

Springer Series in Cognitive and Neural Systems 3

A. Ravishankar Rao
Guillermo A. Cecchi *Editors*

The Relevance of the Time Domain to Neural Network Models

 Springer

The Relevance of the Time Domain to Neural Network Models

Springer Series in Cognitive and Neural Systems

Volume 3

Series Editors

John G. Taylor

King's College, London, UK

Vassilis Cutsuridis

Boston University, Boston, MA, USA

For further volumes:

www.springer.com/series/8572

A. Ravishankar Rao • Guillermo A. Cecchi
Editors

The Relevance of the Time Domain to Neural Network Models

 Springer

Editors

A. Ravishankar Rao
IBM Thomas J. Watson Research Center
1101 Kitchawan Road
Yorktown Heights, NY 10598,
USA
ravirao@us.ibm.com

Guillermo A. Cecchi
Dept. Silicon Technology
IBM Thomas J. Watson Research Center
1101 Kitchawan Road
Yorktown Heights, NY 10598,
USA
gcecchi@us.ibm.com

ISBN 978-1-4614-0723-2

e-ISBN 978-1-4614-0724-9

DOI 10.1007/978-1-4614-0724-9

Springer New York Dordrecht Heidelberg London

Library of Congress Control Number: 2011938345

© Springer Science+Business Media, LLC 2012

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Foreword

What is the relevance of temporal signal structure to the brain? We may gain some insight by comparing the brain to the computer. In the modern computer, signals are binary (have only two possible values), are made to change as quickly as technology permits, and temporal relations between signals are of central importance. The computer is driven by a clock through a quick succession of globally ordered states, while great care and effort is expended to make sure that no signal spills over from one state to the next. Ordered states are defined by commands in a program, each command specifying the setting of a large number of switches. At one time [1], this picture of a digital machine was taken seriously as a model for the brain, switches being identified with neurons. Digital machines are universal, meaning that any conceivable finite process can be realized in them, thus creating the vision that also the processes of the mind could be realized as processes in a physical machine. At the time, this idea was taken as the breakdown of the formerly perceived impenetrable glass wall between mind and matter. Unfortunately, the research program of Artificial Intelligence, which was built on this vision, has not given us intelligence in the machine yet. What is wrong with this vision of the brain as a digital machine? The succession of states in the computer is specified by programs, programs arise in human brains, and thus processes in the computer are imposed on it from outside. The big remaining question regarding the brain is that of the origin of its ordered states and sequences of states.

The role of temporal signal correlations in the brain may well be compared to that in the computer. The purpose of the brain is to coordinate activity in its various parts into ordered states and successions of states, such that things that belong together and form part of a functional whole are activated together. In this task of coordination, the brain is essentially out on its own, with very scant external help, which can in no way be compared to the insight of the computer's programmer. Classical artificial neural network models (important examples being the perceptron and associative memory) tended to grossly underestimate this task of generating and organizing brain states. In these models, time is paced by the presentation of stimuli, the network responding to each input pattern by convergence to a stationary state. This volume concentrates on a different brand of neural network models, in which

the generation of temporal patterns is the focus of interest. As these studies in their turn tend to pay less attention to the solution of functional tasks (beyond the standard problem of segmentation) and concentrate to a large extent on the modeling of brain rhythms that are actually found, it may be of interest if I attempt to give a wider perspective on the functional significance of temporal signal structure.

There are two aspects to the data structure of brain state, that is, to the way neural activity represents cognitive content. Considering neurons as elementary symbols, these aspects are (a) which of these symbols are active in a given psychological moment, and (b) how these symbols are put in relation to each other. If there are several objects in a scene, for example, each to be described by several attributes, a number of neurons will be active to represent the objects and the attributes (aspect (a)), but it is also necessary to represent the information which of the several attributes refer to which of the several objects (aspect (b)). Another example is visual (or more generally, sensory) segmentation: the problem of expressing the subdivision of the sensory field into coherent perceptual objects.

