




Robot Brains

Circuits and Systems for Conscious Machines

Pentti O. Haikonen

 WILEY



TP242.6
H149

Robot Brains

Circuits and Systems for Conscious Machines

Pentti O. Haikonen

Nokia Research Center, Finland



John Wiley & Sons, Ltd

Copyright © 2007 John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England
Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk
Visit our Home Page on www.wiley.com

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770620.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 42 McDougall Street, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 6045 Freemont Blvd, Mississauga, ONT, L5R 4J3

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Anniversary Logo Design: Richard J. Pacifico

Library of Congress Cataloging in Publication Data

Haikonen, Pentti O.

Robot brains : circuits and systems for conscious machines / Pentti O. Haikonen.
p. cm.

Includes bibliographical references and index.

ISBN 978-0-470-06204-3 (cloth)

1. Robotics. 2. Conscious automata. I. Title.

TJ211.15.H33 2007

629.8'9263—dc22

2007029110

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN-13 978-0470-06204-3

Typeset in 11/14pt Times by Integra Software Services Pvt. Ltd, Pondicherry, India

Printed and bound in Great Britain by Antony Rowe Ltd, Chippenham, Wiltshire

This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

Robot Brains

Preface

Should robots understand what they are doing? What is understanding anyhow and could it be realized artificially? These questions relate to the emerging field of machine cognition and consciousness. This is a rather new discipline that is transforming philosophical musings into engineering solutions.

What is machine consciousness? Chrisley, Clowes and Torrance (2005) identify the goals of machine consciousness research as the creation of artefacts that have characteristics typically associated with awareness, self-awareness, emotion, experience, imagination, etc., and the modelling of these in embodied systems such as robots.

The first serious papers about machine consciousness began to appear in the beginning of the 1990s. Taylor (1992, 1997, 1999), Aleksander and Morton (1993) and Aleksander (1996, 2005) began to propose neural mechanisms for consciousness and mind. Duch (1994, 2005) studied a combined neural/symbolic approach to artificial minds. Trehub (1991) and Valiant (1994) presented some algorithms for cognitive processing, but did not address directly the question of machine consciousness. Sloman (2000) outlined the requirements for minds on a more general level. Philosophical approaches were presented, for example, by Dennett (1991) and Baars (1997). Also the author began his work on machine cognition in the same decade and has summarized the foundations of his approach in the book *The Cognitive Approach to Conscious Machines* (Imprint Academic, UK, 2003a).

The common approach to machine cognition has been the traditional program-based way. However, it is known that the ultimate model for cognition, the human brain, is not program based and is definitely not a digital computer. Therefore, there must be another way to realize machine cognition, one that does not depend on programmed computers.

The author visions autonomous robots that perceive and understand the world and act in it in a natural way, without programs and numerical representation of information. This approach considers the cognitive machine as a system that is seamlessly interactive, both internally and externally, in respect to its environment and experience. This approach should result in robots that know and understand what they are doing and why, robots that can plan and imagine their actions and the possible outcome of these. Robots that exhibit properties like these are said to possess machine consciousness, which may or may not have common deeper properties with animal and human consciousness. The creation of these kinds of machines is also expected to illuminate the nature of human consciousness.

This book considers the engineering aspects of nonnumeric cognitive and potentially conscious machines and should be of interest to anyone who wishes to consider the actual realization of these machines and robots. The principles are presented

with the aim of producing dedicated circuits and systems, but obviously everything can be simulated; those who prefer programmed simulations should find it rather easy to implement these principles in the way of computer programs.

I thank the head of the Nokia Research Center, Dr Bob Iannucci, and my superiors, Dr Juha Lehtikoinen and Dr Kari Laurila, for the possibility of working in this interesting field. I want to express my special gratitude to my older brother Dr Terho (Max) Haikonen for reading the manuscript and giving valuable advice. I wish to thank Dr Markku Åberg and Mr Arto Rantala of the Technical Research Centre of Finland, Centre for Microelectronics, who have also perused the manuscript and have given their expert advice. I also wish to thank Mr Tuomo Alamäki for his help and support.

Extra thanks go to my son Pete for the science fiction underground-style comic strip in the last chapter. Finally, my very special thanks go to my ever-young wife Sinikka for her support and patience.

