ij} = -\gamma \frac{\partial E^p}{\partial w_{ij}} \tag{5.10}$$

where γ is a constant of proportionality. Now, we have to determine the partial derivative of E^p:

$$\frac{\partial E^p}{\partial w_{ij}} = \begin{cases} w_{ij} - x_j^p & \text{if unit } i \text{ wins} \\ 0 & \text{otherwise} \end{cases} \tag{5.11}$$

such that

$$\Delta_p w_{ij} = -\gamma(w_{ij} - x_j^p) = \gamma(x_j^p - w_{ij}) \tag{5.12}$$

which is eq. (5.6) written down for one element of $\mathbf{w}_i$.

Therefore, eq. (5.8) is minimised by repeated weight updates using eq. (5.6).

□

An almost identical process of moving cluster centres is used in a large family of conventional clustering algorithms known as *square error clustering methods*, e.g., k-means, FORGY, ISODATA, CLUSTER.

5.1.2 Vector quantisation

Another important use of competitive learning networks is found in *vector quantisation*. A vector quantisation scheme divides the input space in a number of disjoint subspaces and represents each input vector $\mathbf{x}$ by the label of the subspace it falls into (i.e., index k of the winning neuron). The difference with clustering is that we are not so much interested in finding clusters of similar data, but more in quantising the entire input space. The quantisation performed by the competitive learning network is said to 'track the input probability density function:' the density of neurons and thus subspaces is highest in those areas where inputs are most likely to appear, whereas a more coarse quantisation is obtained in those areas where inputs are scarce. An example of tracking the input density is sketched in figure 5.4. Vector quantisation through competitive learning results in a more *fine-grained* discretisation in those areas of the input space where most input have occurred in the past.

5.1.3 Using vector quantisation

In a large number of applications, networks that perform vector quantisation are combined with another type of network in order to perform function approximation. An example of such a network is given in figure 5.5. This network can approximate a function

$$\mathcal{F} : \mathbb{R}^n \rightarrow \mathbb{R}^m$$

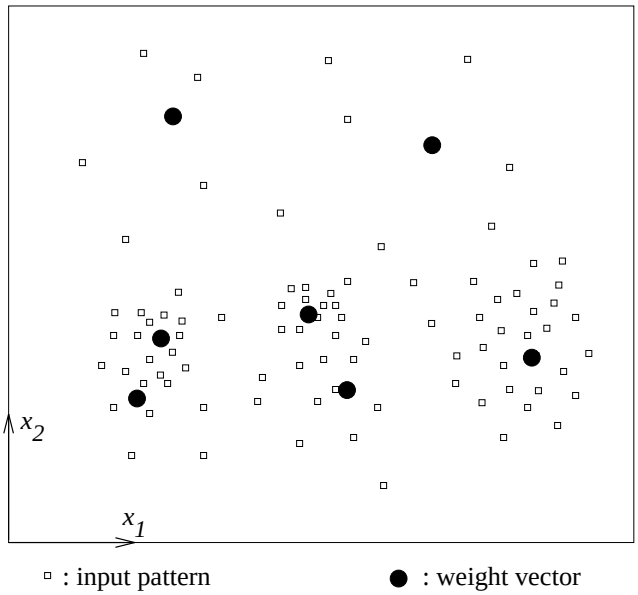

Figure 5.4: This figure visualises the tracking of the input density. The input patterns are drawn from $\mathbb{R}^2$; the weight vectors also lie in $\mathbb{R}^2$. In the areas where inputs are scarce, the upper part of the figure, only few (in this case two) neurons are used to discretise the input space. Thus, the upper part of the input space is divided into two large separate regions. The lower part, however, where many more inputs have occurred, five neurons discretise the input space into five smaller subspaces.

by associating with each neuron i a function value $[w_{1i} \ldots w_{mi}]^T$ which is somehow representative for the function values $\mathcal{F}(\mathbf{x})$ of inputs $\mathbf{x}$ represented by i. This way of approximating a function effectively implements a 'look-up table:' an input $\mathbf{x}$ is assigned to a table entry k with $\|\mathbf{x} - \mathbf{w}_k\| \leq \|\mathbf{x} - \mathbf{w}_i\| \quad \forall i \neq k$, and the function value $[w_{1k} \ldots w_{mk}]^T$ in this table entry is taken as an approximation of $\mathcal{F}(\mathbf{x})$. A well-known example of such a network is the *Counterpropagation* network[8].

Depending on the application, one can choose to perform the vector quantisation before learning the function approximation, or one can choose to learn the quantisation and the approximation layer simultaneously. As an example of the latter, the network presented in figure 5.5 can be supervisedly trained in the following way:

1. present the network with both input $\mathbf{x}$ and function value $\mathbf{d} = \mathcal{F}(\mathbf{x})$;

2. perform the *unsupervised* quantisation step. For each weight vector, calculate the distance from its weight vector to the input pattern and find winner k. Update the weights w_{ij} with equation (5.6);

3. perform the *supervised* approximation step:

$$w_{hk}(t+1) = w_{hk}(t) + \gamma(o_h - w_{hk}(t)). \tag{5.13}$$

   This is simply the δ-rule with $o_h = \sum_i a_i w_{hi} = w_{hk}$ when k is the winning neuron and the desired output is given by $\mathbf{d} = \mathcal{F}(\mathbf{x})$.

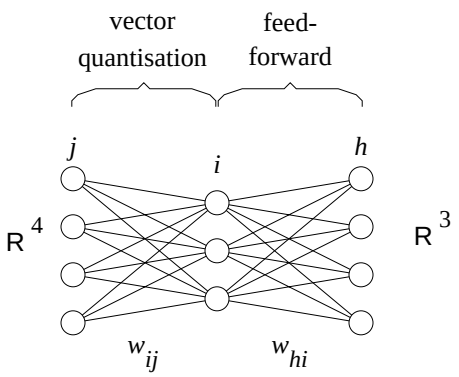

Figure 5.5: A network combining a vector quantisation layer with a 1-layer feed-forward neural network. This network can be used to approximate functions from $\mathbb{R}^4$ to $\mathbb{R}^3$, the input space $\mathbb{R}^4$ is discretised in two disjoint subspaces.

If we define a function $g(\mathbf{x}, k)$ as :

$$g(\mathbf{x}, k) = \begin{cases} 1 & \text{if } k \text{ is winner} \\ 0 & \text{otherwise} \end{cases} \tag{5.14}$$

it can be shown that this learning procedure converges to

$$w_{hi} = \int_{\mathbb{R}^n} o_h g(\mathbf{x}, i) \quad d\mathbf{x}. \tag{5.15}$$

I.e., each table entry converges to the mean function value over all inputs in the subspace represented by that table entry. As we have seen before, the quantisation scheme tracks the input probability density function, which results in a better approximation of the function in those areas where input is most likely to appear.

Not all functions are represented accurately by this combination of quantisation and approximation layers. E.g., a simple identity or combinations of sines and cosines are much better approximated by multilayer back-propagation networks if the activation functions are chosen appropriately. However, if we expect our input to be (a subspace of) a high dimensional input space $\mathbb{R}^n$ and we expect our function $\mathcal{F}$ to be discontinuous at numerous points, the combination of quantisation and approximation is not uncommon and probably very efficient. Of course this combination extends itself much further than the presented combination of the presented single layer competitive learning network and the single layer feed-forward network. The latter could be replaced by a reinforcement learning procedure, which is illustrated in section 5.5. The quantisation layer can be replaced by various other quantisation schemes, such as Kohonen networks (see section 5.2) or octree methods [9]. In fact, various modern statistical function approximation methods (CART, MARS, [10,11]) are based on this very idea, extended with the possibility to have the approximation layer influence the quantisation layer (e.g., to obtain a better or locally more fine-grained quantisation). Recent research [12] also investigates in this direction.

It is an unpleasant habit in neural network literature, to also cover *Learning Vector Quantisation* (LVQ) methods in chapters on unsupervised clustering. Granted that these methods also perform a clustering or quantisation task and use similar learning rules, they are trained *supervisedly* and perform *discriminant analysis* rather than unsupervised clustering. These networks attempt to define 'decision boundaries' in the input space, given a large set of exemplary decisions (the training set); each decision could, e.g., be a correct class label.

A rather large number of slightly different LVQ methods is appearing in recent literature. They are all based on the following basic algorithm:

1. with each output neuron i, a class label (or decision of some other kind) o_i is associated;
2. a learning sample consists of input vector $\mathbf{x}^p$ together with its correct class label o_i^p;
3. using distance measures between weight vectors $\mathbf{w}_i$ and input vector $\mathbf{x}^p$, not only the winner k_1 is determined, but also the second best k_2:

$$\|\mathbf{x}^p - \mathbf{w}_{k_1}\| < \|\mathbf{x}^p - \mathbf{w}_{k_2}\| < \|\mathbf{x}^p - \mathbf{w}_i\| \qquad \forall i \neq k_1, k_2;$$

4. the labels $o_{k_1}^p$, $o_{k_2}^p$ are compared with d^p. The weight update rule given in equation (5.6) is used *selectively* based on this comparison.

An example of the last step is given by the LVQ2 algorithm by Kohonen [13], using the following strategy:

- if $o_{k_1}^p \neq d^p$ and $d^p = o_{k_2}^p$;
- and $\|\mathbf{x}^p - \mathbf{w}_{k_2}\| - \|\mathbf{x}^p - \mathbf{w}_{k_1}\| < \epsilon$;
- then $\mathbf{w}_{k_2}(t+1) = \mathbf{w}_{k_2} + \gamma(\mathbf{x} - \mathbf{w}_{k_2}(t))$
- and $\mathbf{w}_{k_1}(t+1) = \mathbf{w}_{k_1}(t) - \gamma(\mathbf{x} - \mathbf{w}_{k_1}(t))$

I.e., $\mathbf{w}_{k_2}$ with the correct label is moved *towards* the input vector, while $\mathbf{w}_{k_1}$ with the incorrect label is moved *away* from it.

The new LVQ algorithms that are emerging all use different implementations of these different steps, e.g., how to define class labels o_i, how many 'next-best' winners are to be determined, how to adapt the number of output neurons i and how to selectively use the weight update rule.

5.2 Kohonen network

The Kohonen network [14,15] can be seen as an extension to the competitive learning network, although this is chronologically incorrect. Also, the Kohonen network has a different set of applications.

In the Kohonen network, the output units in S are *ordered* in some fashion, often in a two-dimensional grid or array, although this is application-dependent. The ordering, which is chosen by the user[1], determines which output neurons are neighbours.

Now, when learning patterns are presented to the network, the weights to the output units are thus adapted such that the order present in the input space $\mathbb{R}^N$ is preserved in the output,

[1] Of course, variants have been designed which automatically generate the structure of the network [16,17].

i.e., the neurons in S. This means that learning patterns which are near to each other in the input space (where 'near' is determined by the distance measure used in finding the winning unit) must be mapped on output units which are also near to each other, i.e., the same or neighbouring units. Thus, if inputs are uniformly distributed in $\mathbb{R}^N$ and the order must be preserved, the dimensionality of S must be at least N. The mapping, which represents a discretisation of the input space, is said to be *topology preserving*. However, if the inputs are restricted to a subspace of $\mathbb{R}^N$, a Kohonen network can be used of lower dimensionality. For example: data on a two-dimensional manifold in a high dimensional input space can be mapped onto a two-dimensional Kohonen network, which can for example be used for visualisation of the data.

Usually, the learning patterns are random samples from $\mathbb{R}^N$. At time t, a sample $\mathbf{x}(t)$ is generated and presented to the network. Using the same formulas as in section 5.1, the winning unit k is determined. Next, the weights to this winning unit *as well as its neighbours* are adapted using the learning rule

$$\mathbf{w}_i(t+1) = \mathbf{w}_i(t) + \gamma h(i,k)\,(\mathbf{x}(t) - \mathbf{w}_i(t)) \quad \forall i \in S. \tag{5.16}$$

Here, $h(i,k)$ is a decreasing function of the grid-distance between units i and k, such that $h(k,k) = 1$. For example, for $h()$ a Gaussian function can be used, such that (in one dimension!) $h(i,k) = \exp\left(-(i-k)^2\right)$ (see figure 5.6). Due to this collective learning scheme,

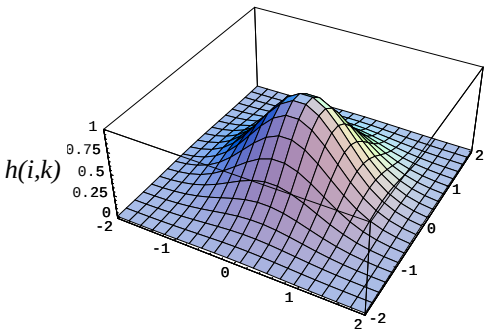

Figure 5.6: Gaussian neuron distance function $h()$. In this case, $h()$ is shown for a two-dimensional grid because it looks nice.

input signals which are near to each other will be mapped on neighbouring neurons. Thus the topology inherently present in the input signals will be preserved in the mapping, such as depicted in figure 5.7.

If the dimensionality of S is less than N, the neurons in the network are 'folded' in the input space, such as depicted in figure 5.8.

The topology-conserving quality of this network has many counterparts in biological brains. The brain is organised in many places so that aspects of the sensory environment are represented in the form of two-dimensional maps. For example, in the visual system, there are several topographic mappings of visual space onto the surface of the visual cortex. There are organised mappings of the body surface onto the cortex in both motor and somatosensory areas, and tonotopic mappings of frequency in the auditory cortex. The use of

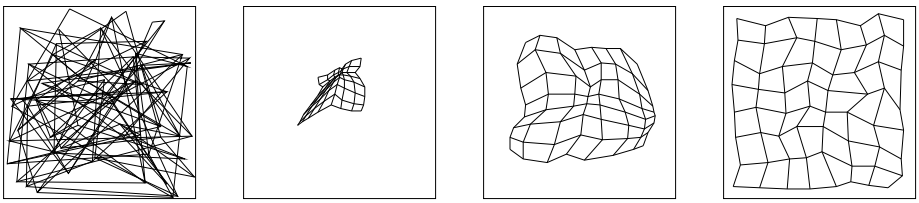

Figure 5.7: A topology-conserving map converging. The weight vectors of a network with two inputs and 8×8 output neurons arranged in a planar grid are shown. A line in each figure connects weight $w_{(i_1,i_2),j}$ with weights $w_{(i_1+1,i_2),j}$ and $w_{(i_1,i_2+1),j}$. The leftmost figure shows the initial weights; the rightmost when the map is almost completely formed.

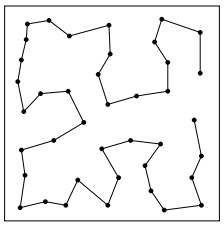

Figure 5.8: The mapping of a two-dimensional input space on a one-dimensional Kohonen network.

topographic representations, where some important aspect of a sensory modality is related to the physical locations of the cells on a surface, is so common that it obviously serves an important information processing function.

It does not come as a surprise, therefore, that already many applications have been devised of the Kohonen topology-conserving maps. Kohonen himself has successfully used the network for phoneme-recognition [18]. Also, the network has been used to merge sensory data from different kinds of sensors, such as auditory and visual, 'looking' at the same scene [19]. Yet another application is in robotics, such as shown in section 7.1.1.

To explain the plausibility of a similar structure in biological networks, Kohonen remarks that the lateral inhibition between the neurons could be obtained via efferent connections between those neurons. In one dimension, those connection strengths form a 'mexican hat' (see figure 5.9).

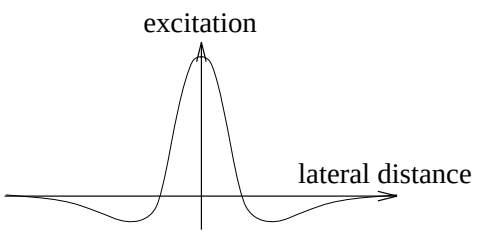

Figure 5.9: Mexican hat. Lateral interaction around the winning neuron as a function of distance: excitation to nearby neurons, inhibition to farther off neurons.

5.3 Principal component networks

5.3.1 Introduction

The networks presented in the previous sections, can be seen as (non linear) vector transformations which map an input vector to a number of binary output elements or neurons. The weights are adjusted in such a way that they could be considered as prototype vectors (vectorial means) for the input patterns for which the competing neuron wins.

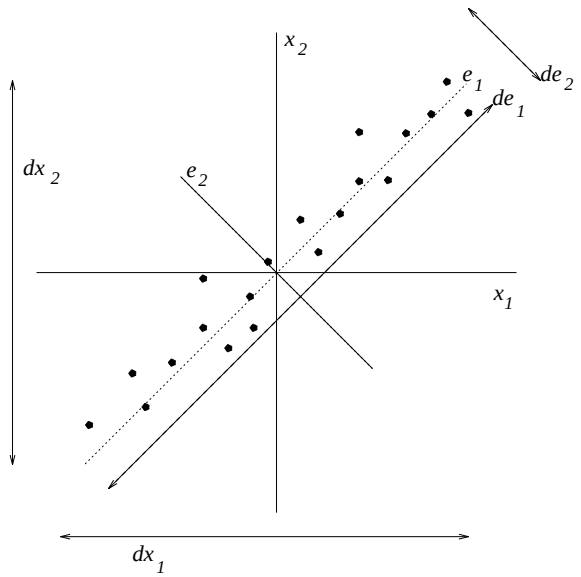

Figure 5.10: Distribution of input samples.

The self-organising transform described in this section rotates the input space in such a way that the values of the output neurons are as uncorrelated as possible and the energy or variances of the patterns is mainly concentrated in a few output neurons. An example is shown in figure 5.10. The two dimensional samples (x_1, x_2) are plotted in the figure. It can be easely seen that x_1 and x_2 are related, such that if we know x_1 we can make a reasonable prediction of x_2 and vica versa since the points are centered around the line $x_1 = x_2$. If we rotate the axis over $\pi/4$ we get the (e_1, e_2) axis as plotted in the figure. Here the conditional prediction has no use because the points have uncorrelated coordinates. Another property of this rotation is that the variance or energy of the transformed patterns is maximised on a lower dimension. This can be intuitively verified by comparing the spreads (d_{x_1}, d_{x_2}) and (d_{e_1}, d_{e_2}) in the figures. After the rotation, the variance of the samples is large along the e_1 axis and small along the e_2 axis.

This transform is very closely related to the eigenvector transformation known from image processing where the image has to be coded or transformed to a lower dimension and reconstructed again by another transform as well as possible (see section 8.3.2).

The next section describes a learning rule which acts as a Hebbian learning rule, but which scales the vector length to unity. In the subsequent section we will see that a linear neuron with a normalised Hebbian learning rule acts as such a transform, extending the theory in the last section to multi-dimensional outputs.

5.3.2 Normalised Hebbian rule

The model considered here consists of one linear(!) neuron with input weights $\mathbf{w}$. The output $o(t)$ of this neuron is given by the usual inner product of its weight $\mathbf{w}$ and the input vector $\mathbf{x}$:

$$o(t) = \mathbf{w}(t)^T \mathbf{x}(t) \tag{5.17}$$

As seen in the previous sections, all models are based on a kind of Hebbian learning. However, the basic Hebbian rule would make the weights grow enormous if there were correlation in the input patterns. This can be overcome by normalising the weight vector to fixed length, typically $\mathbf{1}$, which leads to the following learning rule

$$\mathbf{w}(t+1) = \frac{\mathbf{w}(t) + \gamma o(t)\mathbf{x}(t)}{L\left(\mathbf{w}(t) + \gamma o(t)\mathbf{x}(t)\right)} \tag{5.18}$$

where $L(\cdot)$ denotes a operator which returns the vector length, and γ is a small learning parameter. Compare this learning rule with the normalised learning rule of competitive learning. There the delta rule was normalised, here the standard Hebb rule!

Now the operator which computes the vector length, the norm of the vector, can be approximated by a Taylor expansion around $\gamma = 0$:

$$L\left(\mathbf{w}(t) + \gamma o(t)\mathbf{x}(t)\right) = 1 + \gamma \left.\frac{\partial L}{\partial \gamma}\right|_{\gamma=0} + O(\gamma^2). \tag{5.19}$$

If we substitute this expression for the vector length in equation (5.18), it resolves for small γ to[2]

$$\mathbf{w}(t+1) = \left(\mathbf{w}(t) + \gamma o(t)\mathbf{x}(t)\right)\left(1 - \gamma \left.\frac{\partial L}{\partial \gamma}\right|_{\gamma=0} + O(\gamma^2)\right). \tag{5.20}$$

Since $\left.\frac{\delta L}{\delta \gamma}\right|_{\gamma=0} = o(t)^2$, discarding the higher order terms of γ leads to

$$\mathbf{w}(t+1) = \mathbf{w}(t) + \gamma o(t)\left(\mathbf{x}(t) - o(t)\mathbf{w}(t)\right) \tag{5.21}$$

which is called the 'Oja learning rule' [20]. This learning rule thus modifies the weight in the usual Hebbian sense, the first product terms *is* the Hebb rule $o(t)\mathbf{x}(t)$, but normalises its weight vector directly by the second product term $-o(t)o(t)\mathbf{w}(t)$. What exactly do you expect this learning rule does with the weight vector?

5.3.3 Principal component extractor

Firstly we will recall some old stuff from stochastic theory. Consider an N-dimensional signal $\mathbf{x}(t)$ with

- mean: $\mu = E(\mathbf{x}(t))$.
- correlation matrix: $\mathbf{R} = E((\mathbf{x}(t) - \mu)(\mathbf{x}(t) - \mu)^T)$.

[2] Remembering that $1/(1 + a\gamma) = 1 - a\gamma + O(\gamma^2)$.

In the following we assume the signal mean to be zero, so $\mu = \mathbf{0}$.

From equation (5.21) we see that the expectation of the weights for the Oja learning rule equals

$$E(\mathbf{w}(t+1)|\mathbf{w}(t)) = \mathbf{w}(t) + \gamma \left(\mathbf{R}\mathbf{w}(t) - \left(\mathbf{w}(t)^T \mathbf{R}\mathbf{w}(t)\right) \mathbf{w}(t)\right) \tag{5.22}$$

which has a continuous counterpart as

$$\frac{d}{dt}\mathbf{w}(t) = \mathbf{R}\mathbf{w}(t) - \left(\mathbf{w}(t)^T \mathbf{R}\mathbf{w}(t)\right) \mathbf{w}(t). \tag{5.23}$$

Theorem 1 *Let the eigenvectors* $\mathbf{e}_i$ *of* $\mathbf{R}$ *be ordered with descending associated eigenvalues* λ_i *such that* $\lambda_1 > \lambda_2 > \ldots > \lambda_N$. *With equation (5.23) the weights* $\mathbf{w}(t)$ *will converge to* $\pm\mathbf{e}_1$.

Proof 1 *Since the eigenvectors of* $\mathbf{R}$ *span the N dimensional space, the weight vector can be decomposed as*

$$\mathbf{w}(t) = \sum_i^N \beta_i(t)\mathbf{e}_i. \tag{5.24}$$

Substituting this in the differential equation and concluding the theorem is left as an exercise.