This is generally called the binding problem—the problem of representing relatedness between the symbols represented by neurons. It is now common lore to consider neural signal synchrony as solution to the binding problem: sets of neurons that are relating to each other express this by firing simultaneously. In simple cases, such as the above examples, this seems a perfect solution, as both generation and functional exploitation of signal synchrony are natural to neural networks. Signal synchrony is generated by plausibly existing neural connections. In object-attribute binding, the branching feed-forward connections from the original stimuli to neurons representing objects and attributes can propagate the same signal fluctuations to those neurons as signature of common origin and as expression of relations between attributes and objects. In sensory segmentation, horizontal connections between the neurons in a sensory field, being shaped by spatial closeness and other Gestalt laws, tend to run between neurons responding to the same perceptual object, and these connections thus tend to correlate signals within segments, as has been modelled many times. Functional exploitation, that is, the read-out of signal synchrony, relies on the fact that neurons are coincidence detectors, and thus functional interaction is restricted to sets of signals that are synchronous.

As nice and conceptually coherent the picture engendered by these examples is, it doesn't settle the binding issue, for experimental and for theoretical reasons. It is a disturbing fact that in spite of intensive search and in spite of ample evidence for neural signal synchrony, especially in the form of gamma rhythms (a frequency range from about 35 to 90 hertz), the prediction that signals within sensory segments should be globally correlated has not been confirmed experimentally. This alone raises the question whether there are other mechanisms than signal synchrony by which the brain can express binding, and theory is called upon to work out proposals. (One such proposal for solving the segmentation problem without using temporal binding is described in [2].) And there is more work to do for theory. The above binding examples—attribute-object binding and sensory segmentation—are misleading in their simplicity, reducing the binding issue to the decomposition of the neural state into a few blocks, a view often defended by reference to our inability to

keep simultaneously in mind more than a few chunks of a novel scene (the seven-plus-or-minus-two rule of [3]). On the other hand, we are evidently able to cope with very complex arrays of binding when representing a complex sentence, which necessitates to keep track simultaneously of multiple bindings between semantic, lexical, syntactic and phonetic elements, or when representing a visual scene of familiar structure, which necessitates the simultaneous handling of numerous relations between abstract and concrete patterns and their spatial relationships. Testimony to this complexity are the parsing trees of linguistics or the data structures of computer-based scene analysis (which themselves are all gross simplifications of the reality in our brains). Such complex relational patterns cannot be expressed by signal synchrony within realistic reaction times, given the poor temporal resolution of neural signals (1 to 3 msec, set by response times of neural membranes).

To do justice to the reality of our cognitive apparatus, we need a picture that lets us understand how the neural machinery in our head (or, for that matter, in a mouse's or salamander's head) is able to represent very intricate relational structures, and do so within typical reaction times of small fractions of a second. The called-for mechanisms must not only have high capacity and expressive power, but must in addition be able to store and retrieve relational structures once they have been formed. Finally, a clear picture must be developed for how the brain forms its preferred relational structures and how these preferred structures are to be characterized, for surely they can't be arbitrary.

A foreword is not the place to come forward with the proposal of a new system, but let me just remark that it is my conviction that rapid switching of synapses is part of the mechanism [4], and my laboratory has come to the conclusion that the machinery for storing and retrieving relational structures has the form of connections of a second order, of associative connections between switching synapses [5,6]. It is highly relevant to this book, however, to point out the fundamental significance of the time domain for these structures and processes, whatever they may be in detail.

To say it briefly, temporal signal structure is essential for expressing novel bindings, for laying down relational structures of growing complexity in memory, for reviving relational structures from memory (at a decisively reduced cost in terms of information rate) and for expressing bindings that resist memory storage. The mechanism for generating neural connectivity patterns, and, I claim, also of relational structures in memory, is network self-organization: the network creates structured activity patterns and synapses change in response to signal correlations, thus altering network and activity patterns. This reactive loop between network and activity tends to stabilize certain connectivity patterns, which are characterized by a close correspondence between signal correlations and connections. Network self-organization could perhaps be seen as a sequence of steps, each of which consists in the establishment of a temporal binding pattern followed by plastic change of connections, strengthening those between neurons bound to each other (that is, having correlated signals) while weakening those between neurons that are active but not bound to each other. Even if these individual binding patterns consist merely of one or a few blocks of bound neurons, the result of a sequence of such events can be a very intricate network of relations.