Pentti O. Haikonen

Contents

Preface	ix
1 Introduction	1
1.1 General intelligence and conscious machines	1
1.2 How to model cognition?	3
1.3 The approach of this book	6
2 Information, meaning and representation	9
2.1 Meaning and the nonnumeric brain	9
2.2 Representation of information by signal vectors	11
2.2.1 Single signal and distributed signal representations	11
2.2.2 Representation of graded values	14
2.2.3 Representation of significance	14
2.2.4 Continuous versus pulse train signals	15
3 Associative neural networks	17
3.1 Basic circuits	17
3.1.1 The associative function	17
3.1.2 Basic neuron models	17
3.1.3 The Haikonen associative neuron	19
3.1.4 Threshold functions	20
3.1.5 The linear associator	22
3.2 Nonlinear associators	24
3.2.1 The nonlinear associative neuron group	24
3.2.2 Simple binary associator	26
3.2.3 Associator with continuous weight values	28
3.2.4 Bipolar binary associator	29
3.2.5 Hamming distance binary associator	30
3.2.6 Enhanced Hamming distance binary associator	30
3.2.7 Enhanced simple binary associator	31
3.3 Interference in the association of signals and vectors	32
3.4 Recognition and classification by the associative neuron group	35
3.5 Learning	38
3.5.1 Instant Hebbian learning	38
3.5.2 Correlative Hebbian learning	39
3.6 Match, mismatch and novelty	41
3.7 The associative neuron group and noncomputable functions	42

4	Circuit assemblies	45
4.1	The associative neuron group	45
4.2	The inhibit neuron group	46
4.3	Voltage-to-single signal (V/SS) conversion	46
4.4	Single signal-to-voltage (SS/V) conversion	48
4.5	The 'Winner-Takes-All' (WTA) circuit	49
4.6	The 'Accept-and-Hold' (AH) circuit	51
4.7	Synaptic partitioning	52
4.8	Serial-to-parallel transformation	54
4.9	Parallel-to-serial transformation	56
4.10	Associative Predictors and Sequencers	57
4.11	Timing circuits	61
4.12	Timed sequence circuits	63
4.13	Change direction detection	66
5	Machine perception	69
5.1	General principles	69
5.2	Perception and recognition	70
5.3	Sensors and preprocesses	71
5.4	Perception circuits; the perception/response feedback loop	72
5.4.1	The perception of a single feature	72
5.4.2	The dynamic behaviour of the perception/response feedback loop	74
5.4.3	Selection of signals	76
5.4.4	Perception/response feedback loops for vectors	77
5.4.5	The perception/response feedback loop as predictor	79
5.5	Kinesthetic perception	81
5.6	Haptic perception	82
5.7	Visual perception	84
5.7.1	Seeing the world out there	84
5.7.2	Visual preprocessing	85
5.7.3	Visual attention and gaze direction	86
5.7.4	Gaze direction and visual memory	87
5.7.5	Object recognition	89
5.7.6	Object size estimation	91
5.7.7	Object distance estimation	92
5.7.8	Visual change detection	94
5.7.9	Motion detection	95
5.8	Auditory perception	98
5.8.1	Perceiving auditory scenes	98
5.8.2	The perception of separate sounds	99
5.8.3	Temporal sound pattern recognition	101
5.8.4	Speech recognition	102
5.8.5	Sound direction perception	103
5.8.6	Sound direction detectors	105
5.8.7	Auditory motion detection	111
5.9	Direction sensing	111
5.10	Creation of mental scenes and maps	113

6	Motor actions for robots	117
6.1	Sensorimotor coordination	117
6.2	Basic motor control	117
6.3	Hierarchical associative control	120
6.4	Gaze direction control	122
6.5	Tracking gaze with a robotic arm	126
6.6	Learning motor action sequences	128
6.7	Delayed learning	129
6.8	Moving towards the gaze direction	129
6.9	Task execution	131
6.10	The quest for cognitive robots	134
7	Machine cognition	137
7.1	Perception, cognition, understanding and models	137
7.2	Attention	139
7.3	Making memories	140
7.3.1	Types of memories	140
7.3.2	Short-term memories	140
7.3.3	Long-term memories	142
7.4	The perception of time	143
7.5	Imagination and planning	145
7.6	Deduction and reasoning	146
8	Machine emotions	149
8.1	Introduction	149
8.2	Emotional significance	150
8.3	Pain and pleasure as system reactions	150
8.4	Operation of the emotional soundtrack	152
8.5	Emotional decision making	153
8.6	The system reactions theory of emotions	154
8.6.1	Representational and nonrepresentational modes of operation	154
8.6.2	Emotions as combinations of system reactions	155
8.6.3	The external expressions of emotions	156
8.7	Machine Motivation and willed actions	156
9	Natural language in robot brains	159
9.1	Machine understanding of language	159
9.2	The representation of words	161
9.3	Speech acquisition	161
9.4	The multimodal model of language	163
9.4.1	Overview	163
9.4.2	Vertical grounding of word meaning	165
9.4.3	Horizontal grounding; syntactic sentence comprehension	168
9.4.4	Combined horizontal and vertical grounding	172
9.4.5	Situation models	173
9.4.6	Pronouns in situation models	175
9.5	Inner speech	176

