NEURAL COMPUTING

Research and Applications

Proceedings of the Second Irish Neural Networks Conference
Belfast, Northern Ireland, 25–26 June 1992

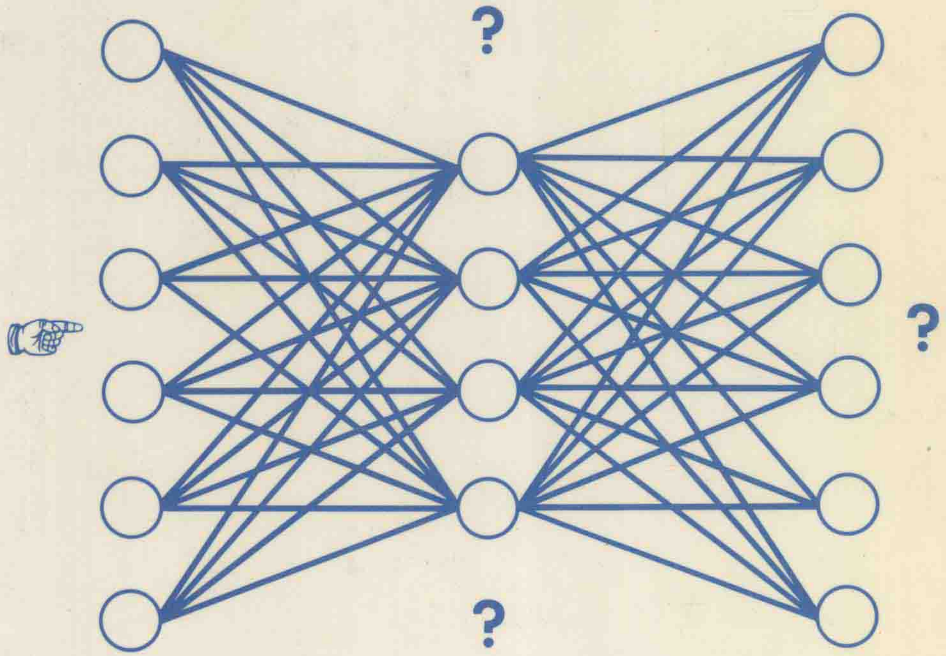

Edited by Gerry Orchard

NEURAL COMPUTING
RESEARCH AND APPLICATIONS

Proceedings of the Second Irish Neural Networks Conference
Belfast, Northern Ireland, 25–26 June 1992

Edited by

Gerry Orchard

Cognitive and Computational Modelling Group
School of Psychology
The Queen's University of Belfast
Northern Ireland

Institute of Physics Publishing
Bristol and Philadelphia

British Library Cataloguing in Publication Data

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

ISBN 0-7503-0259-3

Library of Congress Cataloging-in-Publication Data are available

Published by Institute of Physics Publishing, owned by the Institute of Physics, London. Techno House, Redcliffe Way, Bristol BS1 6NX, England

US Editorial Office: Institute of Physics Publishing, The Public Ledger Building, Suite 1035, Independence Square, Philadelphia, PA 19106, USA

Printed in the Great Britain by Galliard (Printers) Ltd, Great Yarmouth, Norfolk

NEURAL COMPUTING RESEARCH AND APPLICATIONS

Preface

The papers in this volume have been selected from the proceedings of the Second Irish Neural Networks Conference, held at the Queen's University of Belfast during June 1992. The papers are organised into two sections for convenience. The first section includes research on biological and psychological issues and work on the design of neural network architectures and algorithms important for further advances in neural network modelling. The research papers in the second section illustrate the wide range of tasks in industry, commerce, medical diagnosis and psychological modelling where neural nets are either already employed or where opportunities exist for task improvements using neural nets.

The range and depth of treatment of issues varied considerably across the papers presented at the conference with contributions from established workers as well as research students who have recently entered the area. We have attempted to represent this diversity in the papers which follow.