5.3.4 More eigenvectors

In the previous section it was shown that a single neuron's weight converges to the eigenvector of the correlation matrix with maximum eigenvalue, i.e, the weight of the neuron is directed in the direction of highest energy or variance of the input patterns. Here we tackle the question of how to find the remaining eigenvectors of the correlation matrix given the first found eigenvector.

Consider the signal $\mathbf{x}$ which can be decomposed into the basis of eigenvectors $\mathbf{e}_i$ of its correlation matrix $\mathbf{R}$,

$$\mathbf{x} = \sum_i^N \alpha_i \mathbf{e}_i \tag{5.25}$$

If we now subtract the component in the direction of $\mathbf{e}_1$, the direction in which the signal has the most energy, from the signal $\mathbf{x}$

$$\tilde{\mathbf{x}} = \mathbf{x} - \alpha_1 \mathbf{e}_1 \tag{5.26}$$

we are sure that if we decompose $\tilde{\mathbf{x}}$ again into the eigenvector basis, that the coefficient $\alpha_1 = 0$, simply because we just subtracted it. We call $\tilde{\mathbf{x}}$ the *deflation* of $\mathbf{x}$.

If now a second neuron is taught on this signal $\tilde{\mathbf{x}}$, then its weights will be in the direction of the remaining eigenvector with the highest eigenvalue. Since the deflation removed the component in the direction of the first eigenvector, the weight will converge to the remaining eigenvector with maximum eigenvalue. In the previous section we ordered the eigenvalues in magnitude, so according to this definition in the limit we will find $\mathbf{e}_2$. We can continue this strategy and find all the N eigenvectors belonging to the signal $\mathbf{x}$.

We can write the deflation in neural network terms if we see that

$$o = \mathbf{w}^T\mathbf{x} = \mathbf{e}_1^T \sum_i^N \alpha_i \mathbf{e}_i = \alpha_i \tag{5.27}$$

since

$$\mathbf{w} = \mathbf{e}_1 \tag{5.28}$$

So that the deflated vector $\tilde{\mathbf{x}}$ equals

$$\tilde{\mathbf{x}} = \mathbf{x} - o\mathbf{w} \tag{5.29}$$

The term subtracted from the input vector can be interpreted as a kind of a back-projection or expectation. Compare this to the ART described in the next section .

5.4 Adaptive resonance theory

The last unsupervised learning network we discuss differs from the previous networks in that it is recurrent; as with networks in the next chapter, the data is not only fed forward but also back from output to input units.

5.4.1 Background: Adaptive resonance theory

In 1976, Grossberg [3] introduced a model for explaining biological phenomena. The model has three crucial properties:

1. a *normalisation* of the total network activity. Biological systems are usually very adaptive to large changes in their environment. For example, the human eye can adapt itself to large variations in light intensities;

2. *contrast enhancement* of input patterns. The awareness of subtle differences in input patterns can mean a lot in terms of survival. Distinguishing a hiding panther from a resting one makes all the difference in the world. The mechanism used here is contrast enhancement;

3. *short-term memory (STM) storage* of the contrast-enhanced pattern. Before the input pattern can be decoded, it must be stored in the short-term memory. The long-term memory (LTM) implements an arousal mechanism (i.e., the classification), whereas the STM is used to cause gradual changes in the LTM.

The system consists of two layers, $F1$ and $F2$, which are connected to each other via the LTM (see figure 5.11). The input pattern is received at $F1$, whereas classification takes place in $F2$. As mentioned before, the input is not directly classified. First a characterisation takes place by means of extracting features, giving rise to activation in the feature representation field. The expectations, residing in the LTM connections, translate the input pattern to a categorisation in the category representation field. The classification is compared to the expectation of the network, which resides in the LTM weights from $F2$ to $F1$. If there is a match, the expectations are strengthened, otherwise the classification is rejected.

5.4.2 ART1: The simplified neural network model

The ART1 simplified model consists of two layers of binary neurons (with values 1 and 0), called $F1$ (the comparison layer) and $F2$ (the recognition layer) (see figure 5.12). Each neuron in $F1$ is connected to all neurons in $F2$ via the continuous-valued forward long term memory

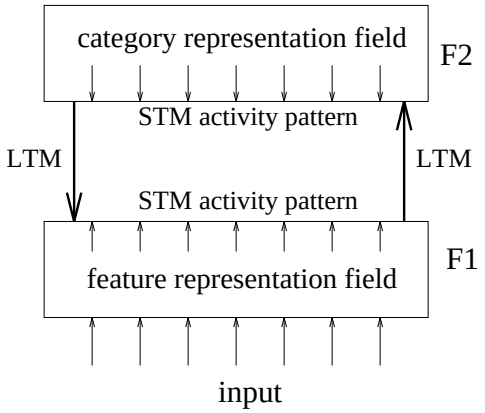

Figure 5.11: The ART architecture.

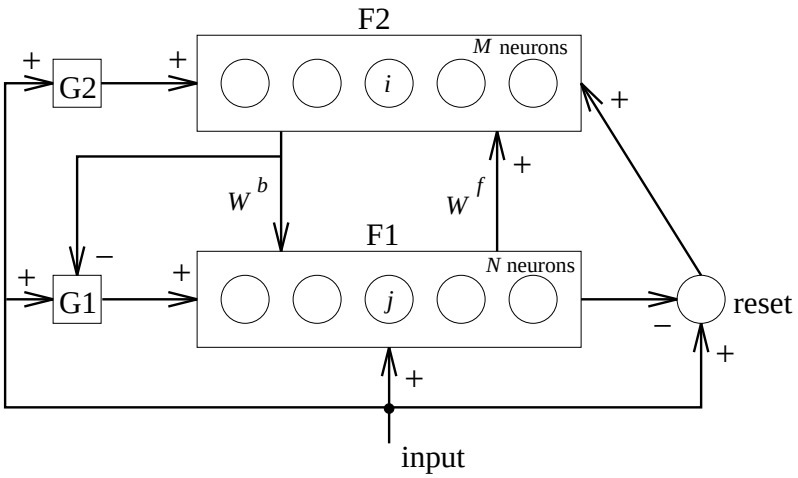

Figure 5.12: The ART1 neural network.

(LTM) W^f, and vice versa via the binary-valued backward LTM W^b. The other modules are gain 1 and 2 ($G1$ and $G2$), and a reset module.

Each neuron in the comparison layer receives three inputs: a component of the input pattern, a component of the feedback pattern, and a gain $G1$. A neuron outputs a 1 if and only if at least three of these inputs are high: the 'two-thirds rule.'

The neurons in the recognition layer each compute the inner product of their incoming (continuous-valued) weights and the pattern sent over these connections. The winning neuron then inhibits all the other neurons via lateral inhibition.

Gain 2 is the logical 'or' of all the elements in the input pattern $\mathbf{x}$.

Gain 1 equals gain 2, except when the feedback pattern from $F2$ contains any 1; then it is forced to zero.

Finally, the reset signal is sent to the active neuron in $F2$ if the input vector $\mathbf{x}$ and the output of $F1$ differ by more than some vigilance level.

Operation

The network starts by clamping the input at $F1$. Because the output of $F2$ is zero, $G1$ and $G2$ are both on and the output of $F1$ matches its input.

The pattern is sent to $F2$, and in $F2$ one neuron becomes active. This signal is then sent back over the backward LTM, which reproduces a binary pattern at $F1$. Gain 1 is inhibited, and only the neurons in $F1$ which receive a 'one' from both $\mathbf{x}$ and $F2$ remain active.

If there is a substantial mismatch between the two patterns, the reset signal will inhibit the neuron in $F2$ and the process is repeated.

Instead of following Carpenter and Grossberg's description of the system using differential equations, we use the notation employed by Lippmann [22]:

1. Initialisation:

$$w_{ji}^b(0) = 1$$
$$w_{ij}^f(0) = \frac{1}{1+N}$$

   where N is the number of neurons in $F1$, M the number of neurons in $F2$, $0 \leq j < N$, and $0 \leq i < M$. Also, choose the vigilance threshold ρ, $0 \leq \rho \leq 1$;

2. Apply the new input pattern $\mathbf{x}$;

3. compute the activation values $\mathbf{a}'$ of the neurons in $F2$:

$$a_i{}' = \sum_{j=0}^{N-1} w_{ij}^f(t)\, x_j; \tag{5.30}$$

4. select the winning neuron k $(0 \leq k < M)$;

5. vigilance test: if

$$\frac{\mathbf{w}_k^b(t) \cdot \mathbf{x}}{\mathbf{x} \cdot \mathbf{x}} > \rho, \tag{5.31}$$

   where $\cdot$ denotes inner product, go to step 7, else go to step 6. Note that $\mathbf{w}_k^b \cdot \mathbf{x}$ essentially is the inner product $\mathbf{x}^* \cdot \mathbf{x}$, which will be large if $\mathbf{x}^*$ and $\mathbf{x}$ are near to each other;

6. neuron k is disabled from further activity. Go to step 3;

7. Set for all l, $0 \leq l < N$:

$$w_{lk}^b(t+1) = w_{lk}^b(t)\, x_l,$$
$$w_{kl}^f(t+1) = \frac{w_{kl}^b(t)\, x_l}{\frac{1}{2} + \sum_{j=0}^{N-1} w_{jk}^b(t)\, x_j};$$

8. re-enable all neurons in $F2$ and go to step 2.

Figure 5.13 shows exemplar behaviour of the network.

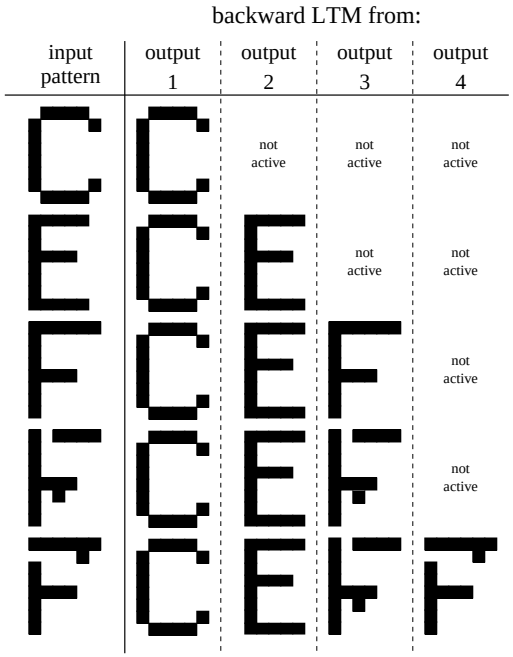

Figure 5.13: An example of the behaviour of the Carpenter Grossberg network for letter patterns. The binary input patterns on the left were applied sequentially. On the right the stored patterns (i.e., the weights of W^b for the first four output units) are shown.

5.4.3 ART1: The original model

In later work, Carpenter and Grossberg [2,23] present several neural network models to incorporate parts of the complete theory. We will only discuss the first model, ART1.

The network incorporates a follow-the-leader clustering algorithm [24]. This algorithm tries to fit each new input pattern in an existing class. If no matching class can be found, i.e., the distance between the new pattern and all existing classes exceeds some threshold, a new class is created containing the new pattern.

The novelty in this approach is that the network is able to adapt to new incoming patterns, while the previous memory is not corrupted. In most neural networks, such as the back-propagation network, all patterns must be taught sequentially; the teaching of a new pattern might corrupt the weights for all previously learned patterns. By changing the structure of the network rather than the weights, ART1 overcomes this problem.

Normalisation

We will refer to a cell in $F1$ or $F2$ with k.

Each cell k in $F1$ or $F2$ receives an input i_k and respond with an activation level a_k.

In order to introduce normalisation in the model, we set $I = \sum i_k$ and let the relative input intensity $\Theta_k = i_k I^{-1}$.

So we have a model in which the change of the response a_k of an input at a certain cell k

- depends inhibitorily on all other inputs and the sensitivity of the cell, i.e., the surroundings of each cell have a negative influence on the cell $-a_k \sum_{l \neq k} i_l$;

- has an excitatory response as far as the input at the cell is concerned $+Bi_k$;
- has an inhibitory response for normalisation $-a_k i_k$;
- has a decay $-Aa_k$.

Here, A and B are constants. The differential equation for the neurons in $F1$ and $F2$ now is

$$\frac{da_k}{dt} = -Aa_k + (B - a_k)i_k - a_k \sum_{l \neq k} i_l, \tag{5.32}$$

with $0 \leq a_k(0) \leq B$ because the inhibitory effect of an input can never exceed the excitatory input.

At equilibrium, when $da_k/dt = 0$, and with $I = \sum i_k$ we have that

$$a_k(A + I) = Bi_k. \tag{5.33}$$

Because of the definition of $\Theta_k = i_k I^{-1}$ we get

$$a_k = \Theta_k \frac{BI}{A + I}. \tag{5.34}$$

Therefore, at equilibrium a_k is proportional to Θ_k, and, since

$$\frac{BI}{A + I} \leq B, \tag{5.35}$$

the total activity $a^{\text{total}} = \sum a_k$ never exceeds B: it is *normalised*.

Contrast enhancement

In order to make $F2$ react better on differences in neuron values in $F1$ (or vice versa), *contrast enhancement* is applied: the contrasts between the neuronal values in a layer are amplified. We can show that eq. (5.32) does not suffice anymore. In order to enhance the contrasts, we chop off all the equal fractions (uniform parts) in $F1$ or $F2$. This can be done by adding an extra inhibitory input proportional to the inputs from the other cells with a factor C:

$$\frac{da_k}{dt} = -Aa_k + (B - a_k)i_k - (a_k + C) \sum_{l \neq k} i_l. \tag{5.36}$$

At equilibrium, when we set $B = (n - 1)C$ where n is the number of neurons, we have

$$a_k = \frac{nCI}{A + I} \left(\Theta_k - \frac{1}{n} \right). \tag{5.37}$$

Now, when an input in which all the i_k are equal is given, then all the a_k are zero: the effect of C is enhancing differences. If we set $B \leq (n - 1)C$ or $C/(B + C) \geq 1/n$, then more of the input shall be chopped off.

Discussion

The description of ART1 continues by defining the differential equations for the LTM. Instead of following Carpenter and Grossberg's description, we will revert to the simplified model as presented by Lippmann [22].

5.5 Reinforcement learning

In the previous chapters a number of supervised training methods have been described in which the weight adjustments are calculated using a set of 'learning samples.' These learning samples can either be given by an external teacher or, in the case that the network is used to control a system which interacts with the world, be derived from the interaction of the system with the environment. However, not always such a set of learning examples is available or can be derived. Often the only information is a signal which is given when a desired goal has been reached (reward) or when the system fails (punishment). This 'reinforcement' signal is characterised by the fact that it is less informative than a set of examples (only a correct-false judgement is given) and that it is often delayed: the success or failure signal which is given at a certain moment is a result of control signals (network outputs) in the past. The question is whether a network can be trained with such poor and delayed training data. A number of algorithms have been developed that make this form of learning possible. Specifically, Barto, Sutton and Anderson [25] have formulated 'reinforcement learning' as a learning strategy which does not need a set of examples provided by a 'teacher.' The system described by Barto explores the space of alternative input-output mappings and uses an evaluative feedback (reinforcement signal) on the consequences of the control signal (network output) on the environment. It has been shown that such reinforcement learning algorithms are implementing an on-line, incremental approximation to the dynamic programming method for optimal control, and are also called 'heuristic' dynamic programming [26].

The basic building blocks in most reinforcement learning schemes are an *Associative Search Element* (ASE) which uses a stochastic method to determine the correct relation between input and output and an *Adaptive Critic Element* (ACE) which learns to give a correct prediction of future reward or punishment (Figure 5.14). The external reinforcement signal r can be generated by a special sensor (for example a collision sensor of a mobile robot) or be derived from the state vector. For example, in control applications, where the state $\mathbf{s}$ of a system should remain in a certain part A of the control space, reinforcement is given by:

$$r = \begin{cases} 0 & \text{if } \mathbf{s} \in A, \\ -1 & \text{otherwise.} \end{cases} \tag{5.38}$$

5.5.1 Associative search

In its most elementary form the ASE gives a binary output value $o(t) \in \{0, 1\}$ as a stochastic function of an input vector. The total input of the ASE is, similar to the neuron presented in chapter 2, the weighted sum of the inputs, with the exception that the bias input in this case is a stochastic variable $\mathcal{N}$ with mean zero normal distribution:

$$i(t) = \sum_{j=1}^{N} w_{Sj} x_j(t) + \mathcal{N}_j. \tag{5.39}$$

The activation function $\mathcal{F}$ is a threshold such that

$$o(t) = a(t) = \begin{cases} 1 & \text{if } i(t) > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{5.40}$$

For updating the weights, a Hebbian type of learning rule is used. However, the update is weighted with the reinforcement signal $r(t)$ and an 'eligibility' e_j is defined instead of the

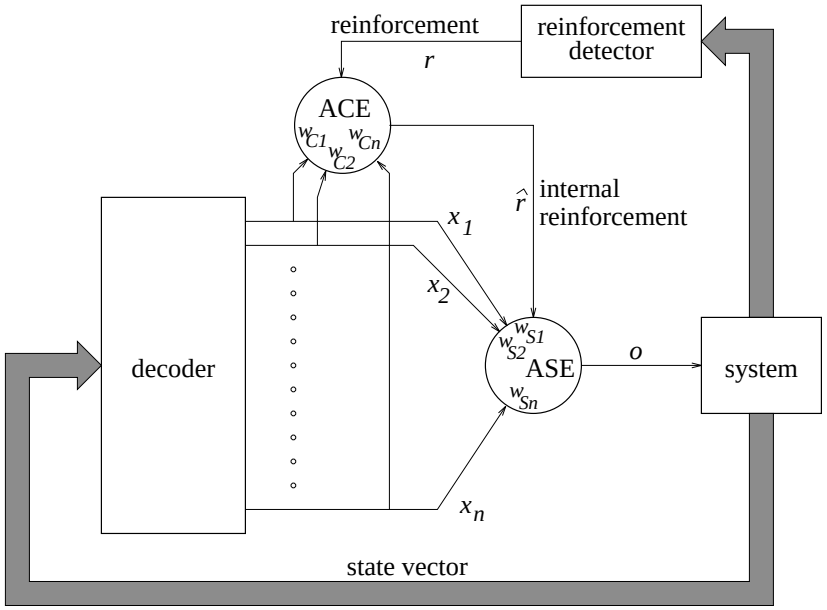

Figure 5.14: Architecture of a reinforcement learning scheme with critic element

product $o(t)x_j(t)$ of input and output:

$$w_{Sj}(t+1) = w_{Sj}(t) + \alpha r(t) e_j(t) \tag{5.41}$$

where α is a learning factor. The eligibility e_j is given by

$$e_j(t+1) = \delta e_j(t) + (1-\delta) o(t) x_j(t) \tag{5.42}$$

with δ the decay rate of the eligibility. The eligibility is a sort of 'memory;' e_j is high if the signals from the input state unit j and the output unit are correlated over some time.

Using $r(t)$ in expression (5.41) has the disadvantage that learning only finds place when there is an external reinforcement signal. Instead of $r(t)$, usually a continuous *internal* reinforcement signal $\hat{r}(t)$ given by the ACE, is used.

Barto and Anandan [27] proved convergence for the case of a *single* binary output unit and a set of *linearly independent* patterns $\mathbf{x}^p$. In control applications, the input vector is the (n-dimensional) state vector $\mathbf{s}$ of the system. In order to obtain a linear independent set of patterns $\mathbf{x}^p$, often a 'decoder' is used, which divides the range of each of the input variables s_i in a number of intervals. The aim is to divide the input (state) space in a number of disjunct subspaces (or 'boxes' as called by Barto). The input vector can therefore only be in one subspace at a time. The decoder converts the input vector into a binary valued vector $\mathbf{x}$, with only one element equal to one, indicating which subspace is currently visited. It has been shown [28] that instead of a-priori quantisation of the input space, a self-organising quantisation, based on methods described in this chapter, results in a better performance.

5.5.2 Adaptive critic

The Adaptive Critic Element (ACE, or 'evaluation network') uses rules from Temporal Difference Learning [29] to define an 'internal reinforcement' signal $\hat{r}$ which is available for each state of the system and no longer incidental (just in case of failure) but continuously present.

The basic idea is that for each state x_j an evaluation value $p(t) = p(x_j(t))$ has to be learned by the 'Adaptive Critic Element' (ACE). This evaluation should be the expectation of the discounted sum of external reinforcement

$$p_{des}(t) = E\left(\sum_{t' \geq t} \gamma^{t'-t} r(t'+1)\right) \tag{5.43}$$

with $0 < \gamma < 1$. From this we can derive that, if $p(t)$ is correctly learned:

$$p(t-1) = r(t) + \gamma p(t). \tag{5.44}$$

An error signal can now be defined for training the ACE:

$$\hat{r}(t) = r(t) + \gamma p(t) - p(t-1). \tag{5.45}$$

$p(t)$ is implemented as a series of 'weights' w_{Cj} to the ACE such that

$$p(t) = w_{Ck} \tag{5.46}$$

if the system is in state k at time t, denoted by $x_k = 1$. The function is learned by adjusting the w_{Cj}'s according to a 'delta-rule' with an error signal δ given by $\hat{r}(t)$:

$$\Delta w_{Cj}(t) = \beta \hat{r}(t) h_j(t). \tag{5.47}$$

β is the learning parameter and $h_j(t)$ indicates the 'trace' of neuron x_j:

$$h_j(t) = \lambda h_j(t-1) + (1-\lambda) x_j(t-1). \tag{5.48}$$

This trace is a low-pass filter or momentum, through which the credit assigned to state j increases while state j is active and decays exponentially after the activity of j has expired.

If $\hat{r}(t)$ is positive, the action u of the system has resulted in a higher evaluation value, whereas a negative $\hat{r}(t)$ indicates a deterioration of the system. $\hat{r}(t)$ can be considered as an *internal reinforcement* signal.

5.5.3 Example: The cartpole system

An example of such a system is the cartpole balancing system (see figure 5.15). Here, a dynamics controller must control the cart in such a way that the pole always stands up straight. The controller applies a 'left' or 'right' force F of fixed magnitude to the cart, which may change direction at discrete time intervals. The model has four state variables:

x the position of the cart on the track,

θ the angle of the pole with the vertical,

$\dot{x}$ the cart velocity, and

$\dot{\theta}$ the angle velocity of the pole.

Furthermore, a set of parameters specify the pole length and mass, cart mass, coefficients of friction between the cart and the track and at the hinge between the pole and the cart, the control force magnitude, and the force due to gravity. The state space is partitioned on the basis of the following quantisation thresholds:

1. x: $\pm 0.8, \pm 2.4$m,
2. θ: $0, \pm 1, \pm 6, \pm 12^\circ$,
3. $\dot{x}$: $\pm 0.5, \pm\infty$ m/s,
4. $\dot{\theta}$: $\pm 50, \pm\infty\circ$/s.

This yields $3 \times 6 \times 3 \times 3 = 162$ regions corresponding to all of the combinations of the intervals. The decoder output is a 162-dimensional vector. A negative reinforcement signal is provided when the state vector gets out of the admissible range: when $x > 2.4$, $x < -2.4$, $\theta > 12^\circ$ or $\theta < -12^\circ$. The system has proved to solve the problem in about 75 learning steps.