10 A cognitive architecture for robot brains	179
10.1 The requirements for cognitive architectures	179
10.2 The Haikonen architecture for robot brains	180
10.3 On hardware requirements	183
11 Machine consciousness	185
11.1 Consciousness in the machine	185
11.1.1 The immateriality of mind	185
11.1.2 The reportability aspect of consciousness	186
11.1.3 Consciousness as internal interaction	188
11.2 Machine perception and qualia	190
11.3 Machine self-consciousness	191
11.3.1 The self as the body	191
11.3.2 The experiencing self	192
11.3.3 Inner speech and consciousness	193
11.3.4 The continuum of the existence of the self	194
11.4 Conscious machines and free will	194
11.5 The ultimate test for machine consciousness	195
11.6 Legal and moral questions	198
Epilogue	201
The dawn of real machine cognition	201
References	203
Index	209

1

Introduction

1.1 GENERAL INTELLIGENCE AND CONSCIOUS MACHINES

Suppose that you were going to see a conscious machine, perhaps a robot. What would you expect to see? Recent advances in robotics have produced animal and humanoid robots that are able to move in very naturalistic ways. Would you find these robots conscious? I do not think so. As impressive as the antics of these artefacts are, their shortcoming is easy to see; the lights may be on, but there is ‘nobody’ at home. The program-controlled microprocessors of these robots do not have the faintest trace of consciousness and the robots themselves do not know what they are doing. These robots are no more aware of their own existence than a cuckoo clock on a good day.

Artificial intelligence (AI) has brought chess-playing computer programs that can beat grand masters and other ‘intelligent’ programs that can execute given specific tasks. However, the intelligence of these programs is not that of the machine; instead it is the intelligence of the programmer who has laid down the rules for the execution of the task. At best and with some goodwill these cases of artificial intelligence might be called ‘specific intelligence’ as they work only for their specific and limited application. In contrast, ‘general intelligence’ would be flexible and applicable over a large number of different problems. Unfortunately, artificial general intelligence has been elusive. Machines do not really understand anything, as they do not utilize meanings.

Robots do not fare well in everyday tasks that humans find easy. Artificial intelligence has not been able to create general intelligence and common sense. Present-day robots are definitely not sentient entities. Many researchers have recognized this and see this as a shortcoming that must be remedied if robots are ever to be as versatile as humans. Robots should be conscious.

Exactly which features and abilities would distinguish a conscious robot from its nonconscious counterpart? It may well be that no matter how ‘conscious’ a robot might be, some philosophers would still pretend to find one or another successful argument against its consciousness.

2 INTRODUCTION

An engineering approach may bypass the philosophical mess by defining what the concept of 'machine consciousness' would involve. However, care should be taken here to consider first what natural consciousness is like so that the cautionary example of artificial intelligence is not repeated; AI is definitely artificial but has somehow managed to exclude intelligence.

Folk psychology describes human consciousness as 'the immaterial feeling of being here'. This is accompanied by the awareness of self, surroundings, personal past, present and expected future, awareness of pain and pleasure, awareness of one's thoughts and mental content. Consciousness is also linked to thinking and imagination, which themselves are often equated to the flow of inner speech and inner imagery. Consciousness is related to self, mind and free will. Consciousness is also seen to allow one to act and execute motions fluently, without any apparent calculations. A seen object can be readily grasped and manipulated. The environment is seen as possibilities for action.

Folk psychology is not science. Thus it is not able to determine whether the above phenomena were caused by consciousness or whether consciousness is the collection of these phenomena or whether these phenomena were even real or having anything to do with consciousness at all. Unfortunately philosophy, while having done much more, has not done much better.

A robot that could understand the world and perceive it as possibilities for action would be a great improvement over existing robots. Likewise, a robot that would communicate with natural language and would even have the flow of natural language inner speech could easily cooperate with humans; it would be like one of us. It can be seen that the machine implementation of the folk psychology hallmarks of consciousness would be beneficial regardless of their actual relationship to 'true' consciousness, whatever that might eventually turn out to be.