Part One: Basic Research

The opening paper is an invited keynote address by John Taylor in which he tackles the central and thorny issue of consciousness. An understanding of conscious attention has been a central aim for cognitive psychology for decades and although at present there is no generally accepted theory, there is agreement that one function of attention is to act as a filter or to select amongst competitive processes at some level within the overall information processing system. Taylor reviews a wide range of neuropsychological and neurophysiological research which suggests that the nucleus reticularis of the thalamus might be the physiological centre for the control of attention. One strength of this multidisciplinary approach is that it brings together work from psychology, neurobiology, neuroanatomy and mathematics and emphasises a role for artificial neural networks in constructing computational models which might provide a firmer grasp of this difficult topic.

In a second invited address, Daniel Amit's paper on attractor neural nets emphasises that empirically derived facts about cortical computation provide a more realistic basis for neural network modelling than an approach which is based on preconceived notions of computation. He suggests that psychologists' relatively uncritical acceptance of simple feedforward networks for cognitive modelling might occlude the possibility that more neurally plausible (and more demanding) approaches may offer a better understanding of the computational processes subserving cognitive function in the long run.

The claim that many artificial neural networks algorithms and architectures are too simple and fail to adequately account for some essential characteristics of real biological systems is taken up by Peter Laming. He argues that the tendency to use the firing rate of a neuron as a measure of activity and then to use a threshold function to modulate onward transmission reduces the flexibility of the net compared with an alternative which uses membrane potentials as a measure of neuron activity. He also points out that neural network modelling has largely ignored the way in which glial cells regulate the neuronal environment and consequently have the capacity to control the responsiveness of networks and subnetworks of neurons.

The biological theme continues with a paper by M G Stewart and co-workers which looks at the time course and location of synaptic modification following a simple learning task in chicks. Their study suggests considerable synaptic plasticity in two loci of the chick forebrain indicated by significant morphological changes in synaptic and dendritic number and they suggest that the different patterns of transient and sustained changes are correlated with short-term and long-term memory formation.

Andrew Schofield and David Foster's paper investigates how backprop nets may be applied to segment visual images consisting of oriented line elements. They compare their orientation detection net with known physiological data and show that some selected units simulate the properties of a variety of cells found in the visual cortex. The authors describe the development of the model into a two stage system capable of discriminating areas of high orientation contrast and go on to compare the results of the model with the performance of humans engaged on the same texture segmentation task.

The paper by R Linggard centres on identifying the basic principles which can account for the operation of a ganglion, a functional unit of highly interconnected neurons, in terms of the action of individual neurons. Working from the assumption that more complex nervous systems have evolved from primitive origins he proposes a model of the neuron that, when connected in a network, permits diverse behaviour.

A psychological perspective on neural net research is introduced with an invited paper by Mike Burton and Vicki Bruce. They outline some of the substantial empirical and theoretical work on the modular systems involved in face recognition and describe the implementation of the theoretical model in an interactive activation (IAC) net. The IAC net, one of the simplest connectionist models, is 'hand-wired' and non-adaptive in the sense that there is no learning algorithm. The authors describe the theoretical developments made possible through this relatively simple implementation and discuss advantages and disadvantages of using an IAC net compared with alternatives such as multi-layer feedforward architectures.

In the following paper, E Stain and R Cowie also discuss the IAC model but this time in reference to connectionist approaches to understanding how the recognition of an individual word is affected by the meaning of the sentence of which it is a part. They review a wide range of recent research which shows evidence for 'top down' sentence context effects on word recognition and they suggest that this evidence provides support for the psychological plausibility of some neural network models.

The G Dunbar and co-workers contribution, which provides neural network models of lexical combination, is concerned with a different level of context effects. They describe an experiment which illustrates the effect of semantic interaction, that is, how the influence of an adjective on the interpretation of a noun varies depending on the noun. They go on to implement several neural network models of the data, discuss novel ways of interpreting weight-state vectors, and conclude with comments on the psychological plausibility of the models.