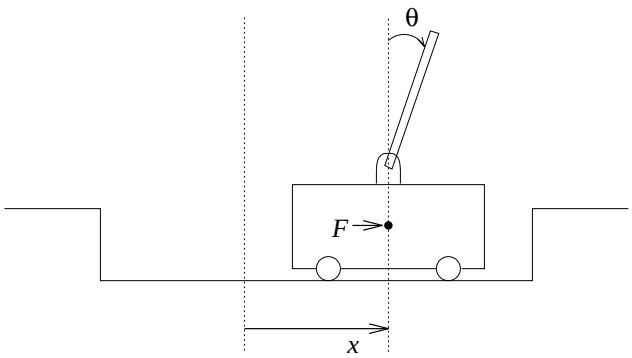

Figure 5.15: The cartpole system.

References

[1] D. E. Rumelhart & D. Zipser, "Feature discovery by competitive learning," *Cognitive Science* 9 (1985), 75–112.

[2] G. A. Carpenter & S. Grossberg, "A massively parallel architecture for a self-organizing neural pattern recognition machine," *Computer Vision, Graphics, and Image Processing* 37 (1987), 54–115.

[3] S. Grossberg, "Adaptive pattern classification and universal recoding I & II," *Biological Cybernetics* 23 (1976), 121–134, 187–202.

[4] K. Fukushima, "Cognitron: A self-organizing multilayered neural network," *Biological Cybernetics* 20 (1975), 121–136.

[5] K. Fukushima, "Neocognitron: A hierarchical neural network capable of visual pattern recognition," *Neural Networks* 1 (1988), 119–130.

[6] R. P. Lippmann, "Review of neural networks for speech recognition," *Neural Computation* 1 (1989), 1–38.

[7] S. C. Ahalt, A. K. Krishnamurthy, P. Chen & D. Melton, "competititve learning algorithms for vector quantization," in *Neural Networks* #3, 1990, 277–290.

[8] R. Hecht-Nielsen, "Counterpropagation networks," *Neural Networks*, 1 (1988).

[9] A. Jansen, P. P. van der Smagt & F. C. A. Groen, "Nested networks for robot control," in *Neural Network Applications*, A. F. Murray, ed., Kluwer Academic Publishers, 1993 (in press).

[10] L. Breiman, J. H. Friedman, R. A. Olshen & C. J. Stone, *Classification and Regression Trees*, Wadsworth and Broks/Cole, 1984.

[11] J. H. Friedman,, "multivariate adaptive regression splines," *Annals of Statistics* 19 (1991), 1–141.

[12] B. E. Rosen, J. M. Goodwin & J. J. Vidal, "Process Control with Adaptive Range Coding," *Biological Cybernetics* 66 (1992), 419–428.

[13] T. Kohonen, *Associative Memory: A System-Theoretical Approach*, Springer-Verlag, 1977.

[14] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biological Cybernetics* 43 (1982), 59–69.

[15] T. Kohonen, *Self-Organization and Associative Memory*, Springer-Verlag, Berlin, 1984.

[16] T. Martinetz & K. Schulten, "A "Neural-gas" network learns topologies," *Proceedings of the 1991 International Conference on Artificial Neural Networks* 1 (1991), 397–402.

[17] B. Fritzke, "Let it grow—self-organizing feature maps with problem dependent cell structure," *Proceedings of the 1991 International Conference on Artificial Neural Networks* 1 (1991), 403–408.

[18] T. Kohonen, M. Makisara & T. Saramaki, "Phonotopic maps—insightful representation of phonological features for speech recognition," in *Proceedings of the 7th IEEE International Conference on Pattern Recognition*, August 1984.

[19] C. Gielen, K. Krommenhoek & J. van Gisbergen, "A procedure for self-organized sensor-fusion in topologically ordered maps," in *Proceedings of the Second International Conference on Autonomous Systems*, T. Kanade, F. C. A. Groen & L. O. Hertzberger, eds., 417–423.

[20] E. Oja, "A simplified neuron model as a principal component analyzer," in *Journal of mathematical biology*, (15), 1982, 267–273.

[21] P. J. H. A. M. van Summeren, *Automatic Classification of Rheumatic Diseases by the ID_3 Induction Algorithm and Two Types of Neural Networks*, Dept. of Mathematics and Computer Science, Vrije Universiteit, Amsterdam, 1990, graduation thesis.

[22] R. P. Lippmann, "An introduction to computing with neural nets," *IEEE Transactions on Acoustics, Speech, and Signal Processing* 2 (April 1987), 4–22.

[23] G. A. Carpenter & S. Grossberg, "ART 2: Self-organization of stable category recognition codes for analog input patterns," *Applied Optics* 26 (1987), 4919–4930.

[24] J. A. Hartigan, *Clustering Algorithms*, John Wiley & Sons, New York, 1975.

[25] A. G. Barto, R. S. Sutton & C. W. Anderson, "Neuronlike adaptive elements that can solve difficult learning problems," *IEEE Transactions on Systems, Man and Cybernetics* 13 (1983), 834–846.

[26] P. W. Werbos, "A menu for designs of reinforcement learning over time," in *Neural Networks for Control*, W. T. Miller III, R. S. Sutton & P. J. Werbos, eds., MIT Press/Bradford, 1990.

[27] A. G. Barto & P. Anandan, "Pattern-recognizing stochastic learning automata," *IEEE Transactions on Systems, Man and Cybernetics* 15 (1985), 360 – 375.

[28] B. J. A. Kröse & J. W. M. van Dam, "Learning to avoid collisions: A reinforcement learning paradigm for mobile robot manipulation," in *Proceedings of IFAC/IFIP/IMACS International Symposium on Artificial Intelligence in Real-Time Control*, IFAC, Laxenburg, Delft, June 1992, 295–300 .

[29] R. S. Sutton, "Temporal credit assignment in reinforcement learning ," University of Massachussetts, Ph.D. Thesis, 1984.

6 Recurrent Networks

Most networks discussed in the previous chapters have in common that they are non-recurrent. All data flows in one direction, and after a finite and known number of steps, depending on the depth of the network, all activation values in the network are known. On the other hand, the weights in the network relax or converge to some desired state by means of a learning rule.

In this chapter, we discuss networks in which the *activation values* undergo a relaxation process but the weights remain unchanged. Such *recurrent networks* have feedback loops from output to input units[1]. When the network is inherently stable, it will evolve to a state in which no neurons change value anymore and the outputs become constant. Otherwise, the network may change activation values *ad infinitum*. Although the latter networks have useful emergent properties, we do not study them here.

Among the early recurrent networks was the *auto-associator* independently described by Anderson [1] and Kohonen [2] in 1977. It consists of a pool of neurons with connections between each unit i and j, $i \neq j$ (see figure 6.1). All connections are weighted.

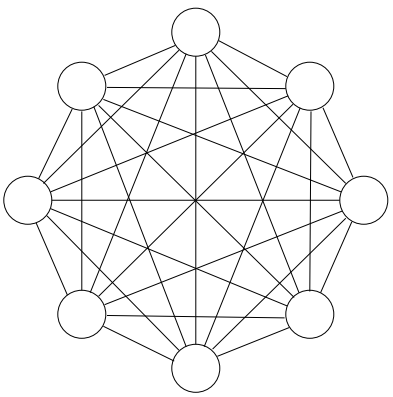

Figure 6.1: The auto-associator network. All neurons are both input and output neurons, i.e., a pattern is clamped, the network iterates to a stable state, and the output of the network consists of the new activation values of the neurons.

In 1982, Hopfield [3] brings together several earlier ideas concerning these networks and presents a complete mathematical analysis based on Ising spin models [4]. It is therefore that this network, which we will describe in this chapter, is generally referred to as the *Hopfield network*.

[1] Or, at least, to units in an 'earlier layer.'

6.1 The Hopfield network

6.1.1 Description

The Hopfield network consists of a set of N interconnected neurons (figure 6.1) which update their activation values asynchronously and independently of other neurons. All neurons are both input and output neurons. The activation values are binary. Originally, Hopfield chose activation values of 1 and 0, but using values $+1$ and -1 presents some advantages discussed below. We will therefore adhere to the latter convention.

The state of the system is given by the activation values[2] $\mathbf{a} = (a_i)$. The net input $i_i(t+1)$ of a neuron i at cycle $t+1$ is a weighted sum

$$i_i(t+1) = \sum_{j \neq i} a_j(t)\, w_{ij} + \theta_i. \tag{6.1}$$

A simple threshold function (figure 2.2) is applied to the net input to obtain the new activation value $a_i(t+1)$ at time $t+1$:

$$a_i(t+1) = \begin{cases} +1 & \text{if } i_i(t+1) > U_i \\ -1 & \text{if } i_i(t+1) < U_i \\ a_i(t) & \text{otherwise,} \end{cases} \tag{6.2}$$

i.e., $a_i(t+1) = \mathrm{sgn}(i_i(t+1))$. For simplicity we henceforth choose $U_i = 0$, but this is of course not essential.

A neuron i in the Hopfield network is called *stable* at time t if, in accordance with equations (6.1) and (6.2),

$$a_i(t) = \mathrm{sgn}(i_i(t-1)). \tag{6.3}$$

A state α is called stable if, when the network is in state α, all neurons are stable. A pattern $\mathbf{x}^p$ is called stable if, when $\mathbf{x}^p$ is clamped, all neurons are stable.

When the extra restriction $w_{ij} = w_{ji}$ is made, the behaviour of the system can be described with an energy function

$$\mathcal{E} = -\tfrac{1}{2} \sum_{i \neq j} \sum a_i a_j w_{ij} - \sum_i \theta_i a_i. \tag{6.4}$$

Theorem 3 *A recurrent network with connections $w_{ij} = w_{ji}$ in which the neurons are updated using rule (6.2) has stable limit points.*

Proof *First, note that the energy expressed in eq. (6.4) is bounded from below, since the a_i are bounded from below and the w_{ij} and θ_i are constant. Secondly, $\mathcal{E}$ is monotonically decreasing when state changes occur, because*

$$\Delta\mathcal{E} = -\Delta a_i \left(\sum_{i \neq j} a_j w_{ij} + \theta_i \right) \tag{6.5}$$

is always negative when a_i changes according to eqs. (6.1) and (6.2).

□

[2] Often, these networks are described using the symbols used by Hopfield: V_i for activation of unit i, T_{ij} for the connection weight between units i and j, and U_i for the external input of unit i. We decided to stick to the more general symbols a_i, w_{ij}, and θ_i.

The advantage of a $+1/-1$ model over a $1/0$ model then is symmetry of the states of the network. For, when some pattern $\mathbf{x}$ is stable, its inverse is stable, too, whereas in the $1/0$ model this is not always true (as an example, the pattern $00\cdots00$ is always stable, but $11\cdots11$ need not be). Similarly, both a pattern and its inverse have the same energy in the $+1/-1$ model.

Removing the restriction of bidirectional connections (i.e., $w_{ij} = w_{ji}$) results in a system that is not guaranteed to settle to a stable state.

6.1.2 Hopfield network as associative memory

A primary application of the Hopfield network is an associative memory. In this case, the weights of the connections between the neurons have to be thus set that the states of the system corresponding with the patterns which are to be stored in the network are stable. These states can be seen as 'dips' in energy space. When the network is cued with a noisy or incomplete test pattern, it will render the incorrect or missing data by iterating to a stable state which is in some sense 'near' to the cued pattern.

The Hebb rule can be used (section 2.3.2) to store P patterns:

$$w_{ij} = \begin{cases} \sum_{p=0}^{P-1} x_i^p x_j^p & \text{if } i \neq j \\ 0 & \text{otherwise,} \end{cases} \tag{6.6}$$

i.e., if x_i^p and x_j^p are equal, w_{ij} is increased, otherwise decreased by one (note that, in the original Hebb rule, weights only increase). It appears, however, that the network gets saturated very quickly, and that about $0.15N$ memories can be stored before recall errors become severe.

There are two problems associated with storing too many patterns:

1. the stored patterns become unstable;
2. spurious stable states appear (i.e., stable states which do not correspond with stored patterns).

The first of these two problems can be solved by an algorithm proposed by Bruce *et al.* [5]:

Algorithm 1 *Given a starting weight matrix* $W = \left[w_{ij}\right]$, *for each pattern* $\mathbf{x}^p$ *to be stored and each element* x_i^p *in* $\mathbf{x}^p$ *define a correction* ϵ_i *such that*

$$\epsilon_i = \begin{cases} 0 & \textit{if } a_i \textit{ is stable and } \mathbf{x}^p \textit{ is clamped;} \\ 1 & \textit{otherwise.} \end{cases} \tag{6.7}$$

Now modify w_{ij} *by* $\Delta w_{ij} = a_i a_j(\epsilon_i + \epsilon_j)$ *if* $i \neq j$. *Repeat this procedure until all patterns are stable.*

Theorem 4 *Algorithm 1 converges in a finite number of steps, provided a solution exists.*

Proof *Exercise.*
□

The second problem stated above can be alleviated by applying the Hebb rule *in reverse* to the spurious stable state, but with a low learning factor [6]. Thus these patterns are weakly *un*stored and will become unstable again.

6.1.3 Neurons with graded response

The network described in section 6.1.1 can be generalised by allowing continuous activation values. Here, the threshold activation function is replaced by a sigmoid. As before, this system can be proved to be stable when a symmetric weight matrix is used [7].

Hopfield networks for optimisation problems

An interesting application of the Hopfield network with graded response arises in a heuristic solution to the NP-complete travelling salesman problem [8]. In this problem, a path of minimal distance must be found between n cities, such that the begin- and end-points are the same.

Hopfield and Tank [9] use a network with $n \times n$ neurons. Each row in the matrix represents a city, whereas each column represents the position in the tour. When the network is settled, each row and each column should have one and only one active neuron, indicating a specific city occupying a specific position in the tour. The neurons are updated using rule (6.2) with a sigmoid activation function between 0 and 1. The activation value $a_{Xi} = 1$ indicates that city X occupies the i^{th} place in the tour.

An energy function describing this problem can be set up as follows. To ensure a correct solution, the following energy must be minimised:

$$\begin{aligned} \mathcal{E} \;=\; & \tfrac{A}{2} \sum_X \sum_i \sum_{j \neq i} a_{Xi} a_{Xj} \\ & + \tfrac{B}{2} \sum_i \sum_X \sum_{X \neq Y} a_{Xi} a_{Yi} \\ & + \tfrac{C}{2} \left(\sum_X \sum_i a_{Xi} - n \right)^2 \end{aligned} \tag{6.8}$$

where A, B, and C are constants. The first and second terms in equation (6.8) are zero if and only if there is a maximum of one active neuron in each row and column, respectively. The last term is zero if and only if there are exactly n active neurons.

To minimise the distance of the tour, an extra term

$$\tfrac{D}{2} \sum_X \sum_{Y \neq X} \sum_i d_{XY} a_{Xi} (a_{Y,i+1} + a_{Y,i-1}) \tag{6.9}$$

is added to the energy, where d_{XY} is the distance between cities X and Y and D is a constant. For convenience, the subscripts are defined modulo n.

The weights are set as follows:

$$\begin{array}{lll} w_{Xi,Yj} = & -A\delta_{XY}(1-\delta_{ij}) & \text{inhibitory connections within each row} \\ & -B\delta_{ij}(1-\delta_{XY}) & \text{inhibitory connections within each column} \\ & -C & \text{global inhibition} \\ & -Dd_{XY}(\delta_{j,i+1}+\delta_{j,i-1}) & \text{data term} \end{array} \tag{6.10}$$

where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise. Finally, each neuron has an external bias input Cn.

Discussion

Although this application is interesting from a theoretical point of view, the applicability is limited. Whereas Hopfield and Tank state that, in a ten city tour, the network converges to a valid solution in 16 out of 20 trials while 50% of the solutions are optimal, other reports show less encouraging results. For example, [10] find that in only 15% of the runs a valid result is obtained, few of which lead to an optimal or near-optimal solution. The main problem is the lack of global information. Since, for an N-city problem, there are $N!$ possible tours, each of which may be traversed in two directions as well as started in N points, the number of different tours is $N!/2N$. Differently put, the N-dimensional hypercube in which the solutions are situated is $2N$ degenerate. The degenerate solutions occur evenly within the hypercube, such that all but one of the final $2N$ configurations are redundant. The competition between the degenerate tours often leads to solutions which are piecewise optimal but globally inefficient.

6.2 Boltzmann machines

The Boltzmann machine, as first described by Ackley, Hinton, and Sejnowski in 1985 [11] is a neural network that can be seen as an extension to Hopfield networks to include hidden units, and with a stochastic instead of deterministic update rule. The weights are still symmetric. The operation of the network is based on the physics principle of *annealing*. This is a process whereby a material is heated and then cooled very, very slowly to a freezing point. As a result, the crystal lattice will be highly ordered, without any impurities, such that the system is in a state of very low energy. In the Boltzmann machine this system is mimicked by changing the deterministic update of equation (6.2) in a stochastic update, in which a neuron becomes active with a probability p,

$$p(a_i \leftarrow +1) = \frac{1}{1 + e^{-\Delta \mathcal{E}_i / T}} \tag{6.11}$$

where T is a parameter comparable with the (synthetic) temperature of the system. This stochastic activation function is not to be confused with neurons having a sigmoid deterministic activation function.

In accordance with a physical system obeying a Boltzmann distribution, the network will eventually reach 'thermal equilibrium' and the relative probability of two global states α and β will follow the Boltzmann distribution

$$\frac{P_\alpha}{P_\beta} = e^{-(\mathcal{E}_\alpha - \mathcal{E}_\beta)/T} \tag{6.12}$$

where P_α is the probability of being in the α^{th} global state, and $\mathcal{E}_\alpha$ is the energy of that state. Note that at thermal equilibrium the units still change state, but the probability of finding the network in any global state remains constant.

At low temperatures there is a strong bias in favour of states with low energy, but the time required to reach equilibrium may be long. At higher temperatures the bias is not so favourable but equilibrium is reached faster. A good way to beat this trade-off is to start at a high temperature and gradually reduce it. At high temperatures, the network will ignore small energy differences and will rapidly approach equilibrium. In doing so, it will perform a search of the coarse overall structure of the space of global states, and will find a good minimum at that coarse level. As the temperature is lowered, it will begin to respond to

smaller energy differences and will find one of the better minima within the coarse-scale minimum it discovered at high temperature.

As multi-layer perceptrons, the Boltzmann machine consists of a non-empty set of visible and a possibly empty set of hidden units. Here, however, the units are binary-valued and are updated stochastically and asynchronously. The simplicity of the Boltzmann distribution leads to a simple learning procedure which adjusts the weights so as to use the hidden units in an optimal way [11]. This algorithm works as follows.

First, the input and output vectors are clamped. The network is then annealed until it approaches thermal equilibrium at a temperature of 0. It then runs for a fixed time at equilibrium and each connection measures the fraction of the time during which both the units it connects are active. This is repeated for all input-output pairs so that each connection can measure $\langle a_i a_j \rangle^{\text{clamped}}$, the expected probability, averaged over all cases, that units i and j are simultaneously active at thermal equilibrium when the input and output vectors are clamped. Similarly, $\langle a_i a_j \rangle^{\text{free}}$ is measured when the output units are not clamped but determined by the network.

In order to determine optimal weights in the network, an error function must be determined. Now, the probability $P^{\text{free}}(\mathbf{a}^p)$ that the visible units are in state $\mathbf{a}^p$ when the system is running freely can be measured. Also, the *desired* probability $P^{\text{clamped}}(\mathbf{a}^p)$ that the visible units are in state $\mathbf{a}^p$ is determined by clamping the visible units and letting the network run. Now, if the weights in the network are correctly set, both probabilities are equal to each other, and the error E in the network must be 0. Otherwise, the error must have a positive value measuring the discrepancy between the network's internal mode and the environment. For this effect, the 'asymmetric divergence' or 'Kullback information' is used:

$$E = \sum_p P^{\text{clamped}}(\mathbf{a}^p) \log \frac{P^{\text{clamped}}(\mathbf{a}^p)}{P^{\text{free}}(\mathbf{a}^p)}, \tag{6.13}$$

Now, in order to minimise E using gradient descent, we must change the weights according to

$$\Delta w_{ij} = -\gamma \frac{\partial E}{\partial w_{ij}}. \tag{6.14}$$

It is not difficult to show that

$$\frac{\partial E}{\partial w_{ij}} = -\frac{1}{T}\left(\langle a_i a_j \rangle^{\text{clamped}} - \langle a_i a_j \rangle^{\text{free}}\right). \tag{6.15}$$

Therefore, each weight is updated by

$$\Delta w_{ij} = \gamma\left(\langle a_i a_j \rangle^{\text{clamped}} - \langle a_i a_j \rangle^{\text{free}}\right). \tag{6.16}$$

References

[1] J. A. Anderson, "Neural models with cognitive implications," in *Basic Processes in Reading Perception and Comprehension Models*, D. LaBerge & S. J. Samuels, eds., Erlbaum, Hillsdale, NJ, 1977, 27–90.

[2] T. Kohonen, *Associative Memory: A System-Theoretical Approach*, Springer-Verlag, 1977.

[3] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the National Academy of Sciences* 79 (1982), 2554–2558.

[4] D. J. Amit, H. Gutfreund & H. Sompolinsky, "Spin-glass models of neural networks," *Physical Review A* 32 (1986), 1007–1018.

[5] A. D. Bruce, A. Canning, B. Forrest, E. Gardner & D. J. Wallace, "Learning and memory properties in fully connected networks," in *AIP Conference Proceedings 151, Neural Networks for Computing*, J. S. Denker, ed., Snowbird, Utah, 1986, 65–70.

[6] J. J. Hopfield, D. I. Feinstein & R. G. Palmer, " 'Unlearning' has a stabilizing effect in collective memories," *Nature* 304 (1983), 159–159.

[7] J. J. Hopfield, "Neurons with graded response have collective computational properties like those of two-state neurons," *Proceedings of the National Academy of Sciences* 81 (May 1984), 3088–3092.

[8] M. R. Garey & D. S. Johnson, *Computers and Intractability*, W. H. Freeman, New York, 1979.

[9] J. J. Hopfield & D. W. Tank, " 'Neural' computation of decisions in optimization problems," *Biological Cybernetics* 52 (1985), 141–152.

[10] G. V. Wilson & G. S. Pawley, "On the stability of the travelling salesman problem algorithm of hopfield and tank," *Biological Cybernetics* 58 (1988), 63–70.

[11] D. H. Ackley, G. E. Hinton & T. J. Sejnowski, "A learning algorithm for Boltzmann machines," *Cognitive Science* 9 (1985), 147–169.

[12] R. A. Hummel & S. W. Zucker, "On the foundations of relaxation labeling processes," *IEEE Transactions on Pattern Analysis and Machine Intelligence* 5 (May 1983), 267–286.

[13] S. Geman & D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6 (1984), 721–724.

Part III

APPLICATIONS

7 Robot Control

An important area of application of neural networks is in the field of robotics. Usually, these networks are designed to direct a manipulator, which is the most important form of the industrial robot, to grasp objects, based on sensor data. Another applications include the steering and path-planning of autonomous robot vehicles.