This can be the starting point for the engineering approach. The folk psychology hallmarks of consciousness should be investigated in terms of cognitive sciences and defined in engineering terms for machine implementation. This would lead to an implementable specification for a 'conscious machine'.

This 'conscious machine' would be equipped with sensors and perception processes as well as some means for a physical response. It would perceive the world in a direct way, as objects and properties out there, just as humans do, and perceive itself as being in the centre of the rest of the world. It would be able to process information, to think, in ways seemingly similar to human thinking, and it would have the flow of inner speech and imagery. It would also be able to observe its own thoughts and introspect its mental content. It would perceive this mental content as immaterial. It would motivate itself and have a will of its own. It would judge the world and its own actions by emotional good-bad criteria. It would be aware of its own existence. It would be able to move and execute actions as freely and readily as humans do. It would emulate the processes of the human brain and cognition. It would appear to be conscious.

This leads to two questions: first, how do the human brain and mind actually work and, second, how could the workings of the brain be best emulated in an artificial

system? Unfortunately the exact answer is not yet known to the first question and it may well be that a definite answer can be found only after we have learned to create successful artificial minds. Thus, models of machine cognition may also help to model human cognition.

1.2 HOW TO MODEL COGNITION?

Presently there are five main approaches to the modelling of cognition that could be used for the development of cognitive machines: the computational approach (artificial intelligence, AI), the artificial neural networks approach, the dynamical systems approach, the quantum approach and the cognitive approach. Neurobiological approaches exist, but these may be better suited for the eventual explanation of the workings of the biological brain.

The computational approach (also known as artificial intelligence, AI) towards thinking machines was initially worded by Turing (1950). A machine would be thinking if the results of the computation were indistinguishable from the results of human thinking. Later on Newell and Simon (1976) presented their Physical Symbol System Hypothesis, which maintained that general intelligent action can be achieved by a physical symbol system and that this system has all the necessary and sufficient means for this purpose. A physical symbol system was here the computer that operates with symbols (binary words) and attached rules that stipulate which symbols are to follow others. Newell and Simon believed that the computer would be able to reproduce human-like general intelligence, a feat that still remains to be seen. However, they realized that this hypothesis was only an empirical generalization and not a theorem that could be formally proven. Very little in the way of empirical proof for this hypothesis exists even today and in the 1970s the situation was not better. Therefore Newell and Simon pretended to see other kinds of proof that were in those days readily available. They proposed that the principal body of evidence for the symbol system hypothesis was negative evidence, namely the absence of specific competing hypotheses; how else could intelligent activity be accomplished by man or machine? However, the absence of evidence is by no means any evidence of absence. This kind of 'proof by ignorance' is too often available in large quantities, yet it is not a logically valid argument. Nevertheless, this issue has not yet been formally settled in one way or another. Today's positive evidence is that it is possible to create world-class chess-playing programs and these can be called 'artificial intelligence'. The negative evidence is that it appears to be next to impossible to create real general intelligence via preprogrammed commands and computations.

The original computational approach can be criticized for the lack of a cognitive foundation. Some recent approaches have tried to remedy this and consider systems that integrate the processes of perception, reaction, deliberation and reasoning (Franklin, 1995, 2003; Sloman, 2000).

There is another argument against the computational view of the brain. It is known that the human brain is slow, yet it is possible to learn to play tennis and other

4 INTRODUCTION

activities that require instant responses. Computations take time. Tennis playing and the like would call for the fastest computers in existence. How could the slow brain manage this if it were to execute computations?

The artificial neural networks approach, also known as connectionism, had its beginnings in the early 1940s when McCulloch and Pitts (1943) proposed that the brain cells, neurons, could be modelled by a simple electronic circuit. This circuit would receive a number of signals, multiply their intensities by the so-called synaptic weight values and sum these modified values together. The circuit would give an output signal if the sum value exceeded a given threshold. It was realized that these artificial neurons could learn and execute basic logic operations if their synaptic weight values were adjusted properly. If these artificial neurons were realized as hardware circuits then no programs would be necessary and biologically plausible artificial replicas of the brain might be possible. Also, neural networks operate in parallel, doing many things simultaneously. Thus the overall operational speed could be fast even if the individual neurons were slow.

However, problems with artificial neural learning led to complicated statistical learning algorithms, ones that could best be implemented as computer programs. Many of today's artificial neural networks are statistical pattern recognition and classification circuits. Therefore they are rather removed from their original biologically inspired idea. Cognition is not mere classification and the human brain is hardly a computer that executes complicated synaptic weight-adjusting algorithms.