This issue of psychological plausibility in neural network modelling is continued in two further papers. The first, by Tony Savage and Roddy Cowie, discusses the connections between learning theory and neural computation and assesses the interaction between the two areas of research. They focus on the role of reinforcement in Pavlovian and instrumental conditioning and examine three neural network models. Their paper suggests that the psychological plausibility of neural network models would be enhanced if they could more sufficiently account for the principles of associative learning in animals, indicated by the work on reinforcement. The second paper, by Gerry Orchard and Elaine Johnston, describes attempts to model the results of a simple human associative learning experiment using a variety of backdrop architectures. They show that the severe interference resulting from sequential learning in backprop makes it inappropriate for modelling paired associate learning and conclude with reservations regarding the application of backprop to models of cognitive function. The point at issue here is that whatever insights might be afforded by the representation of behaviour in a trained net, the backprop procedure is not a plausible mechanism for the acquisition of skills which are sequentially developed over time.

The paper by D Barnes and P Hampson addresses the same issue in work on a different area in psychology. They support the point that a major task for neural network modelling is not simply to simulate behaviours, in their case associated with inference and natural language, but to model their emergence as well. They point to recent research on behaviour analysis which demonstrates the power of human as opposed to non-human learning in drawing inferences, i.e. deriving new information which goes beyond the information provided in the immediate stimulus relationships. Barnes and Hampson describe a three-stage backprop model consisting of an encoder, a relational framing machine and a decoder, and discuss its performance in simulating empirical work on human inferential skills.

John Andrews and Mark Keane's work centres on another sequential behaviour, attention switching. They present two backprop nets aimed at modelling human performance in switching attention between comparative judgement tasks, in this case different common features in visually presented numbers. They compare the performance of a single three-layer net with a multi-net model consisting of a recognition net, a control net and an evaluation net.

The paper by Mark Plumbley outlines the contribution of information theory applied to both supervised and unsupervised learning, with particular reference to perceptual systems. He discusses an approach where the objective is to minimise loss of information as a signal is transmitted through a system and examines the differential effects of noise introduced at input and/or at output on a variety of unsupervised learning algorithms.

The seven papers which follow are concerned with the development of neural net architectures and algorithms. The first of these, from Noel and Amanda Sharkey, examines an issue which is important for both psychology and connectionist modelling, that is, how knowledge of the structure of one task can be transferred to another task, thus forming a basis for adaptive generalisation. They present a method for assessing transfer effects and proceed through a number of simulations to show how prestructuring a net can assist convergence. They compare simulation and empirical results on positive and negative transfer in the paired associate learning paradigm and conclude by pointing out that although prestructuring is an effective method for transferring knowledge between nets, identifying which aspects of prior knowledge are relevant remains a difficult problem.

Backprop is also the focus of the paper by S Namasivayam and J T McMullan who investigate ways of improving on the learning performance of standard nets using a variety of techniques. One of the ways they describe is by increasing connectivity and providing input to units in the topmost layer from any lower layer, a procedure which requires modifying the delta learning rule and increasing redundancy by duplicating the input nodes.

John Mitchell compares the use of the multi-layer feedforward backprop net with the radial basis function (RBF) net for time series prediction. He discusses issues relating to generalisation, learning and performance, concluding that the RBF net is severely limited compared to layered feedforward nets, though the latter suffer shortcomings when dealing with noisy problems. J Kadlec and co-workers also discuss RBF nets in comparison with layered feedforward nets, with an emphasis on modelling non-linear systems based on RBF and linear optimisation.

It is well known that training multi-layer feedforward backprop nets can be lengthy and involve many hundreds or thousands of iterations. In addition, design factors such as the choice of the number of layers and/or the number of units is often a matter of guesswork followed by later manipulation to persuade the net into producing an appropriate coding. Brendan Kiernan argues that such networks are best treated at least initially as a mapping from one binary hypercube onto another, and are therefore an implementation of a binary switching function. In this paper he discusses a perspective where artificial neurons are regarded as devices which divide Euclidean space and he proceeds to outline a way of understanding binary feedforward nets in terms of hypercubes and Karnaugh maps.