In robotics, the major task involves making movements dependent on sensor data. There are four, related, problems to be distinguished [1]:

Forward kinematics. Kinematics is the science of motion which treats motion without regard to the forces which cause it. Within this science one studies the position, velocity, acceleration, and all higher order derivatives of the position variables. A very basic problem in the study of mechanical manipulation is that of *forward kinematics*. This is the static geometrical problem of computing the position and orientation of the end-effector ('hand') of the manipulator. Specifically, given a set of joint angles, the forward kinematic problem is to compute the position and orientation of the tool frame relative to the base frame (see figure 7.1).

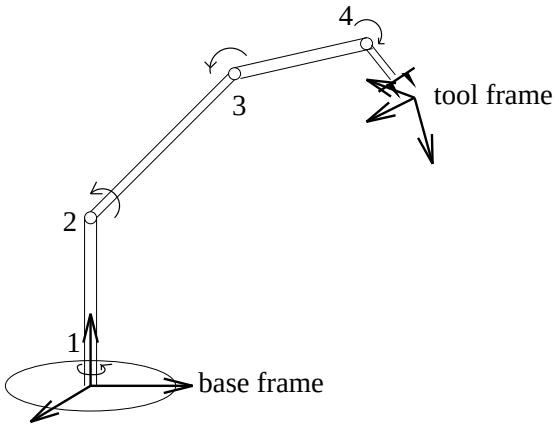

Figure 7.1: An exemplar robot manipulator.

Inverse kinematics. This problem is posed as follows: given the position and orientation of the end-effector of the manipulator, calculate all possible sets of joint angles which could be used to attain this given position and orientation. This is a fundamental problem in the practical use of manipulators.

The inverse kinematic problem is not as simple as the forward one. Because the kinematic equations are nonlinear, their solution is not always easy or even possible in a closed form. Also, the questions of existence of a solution, and of multiple solutions, arise.

Solving this problem is a least requirement for most robot control systems.

Dynamics. Dynamics is a field of study devoted to studying the *forces* required to cause motion. In order to accelerate a manipulator from rest, glide at a constant end-effector velocity, and finally decelerate to a stop, a complex set of torque functions must be applied by the joint actuators.

Trajectory generation. To move a manipulator from here to there in a smooth, controlled fashion each joint must be moved via a smooth function of time. Exactly how to compute these motion functions is the problem of *trajectory generation*.

In the first section of this chapter we will discuss the problems associated with the positioning of the end-effector (in effect, representing the inverse kinematics in combination with sensory transformation). Section 7.2 discusses a network for controlling the dynamics of a robot arm. Finally, section 7.3 describes neural networks for mobile robot control.

7.1 End-effector positioning

The final goal in robot manipulator control is often the positioning of the hand or end-effector in order to be able to, e.g., pick up an object. With the accurate robot arm that are manufactured, this task is often relatively simple, involving the following steps:

1. determine the target coordinates relative to the base of the robot. Typically, when this position is not always the same, this is done with a number of fixed cameras or other sensors which observe the work scene, from the image frame determine the position of the object in that frame, and perform a pre-determined coordinate transformation;
2. with a precise model of the robot (supplied by the manufacturer), calculate the joint angles to reach the target (i.e., the inverse kinematics). This is a relatively simple problem;
3. move the arm (dynamics control) and close the gripper.

The arm motion in point 3 is discussed in section 7.2. Gripper control is not a trivial matter at all, but we will not focus on that.

Involvement of neural networks. So if these parts are relatively simple to solve with a high accuracy, why involve neural networks? The reason is the applicability of robots. When 'traditional' methods are used to control a robot arm, accurate models of the sensors and manipulators (in some cases with unknown parameters which have to be estimated from the system's behaviour; yet still with accurate models as starting point) are required and the system must be calibrated. Also, systems which suffer from wear-and-tear (and which mechanical systems don't?) need frequent recalibration or parameter determination. Finally, the development of more complex (adaptive!) control methods allows the design and use of more flexible (i.e., less rigid) robot systems, both on the sensory and motory side.

7.1.1 Camera–robot coordination is function approximation

The system we focus on in this section is a work floor observed by a fixed cameras, and a robot arm. The visual system must identify the target as well as determine the visual position of the end-effector.

The target position $\mathbf{x}^{\text{target}}$ together with the visual position of the hand $\mathbf{x}^{\text{hand}}$ are input to the neural controller $\mathcal{N}(\cdot)$. This controller then generates a joint position θ for the robot:

$$\theta = \mathcal{N}(\mathbf{x}^{\text{target}}, \mathbf{x}^{\text{hand}}). \tag{7.1}$$

We can compare the neurally generated θ with the optimal θ_0 generated by a fictitious perfect controller $\mathcal{R}(\cdot)$:

$$\theta_0 = \mathcal{R}(\mathbf{x}^{\text{target}}, \mathbf{x}^{\text{hand}}). \tag{7.2}$$

The task of learning is to make the $\mathcal{N}$ generate an output 'close enough' to θ_0.

There are two problems associated with teaching $\mathcal{N}(\cdot)$:

1. generating learning samples which are in accordance with eq. (7.2). This is not trivial, since in useful applications $\mathcal{R}(\cdot)$ is an unknown function. Instead, a form of self-supervised or unsupervised learning is required. Some examples to solve this problem are given below;

2. constructing the mapping $\mathcal{N}(\cdot)$ from the available learning samples. When the (usually randomly drawn) learning samples are available, a neural network uses these samples to represent the whole input space over which the robot is active. This is evidently a form of interpolation, but has the problem that the input space is of a high dimensionality, and the samples are randomly distributed.

We will discuss three fundamentally different approaches to neural networks for robot end-effector positioning. In each of these approaches, a solution will be found for both the learning sample generation and the function representation.

Approach 1: Feed-forward networks

When using a feed-forward system for controlling the manipulator, a self-supervised learning system must be used.

One such a system has been reported by Psaltis, Sideris and Yamamura [2]. Here, the network, which is constrained to two-dimensional positioning of the robot arm, learns by experimentation. Three methods are proposed:

1. *Indirect learning.*
   In indirect learning, a Cartesian target point $\mathbf{x}$ in world coordinates is generated, e.g., by a two cameras looking at an object. This target point is fed into the network, which generates an angle vector θ. The manipulator moves to position θ, and the cameras determine the new position $\mathbf{x}'$ of the end-effector in world coordinates. This $\mathbf{x}'$ again is input to the network, resulting in θ'. The network is then trained on the error $\epsilon_1 = \theta - \theta'$ (see figure 7.2).

   However, minimisation of ϵ_1 does not guarantee minimisation of the overall error $\epsilon = \mathbf{x} - \mathbf{x}'$. For example, the network often settles at a 'solution' that maps all $\mathbf{x}$'s to a single θ (i.e., the mapping $\mathbb{I}$).

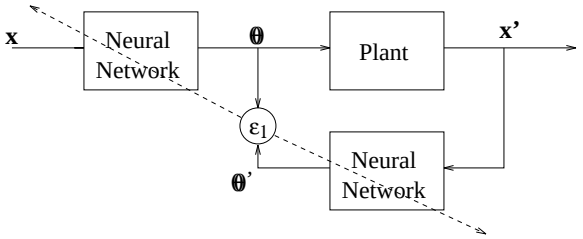

Figure 7.2: Indirect learning system for robotics. In each cycle, the network is used in two different places: first in the forward step, then for feeding back the error.

2. *General learning.*
   The method is basically very much like supervised learning, but here the plant input θ must be provided by the user. Thus the network can directly minimise $|\theta - \theta'|$. The success of this method depends on the interpolation capabilities of the network. Correct choice of θ may pose a problem.

3. *Specialised learning.*
   This method requires knowledge of the Jacobian matrix of the plant. A Jacobian matrix of a multidimensional function F is a matrix of partial derivatives of F, i.e., the multidimensional form of the derivative. For example, if we have $Y = F(X)$, i.e.,

$$\begin{aligned} y_1 &= f_1(x_1, x_2, \ldots, x_n), \\ y_2 &= f_2(x_1, x_2, \ldots, x_n), \\ &\vdots \\ y_m &= f_m(x_1, x_2, \ldots, x_n) \end{aligned}$$

   then

$$\begin{aligned} \delta y_1 &= \frac{\partial f_1}{\partial x_1}\delta x_1 + \frac{\partial f_1}{\partial x_2}\delta x_2 + \ldots + \frac{\partial f_1}{\partial x_n}\delta x_n, \\ \delta y_2 &= \frac{\partial f_2}{\partial x_1}\delta x_1 + \frac{\partial f_2}{\partial x_2}\delta x_2 + \ldots + \frac{\partial f_2}{\partial x_n}\delta x_n, \\ &\vdots \\ \delta y_m &= \frac{\partial f_m}{\partial x_1}\delta x_1 + \frac{\partial f_m}{\partial x_2}\delta x_2 + \ldots + \frac{\partial f_m}{\partial x_n}\delta x_n \end{aligned}$$

   or

$$\delta Y = \frac{\partial F}{\partial X}\delta X. \tag{7.3}$$

   Eq. (7.3) is also written as

$$\delta Y = J(X)\delta X \tag{7.4}$$

   where J is the Jacobian matrix of F. So, the Jacobian matrix can be used to calculate the change in the function when its parameters change.

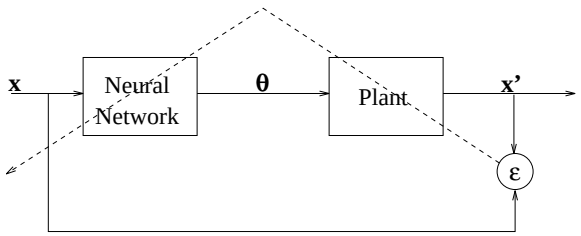

Figure 7.3: The system used for specialised learning.

Now, in this case we have

$$J_{ij} = \left[\frac{\partial P_i}{\partial \theta_j}\right] \tag{7.5}$$

where $P_i(\theta)$ the ith element of the plant output for input θ. The learning rule applied here regards the plant as an additional and unmodifiable layer in the neural network. The total error $\epsilon = \mathbf{x} - \mathbf{x}'$ is propagated back through the plant by calculating the δ_j as in eq. (4.14):

$$\delta_j = \mathcal{F}'(i_j) \sum_i \delta_i \frac{\partial P_i(\theta)}{\partial \theta_j},$$
$$\delta_i = x_i - x'_i,$$

where i iterates over the outputs of the plant. When the plant is an unknown function, $\frac{\partial P_i(\theta)}{\partial \theta_j}$ can be approximated by

$$\frac{\partial P_i(\theta)}{\partial \theta_j} \approx \frac{P_i(\theta + h\theta_j \mathbf{e}_j) - P_i(\theta)}{h} \tag{7.6}$$

where $\mathbf{e}_j$ is used to change the scalar θ_j into a vector. This approximate derivative can be measured by slightly changing the input to the plant and measuring the changes in the output.

A somewhat similar approach is taken in [3] and [4]. Again a two-layer feed-forward network is trained with back-propagation. However, instead of calculating a desired output vector the input vector which should have invoked the current output vector is reconstructed, and back-propagation is applied to this new input vector and the existing output vector.

The configuration used consists of a monocular manipulator which has to grasp objects. Due to the fact that the camera is situated in the hand of the robot, the task is to move the hand such that the object is in the centre of the image and has some predetermined size (in a later article, a biologically inspired system is proposed [5] in which the visual flow-field is used to account for the monocularity of the system, such that the dimensions of the object need not to be known anymore to the system).

One step towards the target consists of the following operations:

1. measure the distance from the current position to the target position in camera domain, $\mathbf{x}$;

2. use this distance, together with the current state θ of the robot, as input for the neural network. The network then generates a joint displacement vector $\Delta\theta$;

3. send $\Delta\theta$ to the manipulator;

4. again measure the distance from the current position to the target position in camera domain, $\mathbf{x}'$;

5. calculate the move made by the manipulator in visual domain, $\mathbf{x} - {}^{t+1}_{t}R\mathbf{x}'$, where ${}^{t+1}_{t}R$ is the rotation matrix of the second camera image with respect to the first camera image;

6. teach the learning pair $(\mathbf{x} - {}^{t+1}_{t}R\mathbf{x}', \theta; \Delta\theta)$ to the network.

This system has shown to learn correct behaviour in only tens of iterations, and to be very adaptive to changes in the sensor or manipulator [4].

By using a feed-forward network, the available learning samples are approximated by a single, smooth function consisting of a summation of sigmoid functions. As mentioned in section 4, a feed-forward network with one layer of sigmoid units is capable of representing practically any function. But how are the optimal weights determined in finite time to obtain this optimal representation? Experiments have shown that, although a reasonable representation can be obtained in a short period of time, an accurate representation of the function that governs the learning samples is often not feasible or extremely difficult [6]. The reason for this is the global character of the approximation obtained with a feed-forward network with sigmoid units: every weight in the network has a *global* effect on the final approximation that is obtained.

Building local representations is the obvious way out: every part of the network is responsible for a small subspace of the total input space. Thus accuracy is obtained locally (Keep It Small & Simple). This is typically obtained with a Kohonen network.

Approach 2: Topology conserving maps

Ritter, Martinetz, and Schulten [7,8] describe the use of a Kohonen-like network for robot control. We will only describe the kinematics part, since it is the most interesting and straightforward.

The system described by Ritter *et al.* consists of a robot manipulator with three degrees of freedom (orientation of the end-effector is not included) which has to grab objects in 3D-space. The system is observed by two fixed cameras which output their (x, y) coordinates of the object and the end effector (see figure 7.4).

Each run consists of two movements. In the gross move, the observed location of the object $\mathbf{x}$ (a four-component vector) is input to the network. As with the Kohonen network, the neuron j with highest activation value is selected as winner, because its weight vector $\mathbf{w}_j$ is nearest to $\mathbf{x}$. The neurons, which are arranged in a 3-dimensional lattice, correspond in a $1-1$ fashion with subregions of the 3D workspace of the robot, i.e., the neuronal lattice is a *discrete representation* of the workspace. With each neuron a vector θ and Jacobian matrix A are associated. During gross move θ_j is fed to the robot which makes its move, resulting in retinal coordinates $\mathbf{x}_i$ of the end-effector. To correct for the discretisation of the working space, an additional move is made which is dependent of the distance between the neuron

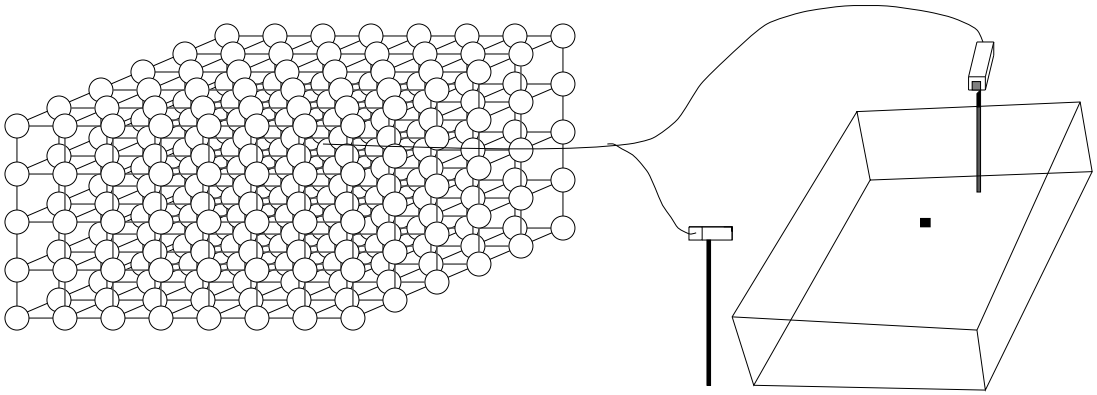

Figure 7.4: A Kohonen network merging the output of two cameras.

and the object in space $\mathbf{w}_j - \mathbf{x}$; this small displacement in Cartesian space is translated to an angle change using the Jacobian A_j:

$$\theta^{\text{final}} = \theta_j + A_j(\mathbf{x} - \mathbf{w}_j) \tag{7.7}$$

which is a first-order Taylor expansion of θ^{final}. The final retinal coordinates of the end-effector after this fine move are in $\mathbf{x}_f$.

Learning proceeds as follows: when an improved estimate $(\theta, A)^*$ has been found, the following adaptations are made *for all neurons* i:

$$\mathbf{w}_i^{\text{new}} = \mathbf{w}_i^{\text{old}} + \gamma(t)\, h_{ij}(t)\, \left(\mathbf{x} - \mathbf{w}_i^{\text{old}}\right),$$
$$(\theta, A)_i^{\text{new}} = (\theta, A)_i^{\text{old}} + \gamma'(t)\, h'_{ij}(t)\, \left((\theta, A)_i^* - (\theta, A)_i^{\text{old}}\right).$$

If $h_{ij}(t) = h'_{ij}(t) = \delta_{ij}$, this is similar to perceptron learning. Here, as with the Kohonen learning rule, a distance function is used such that $h_{ij}(t)$ and $h'_{ij}(t)$ are Gaussians depending on the distance between neurons i and j with a maximum at $i = j$ (cf. eq. (5.6)).

An improved estimate $(\theta, A)^*$ is obtained as follows.

$$\theta^* = \theta_j + A_j(\mathbf{x} - \mathbf{x}_f), \tag{7.8}$$

$$\begin{aligned} A^* &= A_j + A_j(\mathbf{x} - \mathbf{w}_j - \mathbf{x}_f + \mathbf{x}_i) \times \frac{(\mathbf{x}_f - \mathbf{x}_i)^T}{\|\mathbf{x}_f - \mathbf{x}_i\|^2} \\ &= A_j + (\Delta\theta - A_j\Delta\mathbf{x})\frac{\Delta\mathbf{x}^T}{\|\Delta\mathbf{x}\|^2}. \end{aligned} \tag{7.9}$$

In eq. (7.8), the final error $\mathbf{x} - \mathbf{x}_f$ in Cartesian space is translated to an error in joint space via multiplication by A_j. This error is then added to θ_j to constitute the improved estimate θ^* (steepest descent minimisation of error).

In eq. (7.9), $\Delta\mathbf{x} = \mathbf{x}_f - \mathbf{x}_i$, i.e., the change in retinal coordinates of the end-effector due to the fine movement, and $\Delta\theta = A_j(\mathbf{x} - \mathbf{w}_j)$, i.e., the related joint angles during fine movement. Thus eq. (7.9) can be recognised as an error-correction rule of the Widrow-Hoff type for Jacobians A.

It appears that after 6,000 iterations the system approaches correct behaviour, and that after 30,000 learning steps no noteworthy deviation is present.

7.2 Robot arm dynamics

While end-effector positioning via sensor–robot coordination is an important problem to solve, the robot itself will not move without dynamic control of its limbs.

Again, accurate control with non-adaptive controllers is possible only when accurate models of the robot are available, and the robot is not too susceptible to wear-and-tear. This requirement has led to the current-day robots that are used in many factories. But the application of neural networks in this field changes these requirements.

One of the first neural networks which succeeded in doing dynamic control of a robot arm was presented by Kawato, Furukawa, and Suzuki [9]. They describe a neural network which generates motor commands from a desired trajectory in joint angles. Their system does not include the trajectory generation or the transformation of visual coordinates to body coordinates.

The network is extremely simple. In fact, the system is a feed-forward network, but by carefully choosing the basis functions, the network can be restricted to one learning layer such that finding the optimal is a trivial task. In this case, the basis functions are thus chosen that the function that is approximated is a linear combination of those basis functions. This approach is similar to that presented in section 4.4.

Dynamics model. The manipulator used consists of three joints as the manipulator in figure 7.1 without wrist joint. The desired trajectory $\theta_d(t)$, which is generated by another subsystem, is fed into the inverse-dynamics model (figure 7.5). The error between $\theta_d(t)$ and $\theta(t)$ is fed into the neural model.

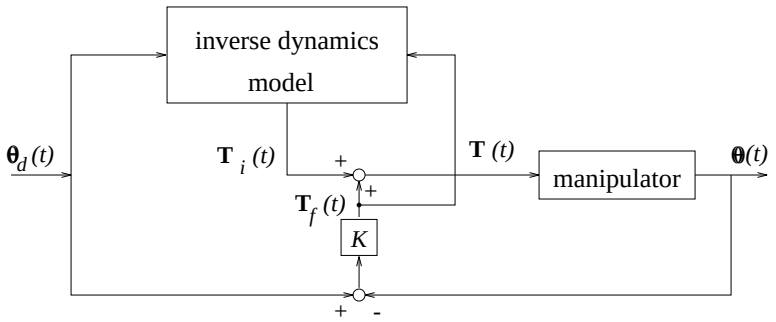

Figure 7.5: The neural model proposed by Kawato *et al.*

The neural model, which is shown in figure 7.6, consists of three perceptrons, each one feeding in one joint of the manipulator. The desired trajectory $\theta_d = (\theta_{d1}, \theta_{d2}, \theta_{d3})$ is fed into 13 nonlinear subsystems. The resulting signals are weighted and summed, such that

$$T_{ik}(t) = \sum_{l=1}^{13} w_{lk} x_{lk}, \qquad (k = 1, 2, 3), \tag{7.10}$$

with

$$\begin{aligned} x_{l1} &= f_l(\theta_{d1}(t), \theta_{d2}(t), \theta_{d3}(t)), \\ x_{l2} = x_{l3} &= g_l(\theta_{d1}(t), \theta_{d2}(t), \theta_{d3}(t)), \end{aligned}$$

and f_l and g_l as in table 7.1.

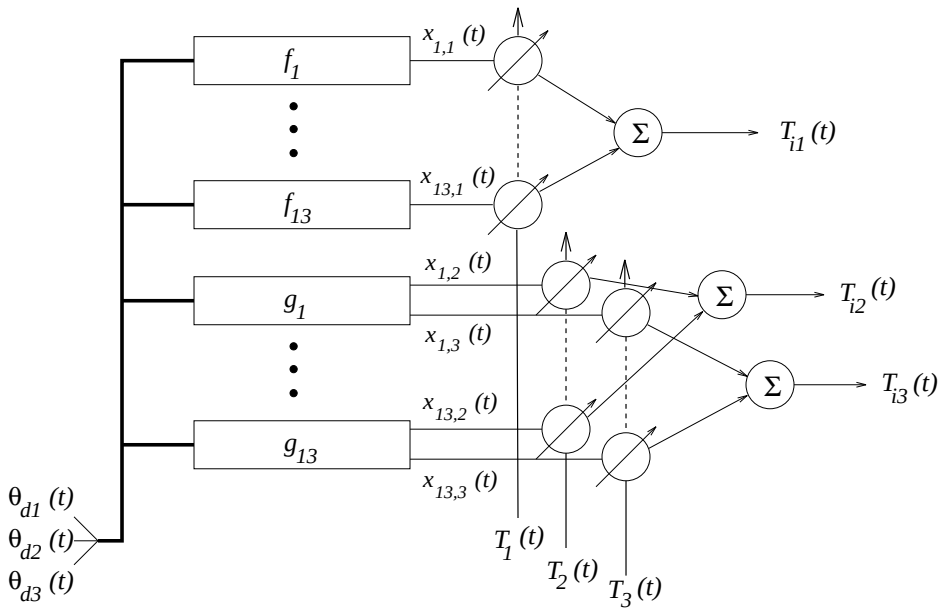

Figure 7.6: The neural network used by Kawato *et al.* There are three neurons, one per joint in the robot arm. Each neuron feeds from thirteen nonlinear subsystems. The upper neuron is connected to the rotary base joint (cf. joint 1 in figure 7.1), the other two neurons to joints 2 and 3.