The human brain has some 10^{11} neurons and each neuron may have tens of thousands of synaptic inputs and input weights. Many artificial neural networks learn by tweaking the synaptic weight values against each other when thousands of training examples are presented. Where in the brain would reside the computing process that would execute synaptic weight adjusting algorithms? Where would these algorithms have come from? The evolutionary feasibility of these kinds of algorithms can be seriously doubted. Complicated algorithms do not evolve via trial and error either. Moreover, humans are able to learn with a few examples only, instead of having training sessions with thousands or hundreds of thousands of examples. It is obvious that the mainstream neural networks approach is not a very plausible candidate for machine cognition although the human brain is a neural network.

Dynamical systems were proposed as a model for cognition by Ashby (1952) already in the 1950s and have been developed further by contemporary researchers (for example Thelen and Smith, 1994; Gelder, 1998, 1999; Port, 2000; Wallace, 2005). According to this approach the brain is considered as a complex system with dynamical interactions with its environment. Gelder and Port (1995) define a dynamical system as a set of quantitative variables, which change simultaneously and interdependently over quantitative time in accordance with some set of equations. Obviously the brain is indeed a large system of neuron activity variables that change over time. Accordingly the brain can be modelled as a dynamical system if the neuron activity can be quantified and if a suitable set of, say, differential equations can be formulated. The dynamical hypothesis sees the brain as comparable to analog

feedback control systems with continuous parameter values. No inner representations are assumed or even accepted. However, the dynamical systems approach seems to have problems in explaining phenomena like 'inner speech'. A would-be designer of an artificial brain would find it difficult to see what kind of system dynamics would be necessary for a specific linguistically expressed thought. The dynamical systems approach has been criticized, for instance by Eliasmith (1996, 1997), who argues that the low dimensional systems of differential equations, which must rely on collective parameters, do not model cognition easily and the dynamicists have a difficult time keeping arbitrariness from permeating their models. Eliasmith laments that there seems to be no clear ways of justifying parameter settings, choosing equations, interpreting data or creating system boundaries. Furthermore, the collective parameter models make the interpretation of the dynamic system's behaviour difficult, as it is not easy to see or determine the meaning of any particular parameter in the model. Obviously these issues would translate into engineering problems for a designer of dynamical systems.

The quantum approach maintains that the brain is ultimately governed by quantum processes, which execute nonalgorithmic computations or act as a mediator between the brain and an assumed more-or-less immaterial 'self' or even 'conscious energy field' (for example Herbert, 1993; Hameroff, 1994; Penrose, 1989; Eccles, 1994). The quantum approach is supposed to solve problems like the apparently nonalgorithmic nature of thought, free will, the coherence of conscious experience, telepathy, telekinesis, the immortality of the soul and others. From an engineering point of view even the most practical propositions of the quantum approach are presently highly impractical in terms of actual implementation. Then there are some proposals that are hardly distinguishable from wishful fabrications of fairy tales. Here the quantum approach is not pursued.

The cognitive approach maintains that conscious machines can be built because one example already exists, namely the human brain. Therefore a cognitive machine should emulate the cognitive processes of the brain and mind, instead of merely trying to reproduce the results of the thinking processes. Accordingly the results of neurosciences and cognitive psychology should be evaluated and implemented in the design if deemed essential. However, this approach does not necessarily involve the simulation or emulation of the biological neuron as such, instead, what is to be produced is the abstracted information processing function of the neuron.

A cognitive machine would be an embodied physical entity that would interact with the environment. Cognitive robots would be obvious applications of machine cognition and there have been some early attempts towards that direction. Holland seeks to provide robots with some kind of consciousness via internal models (Holland and Goodman, 2003; Holland, 2004). Kawamura has been developing a cognitive robot with a sense of self (Kawamura, 2005; Kawamura *et al.*, 2005). There are also others. Grand presents an experimentalist's approach towards cognitive robots in his book (Grand, 2003).

A cognitive machine would be a complete system with processes like perception, attention, inner speech, imagination, emotions as well as pain and pleasure. Various

technical approaches can be envisioned, namely indirect ones with programs, hybrid systems that combine programs and neural networks, and direct ones that are based on dedicated neural cognitive architectures. The operation of these dedicated neural cognitive architectures would combine neural, symbolic and dynamic elements. However, the neural elements here would not be those of the traditional neural networks; no statistical learning with thousands of examples would be implied, no backpropagation or other weight-adjusting algorithms are used. Instead the networks would be associative in a way that allows the symbolic use of the neural signal arrays (vectors). The ‘symbolic’ here does not refer to the meaning-free symbol manipulation system of AI; instead it refers to the human way of using symbols with meanings. It is assumed that these cognitive machines would eventually be conscious, or at least they would reproduce most of the folk psychology hallmarks of consciousness (Haikonen, 2003a, 2005a). The engineering aspects of the direct cognitive approach are pursued in this book.