D Ellison investigate some of the properties of the Albus perceptron (CMAC), an adaptive system based on partly on the perceptron and influenced by the processing structure of the cerebellum. He describes the advantages of the CMAC compared to the classic perceptron and presents the results of a number of experiments which show the effect of varying the generalisation width on convergence rates for different presentation orders of individual training patterns.

E Luk and A D Fagan engage the problem of reducing the intersymbol interference in a non-linear digital communications channel. They compare the results of a transverse equaliser, considered as a special case of a multi-layer perceptron, with a standard MLP-based system, and discuss further research on modifying the standard sigmoidal activation function to improve performance.

Part Two: Applications Research

The first three papers in this section, which describe applications of neural networks in engineering, are introduced by an invited paper from G W Irwin. He outlines the advantages of neural networks, describes areas of current research and provides a viewpoint on the potential benefits, as well as some of the limitations, of using neural networks in control system engineering, particularly auto-tuning of PID controllers, non-linear plant identification and model-based adaptive control. One of the points he makes, which resonates with similar problems found in other areas of neural net research, is the difficulty of identifying the parameters distributively represented in a trained neural net and relating these to the physical system which it is modelling. Other difficulties in net design such as the number of hidden units and the method for input coding appear to be as problematic for engineers as for workers engaged in psychological and biological modelling. Given these problem areas, Irwin stresses the need to base the development of neural net models firmly on existing knowledge and theoretical results, a point also made elsewhere in this volume.

Ali Zalzala, in this paper on the learning control of robotic systems, minimises these problems of optimising net parameters (e.g. number of hidden units, input coding) in his adaptation of a basic multi-layer backprop net to provide a more effective system for controlling a robotic arm. He achieves this by modifying the net in order to provide some initial knowledge of the manipulator constraints. Although Zalzala emphasises that this is a modification best suited to engineering and robotics applications, his system may be useful for behavioural scientists interested in biological movement control.

The paper by Rene Biewald tackles the problem of controlling robot vehicle navigation. His approach is to employ the type of constrained symbolic (rather than spatially accurate) world knowledge which covers path control, obstacle avoidance, localisation and global navigation within the semi-regular environment common in industrial settings. The system he describes employs a neural network controller which uses ultrasonic sensory data to incorporate obstacle avoidance into various local navigation strategies.

Interpreting handwriting is a classically difficult problem for humans and even worse for any automated system. It could easily be argued that the only way of reading the wide variety of handwriting styles that exist is through a continual process of learning through experience to recognise patterns which may be seriously distorted away from a given or derived prototype. It is an important area for neural network research with applications particularly relevant to banks, postal services, insurance companies, credit card agencies and so on, all of whom have to spend a great deal of time and money in reading and authenticating handwritten symbols. The papers by D K R McCormack and B M Brown on signature verification, and by Michal Morciniec on digit classification, tackle the problem using different versions of backpropagation.

Speech and speaker recognition, another area with widespread applications, is the subject of the paper by P Linford and co-workers which demonstrates effective collaboration between academic and industrial researchers. They investigate the performance of a system where short frames of a speech signai are converted to spectral parameters which are then vector quantised using a Kohonen net. The considerable overlap in the frames encodes frame sequence and allows words to be defined as a list, or histogram, of the frame types which they contain. In determining the optimum frame duration, they found that the shortest duration of 6·4 mS produced best performance, a value smaller than that of conscious recognition. They conclude by discussing the implications of the results in terms of human speech perception.

E L Hines and co-workers describe the application of a variety of neural nets to the Warwick Electronic Nose, an artificial system based on the structure of human olfactory receptors. The standard backprop procedure was found to suffer the usual problems in selecting appropriate net parameters, avoidance of local minima and slow convergence. They compare the results of using standard backprop and an adaptive backpropagation algorithm (ABPA) using inputs normalised to ranges 0·1 and −1·1. They show that the ABPA with inputs normalised to the range −1·1 produce a significant reduction in training time, and they discuss directions for further improvements.