So far, network self-organization has been mostly applied to the generation of static networks, as illustrated by models of the ontogenesis of the visual system with its retinotopic connection patterns and columnar arrangements of sensory features (orientation, motion, stereo, color; for an example see [7]). If, however, synapses are allowed to switch on a fast time scale, a given set of neurons can support a number of alternate connectivity patterns, to be activated at different times. An important application of this could be neighborhood-preserving fiber projections corresponding to different transformation parameters to solve the problem of, for example, position-invariant pattern recognition [6]. For a model for how such alternate relational networks and their control structures could be generated by network self-organization, see [8].

Whereas the capacity of short-term memory is severely limited, as by Miller's seven-plus-or-minus-two rule, the capacity of long-term memory is generally held as virtually unlimited. The price to be paid is the laborious process of transferring short-term memory into long-term memory. Maybe this process is laborious because it necessitates the establishment of a new permanent relational network with the help of quite a number of consecutive activity binding patterns, as mentioned above.

Let me come back to our comparison between computer and brain. McCulloch and Pitts identified neurons with what in modern parlance are the logic gates—or bistable elements, or bits—of a digital machine. The bits of the computer can actually play the role of elements of pattern representations, analogous to the interpretation of neurons as elementary symbols. Many of them do, however, control switches (hence the name gate). Maybe it is time to reinterpret McCulloch and Pitts networks correspondingly, taking some of the “neurons” as elementary symbols, as is customary, but taking others as switches that can be opened and closed, an idea expressed already in [9].

The computer makes extensive use of temporal binding. All the bit settings in a given state are related to each other in the sense of forming one coherent functional state as specified in a program command. All signals necessary to constitute a state must have arrived at their target before the computer clock triggers the next state. The computer can afford this tight regime as its signals and pathways by now have a bandwidth of more than a gigahertz. In the brain, where the signal bandwidth is less than one kilohertz, a state comes into existence as the result of signals arriving without precise synchronization, so that the transition from one state to the next is a smooth and gradual affair.

The greatest step to be taken to transition from the computer to the brain is to find an explanation for the origin of states. As has been said above, whereas in the computer the switch settings essential for state organization are programmer-imposed, brain states must be self-organized. The gradual affair of brain state establishment may not just be a weakness but may be essential to this self-organization. If the brain has mechanisms to assess a state's level of self-consistency or completeness, it can iterate as long as it takes to establish a valid state. This complexity is the price the brain has to pay to be capable of programming itself as it goes along. If the state leaves behind a permanent trace that makes it easier to establish it, or parts of it, later again, and this self-programming may, after extensive exercise, install the equivalent of complex algorithms.

Unfortunately, our neural models are still very weak relative to this goal of brain state organization. This may be responsible for one great shortcoming of current neural network models and of related approaches—their inability to scale up in terms of numbers of elements or of functional sophistication to anything like the brains of even small animals. The difficulty is that larger systems cannot be made to converge to definite structures under the influence of training input. The solution to this problem must lie in decisive reduction of the systems' number of internal degrees of freedom, to be achieved by network self-organization (the one gigabyte of human genetic information not being enough to code for the petabyte needed to note down the wiring diagram of the human cortex). As an essential ingredient of any theory of network self-organization will be a clear understanding of the way in which temporal signal structure is shaped by a given network, the contents of this book seems to be highly relevant to neural network models of the coming decade.