The feedback torque $\mathbf{T}_f(t)$ in figure 7.5 consists of

$$T_{fk}(t) = K_{pk}(\theta_{dk}(t) - \theta_k(t)) + K_{vk}\frac{d\theta_k(t)}{dt}, \qquad (k = 1, 2, 3),$$
$$K_{vk} = 0 \quad \text{unless } |\theta_k(t) - \theta_{dk}(\text{objective point})| < \varepsilon.$$

The feedback gains $\mathbf{K}_p$ and $\mathbf{K}_v$ were computed as $(517.2, 746.0, 191.4)^T$ and $(16.2, 37.2, 8.4)^T$.

Next, the weights adapt using the delta rule

$$\gamma\frac{dw_{ik}}{dt} = x_{ik}T_1 = x_{ik}(T_{fk} - T_{ik}), \qquad (k = 1, 2, 3). \tag{7.11}$$

A desired move pattern is shown in figure 7.7. After 20 minutes of learning the feedback torques are nearly zero such that the system has successfully learned the transformation. Although the applied patterns are very dedicated, training with a repetitive pattern $\sin(\omega_k t)$, with $\omega_1 : \omega_2 : \omega_3 = 1 : \sqrt{2} : \sqrt{3}$ is also successful.

The usefulness of neural algorithms is demonstrated by the fact that novel robot architectures, which no longer need a very rigid structure to simplify the controller, are now constructed. For example, several groups [11,12] report on work with a pneumatic musculo-skeletal robot arm, with rubber actuators replacing the DC motors. The very complex dynamics and environmental temperature dependency of this arm make the use of non-adaptive algorithms impossible, where neural networks succeed.

l	$f_l(\theta_1,\theta_2,\theta_3)$	$g_l(\theta_1,\theta_2,\theta_3)$
1	$\ddot{\theta}_1$	$\ddot{\theta}_2$
2	$\ddot{\theta}_1 \sin^2\theta_2$	$\ddot{\theta}_3$
3	$\ddot{\theta}_1 \cos^2\theta_2$	$\ddot{\theta}_2 \cos\theta_3$
4	$\ddot{\theta}_1 \sin^2(\theta_2+\theta_3)$	$\ddot{\theta}_3 \cos\theta_3$
5	$\ddot{\theta}_1 \cos^2(\theta_2+\theta_3)$	$\dot{\theta}_1^2 \sin\theta_2 \cos\theta_2$
6	$\ddot{\theta}_1 \sin\theta_2 \sin(\theta_2+\theta_3)$	$\dot{\theta}_1^2 \sin(\theta_2+\theta_3) \cos(\theta_2+\theta_3)$
7	$\dot{\theta}_1\dot{\theta}_2 \sin\theta_2 \cos\theta_2$	$\dot{\theta}_1^2 \sin\theta_2 \cos(\theta_2+\theta_3)$
8	$\dot{\theta}_1\dot{\theta}_2 \sin(\theta_2+\theta_3) \cos(\theta_2+\theta_3)$	$\dot{\theta}_1^2 \cos\theta_2 \sin(\theta_2+\theta_3)$
9	$\dot{\theta}_1\dot{\theta}_2 \sin\theta_2 \cos(\theta_2+\theta_3)$	$\dot{\theta}_2^2 \sin\theta_3$
10	$\dot{\theta}_1\dot{\theta}_2 \cos\theta_2 \sin(\theta_2+\theta_3)$	$\dot{\theta}_3^2 \sin\theta_3$
11	$\dot{\theta}_1\dot{\theta}_3 \sin(\theta_2+\theta_3) \cos(\theta_2+\theta_3)$	$\dot{\theta}_2\dot{\theta}_3 \sin\theta_3$
12	$\dot{\theta}_1\dot{\theta}_3 \sin\theta_2 \cos(\theta_2+\theta_3)$	$\dot{\theta}_2$
13	$\dot{\theta}_1$	$\dot{\theta}_3$

Table 7.1: Nonlinear transformations used in the Kawato model.

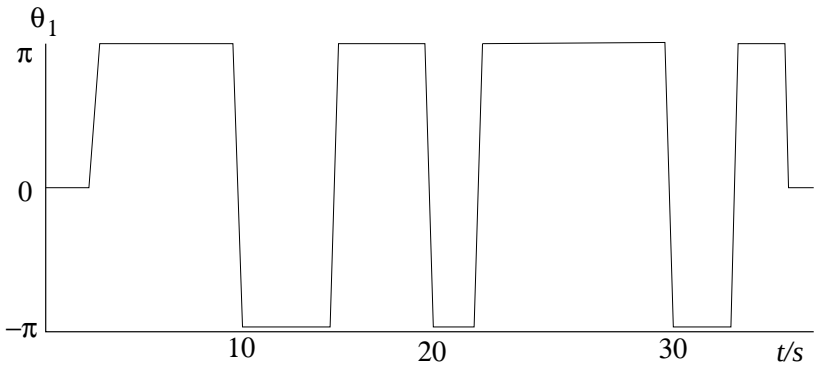

Figure 7.7: The desired joint pattern for joints 1. Joints 2 and 3 have similar time patterns.

7.3 Mobile robots

In the previous sections some applications of neural networks on robot arms were discussed. In this section we focus on mobile robots. Basically, the control of a robot arm and the control of a mobile robot is very similar: the (hierarchical) controller first plans a path, the path is transformed from Cartesian (world) domain to the joint or wheel domain using the inverse kinematics of the system and finally a dynamic controller takes care of the mapping from set-points in this domain to actuator signals. However, in practice the problems with mobile robots occur more with path-planning and navigation than with the dynamics of the system. Two examples will be given.

7.3.1 Model based navigation

Jorgensen [13] describes a neural approach for path-planning. Robot path-planning techniques can be divided into two categories. The first, called *local planning* relies on information available from the current 'viewpoint' of the robot. This planning is important, since it is able to deal with fast changes in the environment. Unfortunately, by itself local data is generally not adequate since occlusion in the line of sight can cause the robot to wander into dead end corridors or choose non-optimal routes of travel. The second situation is called

global path-planning, in which case the system uses global knowledge from a topographic map previously stored into memory. Although global planning permits optimal paths to be generated, it has its weakness. Missing knowledge or incorrectly selected maps can invalidate a global path to an extent that it becomes useless. A possible third, 'anticipatory' planning combined both strategies: the local information is constantly used to give a best guess what the global environment may contain.

Jorgensen investigates two issues associated with neural network applications in unstructured or changing environments. First, can neural networks be used in conjunction with direct sensor readings to associatively approximate global terrain features not observable from a single robot perspective. Secondly, is a neural network fast enough to be useful in path relaxation planning, where the robot is required to optimise motion and situation sensitive constraints.

For the first problem, the system had to store a number of possible sensor maps of the environment. The robot was positioned in eight positions in each room and 180° sonar scans were made from each position. Based on these data, for each room a map was made. To be able to represent these maps in a neural network, the map was divided into 32×32 grid elements, which could be projected onto the 32×32 nodes neural network. The maps of the different rooms were 'stored' in a Hopfield type of network. In the operational phase, the robot wanders around, and enters an unknown room. It makes one scan with the sonar, which provides a partial representation of the room map (see figure 7.8). This pattern is clamped onto the network, which will regenerate the best fitting pattern. With this information a global path-planner can be used. The results which are presented in the paper are not very encouraging. With a network of 32×32 neurons, the total number of weights is 1024 squared, which costs more than 1 Mbyte of storage if only one byte per weight is used. Also the speed of the recall is low: Jorgensen mentions a recall time of more than two and a half hour on an IBM AT, which is used on board of the robot.

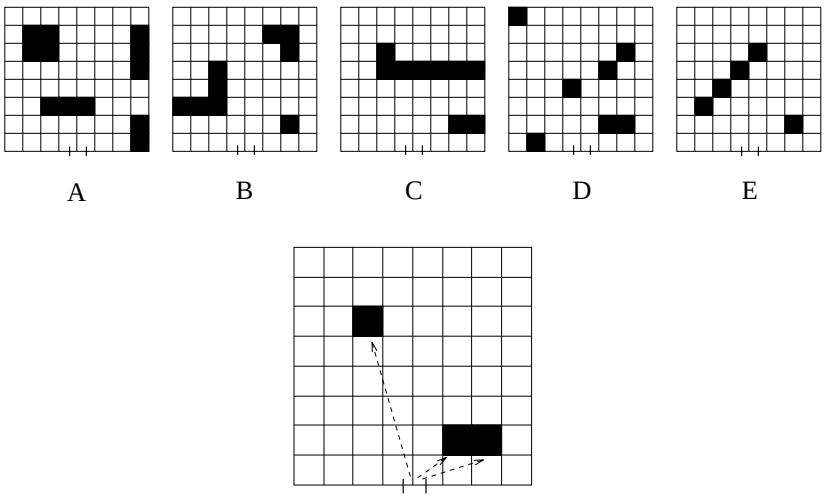

Figure 7.8: Schematic representation of the stored rooms, and the partial information which is available from a single sonar scan.

Also the use of a simulated annealing paradigm for path planning is not proving to be an effective approach. The large number of settling trials (> 1000) is far too slow for real time, when the same functions could be better served by the use of a potential field approach or distance transform.

7.3.2 Sensor based control

Very similar to the sensor based control for the robot arm, as described in the previous sections, a mobile robot can be controlled directly using the sensor data. Such an application has been developed at Carnegy-Mellon by Touretzky and Pomerleau. The goal of their network is to drive a vehicle along a winding road. The network receives two type of sensor inputs from the sensory system. One is a 30×32 (see figure 7.9) pixel image from a camera mounted on the roof of the vehicle, where each pixel corresponds to an input unit of the network. The other input is an 8×32 pixel image from a laser range finder. The activation levels of units in the range finder's retina represent the distance to the corresponding objects.

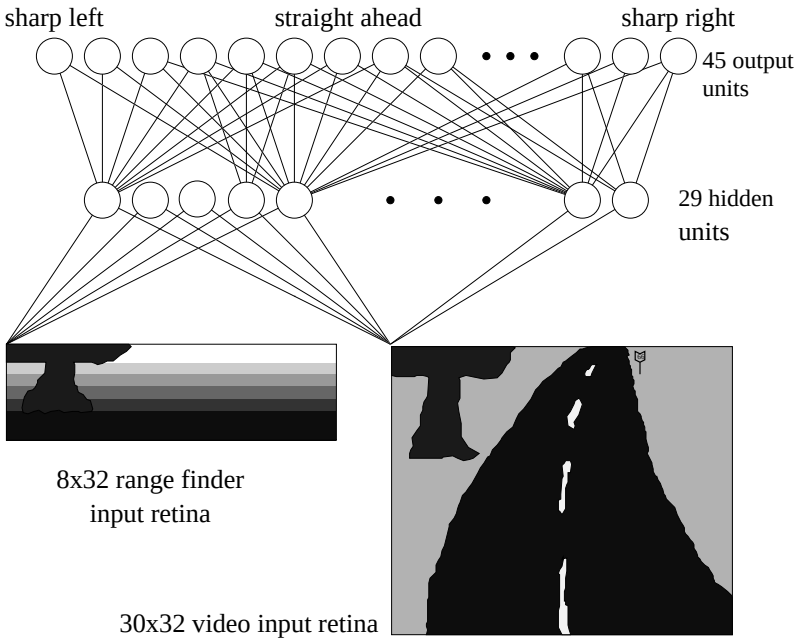

Figure 7.9: The structure of the network for the autonomous land vehicle.

The network was trained by presenting it samples with as inputs a wide variety of road images taken under different viewing angles and lighting conditions. 1,200 Images were presented, 40 times each while the weights were adjusted using the back-propagation principle. The authors claim that once the network is trained, the vehicle can accurately drive (at about 5 km/hour) along '... a path though a wooded area adjoining the Carnegie Mellon campus, under a variety of weather and lighting conditions.' The speed is nearly twice as high as a non-neural algorithm running on the same vehicle.

Although these results show that neural approaches can be possible solutions for the sensor based control problem, there still are serious shortcomings. In simulations in our own laboratory, we found that networks trained with examples which are provided by human operators are not always able to find a correct approximation of the human behaviour. This is the case if the human operator uses other information than the network's input to generate the steering signal. Also the learning of in particular back-propagation networks is dependent on the sequence of samples, and, for all supervised training methods, depends on the distribution of the training samples.

References

[1] J. J. Craig, *Introduction to Robotics* (2nd ed.), Addison-Wesley Publishing Company, 1989.

[2] D. Psaltis, A. Sideris & A. A. Yamamura, "A multilayer neural network controller," *IEEE Control Systems Magazine* 8 (April 1988), 17–21.

[3] B. J. A. Kröse, M. J. van der Korst & F. C. A. Groen, "Learning strategies for a vision based neural controller for a robot arm," in *IEEE International Workshop on Intelligent Motor Control*, O. Kaynak, ed., Istanbul, August 1990, 199–203.

[4] P. P. van der Smagt & B. J. A. Kröse, "A real-time learning neural robot controller," in *Proceedings of the 1991 International Conference on Artificial Neural Networks*, Espoo, Finland, June 1991, 351–356.

[5] P. van der Smagt, B. J. A. Kröse & F. C. A. Groen, "Using time-to-contact to guide a robot manipulator," in *Proceedings of the 1992 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Raleigh, N. C., June 1992, 177–182.

[6] A. Jansen, P. P. van der Smagt & F. C. A. Groen, "Nested networks for robot control," in *Neural Network Applications*, A. F. Murray, ed., Kluwer Academic Publishers, 1993 (in press).

[7] H. J. Ritter, T. M. Martinetz & K. J. Schulten, "Topology-conserving maps for learning visuo-motor-coordination," *Neural Networks* 2 (1989), 159–168.

[8] H. Ritter, T. Martinetz & K. Schulten, *Neural Computation and Self-Organizing Maps*, Addison-Wesley, 1992, *Neuronale Netze*, Addison-Wesley, 1991.

[9] M. Kawato, K. Furukawa & R. Suzuki, "A hierarchical neural-network model for control and learning of voluntary movement," *Biological Cybernetics* 57 (1987), 169–185.

[10] W. T. Miller III, "Real-time application of neural networks for sensor-based control of robots with vision," *IEEE Transactions on Systems, Man, and Cybernetics* 19 (August 1989), 825–831.

[11] M. Katayama & M. Kawato, "A parallel-hierarchical neural network model for motor control of a musculo-skeletal system," ATR Auditory and Visual Perception Research Laboratories, TR–A–0145, April 1992.

[12] T. Hesselroth, K. Sarkar, P. van der Smagt & K. Schulten, "Neural network control of a pneumatic robot arm," Theoretical Biophysics, University of Illinois at Urbana/Champaign, UIUC–BI–TB–92–15 (*IEEE Systems, Man, and Cybernetics*, to be published), August 1992.

[13] C. C. Jorgensen, "Neural network representation of sensor graphs in autonomous robot path planning," in *IEEE First International Conference on Neural Networks* #IV, San Diego, June 1987, 507–515.

8 Vision

8.1 Introduction

In this chapter we illustrate some applications of neural networks which deal with visual information processing. In the neural literature we find roughly two types of problems: the modelling of biological vision systems and the use of artificial neural networks for machine vision. We will focus on the latter.

The primary goal of machine vision is to obtain information about the environment by processing data from one or multiple two-dimensional arrays of intensity values ('images'), which are projections of this environment on the system. This information can be of different nature:

- recognition: the classification of the input data in one of a number of possible classes;
- geometric information about the environment, which is important for autonomous systems;
- compression of the image for storage and transmission.

Often a distinction is made between low level (or early) vision, intermediate level vision and high level vision. Typical low-level operations include image filtering, isolated feature detection and consistency calculations. At a higher level segmentation can be carried out, as well as the calculation of invariants. The high level vision modules organise and control the flow of information from these modules and combine this information with high level knowledge for analysis.

Computer vision already has a long tradition of research, and many algorithms for image processing and pattern recognition have been developed. There appear to be two computational paradigms that are easily adapted to massive parallelism: local calculations and neighbourhood functions. Calculations that are strictly localised to one area of an image are obviously easy to compute in parallel. Examples are filters and edge detectors in early vision. A cascade of these local calculations can be implemented in a feed-forward network.

The first section describes feed-forward networks for vision. Section 8.3 shows how back-propagation can be used for image compression. In the same section, it is shown that the PCA neuron is ideally suited for image compression. Finally, sections 8.4 and 8.5 describe the cognitron for optical character recognition, and relaxation networks for calculating depth from stereo images.

8.2 Feed-forward types of networks

The early feed-forward networks as the perceptron and the adaline were essentially designed to be be visual pattern classifiers. In principle a multi-layer feed-forward network is able to learn to classify all possible input patterns correctly, but an enormous amount of connections is needed (for the perceptron, Minsky showed that many problems can only be solved if each hidden unit is connected to all inputs). The question is whether such systems can still be regarded as 'vision' systems. No use is made of the spatial relationships in the input patterns and the problem of classifying a set of 'real world' images is the same as the problem of classifying a set of artificial random dot patterns which are, according to Smeulders [Smeulders needs no reference] no 'images.' For that reason, most successful neural vision applications combine self-organising techniques with a feed-forward architecture, such as for example the neocognitron [1], described in section 8.4. The neocognitron performs the mapping from input data to output data by a layered structure in which at each stage increasingly complex features are extracted. The lower layers extract local features such as a line at a particular orientation and the higher layers aim to extract more global features.

Also there is the problem of translation invariance: the system has to classify a pattern correctly independent of the location on the 'retina.' However, a standard feed-forward network considers an input pattern which is translated as a totally 'new' pattern. Several attempts have been described to overcome this problem, one of the more exotic ones by Widrow [2] as a layered structure of adalines.

8.3 Self-organising networks for image compression

In image compression one wants to reduce the number of bits required to store or transmit an image. We can either require a perfect reconstruction of the original or we can accept a small deterioration of the image. The former is called a lossless coding and the latter a lossy coding. In this section we will consider lossy coding of images with neural networks.

The basic idea behind compression is that an n-dimensional stochastic vector $\mathbf{n}$, (part of) the image, is transformed into an m-dimensional stochastic vector

$$\mathbf{m} = \mathbf{T}\mathbf{n}. \tag{8.1}$$

After transmission or storage of this vector $\tilde{\mathbf{m}}$, a discrete version of $\mathbf{m}$, we can make a reconstruction of $\mathbf{n}$ by some sort of inverse transform $\tilde{\mathbf{T}}$ so that the reconstructed signal equals

$$\tilde{\mathbf{n}} = \tilde{\mathbf{T}}\tilde{\mathbf{n}}. \tag{8.2}$$

The error of the compression and reconstruction stage together can be given as

$$\epsilon = E\left[\|\mathbf{n} - \tilde{\mathbf{n}}\|\right]. \tag{8.3}$$

There is a trade-off between the dimensionality of $\mathbf{m}$ and the error ϵ. As one decreases the dimensionality of $\mathbf{m}$ the error increases and vice versa, i.e., a better compression leads to a higher deterioration of the image. The basic problem of compression is finding $\mathbf{T}$ and $\tilde{\mathbf{T}}$ such that the information in $\mathbf{m}$ is as compact as possible with acceptable error ϵ. The definition of acceptable depends on the application area.

The cautious reader has already concluded that dimension reduction is in itself not enough to obtain a compression of the data. The main importance is that some aspects of an image

are more important for the reconstruction then others. For example, the mean grey level and generally the low frequency components of the image are very important, so we should code these features with high precision. Other, like high frequency components, are much less important so these can be coarse-coded. So, when we reduce the dimension of the data, we are actually trying to concentrate the information of the data in a few numbers (the low frequency components) which can be coded with precision, while throwing the rest away (the high frequency components). In this section we will consider coding an image of 256×256 pixels. It is a bit tedious to transform the whole image directly by the network. This requires a huge amount of neurons. Because the statistical description over parts of the image is supposed to be stationairy, we can break the image into 1024 blocks of size 8×8, which is large enough to entail a local statistical description and small enough to be managed. These blocks can then be coded separately, stored or transmitted, where after a reconstruction of the whole image can be made based on these coded 8×8 blocks.

8.3.1 Back-propagation

The process above can be interpreted as a 2-layer neural network. The inputs to the network are the 8×8 patters and the desired outputs are the same 8×8 patterns as presented on the input units. This type of network is called an auto-associator.

After training with a gradient search method, minimising ϵ, the weights between the first and second layer can be seen as the coding matrix $\mathbf{T}$ and between the second and third as the reconstruction matrix $\tilde{\mathbf{T}}$.

If the number of hidden units is smaller then the number of input (output) units, a compression is obtained, in other words we are trying to squeeze the information through a smaller channel namely the hidden layer.

This network has been used for the recognition of human faces by Cotrell [cotrell]. He uses an input and output layer of 64×64 units (!) on which he presented the whole face at once. The hidden layer, which consisted of 64 units, was classified with another *network* by means of a delta rule. Is this complex network invariant to translations in the input?

8.3.2 Linear networks

It is known from statistics that the optimal transform from an n-dimensional to an m-dimensional stochastic vector, optimal in the sense that ϵ contains the lowest energy possible, equals the concatination of the first m eigenvectors of the correlation matrix $\mathbf{R}$ of $\mathbf{N}$. So if $(\mathbf{e}_1, \mathbf{e}_2, .., \mathbf{e}_n)$ are the eigenvectors of $\mathbf{R}$, ordered in decreasing corresponding eigenvalue, the transformation matrix is given as $\mathbf{T} = [\mathbf{e}_1 \mathbf{e}_2 \ldots \mathbf{e}_2]^T$.

In section 5.3.1 a linear neuron with a normalised Hebbian learning rule was able to learn the eigenvectors of the correlation matrix of the input patterns. The definition of the optimal transform given above, suits exactly in the PCA network we have described.

So we end up with a $64 \times m \times 64$ network, where m is the desired number of hidden units which is coupled to the total error ϵ. Since the eigenvalues are ordered in decreasing values, which are the outputs of the hidden units, the hidden units are ordered in importance for the reconstruction.

Sanger [3] used this implementation for image compression. The test image is shown in figure 8.1. It is 256×256 with 8 bits/pixel.

Figure 8.1: Input image for the network. The image is divided into 8×8 blocks which are fed to the network.

After training the image four times, thus generating 4×1024 learning patterns of size 8×8, the weights of the network converge into figure 8.2.