The application of pattern recognition and classification nets to assist with medical diagnosis has become more widespread in recent years. The paper by L Kilmartin and co-workers compares the performance of a multi-layer perceptron net, a nearest neighbour classifier and a Bayes classifier in the diagnosis of liver disorders. The MLP net significantly outperformed the other system and, given the same level of information, approached the level of diagnostic ability of a human physician. S J McKenna and co-workers examine a modularized two-stage neural network approach to cervical cell screening. In the first, unsupervised, learning stage cell images are Fourier transformed and inputted to a variable number of competitive modules. The outputs were then inputted to a feed forward net using backprop or conjugate gradient descent for binary classification. The authors compare the effects of different methods for representing the frequency domain image and different supervised learning architectures on the overall efficiency of the system. Both of these papers provide neural net systems which have obvious further applications within medical diagnosis, in addition to their contributions to the specific diagnostic tasks.

Christopher Burdorf and John Fitch's paper outlines the problems involved in trying to predict fluctuations in commodity market prices. Although the chaotic nature of dynamically changing interaction commodity markets makes accurate prediction of prices an unrealisable goal, even for neural networks, the neural net developed by Burdorf and Fitch provides a model which enables a better understanding of the processes involved.

Using portable laptop and notebook computers away from the office, at home or in transit, requires a high level of personal organisation to ensure that all the necessary software and data files are installed. Forgetting just one file might mean carrying around a piece of expensive and temporarily useless computer hardware with great detriment to efficiency and personal embarrassment if one is about to give a presentation. Russell Beale and co-workers' paper reviews the recent growth in what they term 'distributed working', and outlines a neural net approach to assist users by modelling their individual working habits in order to identify which items need to be downloaded onto the portable from the base system.

The final paper, by John Taylor, presents concluding comments on the conference in which he provides an overview of the considerable range of material presented. He discusses overall trends in current neural network research and suggests that one way forward for the development of artificial neural networks is to look again at the most effective information handlers—ourselves and other animals—for inspiration.

Gerry Orchard
Conference Organiser

Acknowledgements

We would like to thank the following individuals and organisations whose support ensured the success of the Second Irish Neural Networks Conference.

- The Secretary of State for Northern Ireland and the Northern Ireland Office, particularly the Department of Education and the Department of Health and Social Security, for funding guest speakers and providing a conference reception in the Parliament Buildings at Stormont.
- The Industrial Research and Technology Unit (IRTU) for sponsoring the publication of this volume.
- The Vice-Chancellor of Queen's University, the Dean of the Faculty of Science, and the Distinguished Scholars Fund for sponsoring guest speakers.
- The Director of the School of Psychology, Ken Brown, for sponsoring the conference and making available technical and administrative facilities, and the Head of Research Division, Roddy Cowie for encouragement and assistance.
- The guest speakers John Taylor, Daniel Amit, George Irwin, Mike Burton and Vicki Bruce for presenting such excellent keynote addresses, and to the session Chairpersons for their adroit regulation of the presentations.
- Very special thanks to Jennifer Cleland, the Conference Administrator, whose efficient management and administration ensured the smooth running of the conference itself and who was also responsible for the preparation and copy editing of this volume.
- Finally, with thanks to all contributors for their presentations and to the participants, all of whom collectively generated the stimulating, enthusiastic and creative atmosphere characteristic of neural network research.

Contents

ix **Preface**

xv **Acknowledgements**

Part One. Neural computing: Basic research

1 Modelling the mind
J G Taylor

21 Cognitive neuro-physiology: An empirical basis for neutral modelling and cognitive psychology
D J Amit

29 Modelling membrane potentials is more flexible than spikes
P R Laming

43 Synaptic correlates of short- and long-term memory formation in the chick forebrain following one-trial passive avoidance training
M G Stewart, T P Doubell, A Hunter and M Lowndes

53 A modular network model for texture segmentation based on orientation contrast
A J Schofield and D H Foster

63 A neural mechanism for diverse behaviour
R Linggard

73 An interactive activation model of human face recognition
A M Burton and V Bruce

85 Why connectionist networks provide natural models of the way sentence context affects identification of a word
E P Strain and R I D Cowie