References

1. McCulloch WS, Pitts W (1943) A logical calculus of ideas immanent in nervous activity. *Bull Math Biophys* 5:115–133
2. Wersing H, Steil JJ, Ritter HJ (2001) A competitive layer model for feature binding and sensory segmentation. *Neural Comput* 13(2):357–387. <http://ni.www.techfak.uni-bielefeld.de/files/WersingSteilRitter2001-ACL.pdf>
3. Miller GA (1956) The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychol Rev* 63:81–97
4. von der Malsburg C (1981) The correlation theory of brain function. Internal Report, 81-2, Max-Planck-Institut für Biophysikalische Chemie, Göttingen, Reprinted in Domany E, van Hemmen JL, Schulten K (eds) *Models of neural networks II*, Chap 2. Springer, Berlin, pp 95–119
5. Lücke J (2005) Information processing and learning in networks of cortical columns. Shaker Verlag, Dissertation
6. Wolfrum P, Wolff C, Lücke J, von der Malsburg C (2008) A recurrent dynamic model for correspondence-based face recognition. *J Vis* 8(7):1–18. <http://journalofvision.org/8/7/34/>, <http://journalofvision.org/8/7/34/Wolfrum-2008-jov-8-7-34.pdf>
7. Grabska-Barwinska A, von der Malsburg C (2009) Establishment of a Scaffold for orientation maps in primary visual cortex of higher mammals. *J Neurosci* 28:249–257. <http://www.jneurosci.org/cgi/content/full/28/1/249>
8. Bergmann U, von der Malsburg C (2010) A bilinear model for consistent topographic representations. In: *Proceedings of ICANN, Part III, LNCS*, vol 6354
9. Sejnowski TJ (1981) Skeleton filters in the brain. In: Hinton GE, Anderson JA (eds) *Parallel models of associative memory*. Lawrence Erlbaum, Hillsdale, pp 189–212

Acknowledgements

We are delighted to bring out a book dedicated to understanding the role of timing information in brain function. This has proven to be a daunting challenge. However, with the aid of advanced neuroscientific measurement techniques, more sophisticated mathematical modeling techniques, increased computational power and fast hardware implementations, we are making rapid progress.

We are very grateful to the contributing authors of the various chapters in the book for their valuable insights. We are particularly delighted to receive a Foreword written by Dr. Christoph von der Malsburg, a pioneer in this field.

We appreciate the efficient publication services provided by Ann Avouris and her staff at Springer. We are also grateful to the management at IBM Research, specifically Dr. Charles Peck and Dr. Ajay Royyuru in the Computational Biology Center for their support of this publication project.

A.R. Rao
G.A. Cecchi

Contents

1	Introduction	1
	Guillermo Cecchi and A. Ravishankar Rao	
2	Adaptation and Contraction Theory for the Synchronization of Complex Neural Networks	9
	Pietro DeLellis, Mario di Bernardo, and Giovanni Russo	
3	Temporal Coding Is Not Only About Cooperation—It Is Also About Competition	33
	Thomas Burwick	
4	Using Non-oscillatory Dynamics to Disambiguate Pattern Mixtures .	57
	Tsvi Achler	
5	Functional Constraints on Network Topology via Generalized Sparse Representations	75
	A. Ravishankar Rao and Guillermo A. Cecchi	
6	Evolution of Time in Neural Networks: From the Present to the Past, and Forward to the Future	99
	Ji Ryang Chung, Jaerock Kwon, Timothy A. Mann, and Yoonsuck Choe	
7	Synchronization of Coupled Pulse-Type Hardware Neuron Models for CPG Model	117
	Ken Saito, Akihiro Matsuda, Katsutoshi Saeki, Fumio Uchikoba, and Yoshifumi Sekine	
8	A Universal Abstract-Time Platform for Real-Time Neural Networks	135
	Alexander D. Rast, M. Mukaram Khan, Xin Jin, Luis A. Plana, and Steve B. Furber	
9	Solving Complex Control Tasks via Simple Rule(s): Using Chaotic Dynamics in a Recurrent Neural Network Model	159
	Yongtao Li and Shigetoshi Nara	

10 Time Scale Analysis of Neuronal Ensemble Data Used to Feed Neural Network Models 179
N.A.P. Vasconcelos, W. Blanco, J. Faber, H.M. Gomes, T.M. Barros,
and S. Ribeiro

11 Simultaneous EEG-fMRI: Integrating Spatial and Temporal Resolution 199
Marcio Junior Sturzbecher and Draulio Barros de Araujo

Index 219

Contributors

Tsvi Achler Siebel Center, University of Illinois Urbana-Champaign, 201 N. Goodwin Ave, Urbana, IL 61801, USA, achler@illinois.edu

Draulio Barros de Araujo Department of Physics, FFCLRP, University of Sao Paulo, Ribeirao Preto, Brazil, draulio@neuro.ufrn.br; Brain Institute, Federal University of Rio Grande do Norte, Natal, Brazil; Onofre Lopes University Hospital, Federal University of Rio Grande do Norte, Natal, Brazil