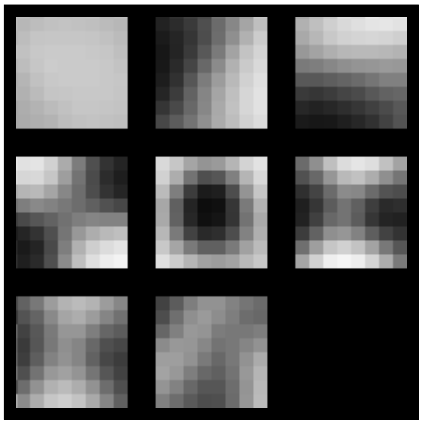

ordering of the neurons:

neuron 0	neuron 1	neuron 2
neuron 3	neuron 4	neuron 5
neuron 6	neuron 7	(unused)

Figure 8.2: Weights of the PCA network. The final weights of the network trained on the test image. For each neuron, an 8×8 rectangle is shown, in which the grey level of each of the elements represents the value of the weight. Dark indicates a large weight, light a small weight.

8.3.3 Principal components as features

If parts of the image are very characteristic for the scene, like corners, lines, shades etc., one speaks of features of the image. The extraction of features can make the image understanding task on a higher level much easer. If the image analysis is based on features it is very important that the features are tolerant of noise, distortion etc.

From an image compression viewpoint it would be smart to code these features with as little bits as possible, just because the definition of features was that they occur frequently in the image.

So one can ask oneself if the two described compression methods also extract features from the image. Indeed this is true and can most easily be seen in fig. 8.2. It might not be clear directly, but one can see that the weights are converged to:

- neuron 0: the mean grey level;
- neuron 1 and neuron 2: the first order gradients of the image;
- neuron 3 ... neuron 5: second orders derivates of the image.

The features extracted by the principal compentent network are the gradients of the image.

8.4 The cognitron and neocognitron

Yet another type of unsupervised learning is found in the *cognitron*, introduced by Fukushima as early as 1975 [4]. This network, with primary applications in pattern recognition, was improved at a later stage to incorporate scale, rotation, and translation invariance resulting in the *neocognitron* [1], which we will not discuss here.

8.4.1 Description of the cells

Central in the cognitron is the type of neuron used. Whereas the Hebb synapse (unit i, say), which is used in the perceptron model, increases an incoming weight (w_{ij}) if and only if the incoming signal (a_j) is high and a control input is high, the synapse introduced by Fukushima increases (the absolute value of) its weight ($|w_{ij}|$) only if it has positive input a_j and a maximum activation value $a_i = \max(a_{k_1}, a_{k_2}, \ldots, a_{k_n})$, where $k_1, k_2, \ldots, k_n$ are all 'neighbours' of i. Note that this learning scheme is competitive and unsupervised, and the same type of neuron has, at a later stage, been used in the competitive learning network (section 5.1) as well as in other unsupervised networks.

Fukushima distinguishes between *excitatory* inputs and *inhibitory* inputs. The output of an excitatory cell u is given by[1]

$$u(i) = \mathcal{F}\left[\frac{1+e}{1+h} - 1\right] = \mathcal{F}\left[\frac{e-h}{1+h}\right], \tag{8.4}$$

where e is the excitatory input from u-cells and h the inhibitory input from v-cells. The activation function is

$$\mathcal{F}(x) = \begin{cases} x & \text{if } x \geq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{8.5}$$

When the inhibitory input is small, i.e., $h \ll 1$, $u(i)$ can be approximated by $u(i) = e - h$, which agrees with the formula for a conventional linear threshold element (with a threshold of zero).

When both the excitatory and inhibitory inputs increase in proportion, i.e.,

$$e = \epsilon x, \qquad h = \eta x \tag{8.6}$$

(ϵ, η constants) and $\epsilon > \eta$, then eq. (8.4) can be transformed into

$$u(i) = \frac{(\epsilon - \eta)x}{1 + \eta x} = \frac{\epsilon - \eta}{2\eta}\Big(1 + \tanh(\tfrac{1}{2}\log \eta x)\Big) \tag{8.7}$$

i.e., a squashing function as in figure 2.2.

[1] Here our notational system fails. We adhere to Fukushima's symbols.

8.4.2 Structure of the cognitron

The basic structure of the cognitron is depicted in figure 8.3.

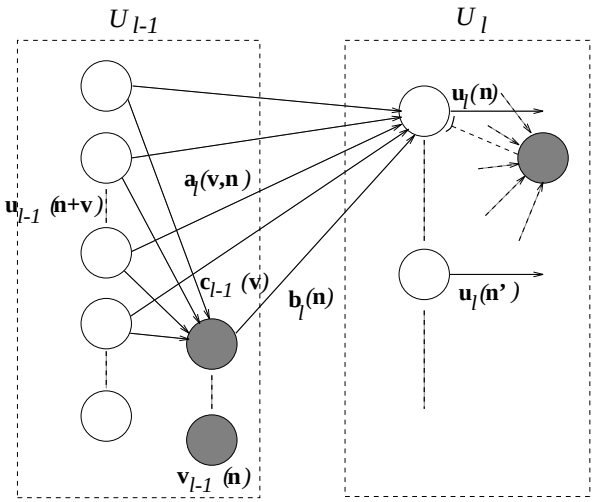

Figure 8.3: The basic structure of the cognitron.

The cognitron has a multi-layered structure. The l-th layer U_l consists of excitatory neurons $u_l(\mathbf{n})$ and inhibitory neurons $v_l(\mathbf{n})$, where $\mathbf{n} = (n_x, n_y)$ is a two-dimensional location of the cell.

A cell $u_l(\mathbf{n})$ receives inputs via modifiable connections $a_l(\mathbf{v}, \mathbf{n})$ from neurons $u_{l-1}(\mathbf{n} + \mathbf{v})$ and connections $b_l(\mathbf{n})$ from neurons $v_{l-1}(\mathbf{n})$, where $\mathbf{v}$ is in the connectable area (cf. area of attention) of the neuron. Furthermore, an inhibitory cell $v_{l-1}(\mathbf{n})$ receives inputs via fixed excitatory connections $c_{l-1}(\mathbf{v})$ from the neighbouring cells $u_{l-1}(\mathbf{n} + \mathbf{v})$, and yields an output equal to its weighted input:

$$v_{l-1}(\mathbf{n}) = \sum_{\mathbf{v}} c_{l-1}(\mathbf{v})\, u_{l-1}(\mathbf{n} + \mathbf{v}). \tag{8.8}$$

where $\sum_{\mathbf{v}} c_{l-1}(\mathbf{v}) = 1$ and are fixed.

It can be shown that the growing of the synapses (i.e., modification of the a and b weights) ensures that, if an excitatory neuron has a relatively large response, the excitatory synapses grow faster than the inhibitory synapses, and vice versa.

Receptive region

For each cell in the cascaded layers described above a connectable area must be established. A connection scheme as in figure 8.4 is used: a neuron in layer U_l connects to a small region in layer U_{l-1}.

If the connection region of a neuron is constant in all layers, a too large number of layers is needed to cover the whole input layer. On the other hand, increasing the region in later layers results in so much overlap that the output neurons have near identical connectable areas and thus all react similarly. This again can be prevented by increasing the size of the vicinity area in which neurons compete, but then only one neuron in the output layer will react to some input stimulus. This is in contradiction with the behaviour of biological brains.

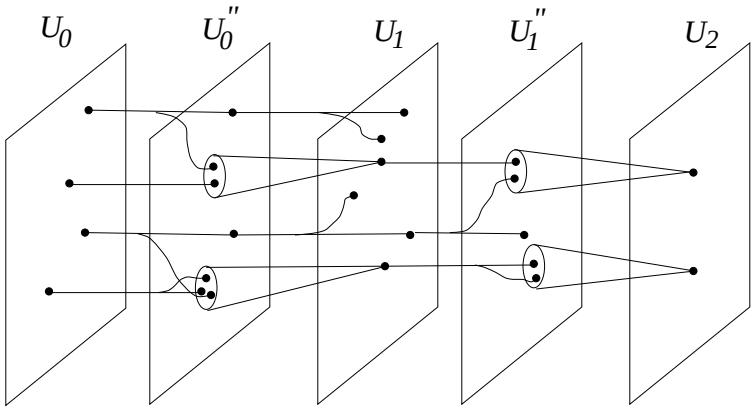

Figure 8.4: Cognitron receptive regions.

A solution is to distribute the connections probabilistically such that connections with a large deviation are less numerous.

8.4.3 Simulation results

In order to illustrate the working of the network, a simulation has been run with a four-layered network with 16×16 neurons in each layer. The network is trained with four learning patterns, consisting of a vertical, a horizontal, and two diagonal lines. Figure 8.5 shows the activation levels in the layers in the first two learning iterations.

After 20 learning iterations, the learning is halted and the activation values of the neurons in layer 4 are fed back to the input neurons; also, the maximum output neuron alone is fed back, and thus the input pattern is 'recognised' (see figure 8.6).

8.5 Relaxation types of networks

As demonstrated by the Hopfield network, a relaxation process in a connectionist network can provide a powerful mechanism for solving some difficult optimisation problems. Many vision problems can be considered as optimisation problems, and are potential candidates for an implementation in a Hopfield-like network. A few examples that are found in the literature will be mentioned here.

8.5.1 Depth from stereo

By observing a scene with two cameras one can retrieve depth information out of the images by finding the pairs of pixels in the images that belong to the same point of the scene. The calculation of the depth is relatively easy; finding the correspondences is the main problem. One solution is to find features such as corners and edges and match those, reducing the computational complexity of the matching. Marr [5] showed that the correspondence problem can be solved correctly when taking into account the physical constraints underlying the process. Three matching criteria were defined:

- Compatibility: Two descriptive elements can only match if they arise from the same physical marking (corners can only match with corners, 'blobs' with 'blobs,' etc.);

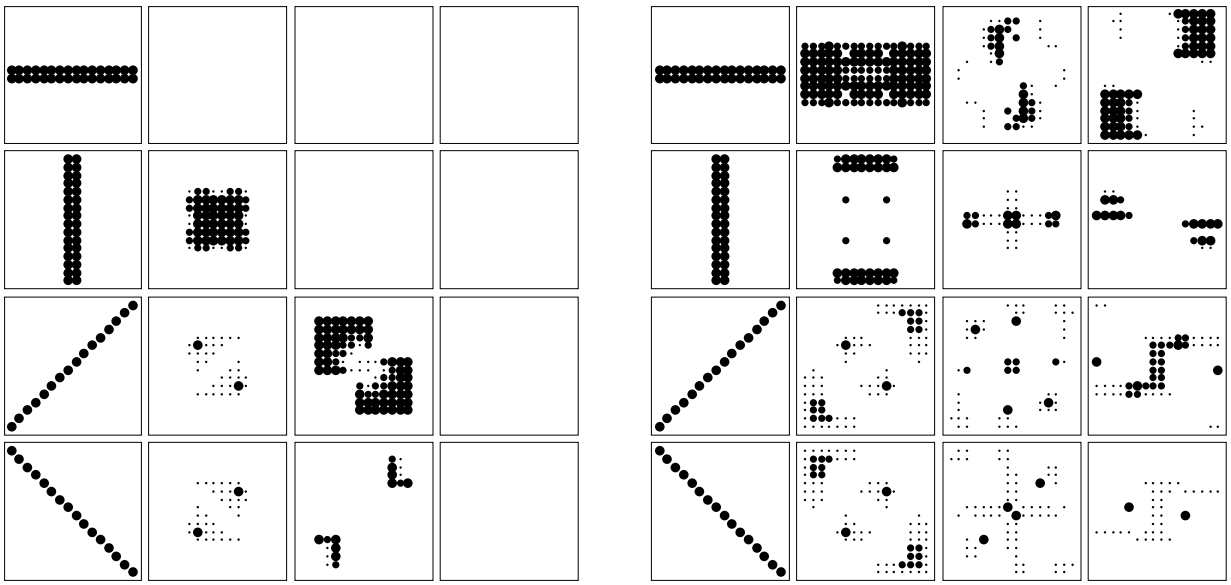

a. b.

Figure 8.5: Two learning iterations in the cognitron.
Four learning patterns (one in each row) are shown in iteration 1 (a.) and 2 (b.). Each column in a. and b. shows one layer in the network. The activation level of each neuron is shown by a circle. A large circle means a high activation. In the first iteration (a.), a structure is already developing in the second layer of the network. In the second iteration, the second layer can distinguish between the four patterns.

- Uniqueness: Almost always a descriptive element from the left image corresponds to exactly one element in the right image and vice versa;
- Continuity: The disparity of the matches varies smoothly almost everywhere over the image.

Marr's 'cooperative' algorithm (also a 'non-cooperative' or local algorithm has been described [5]) is able to calculate the disparity map from which the depth can be reconstructed. This algorithm is some kind of neural network, consisting of neurons $N(x, y; d)$, where neuron $N(x, y; d)$ represents the hypothesis that pixel (x, y) in the left image corresponds with pixel $(x + d, y)$ in the right image. The update function is

$$N^{t+1}(x, y; d) = \sigma \left(\sum_{\substack{x', y'; d' \in \\ S(x,y;d)}} N^t(x', y'; d') - \epsilon \sum_{\substack{x', y'; d' \in \\ O(x,y;d)}} N^t(x', y'; d') + N^0(x, y; d) \right). \tag{8.9}$$

Here, ϵ is an inhibition constant, σ is a threshold function, $S(x, y; d)$ is the local excitatory neighbourhood, and $O(x, y; d)$ is the local inhibitory neighbourhood, which are chosen as follows:

$$S(x, y; d) = \{ (r, s, t) \mid (r = x \vee r = x - d) \wedge s = y \}, \tag{8.10}$$

$$O(x, y; d) = \{ (r, s, t) \mid d = t \wedge \| (r, s) - (x, y) \| \leq w \}. \tag{8.11}$$

The network is loaded with the cross correlation of the images at first: $N^0(x, y; d) = I_l(x, y)I_r(x + d, y)$, where I_l and I_r are the intensity matrices of the left and right image respectively. This network state represents all possible matches of pixels. Then the set of possible matches is reduced by recursive application of the update function until the state of the network is stable.

The algorithm converges in about ten iterations. Then the disparity of a pixel (x, y) is displayed by the firing neuron in the set $\{N(r, s; d) \mid r = x, s = y\}$. In each of these sets there should be exactly one neuron firing, but if the algorithm could not compute the exact disparity, for instance at hidden contours, there may be zero or more than one neurons firing. Improvements on this algorithm have been described [6].

8.5.2 Image restoration and image segmentation

The restoration of degraded images is a branch of digital picture processing closely related to image segmentation and boundary finding. An analysis of the major applications and procedures may be found in [7]. An algorithm which is based on the minimisation of an energy function and can very well be parallelised is given by Geman and Geman [8]. Their approach is based on stochastic modelling, in which image samples are considered to be generated by a random process that changes its statistical properties from region to region. The random process that that generates the image samples is a two-dimensional analogue of a Markov process, called a Markov random field. Image segmentation is then considered as a statistical estimation problem in which the system calculates the optimal estimate of the region boundaries for the input image. Simultaneously estimation of the region properties and boundary properties has to be performed, resulting in a set of nonlinear estimation equations that define the optimal estimate of the regions. The system must find the maximum a posteriori probability estimate of the image segmentation. Geman and Geman showed that the problem can be recast into the minimisation of an energy function, which, in turn, can be solved approximately by optimisation techniques such as simulated annealing. The interesting point is that simulated annealing can be implemented using a network with local connections, in which the network iterates into a global solution using these local operations.

8.5.3 Silicon retina

Mead and his co-workers [9] have developed an analogue VLSI vision preprocessing chip modelled after the retina. The design not only replicates many of the important functions of the first stages of retinal processing, but it does so by replicating in a detailed way both the structure and dynamics of the constituent biological units. The logarithmic compression from photon input to output signal is accomplished by analogue circuits, while similarly space and time averaging and temporal differentiation are accomplished by analogue processes and a resistive network (see section 10.2.1).

References

[1] K. Fukushima, "Neocognitron: A hierarchical neural network capable of visual pattern recognition," *Neural Networks* 1 (1988), 119–130.

[2] B. Widrow, R. G. Winter & R. A. Baxter, "Layered neural nets for pattern recognition," *IEEE Transactions on Acoustics, Speech, and Signal Processing* 36 (1988), 1109–1117.

[3] T. D. Sanger, "Optimal unsupervised learning in a single-layer linear feedforward neural network," in *Neural Networks*, (2), 459–473.

[4] K. Fukushima, "Cognitron: A self-organizing multilayered neural network," *Biological Cybernetics* 20 (1975), 121–136.

[5] D. Marr, *Vision*, W. H. Freeman, San Francisco, 1982.

[6] H. Braun & T. Müller, "Enhancing Marr's cooperative algorithm," in *Proceedings of the International Conference on Neural Networks* #1, Paris, July 1990, 38–41.

[7] A. Rosenfeld & A. C. Kak, *Digital Picture Processing*, Academic Press, New York, 1982.

[8] S. Geman & D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6 (1984), 721–724.

[9] C. Mead, *Analog VLSI and Neural Systems*, Addison-Wesley, Reading, MA, 1989.

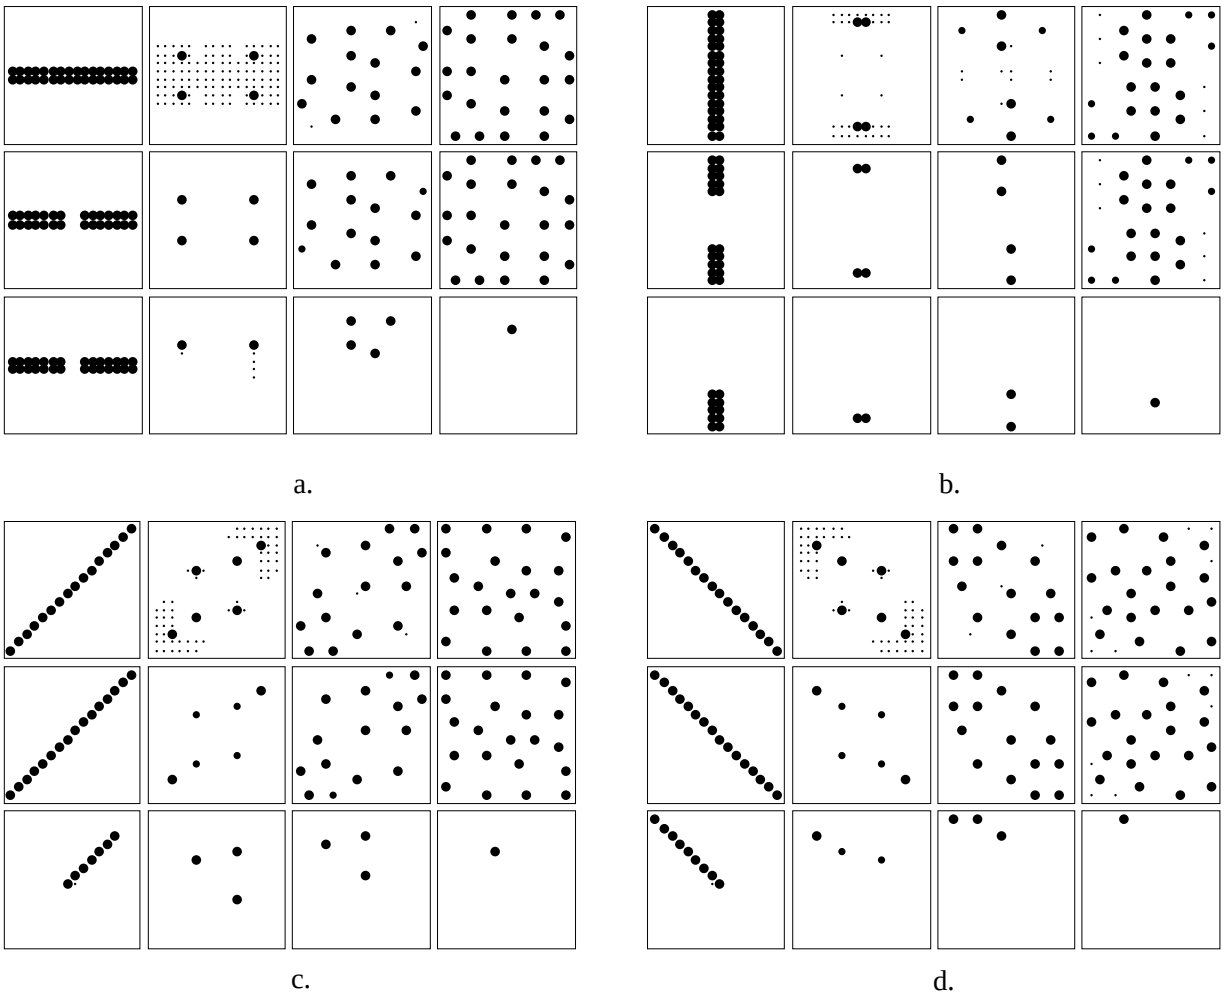

Figure 8.6: Feeding back activation values in the cognitron.

The four learning patterns are now successively applied to the network (row 1 of figures a–d). Next, the activation values of the neurons in layer 4 are fed back to the input (row 2 of figures a–d). Finally, all the neurons except the most active in layer 4 are set to 0, and the resulting activation values are again fed back (row 3 of figures a–d). After as little as 20 iterations, the network has shown to be rather robust.

Part IV

IMPLEMENTATIONS

Implementation of neural networks can be divided into three categories:

- software simulation;
- (hardware) emulation[2];
- hardware implementation.

The distinction between the former two categories is not clear-cut. We will use the term *simulation* to describe software packages which can run on a variety of host machines (e.g., PYGMALION, the Rochester Connectionist Simulator, NeuralWare, Nestor, etc.). Implementation of neural networks on general-purpose multi-processor machines such as the Connection Machine, the Warp, transputers, etc., will be referred to as *emulation*. *Hardware implementation* will be reserved for neuro-chips and the like which are specifically designed to run neural networks.

To evaluate and provide a taxonomy of the neural network simulators and emulators discussed, we will use the descriptors of table 8.1 (cf. [2]).

1. **Equation type**: many networks are defined by the type of equation describing their operation. For example, Grossberg's ART (cf. section 5.4) is described by the differential equation

$$\frac{dx_i}{dt} = -Ax_i + (B - x_i)I_i - x_i \sum_{j \neq i} I_j, \tag{8.12}$$

   in which $-Ax_i$ is a decay term, $+BI_i$ is an external input, $-x_iI_i$ is a normalisation term, and $-x_i \sum_{j \neq i} I_j$ is a neighbour shut-off term for competition. Although differential equations are very powerful, they require a high degree of flexibility in the software and hardware and are thus difficult to implement on special-purpose machines. Other types of equations are, e.g., difference equations as used in the description of Kohonen's topological maps (see section 5.2), and optimisation equations as used in back-propagation networks.
2. **Connection topology**: the design of most general purpose computers includes random access memory (RAM) such that each memory position can be accessed with uniform speed. Such designs always present a trade-off between size of memory and speed of access. The topology of neural networks can be matched in a hardware design with fast local interconnections instead of global access. Most networks are more or less local in their interconnections, and a global RAM is unnecessary.
3. **Processing schema**: although most artificial neural networks use a synchronous update, i.e., the output of the network depends on the previous state of the network, asynchronous update, in which components or blocks of components can be updated one by one, can be implemented much more efficiently. Also, continuous update is a possibility encountered in some implementations.
4. **Synaptic transmission mode**: most artificial neural networks have a transmission mode based on the neuronal activation values multiplied by synaptic weights. In these models, the propagation time from one neuron to another is neglected. On the other hand, biological neurons output a series of pulses in which the frequency determines the neuron output, such that propagation times are an essential part of the model. Currently, models arise which make use of temporal synaptic transmission [3,4].