95 Semantic interaction: A connectionist model of lexical combination
G Dunbar, M Kempen and N Maessen

105 Artificial Rewards: A comparative review of reinforcement in learning theory and neural computing
A J Savage and R I D Cowie

119 Sequential learning and interference in backpropagation models of paired associate learning
G A Orchard and E Johnston

129 Learning to learn: The contribution of behaviour analysis to connectionist models of inferential skill in humans
D Barnes and P J Hampson

139 A cognitive model of attention switching
J Andrews and M T Keane

145 Information theory and neural network learning algorithms
M D Plumbley

157 Prestructured neural nets and the transfer of knowledge
N E Sharkey and A J C Sharkey

165 The effect of all-connective backpropagation algorithm with redundancy on the learning characteristic of a neural network
S Namasivayam and J T McMullan

175 Comparing feedforward neural network models for time series prediction
J Mitchell

183 A nonlinear systolic filter with radial basis function estimation
J Kadlec, F M F Gaston and G W Irwin

191 Hypercube cuts and Karnaugh maps: Towards an understanding of binary feedforward neural networks
B Kiernan

199 Generalisation and convergence in the multi-dimensional Albus Perceptron (CMAC)
D Ellison

209 A neural network approach to equalization of a non-linear channel
E Luk and A D Fagan

Part Two. Neural computing: Applications research

215 Neural networks in control systems engineering
G W Irwin

227 A neural network methodology for engineering applications: The learning control of robotic systems
A M S Zalzala

235 A neural controller for navigation for non-holonomic mobile robots using sensory information
R Biewald

243 Handwritten signature verification using the backpropagation neural network
D K R McCormack and B M Brown

253 Hand-written digit classification: explorations in connectionism
M Morciniec

261 Speech recognition using a neural network VQ system
P Linford, R Linggard and W Millar

267 Improved rate of convergence in a MLP based electronic nose
E L Hines, J W Gardner, W W Fung and A A Fekadu

273 A neural network based diagnostic tool for liver disorders
L Kilmartin, E Ambikairajah and S M Lavelle

281 Using a two-stage artificial neural network to detect abnormal cervical cells from their frequency domain image
S J McKenna, I W Ricketts, A Y Cairns and K A Hussein

289 Using chaos to model commodity market price fluctuations in neural networks
C Burdorf and J Fitch

297 Neural network task identification for distributed working support
R Beale, A Dix and J Finlay

307 Concluding comments
J G Taylor

313 Author Index

Paper presented at Neural Computing Research and Applications: Part One
Queen's University of Belfast, Northern Ireland, 25–26 June 1992

Modelling the mind

John G Taylor

Centre for Neural Networks, Dept of Mathematics, King's College, London, WC2R 2LS, UK

Abstract: A particular mid-brain structure, the nucleus reticularis of the thalamus (NRT) is proposed to act as a global gateway for the control of attention. This conjecture is supported by neuropsychological and neurophysiological data, and further bolstered by a set of analytic models based on simplified circuitry of the system. This model, the so called 'Conscious I', from the analogy of the NRT and the retina, leads to insights into the nature of endogenous and exogenous attention and to a new approach to consciousness. This latter, based on the ' 'Relational Mind' model, requires especially clues given by, and roles assigned to, short-term memory and imagery.

1. Introduction

Consciousness has still to be explained convincingly. Yet there is increasing support, both from psychology and neurobiology, that some underlying neural explanation is possible. If attention is considered as the gateway to consciousness, then the neural basis for the latter is becoming clear through that for the former (Posner & Petersen 1990). It involves both cortical and mid-brain structures, the latter particularly involving the pulvinar nucleus of the thalamus (La Berge 1990).