T.M. Barros Edmond and Lily Safra International Institute of Neuroscience of Natal (ELS-IINN), Rua Professor Francisco Luciano de Oliveira 2460, Bairro Candelária, Natal, RN, Brazil

W. Blanco Brain Institute, Federal University of Rio Grande do Norte (UFRN), Natal, RN, 59078-450, Brazil

Thomas Burwick Frankfurt Institute for Advanced Studies (FIAS), Goethe-Universität, Ruth-Moufang-Str. 1, 60438 Frankfurt am Main, Germany, burwick@fias.uni-frankfurt.de

Guillermo Cecchi IBM Research, Yorktown Heights, NY 10598, USA, gcecchi@us.ibm.com

Yoonsuck Choe Department of Computer Science and Engineering, Texas A&M University, 3112 TAMU, College Station, TX 77843-3112, USA, choe@cs.tamu.edu

Ji Ryang Chung Department of Computer Science and Engineering, Texas A&M University, 3112 TAMU, College Station, TX 77843-3112, USA, jchung@cse.tamu.edu

Pietro DeLellis Department of Systems and Computer Engineering, University of Naples Federico II, Naples, Italy, pietro.delellis@unina.it

Mario di Bernardo Department of Systems and Computer Engineering, University of Naples Federico II, Naples, Italy, mario.dibernardo@unina.it; Department of Engineering Mathematics, University of Bristol, Bristol, UK, m.dibernardo@bristol.ac.uk

J. Faber Fondation Nanosciences & CEA/LETI/CLIMATEC, Grenoble, 38000, France

Steve B. Furber School of Computer Science, University of Manchester, Manchester, UK M13 9PL, steve.furber@manchester.ac.uk

H.M. Gomes Department of Systems and Computation, Federal University of Campina Grande (UFCG), Campina Grande, PB, 58249-900, Brazil

Xin Jin School of Computer Science, University of Manchester, Manchester, UK M13 9PL, jinx@cs.man.ac.uk

M. Mukaram Khan School of Computer Science, University of Manchester, Manchester, UK M13 9PL, khanm@cs.man.ac.uk

Jaeroock Kwon Department of Electrical and Computer Engineering, Kettering University, 1700 W. University Avenue, Flint, MI 48504, USA, jkwon@kettering.edu

Yongtao Li Department of Electrical & Electronic Engineering, Graduate School of Natural Science and Technology, Okayama University, Okayama, Japan, yongtaoli@es.hokudai.ac.jp

Timothy A. Mann Department of Computer Science and Engineering, Texas A&M University, 3112 TAMU, College Station, TX 77843-3112, USA, [mann@cse.tamu.edu](mailto:manna@cse.tamu.edu)

Akihiro Matsuda College of Science and Technology, Nihon University, 7-24-1 Narashinodai, Funabashi-shi, Chiba, 274-8501 Japan

Shigetoshi Nara Department of Electrical & Electronic Engineering, Graduate School of Natural Science and Technology, Okayama University, Okayama, Japan, nara@chaos.elec.okayama-u.ac.jp

Luis A. Plana School of Computer Science, University of Manchester, Manchester, UK M13 9PL, plana@cs.man.ac.uk

A. Ravishankar Rao IBM Research, Yorktown Heights, NY 10598, USA, ravirao@us.ibm.com

Alexander D. Rast School of Computer Science, University of Manchester, Manchester, UK M13 9PL, rasta@cs.man.ac.uk

S. Ribeiro Brain Institute, Federal University of Rio Grande do Norte (UFRN), Natal, RN, 59078-450, Brazil, sidartaribeiro@neuro.ufrn.br; Neuroscience Graduate Program, Federal University of Rio Grande do Norte (UFRN), Natal, RN, 59078-450, Brazil; Psychobiology Graduate Program, Federal University of Rio Grande do Norte (UFRN), Natal, RN, 59078-450, Brazil

Giovanni Russo Department of Systems and Computer Engineering, University of Naples Federico II, Naples, Italy, giovanni.russo2@unina.it