Table 8.1: A possible taxonomy.

[2] The term *emulation* (see, e.g., [1] for a good introduction) in computer design means running one computer to execute instructions specific to another computer. It is often used to provide the user with a machine which is seemingly compatible with earlier models.

The following chapters describe general-purpose hardware which can be used for neural network applications, and neuro-chips and other dedicated hardware.

References

[1] E. G. Mallach, "Emulator architecture," *Computer* 8 (August 1975), 24–32.

[2] AFCEA International Press, *DARPA Neural Network Study*, 1988.

[3] A. F. Murray, "Pulse arithmetic in VLSI neural networks," *IEEE Micro* 9 (December 1989), 64–74.

[4] M. S. Tomlinson Jr. & D. J. Walker, "DNNA: A digital neural network architecture," in *Proceedings of the International Conference on Neural Networks* #II, Paris, July 1990, 589–592.

9 General Purpose Hardware

Parallel computers [1] can be divided into several categories. One important aspect is the *granularity* of the parallelism. Broadly speaking, the granularity ranges from *coarse-grain* parallelism, typically up to ten processors, to *fine-grain* parallelism, up to thousands or millions of processors.

Both fine-grain and coarse-grain parallelism is in use for emulation of neural networks. The former model, in which one or more processors can be used for each neuron, corresponds with table 8.1's type 2, whereas the second corresponds with type 1. We will discuss one model of both types of architectures: the (extremely) fine-grain Connection Machine and coarse-grain Systolic arrays, viz. the Warp computer. A more complete discussion should also include transputers which are very popular nowadays due to their very high performance/price ratio [2–4]. In this case, descriptor 1 of table 8.1 is most applicable.

Besides the granularity, the computers can be categorised by their operation. The most widely used categorisation is by Flynn [5] (see table 9.1). It distinguishes two types of parallel computers: SIMD (Single Instruction, Multiple Data) and MIMD (Multiple Instruction, Multiple Data). The former type consists of a number of processors which execute the same instructions but on different data, whereas the latter has a separate program for each processor. Fine-grain computers are usually SIMD, while coarse grain computers tend to be MIMD (also in correspondence with table 8.1, entries 1 and 2).

		Number of Data Streams	
		single	multiple
Number of Instruction Streams	single	SISD (von Neumann)	SIMD (vector, array)
	multiple	MISD (pipeline?)	MIMD (multiple micros)

Table 9.1: Flynn's classification [5].

Table 9.2 shows a comparison of several types of hardware for neural network simulation. The speed entry, measured in interconnects per second, is an important measure which is often used to compare neural network simulators. It measures the number of multiply-and-add operations that can be performed per second. However, the comparison is not 100% honest: it does not always include the time needed to fetch the data on which the operations are to be performed, and may also ignore other functions required by some algorithms such

as the computation of a sigmoid function. Also, the speed is of course dependent of the algorithm used.

	HARDWARE	WORD LENGTH	STORAGE (K Intcnts)	SPEED (K Int/s)	COST (K$)	SPEED / COST
WORKSTATIONS						
Micro/Mini	PC/AT	16	100	25	5	5.0
Computers	Sun 3	32	250	250	20	12.5
	VAX	32	100	100	300	0.33
	Symbolics	32	32,000	35	100	0.35
Attached	ANZA	8–32	500	45	10	4.5
Processors	$\Delta - 1$	32	1,000	10,000	15	667
	Transputer	16	2,000	3,000	4	750
Bus-oriented	Mark III, IV	16	1,000	500	75	6.7
	MX/1–16	16	50,000	120,000	300	400
MASSIVELY	CM–2 (64K)	32	64,000	13,000	2,000	6.5
PARALLEL	Warp (10)	32	320	17,000	300	56.7
	Warp (20)			32,000		
	Butterfly (64)	32	60,000	8,000	500	16
SUPER-COMPUTERS	Cray XMP	64	2,000	50,000	4,000	12.5

Table 9.2: Hardware machines for neural network simulation.
Note that the reported speeds are measured for small networks.

9.1 The Connection Machine

9.1.1 Architecture

One of the most outstanding fine-grain SIMD parallel machines is Daniel Hillis' Connection Machine [6,7], originally developed at MIT and later built at Thinking Machines Corporation. The original model, the CM–1, consists of 64K (65,536) one-bit processors, divided up into four units of 16K processors each. The units are connected via a cross-bar switch (the *nexus*) to up to four front-end computers (see figure 9.1). The large number of extremely simple processors make the machine a *data parallel* computer, and can be best envisaged as active memory.

Each processor chip contains 16 processors, a control unit, and a router. It is connected to a memory chip which contains 4K bits of memory per processor. Each processor consists of a one-bit ALU with three inputs and two outputs, and a set of registers. The control unit decodes incoming instructions broadcast by the front-end computers (which can be DEX VAXes or Symbolics Lisp machines). At any time, a processor may be either listening to the incoming instruction or not.

The router implements the communication algorithm: each router is connected to its nearest neighbours via a two-dimensional grid (the *NEWS* grid) for fast neighbour communication; also, the chips are connected via a Boolean 12-cube, i.e., chips i and j are connected if and only if $|i - j| = 2^k$ for some integer k. Thus at most 12 hops are needed to deliver a message. So there are 4,096 routers connected by 24,576 bidirectional wires.

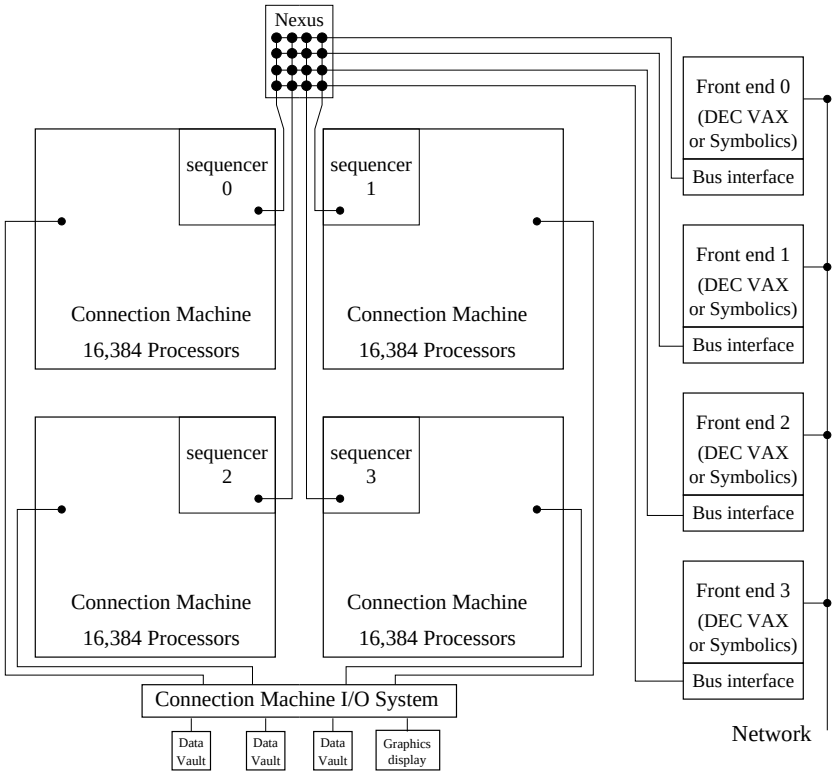

Figure 9.1: The Connection Machine system organisation.

By slicing the memory of a processor, the CM can also implement virtual processors.

The CM–2 differs from the CM–1 in that it has 64K bits instead of 4K bits memory per processor, and an improved I/O system.

9.1.2 Applicability to neural networks

There have been a few researchers trying to implement neural networks on the Connection Machine [8,9]. Even though the Connection Machine has a topology which matches the topology of most artificial neural networks very well, the *relatively* slow message passing system makes the machine not very useful as a general-purpose neural network simulator. It appears that the Connection Machine suffers from a dramatic decrease in throughput due to communication delays [10]. Furthermore, the cost/speed ratio (see table 9.2) is very bad compared to, e.g., a transputer board. As an effect, the Connection Machine is not widely used for neural network simulation.

One possible implementation is given in [8]. Here, a back-propagation network is implemented by allocating one processor per unit and one per outgoing weight and one per incoming weight. The processors are thus arranged that each processor for a unit is immediately followed by the processors for its outgoing weights and preceded by those for its incoming weights. The feed-forward step is performed by first clamping input units and next executing a *copy-scan* operation by moving those activation values to the next k processors (the outgoing weight processors). The weights then multiply themselves with the activation values and perform a *send* operation in which the resulting values are sent to the processors

allocated for incoming weights. A *plus-scan* then sums these values to the next layer of units in the network. The feedback step is executed similarly. Both the feed-forward and feedback steps can be interleaved and pipelined such that no layer is ever idle. For example, for the feed-forward step, a new pattern $\mathbf{x}^p$ is clamped on the input layer while the next layer is computing on $\mathbf{x}^{p-1}$, etc.

To prevent inefficient use of processors, one weight could also be represented by one processor.

9.2 Systolic arrays

Systolic arrays [11] take the advantage of laying out algorithms in two dimensions. The design favours compute-bound as opposed to I/O-bound operations. The name *systolic* is derived from the analogy of pumping blood through a heart and feeding data through a systolic array.

A typical use is depicted in figure 9.2. Here, two band matrices A and B are multiplied and added to C, resulting in an output $C + AB$. Essential in the design is the reuse of data elements, instead of referencing the memory each time the element is needed.

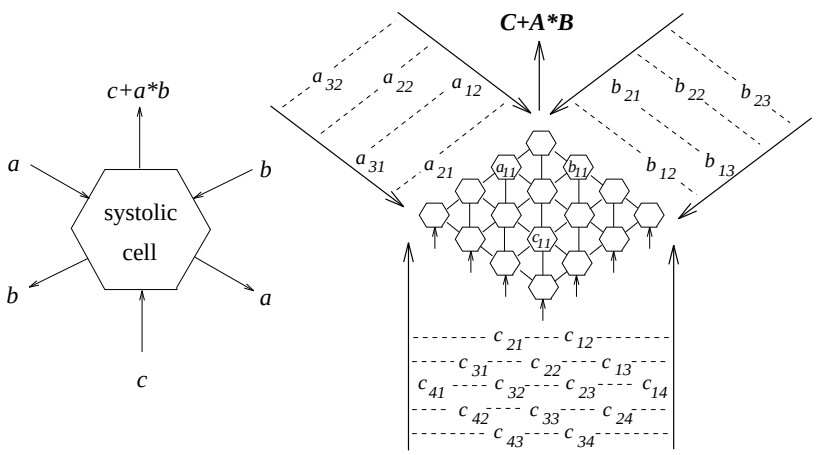

Figure 9.2: Typical use of a systolic array.

The Warp computer, developed at Carnegie Mellon University, has been used for simulating artificial neural networks [12] (see table 9.2). It is a system with ten or more programmable one-dimensional systolic arrays. Two data streams, one of which is bi-directional, flow through the processors (see figure 9.3). To implement a matrix product $W\mathbf{x} + \theta$, the W is not a stream as in figure 9.2 but stored in the memory of the processors.

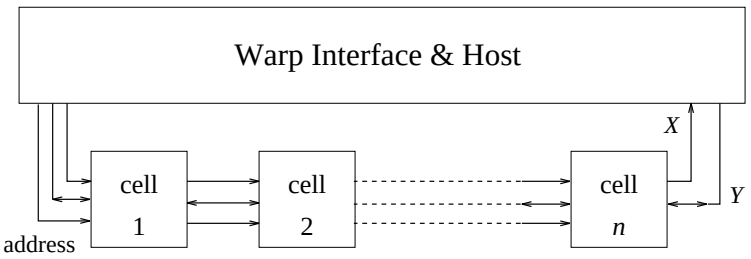

Figure 9.3: The Warp system architecture.

References

[1] G. S. Almasi & A. Gottlieb, *Highly Parallel Computing*, The Benjamin/Cummings Publishing Company Inc., 1989.

[2] John A. Board, Jr., ed., *Transputer Research and Applications 2: Proceedings of the Second Conference on the North American Transputer Users Group*, IOS Press, Durham, NC, 1989.

[3] IOS Press, *Parallel Programming of Transputer Based Machines: Proceedings of the 7th Occam User Group Technical Meeting*, Grenoble, France, 1987.

[4] R. Eckmiller, G. Hartmann & G. Hauske, eds., *Parallel Processing in Neural Systems and Computers*, Elsevier Science Publishers B.V., 1990.

[5] M. J. Flynn, "Some computer organizations and their effectiveness," *IEEE Transactions on Computers* C–21 (1972), 948–960.

[6] W. D. Hillis, *The Connection Machine*, The MIT Press, 1985.

[7] Thinking Machines Corporation, "Connection Machine model CM–2 technical summary," HA87–4, 1987.

[8] G. Blelloch & C. R. Rosenberg, "Network learning on the Connection Machine," in *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, Milan, Italy, August 1987, 323–326.

[9] A. Singer, "Implementations of artificial neural networks on the Connection Machine," Thinking Machines Corporation, RL90–2, Cambridge, MA, January 1990.

[10] R. A. Hummel, 1989, 1990, Personal Communication.

[11] H. T. Kung & C. E. Leierson, "Systolic arrays (for VLSI)," in *Sparse Matrix Proceedings, 1978*, Academic Press, 1979, also in: *Algorithms for VLSI Processor Arrays*, C. Mead & L. Conway, eds., Addison-Wesley, 1980, 271-292.

[12] D. A. Pomerleau, G. L. Gusciora, D. S. Touretzky & H. T. Kung, "Neural network simulation at Warp speed: How we got 17 million connections per second," in *IEEE Second International Conference on Neural Networks* #II, 1988, 143–150.

10 Dedicated Neuro-Hardware

Recently, many neuro-chips have been designed and built. Although many techniques, such as digital and analogue electronics, optical computers, chemical implementation, and bio-chips, are investigated for implementing neuro-computers, only digital and analogue electronics, and in a lesser degree optical implementations, are at present feasible techniques. We will therefore concentrate on such implementations.

10.1 General issues

10.1.1 Connectivity constraints

Connectivity within a chip

A major problem with neuro-chips always is the connectivity. A single integrated circuit is, in current-day technology, planar with limited possibility for cross-over connections. This poses a problem. Whereas connectivity to nearest neighbour can be implemented without problems, connectivity to the second nearest neighbour results in a cross-over of four which is already problematic. On the other hand, full connectivity between a *set* of input and output units can be easily attained when the input and output neurons are situated near two edges of the chip (see figure 10.1). Note that the number of neurons in the chip grows linearly with the size of the chip, whereas in the earlier layout, the dependence is quadratic.

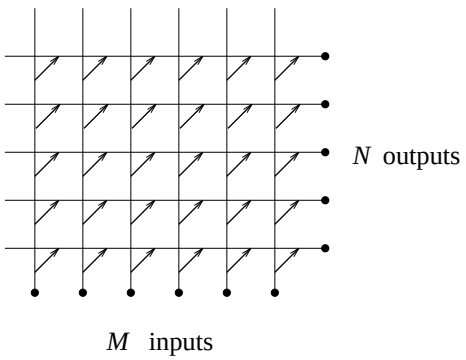

Figure 10.1: Connections between M input and N output neurons.

Connectivity between chips

To build large or layered ANN's, the neuro-chips have to be connected together. When only few neurons have to be connected together, or the chips can be placed in subsequent rows in feed-forward types of networks, this is no problem. But in other cases, when large numbers of neurons in one chip have to be connected to neurons in other chips, there are a number of problems:

- designing chip packages with a very large number of input or output leads;
- fan-out of chips: each chip can ordinarily only send signals two a small number of other chips. Amplifiers are needed, which are costly in power dissipation and chip area;
- wiring.

A possible solution would be using optical interconnections. In this case, an external light source would reflect light on one set of neurons, which would reflect part of this light using deformable mirror spatial light modulator technology on to another set of neurons. Also under development are three-dimensional integrated circuits.

10.1.2 Analogue vs. digital

Due to the similarity between artificial and biological neural networks, analogue hardware seems a good choice for implementing artificial neural networks, resulting in cheaper implementations which operate at higher speed. On the other hand, digital approaches offer far greater flexibility and, not to be neglected, arbitrarily high accuracy. Also, digital chips can be designed without the need of very advanced knowledge of the circuitry using CAD/CAM systems, whereas the design of analogue chips requires good theoretical knowledge of transistor physics as well as experience.

An advantage that analogue implementations have over digital neural networks is that they closely match the physical laws present in neural networks (table 8.1, point 1). First of all, weights in a neural network can be coded by one single analogue element (e.g., a resistor) where several digital elements are needed[1]. Secondly, very simple rules as Kirchoff's laws[2] can be used to carry out the addition of input signals. As another example, Boltzmann machines (section 6.2) can be easily implemented by amplifying the natural noise present in analogue devices.

10.1.3 Optics

As mentioned above, optics could be very well used to interconnect several (layers of) neurons. One can distinguish two approaches. One is to store weights in a planar transmissive or reflective device (e.g., a spatial light modulator) and use lenses and fixed holograms for interconnection. Figure 10.2 shows an implementation of optical matrix multiplication. When N is the linear size of the optical array divided by wavelength of the light used, the array has capacity for N^2 weights, so it can fully connect N neurons with N neurons [2].

[1] On the other hand, the opposite can be found when considering the *size* of the element, especially when high accuracy is needed. However, once artificial neural networks have outgrown rules like back-propagation, high accuracy might not be needed.

[2] The Kirchoff laws state that for two resistors R_1 and R_2 (1) in series, the total resistance can be calculated using $R = R_1 + R_2$, and (2) in parallel, the total resistance can be found using $1/R = 1/R_1 + 1/R_2$ [1].

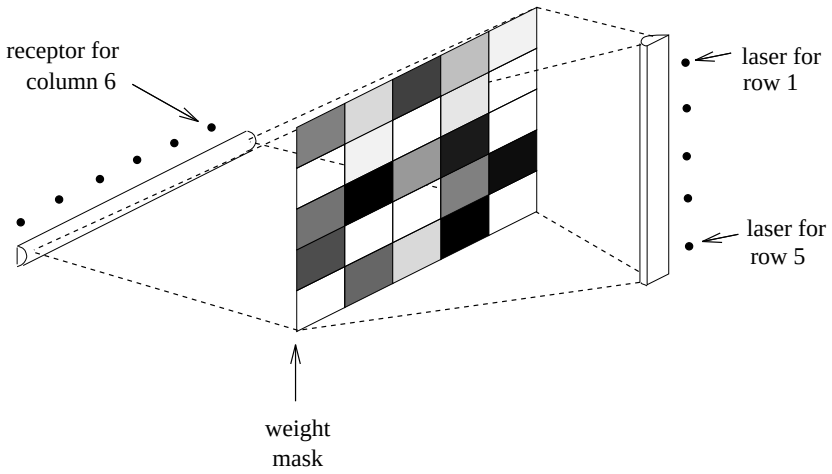

Figure 10.2: Optical implementation of matrix multiplication.

A second approach uses volume holographic correlators, offering connectivity between two areas of N^2 neurons for a total of N^4 connections[3]. A possible use of such volume holograms in an all-optical network would be to use the system for image completion [3]. A number of images could be stored in the hologram. The input pattern is correlated with each of them, resulting in output patterns with a brightness varying with the degree of correlation. The images are fed into a threshold device which will conduct the image with highest brightness better than others. This enhancement can be repeated for several loops.

10.1.4 Learning vs. non-learning

It is generally agreed that the major forte of neural networks is their ability to *learn*. Whereas a network with fixed, pre-computed, weight values could have its merit in industrial applications, on-line adaptivity remains a design goal for most neural systems.

With respect to learning, we can distinguish between the following levels:

1. **fixed weights**: the design of the network determines the weights. Examples are the retina and cochlea chips of Carver Mead's group discussed below (cf. a ROM (Read-Only Memory) in computer design);

2. **pre-programmed weights**: the weights in the network can be set only once, when the chip is installed. Many optical implementations fall in this category (cf. PROM (Programmable ROM));

3. **programmable weights**: the weights can be set more than once by an external device (cf. EPROM (Erasable PROM) or EEPROM (Electrically Erasable PROM));

4. **on-site adapting weights**: the learning mechanism is incorporated in the network (cf. RAM (Random Access Memory)).

[3] Well ... not exactly. Due to diffraction, the total number of independent connections that can be stored in an ideal medium is N^3, i.e., the volume of the hologram divided by the cube of the wavelength. So, in fact $N^{3/2}$ neurons can be connected with $N^{3/2}$ neurons.

10.2 Implementation examples

10.2.1 Carver Mead's silicon retina

The chips devised by Carver Mead's group at Caltech [4] are heavily inspired by biological neural networks. Mead attempts to build analogue neural chips which match biological neurons as closely as possible, including extremely low power consumption, fully analogue hardware, and operation in continuous time (table 8.1, point 3). One example of such a chip is the *Silicon Retina* [5].

Retinal structure

The off-center retinal structure can be described as follows. Light is transduced to electrical signals by photo-receptors which have a primary pathway through the *triad synapses* to the *bipolar cells*. The bipolar cells are connected to the *retinal ganglion* cells which are the output cells of the retina. The *horizontal cells*, which are also connected via the triad synapses to the photo-receptors, are situated directly below the photo-receptors and have synapses connected to the axons leading to the bipolar cells.

The system can be described in terms of the triad synapse's three elements:

1. the photo-receptor outputs the logarithm of the intensity of the light;
2. the horizontal cells form a network which averages the photo-receptor over space and time;
3. the output of the bipolar cell is proportional to the difference between the photo-receptor output and the horizontal cell output.

The photo-receptor

The photo-receptor circuit outputs a voltage which is proportional to the logarithm of the intensity of the incoming light. There are two important consequences:

1. several orders of magnitude of intensity can be handled in a moderate signal level range;
2. the voltage difference between two points is proportional to the contrast ratio of their illuminance.

The photo-receptor can be implemented using a photo-detector, two FET's[4] connected in series[5] and one transistor (see figure 10.3). The lowest photo-current is about $10^{-14}A$ or 10^5 photons per second, corresponding with a moonlit scene.

Horizontal resistive layer

Each photo-receptor is connected to its six neighbours via resistors forming a hexagonal array. The voltage at every node in the network is a spatially weighted average of the photo-receptor inputs, such that farther away inputs have less influence (see figure 10.4(a)).

[4] Field Effect Transistor

[5] A detailed description of the electronics involved is out of place here. However, we will provide figures where useful. See [4] for an in-depth study.