This implies that the search for a purely cerebral origin of consciousness is not expected to be successful, in spite of the interesting insights to which it may lead (Dimond 1976, Crick & Koch 1990). The difficulty any solely cortical model is faced with is that of uniqueness. How can many competing activities in different cerebral areas meld together effectively to give the single stream of consciousness? This question is particularly relevant as the nature of cortical modularity and connectivity becomes clearer (Kaas 1989). It seems ever more problematic that the multifarious cortical areas could ever be combined to give any sensible coherence to consciousness. No wonder that the phrase 'multiple pandemoniums' was used in a recently claimed explanation of consciousness (Dennett 1991).

The idea of filtering (Broadbent 1958), of feature selection (Triesman & Gelade 1980) or of stimulus similarity (Duncan & Humphreys 1989) all argue for some sort of competitive activity at the basis of attentional processing. The thesis of this paper is that such competition arises from conjoint activity of cortex, thalamus and its associated reticular nucleus (NRT). In particular it is the special structure and mode of action of the latter which leads it to possess global control over its thalamic sources of inputs, and thence over more general cortical activity. At the same time cortical and non-cortical sources of thalamic input themselves exert modifying effects on the NRT, so altering attention. It is the sheet-like connected structure of the NRT, together with its inhibitory action, which is conjectured to lead to the strong correlation between distant cortical activities. This latter relationship leads to a neurobiological implementation of a recent relational theory of consciousness (Taylor 1991).

The ideas of this paper are meant to form a bridge between the neurobiological and psychological approaches to attention and consciousness. They lead to tentative extensions of the model to higher cognitive functions such as self-consciousness, reasoning and planning, as well as to emotional and goal related features. The model also gives a functional framework for neuronatomical data on thalamic, NRT and cortical connectivity on the basis of a mathematical model of the neural systems being considered, so leading the way to the emergence of a mathematics of consciousness.

The paper is organised in the following manner. In the next section we review work related to our model that is not readily accessible elsewhere, that is, on the structure of NRT on neural network modelling, and on the relational approach to consciousness. In the following section the basic features of the model are presented. This is based on a more detailed mathematical analysis, which is presented elsewhere (Taylor 1992d). Further developments of the model are outlined in section 4, in particular associated with possible species differences in feedback control and consciousness, on the nature of exogenous attention, on the manner in which orientation to a possible target might be achieved, on cortical transformations associated with feature detection and object recognition, and on short-term memory and imagery. The paper concludes with a discussion.

2. Related work

This paper is an interdisciplinary attack on the problem of consciousness, involving the fields of psychology, neurobiology, neuroanatomy and mathematics. It seems therefore of value to provide a brief background to certain of these topics so as to make the paper more comprehensible. At the same time relation to the work of others can be delineated.

Excellent recent reviews on filtering in attention have already been given (La Berge 1990, Sandon 1990, see Posner & Petersen [1990] for neurobiological aspects of attention). We thus limit our discussion here to the three topics of the neuroanatomy of the thalamic-NRT-cortical system, the mathematical modelling of neural activity to be used here, and a description of the relational approach to mind (a general guide to the research in this paper).

2.1. Nucleus Reticularis of the Thalamus

The NRT, present in all mammals, is a thin, curved sheet of cells so situated between thalamic relay cells and their cortical target sites as to be highly suggestive of its possible controlling influence on cortical input and activity. It is pierced by thalamo-cortical and returning cortico-thalamic fibres in a roughly topographic fashion, especially for inputs to and from primary sensory cortical areas. The main organisational principle (Jones 1975) is that NRT can be considered as a series of overlapping sectors, each related to a particular dorsal thalamic nucleus (or nuclei). The axons penetrating NRT give off excitatory collaterals to it, whilst NRT cells themselves only feed back inhibitorily in a roughly topographic fashion to the thalamic relay cells from which they have received their collaterals. The main NRT cell neurotransmitter is GABA, which is a well-known inhibitor. The NRT structure eminently qualifies it to be some sort of integrative filter modulating thalamic and cortical activity, the filter itself controlled by cortical, mid-brain and some brain stem structures (which also have inputs to NRT). That NRT performs a control function has been remarked on briefly in numerous papers. For example, as noted by Schiebel (1980): "Situated like a thin nuclear sheet draped over the lateral and anterior surfaces of the thalamus, it has been likened to the screen grid interposed between cathode and anode in the triode or pentode vacuum tube".