1.3 THE APPROACH OF THIS BOOK

This book outlines an engineering approach towards cognitive and conscious machines. These machines would perceive the world directly, as objects and properties out there, and integrate these percepts seamlessly into their cognitive processes. They would have the flow of inner speech and imagery. They would observe and introspect their mental content and perceive this mental content as immaterial. They would judge the world and their own actions by emotional good–bad criteria. They would be self-motivated agents with a will of their own. They would be able to move and execute actions as freely and readily as humans do. They would produce the hallmarks of consciousness.

The requirements for the direct perception of the world and the seamless integration would seem to call for a special way of perception and information representation. Here the traditional way of information representation of the artificial intelligence, the representation of information as statements in a formal language, is rejected. Instead, associative nonnumeric neural networks, distributed representations and cognitive architectures are investigated as the solution.

The book chapters proceed in logical order and build on previously presented matter; thus they should be studied sequentially. Some matters may appear trivial to advanced readers. However, in that case the reader should pay attention as apparently familiar matters may also have novel twists and meanings in this context. The last chapter, ‘Machine Consciousness’, should also be accessible to enlightened nonengineer readers who are interested in machine consciousness research. The contents of the book chapters are summarized here.

Chapter 1, ‘Introduction’, is the introduction in the present chapter.

Chapter 2, ‘Information, Meaning and Representation’, notes that humans operate with meanings while computers normally do not. It is argued that the brain is not a digital computer and therefore a cognitive robot should not be one either. Instead of digital representation of information a nonnumeric information representation by

large signal arrays, vectors, is proposed. This method is used in the approach of this book.

Chapter 3, 'Associative Neural Networks', presents neural networks that can be used for the associative processing of signal vectors. With these networks vectors can be associated with each other and can be evoked by each other. The neural network that executes this operation is called an associator. The limitations of the traditional neural associator are pointed out and enhanced associators that remedy these limitations are introduced. These associators are based on artificial associative neurons and here one specific execution of these, the Haikonen associative neuron, is presented. The proposed associators are not statistical neural networks and do not utilize learning algorithms such as 'backpropagation'.

Chapter 4, 'Circuit Assemblies', presents some example realizations and definitions for the basic circuit assemblies that are used as building blocks in the systems that are introduced later on in this book. These circuits include the basic associative neuron group, a 'Winner-Takes-All' circuit, an 'Accept-and-Hold' circuit, associative predictors and sequencers, timed sequence circuits and others.

Chapter 5, 'Machine Perception', describes the principles and circuits that would allow direct perception of the world, apparently as such in the machine. The associative neuron-based perception/response feedback loop is introduced as the fundamental building block that is applicable to every sensory modality, such as kinesthetic, haptic, visual and auditory perception. The purpose of this chapter is not to give an exhaustive treatment on the issues of perception, which would be many; instead the purpose is to present the basic ideas that are necessary for the understanding of the cognitive system that is outlined in the following chapters.

Chapter 6, 'Motor Actions for Robots', describes the integration of motor systems into the associative system. According to this approach a robot will be able to execute motor actions in the directly perceived environment without any numeric computations. Hierarchical control loops allow the planning and evocation of action in associative terms and the control of the actual execution in motor control terms. Simple control examples are described.

Chapter 7, 'Machine Cognition', argues that reactive processes are not sufficient for an autonomous robot and deliberative processes are needed. These processes call for higher-level cognition and the use of mental models. The issues of understanding, memories, perception of time, imagination and reasoning in the cognitive machine are considered.

Chapter 8, 'Machine Emotions', describes how the functional equivalents of pain, pleasure, good and bad can be evolved into a system of values, motivations, attention and learning control, and emotional decision making. The Haikonen System Reactions Theory of Emotions is applied to the cognitive machine.

Chapter 9, 'Natural Language in Robot Brains', outlines how natural language can be used and understood by the machine using the associative neural networks. The author's 'multimodal model of language' that seamlessly integrates language, multisensory perception and motor actions is utilized here. The grounding of word meaning is explained. The robot utilizes natural language also as 'inner