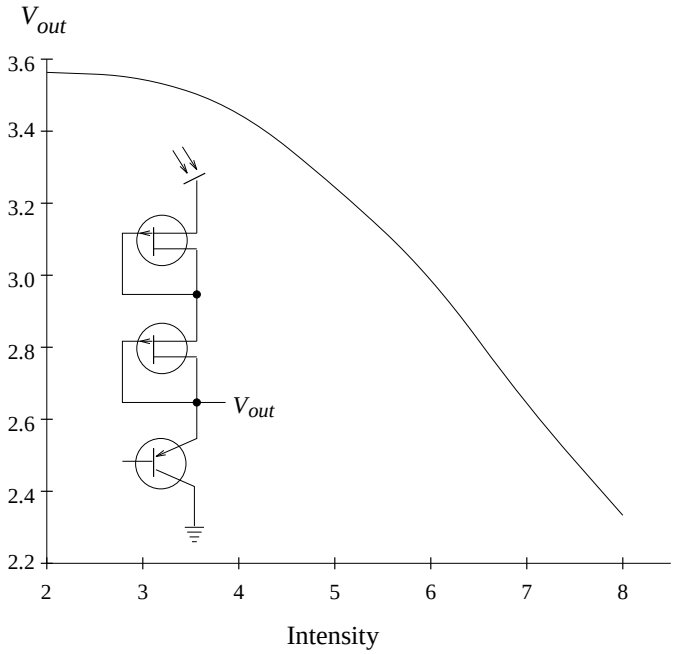

Figure 10.3: The photo-receptor used by Mead. To prevent current being drawn from the photo-receptor, the output is only connected to the gate of the transistor.

Bipolar cell

The output of the bipolar cell is proportional to the difference between the photo-receptor output and the voltage of the horizontal resistive layer. The architecture is shown in figure 10.4(b). It consists of two elements: a wide-range amplifier which drives the resistive network towards the photo-receptor output, and an amplifier sensing the voltage difference between the photo-receptor output and the network potential.

Implementation

A chip was built containing 48×48 pixels. The output of every pixel can be accessed independently by providing the chip with the horizontal and vertical address of the pixel. The selectors can be run in two modes: static probe or serial access. In the first mode, a single row and column are addressed and the output of a single pixel is observed as a function of time. In the second mode, both vertical and horizontal shift registers are clocked to provide a serial scan of the processed image for display on a television display.

Performance

Several experiments show that the silicon retina performs similarly as biological retina [5]. Similarities are shown between sensitivity for intensities; time responses for a single output when flashes of light are input; response to contrast edges.

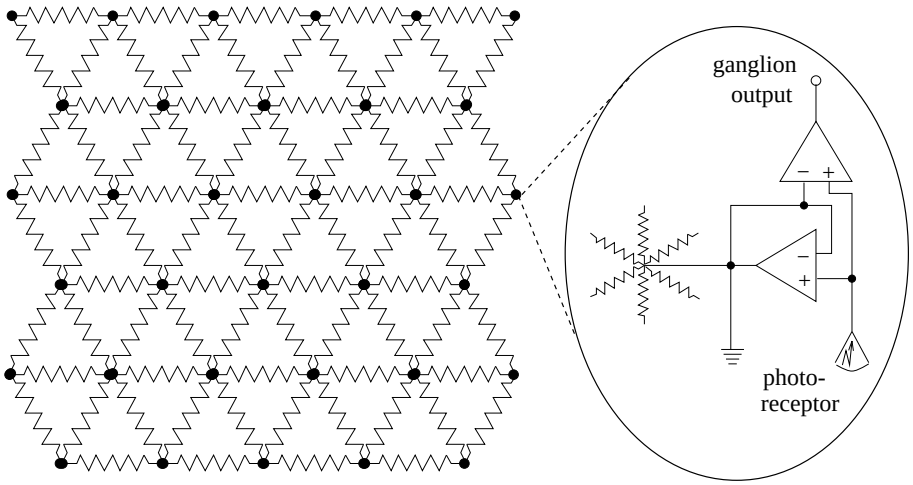

Figure 10.4: The resistive layer (a) and, enlarged, a single node (b).

10.2.2 LEP's LNeuro chip

A radically different approach is the LNeuro chip developed at the Laboratoires d'Electronique Philips (LEP) in France [6,7]. Whereas most neuro-chips implement Hopfield networks (section 6.1) or, in some cases, Kohonen networks (section 5.2) (due to the fact that these networks have *local* learning rules), these digital neuro-chips can be configured to incorporate any learning rule and network topology.

Architecture

The LNeuro chip, depicted in figure 10.5, consists of an multiply-and-add or *relaxation* part, and a learning part. The LNeuro 1.0 has a parallelism of 16. The weights w_{ij} are 8 bits long in the relaxation phase. and 16 bit in the learning phase.

Multiply-and-add

The multiply-and-add in fact performs a matrix multiplication

$$a_i(t+1) = \mathcal{F}\left(\sum_j w_{ij} a_j(t)\right). \tag{10.1}$$

The input activations a_j are kept in the *neural state registers*. For each neural state there are two registers. These can be used to implement synchronous or asynchronous update. In the former mode, the computed state of neurons wait in registers until all states are known; then the whole register is written into the register used for the calculations. In asynchronous mode, however, every new state is directly written into the register used for the next calculation.

The arithmetical logical unit (ALU) has an external input to allow for accumulation of external partial products. This can be used to construct larger, structured, or higher-precision networks.

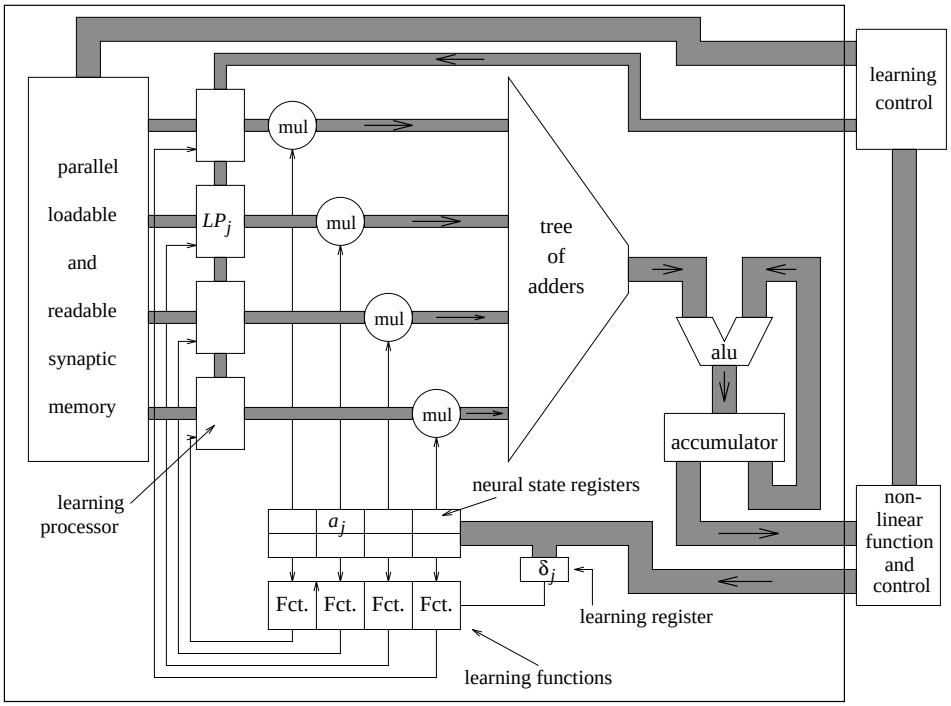

Figure 10.5: The LNeuro chip. For clarity, only four neurons are drawn.

The neural states (a_j) are coded in one to eight bits, whereas either eight or sixteen bits can be used for the weights which are kept in a RAM. In order to save silicon area, the multiplications $w_{ij}a_j$ are serialised over the bits of a_j, replacing N eight by eight bit parallel multipliers by N eight bit AND gates. The partial products are saved and added in the *tree of adders*.

The computation thus increases linearly with the number of neurons (instead of quadratic in simulation on serial machines).

The activation function is, for reasons of flexibility, kept off-chip. The results of the weighted sum calculation go off-chip serially (i.e., bit by bit), and the result must be written back to the neural state registers.

Finally, a column of latches is included to temporarily store memory values, such that during a multiply of the weight with several bits the memory can be freely accessed. These latches in fact take part in the learning mechanism described below.

Learning

The remaining parts in the chip are dedicated to the learning mechanism. The learning mechanism is designed to implement the Hebbian learning rule [8]

$$w_{ij} \leftarrow w_{ij} + \delta_i a_j \tag{10.2}$$

where δ_i is a scalar which only depends on the output neuron i. To simplify the circuitry, eq. (10.2) is simplified to

$$w_{ij} \leftarrow w_{ij} + g(a_i, a_j)\delta_i \tag{10.3}$$

where $g(a_i, a_j)$ can have value -1, 0, or $+1$. In effect, eq. (10.3) either increments or decrements the w_{ij} with δ_i, or keeps w_{ij} unchanged. Thus eq. (10.2) can be simulated by executing eq. (10.3) several times over the same set of weights.

The weights $\mathbf{w}_i$ related to the output neuron i are all modified in parallel. A learning step proceeds as follows. Every *learning processor* (see figure 10.5) LP_j loads the weight w_{ij} from the synaptic memory, the δ_i from the *learning register*, and the neural state a_j. Next, they all modify their weights in parallel using eq. (10.3) and write the adapted weights back to the synaptic memory, also in parallel.

References

[1] R. P. Feynman, R. B. Leighton & M. Sands, *The Feynman Lectures on Physics*, Addison-Wesley Publishing Company, Reading (MA), Menlo Park (CA), London, Sidney, Manila, 1983.

[2] N. H. Fahrat, D. Psaltis, A. Prata & E. Paek, "Optical implementation of the Hopfield model," *Applied Optics* 24 (1985), 1469–1475.

[3] Y. A. Abu-Mostafa & D. Psaltis, "Optical neural computers," *Scientific American* (March 1987).

[4] C. Mead, *Analog VLSI and Neural Systems*, Addison-Wesley, Reading, MA, 1989.

[5] C. A. Mead & M. A. Mahowald, "A silicon model of early visual processing," *Neural Networks* 1 (1988), 91–97.

[6] J. B. Theeten, M. Duranton, N. Mauduit & J. A. Sirat, "The LNeuro chip: A digital VLSI with on-chip learning mechanism," in *Proceedings of the International Conference on Neural Networks* #I, Paris, July 1990, 593–596.

[7] M. Duranton & J. A. Sirat, "Learning on VLSI: A general purpose digital neurochip," in *Proceedings of the Third International Joint Conference on Neural Networks*, Washington, June 1989.

[8] D. O. Hebb, *The Organization of Behaviour*, Wiley, New York, 1949.

Author Index

A

Abu-Mostafa, Y. A. 117
Ackley, D. H. 73, 74
Ahalt, S. C. 48
Almasi, G. S. 109
Almeida, L. B. 41
Amit, D. J. 69
Anandan, P. 64
Anderson, C. W. 63
Anderson, J. A. 9, 13, 18, 69

B

Ballard, D. H. 16
Barto, A. G. 63, 64
Blelloch, G. 111
Board, J. A. 109
Boomgaard, R. van den 41
Braun, H. 101
Breiman, L 51
Bruce, A. D. 71
Bulirsch, R. 40

C

Canning, A. 71
Carpenter, G. A. 45, 61, 62
Craig, J. J. 79
Cybenko, G. 33

D

Dam, J. W. M. 64
DARPA 107
Dastani, M. M. 38, 39
Duranton, M. 120

E

Eckmiller, R. 109

F

Fahrat, N. H. 116
Feinstein, D. I. 71
Feldman, J. A. 16
Feynman, R. P. 116
Flannery, B. P. 39
Flynn, M. J. 109
Forrest, B. 71
Friedman, J. H. 51
Fritzke, B. 52
Fukushima, K. 13, 45, 94, 97
Funahashi, K.-I. 33
Furukawa, K. 86

G

Gardner, E. 71
Garey, M. R. 72
Geman, D. 101
Geman, S. 101
Gielen, C. 54
Gisbergen, J. van 54
Gorman, R. P. 42
Gottlieb, A. 109
Groen, F. C. A. 51, 83, 84

Grossberg, S. 13, 45, 58, 61, 62, 107
Gusciora, G. L. 112
Gutfreund, H. 69

H

Hartigan, J. A. 61
Hartman, E. J. 33
Hartmann, G. 109
Hauske, G. 109
Hebb, D. O. 18, 121
Hecht-Nielsen, R. 50
Hertz, J. 9
Hesselroth, T. 87
Hestenes, M. R. 40
Hillis, W. D. 110
Hinton, G. E. 33, 73, 74
Hoff, M. E. 25, 31
Hopfield, J. J. 18, 69, 71, 72
Hornik, K. 33
Hummel, R. A. 111

J

Jansen, A. 51, 84
Johnson, D. S. 72
Jorgensen, C. C. 88
Josin, G. 42

K

Kak, A. C. 101
Katayama, M. 87
Kawato, M. 86, 87
Keeler, J. D. 33
Kohonen, T. 13, 18, 45, 52, 54, 69, 107
Korst, M. van der 83
Kowalski, J. M. 33
Kröse, B. J. A. 64, 83, 84
Krogh, A. 9
Krommenhoek, K. 54
Kung, H. T. 112

L

Le Cun, Y. 33
Leierson, C. E. 112
Leighton, R. B. 116
Lippmann, R. P. 46, 60, 62

M

Müller, T. 101
Mahowald, M. A. 118, 119
Makisara, M. 54
Mallach, E. G. 107
Marr, D. 99, 100
Martinetz, T. M. 9, 52, 84
Mauduit, N. 120
McClelland, J. L. 9, 15
McCulloch, W. S. 13
Mead, C. A. 101, 118, 119
Mel, B. W. 16
Minsky, M. 13, 26–29, 33, 94
Murray, A. F. 107

O

Oja, E. 56

P

Paek, E. 116
Palmer, R. G. 9, 71
Papert, S. 13, 26–29, 33
Parker, D. B. 33
Pawley, G. S. 73
Pearlmutter, B. A. 18
Pitts, W. 13
Polak, E. 40
Pomerleau, D. A. 112
Powell, M. J. D. 40, 41
Prata, A. 116
Press, W. H. 39
Psaltis, D. 81, 116, 117

R

Ritter, H. J. 9, 84
Rosen, B.E. 51
Rosenberg, C. R. 42, 111
Rosenblatt, F. 26, 27
Rosenfeld, A. 101
Rosenfeld, E. 9
Rumelhart, D. E. 9, 15, 33, 45

S

Sands, M. 116
Sanger, T. D. 95
Saramaki, T. 54
Sarkar, K. 87
Schulten, K. J. 9, 52, 84, 87
Sejnowski, T. J. 42, 73, 74
Sideris, A. 81
Silva, F. M. 41
Singer, A. 111
Sirat, J. A. 120
Smagt, P. P. van der 51, 83, 84, 87
Smeulders, A. 41
Sompolinsky, H. 69
Stiefel, E. 40
Stinchcombe, M. 33
Stoer, J. 40
Summeren, P. J. H. A. M. van 58
Sutton, R. S. 63, 64
Suzuki, R. 86

T

Tank, D. W. 72
Teukolsky, S. A. 39
Theeten, J. B. 120
Tomlinson, M. S. 107
Touretzky, D. S. 112

V

Vetterling, W. T. 39

W

Walker, D. J. 107
Wallace, D. J. 71
Wasserman, P. D. 9
Werbos, P. 33
Werbos, P. W. 63
White, H. 33
Widrow, B. 25, 26, 31, 94
Williams, R. J. 33
Wilson, G. V. 73
Winter, R. 94

Y

Yamamura, A. A. 81

Z

Zipser, D. 45

Subject Index

A

ACE 63
activation function 17, 20
  hard limiting, 17
  Heaviside, 27
  linear, 17
  nonlinear, 34
  semi-linear, 17
  sgn, 27
  sigmoid, 17, 36, 39
    derivative of, 36
  threshold, 70
adaline 18, 25, 26
  vision, 94
adaptive critic element 63
adaptive critic element, *see* ACE
adaptive resonance theory, *see* ART
analogue implementation 115, 116, 118
annealing 73
ART 45, 58, 107
ASE 63
  activation function, 63
  bias, 63
associative learning 18
associative memory 71
  instability of stored patterns, 71
  spurious stable states, 71
associative search element 63
associative search element, *see* ASE
asymmetric divergence 74
asynchronous update 16, 70
auto-associator 69
  vision, 95

B

back-propagation 33, 38, 42, 107, 116
  advanced training algorithms, 39
  conjugate gradient, 39
  derivation of, 35
  discovery of, 13
  gradient descent, 35, 37, 39
  implementation on Connection Machine, 111
  learning by pattern, 37
  learning rate, 37
  local minima, 39
  momentum, 37
  network paralysis, 38
  oscillation in, 37
  vision, 95
bias 20
bio-chips 115
bipolar cells
  retina, 118
  silicon retina, 119
Boltzmann distribution 73
Boltzmann machine 73, 74, 116

C

Carnegie Mellon University 112
CART 51
cartpole 65
chemical implementation 115
CLUSTER 49
clustering 45, 46
coarse-grain parallelism 109
coding
  lossless, 94
  lossy, 94

cognitron 45, 97
competitive learning 45, 46
    error-function, 48
    frequency sensitive, 48
conjugate directions 40
conjugate gradient 39
Connection Machine 107, 109–111
    architecture, 110
    communication, 110
    NEWS grid, 110
    nexus, 110
connectionist models 13
connectivity
    constraints, 115
    optical, 116
convergence
    steepest descent, 41
cooperative algorithm 100
correlation matrix 56
counterpropagation 49
counterpropagation network 50

D

DARPA neural network study 9
decoder 64
deflation 57
delta rule 19, 26, 31, 32
    generalised, 33–36
digital implementation 115, 116
dimensionality reduction 45
discovery vs. creation 13
discriminant analysis 52
distance measure 47
dynamic programming 63
dynamics
    in neural networks, 18
    robotics, 80, 86

E

EEPROM 117
eigenvector transformation 55
eligibility 63
emulation 107, 109
energy 20
    Hopfield network, 70, 71
    travelling salesman problem, 72
EPROM 117
error 20
    back-propagation, 34
    competitive learning, 48
    perceptron, 31
    quadratic, 34
excitation 16
exclusive-or, *see* XOR
external input 20
eye 58

F

face recognition 95
feature extraction 45, 96
feed-forward network 18, 20, 33, 35, 41, 42
FET 118
fine-grain parallelism 109
five, *see* three
FORGY 49
forward kinematics 79

G

Gaussian 85
generalised delta rule 33–36
gradient descent 31, 35, 37, 39
granularity of parallelism 109

H

hard limiting activation function 17
Heaviside 27
Hebb rule 18, 30, 55, 71, 121
    normalised, 56
Hessian 40
high level vision 93
holographic correlators 117
Hopfield network 69, 70, 89, 120

as associative memory, 71
instability of stored patterns, 71
spurious stable states, 71
energy, 70, 71
graded response neurons, 72
optimisation, 72
stable limit points, 70
stable neuron in, 70
stable pattern in, 70
stable state in, 70
stable storage algorithm, 71
stochastic update, 73
symmetry, 71
un-learning, 71
horizontal cells
retina, 118
silicon retina, 118

I

image compression 94
back-propagation, 95
PCA, 95
self-organising networks, 94
implementation 107
analogue, 115, 116, 118
chemical, 115
connectivity constraints, 115
digital, 115, 116
on Connection Machine, 111
optical, 115, 116
silicon retina, 119
information gathering 15
inhibition 16
instability of stored patterns 71
intermediate level vision 93
inverse kinematics 79
Ising spin model 69
ISODATA 49

J

Jacobian matrix 82, 84, 85

K

k-means clustering 49
Kirchoff laws 116
KISS 84
Kohonen network 52, 120
3-dimensional, 84
for robot control, 84
Kullback information 74

L

leaky learning 48
learn vs. store 20
learning 18, 20, 117
associative, 18
LNeuro, 121
self-supervised, 18, 81
supervised, 18
unsupervised, 18, 81
learning rate 19, 20
back-propagation, 37
learning vector quantisation 52
LEP 120
linear activation function 17
linear convergence 41
linear networks
vision, 95
linear threshold element 26
LMS, *see* delta rule
LNeuro 120, 121
activation function, 121
ALU, 120
learning, 121
RAM, 121
local minima
back-propagation, 39
look-up table 16, 50
lossless coding 94
lossy coding 94
low level vision 93
LVQ2 52

M

Markov random field 101
MARS 51
mean vector 56
MIMD 109
MIT 110
mobile robots 88
momentum 37
multi-layer perceptron 30, 74

N

neocognitron 94, 97
Nestor 107
NETtalk 42
network paralysis
    back-propagation, 38
NeuralWare 107
neuro-computers 115
nexus 110
non-cooperative algorithm 100
non-recurrent 69
none 45
normalisation 56
notation 19

O

octree methods 51
offset 20
Oja learning rule 56, 57
optical implementation 115, 116
optimisation 72
oscillation in back-propagation 37
output vs. activation of a unit 20

P

panther
    hiding, 58
    resting, 58
parallel distributed processing 13, 15
parallelism
    coarse-grain, 109
    fine-grain, 109
PCA 55
    image compression, 95
PDP 13, 15
perceptron 13, 18, 25–31
    convergence theorem, 30
    error, 31
    learning rule, 30
    threshold, 30
    vision, 94
photo-receptor
    retina, 118
    silicon retina, 118, 119
positive definite 40
Principal components 55
PROM 117
prototype vectors 55
PYGMALION 107

R

RAM 107, 117
recurrent networks 18, 20, 69
reinforcement learning 45, 63
relaxation 18
resistor 116
retina 94
    bipolar cells, 118
    horizontal cells, 118
    photo-receptor, 118
    retinal ganglion, 118
    structure, 118
    triad synapses, 118
retinal ganglion 118
robotics 79
    dynamics, 80, 86
    forward kinematics, 79
    inverse kinematics, 79
    trajectory generation, 80
Rochester Connectionist Simulator 107
ROM 117

S

self-organisation 18, 45

self-organising networks 45
    image compression, 94
    vision, 94
self-supervised learning 18, 81
semi-linear activation function 17
sgn function 27
sigma unit 16
sigma-pi unit 16
sigmoid activation function 17, 36, 39
    derivative of, 36
silicon retina 101, 118
    bipolar cells, 119
    horizontal cells, 118
    implementation, 119
    photo-receptor, 118, 119
SIMD 109, 110
simulated annealing 73
simulation 107
    taxonomy, 107
spurious stable states 71
stable limit points 70
stable neuron 70
stable pattern 70
stable state 70
stable storage algorithm 71
steepest descent 85
    convergence, 41
stochastic update 17, 73
storage 20
supervised learning 18
synchronous update 16
systolic 112
systolic arrays 109, 112

T

target 31
temperature 73
temporal difference learning 64
terminology 20
Thinking Machines Corporation 110
threshold 20
topologies 17
topology-conserving map 45, 53, 107
training 18
trajectory generation 80
transistor 118, 119
transputer 107, 109
travelling salesman problem 72
    energy, 72
triad synapses 118
TSP, *see* travelling salesman problem

U

unsupervised learning 18, 45, 46, 81
update of a unit 15
    asynchronous, 16, 70
    stochastic, 73
    synchronous, 16

V

vector quantisation 45, 46, 49
vision 93
    high level, 93
    intermediate level, 93
    low level, 93

W

Warp 107, 109, 112
Widrow-Hoff rule 19
winner-take-all 46

X

XOR problem 27–29