There are various features of the detailed circuitry and functionality which must be properly account for in any serious modelling. Firstly the nature of the inter-connectivity on the NRT sheet is itself species-specific. Thus in the rat, fine structure analysis shows only axo-dendritic synapses of presumed excitatory or inhibitory type (O'Hara & Lieberman 1985), but that of the cat and monkey have very clear dendro-dendritic synapses (O'Hara 1988, Deschenes et al. 1985), even some of these being reciprocal (Deschenes et al. 1985). It should be added that these dendro-dendritic synapses occur near cell bodies, and not on their distal dendrites.

Secondly the exact nature of the inhibitory effect of the GABA-ergic NRT cells is unclear. Thus evidence has been presented (McCormick & Prince 1986, Spreafico et al. 1988) that local application of GABA to NRT neurons cause depolarization of cell membrane rather than hyperpolarisation. However direct evidence of lateral inhibition of NRT neurons on each other was earlier provided by Ahlson and Lindstrom (1982). Moreover, the equilibrium potential E_{Cl} of the chloride conductance (the ion channel by which GABA influences the membrane potential) is about -65mV (McCormick & Prince 1986) whilst the average membrane potential is about -56mV (Avanzini et al. 1989). Thus the effect of GABA on neurons at or above their resting potential will be expected to be inhibitory. It is interesting to note that this effect will become excitatory if the membrane potential goes below E_{Cl}. This may be a useful control mechanism to prevent NRT neurons becoming so hyperpolarised as to be in their bursting state. Throughout this paper we will assume the action of GABA on NRT neurons is inhibitory, and that the neurons have membrane potentials always above E_{Cl} (although this latter will not be explicitly discussed in the modelling). The effect of GABA ergic feedback in thalamic relay cells and interneurons will be considered later.

Thirdly, the continuous nature of the NRT sheet is reasonably well supported by anatomical and cytoarchitectonic evidence, except for that part of it adjacent to the visual input (through the lateral geniculate LGN) termed the perigeniculate nucleus (PGN). Presently weight of opinion is in favour of complete connectivity between PGN and NRT. However, we will argue later (section 4.3) that may be incorrect. Fourthly, there are at least two modes of action of thalamic relay and NRT neurons: (i) relay-like behaviour, defined by tonic firing in response to inputs and (ii) aphasic bursting discharge behaviour, with the ability to maintain rhythmic burst discharges at about 6 to 8 Hz. There is great relevance of the second mode (ii) in thalamic and cerebral spindling (sequences of rhythmic bursting) activity in sleep and in the sleep-wake transition (Avanzini et al. 1989). Since we are mainly interested here in the waking state, only mode (i) of the NRT (and thalamic) neurons will be considered explicitly. Fifthly, there is an important species-specific difference in the NRT feedback to thalamus. In the rat only LGN has inhibitory interneurons, but other thalamic nuclei with output to NRT are known to possess very few non-relay cells (Barbaresi et al. 1986, Harris & Hendrickson 1987). In cat and monkey, however, such interneurons are apparently widespread. Inhibitory NRT feedback on these latter interneurons was used as an important part of the model of Steriade, Domich and Oakson (1986) and later by La Berge (1990). This model assumed disinhibition of interneuronal input control by inhibitory NRT feedback. However in rat, without such interneurons, one might expect NRT inhibition to feed back directly on to input thalamic relay cells, so having just the opposite effect on thalamic inputs. We will discuss this problem when we turn to the details of our model, and in particular consider it in the context of stability. Sixthly, there are also interesting features of the global wiring diagram in which NRT is concerned. Thus NRT is claimed not to be connected at all to the anterior thalamic nuclei in cats (Jones 1975, Pare et al. 1987), but there is apparently such connectivity in rats (Ohara & Lieberman 1982). The question of how important NRT inputs are to conscious processing will be discussed later in the paper.