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

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

Printed on acid-free paper

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 stan-
dard 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 mo-
ment, 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 re-
fer 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 re-
latedness 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 neu-
rons 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 propos-
als. (One such proposal for solving the segmentation problem without using tem-
poral 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

Frankfurt, Germany   Christoph von der Malsburg

Acknowledgements

Contents

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/CLINATEC, 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, jinxa@cs.man.ac.uk

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

Jaerock 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

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

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

Ken Saito College of Science and Technology, Nihon University, 7-24-1 Narashinodai, Funabashi-shi, Chiba, 274-8501 Japan, kensaito@eme.cst.nihon-u.ac.jp

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

Marcio Junior Sturzbecher Department of Physics, FFCLRP, University of Sao Paulo, Ribeirao Preto, Brazil

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

N.A.P. Vasconcelos Brain Institute, Federal University of Rio Grande do Norte (UFRN), Natal, RN, 59078-450, Brazil; Department of Systems and Computation, Federal University of Campina Grande (UFCG), Campina Grande, PB, 58249-900, Brazil; Edmond and Lily Safra International Institute of Neuroscience of Natal (ELS-IINN), Rua Professor Francisco Luciano de Oliveira 2460, Bairro Candelária, Natal, RN, Brazil; Faculdade Natalense para o Desenvolvimento do Rio Grande do Norte (FARN), Natal, RN 59014-540, Brazil; Faculdade de Natal, Natal, RN, 59064-740, Brazil

Chapter 1
Introduction

Guillermo Cecchi and A. Ravishankar Rao

Abstract The field of neural modeling uses neuroscientific data and measurements to build computational abstractions that represent the functioning of a neural system. The timing of various neural signals conveys important information about the sensory world, and also about the relationships between activities occurring in different parts of a brain. Both theoretical and experimental advances are required to effectively understand and model such complex interactions within a neural system. This book aims to develop a unified understanding of temporal interactions in neural systems, including their representation, role and function. We present three different research perspectives arising from theoretical, engineering and experimental approaches.

A significant amount of effort in neural modeling is directed towards understanding the representation of external objects in the brain, prominently in primary and associative cortical areas, and along the pathways that process sensory information from the periphery. There is also a rapidly growing interest in modeling the intrinsically generated activity in the brain represented by the default mode state, the emergent behavior that gives rise to critical phenomena such as neural avalanches, and the self-generated activity required to drive behavior. Time plays a critical role in these intended modeling domains, from the mundane yet exquisite discriminations the mammalian auditory system achieves in echolocation and voice recognition, to the precise timing involved in high-end activities such as competitive sports or professional music performance.

The effective incorporation of time in neural network models, however, has been a challenging task. Inspired by early experimental observations of oscillatory activity in electro-encephalogram recordings, and more recently in magneto-encephalogram observations [9], many theoretical efforts have been focused on the emergence and functionality of oscillations and synchronized activity of neural pop-

G. Cecchi (✉) · A.R. Rao
IBM Research, Yorktown Heights, NY 10598, USA
e-mail: gcecchi@us.ibm.com

A.R. Rao
e-mail: ravirao@us.ibm.com

ulations. In fact, the phenomenon of synchronization of ensembles of oscillators has been long recognized as an essential feature of biological systems. The pioneering work of Winfree and Kuramoto [10, 15] laid the foundation for a theoretical analysis of oscillator networks with relatively simple configurations. In neuroscience, similarly, Wilson and Cowan provided a framework to analyze the conditions under which ensembles of locally interacting neurons can give rise to oscillations. These early ideas have permitted researchers to find a conceptual link between neural modeling and the variety of complex perceptual and cognitive brain states that require local or long-distance coordination of neural ensembles, such as perceptual binding of real and illusory contours [5], face recognition [12], sensorimotor tasks [8], and attention [4], for which experimental evidence of involvement of oscillations has been documented.

Despite this wealth of observations, theoretical frameworks to conceptualize the full functional implications of oscillatory networks remain scattered and disconnected. Moreover, most modeling efforts are descriptive or heuristic, and tailored to specific, constrained situations. The purpose of this book is to provide a rallying point to develop a unified view of how the time domain can be effectively employed in neural network models. We will concentrate on three broad lines of research that run in parallel, but have enough sleepers to bind them together.

A first direction to consider is the utilization of ensembles of oscillators with the purpose of achieving specific, well-designed computational tasks. This is exemplified by the use of synchronization between oscillators as a means to solve the binding problem, that is, how to carve out all the units that, across different levels of abstraction, contribute to the emergence of a perceptual or cognitive object, without including other unrelated active units and without requiring a combinatorial, exploding number of connections. This issue, with implications for several related problems in neuroscience and signal processing such as perceptual integration and image segmentation, has been the focus of intense research in recent years, including approaches from biological, engineering and mathematical perspectives that range from physiological plausibility to adaptive network topologies for synchronization.

Motivated mostly by the need to implement practical applications for signal processing, a second line of interest is the development of dedicated hardware solutions for fast, real-time simulations of large-scale networks of oscillators. The implications of this type of research are enormous, and certainly go beyond the application to signal processing, as a successful hardware simulator could in principle form the basis on which to expand and multiply recent advances in the field of neural implants, neuro-prosthetics and brain-machine interfaces.

Finally, the advancement of electro-physiology techniques, including the development of new ones, together with practically inexhaustible computational power, has allowed for increased availability of experimental data to match the different levels of abstraction, be they multi-electrode recordings, electro- and magneto-encephalo- and cortico-grams (EEG, MEG, ECoG), functional magnetic resonance (fMRI), and ultra-fast optical imaging [1, 2, 6, 7, 11]. As a consequence, a third and in fact enormously active direction of research is the extraction of temporal patterns from the various sources of brain data, with objectives that range from the theoretically motivated to purely clinical studies.

The chapters have been organized accordingly, reflecting these three research approaches: theoretical, engineering and experimental. Following is a brief description of each contribution in relation to the objectives of the book.

1.1 Theoretical Background

While the synchronization of relaxation oscillators in simple network topologies is well understood, extensions to less constrained conditions are still an open area of research. In Chap. 2, De Lellis, di Bernardo and Russo investigate the conditions that allow synchronization in dynamical systems consisting of networks of units with arbitrary dynamics, topologies and coupling functions. They explore two parallel approaches. First, they focus on network plasticity: under very general assumptions on the vector field that defines the dynamics of the neural units (the QUAD condition, a generalization of Lipschitz's condition for uniqueness of solution of ordinary differential equations), it is possible to derive conditions for the dynamics of weight adaptation that ensure synchronization. Interestingly, the resulting update rule looks like a generalization of Hebbian learning, providing a plausible path for physiological interpretation. Alternatively, the authors ask whether it is possible to specify the node dynamics such that under simple coupling functions, and without the need of topology or weight adaptation, synchronization will be guaranteed. They answer this question through the use of contraction theory, which defines exponentially decaying bounds for the divergence of nearby trajectories of a dynamical system, if the norm-induced measure of the Jacobian is appropriately bound over the entire phase space. With this result, they show that the simple addition of self-inhibitory feedback to the units results in asymptotically synchronizing networks, for a large and simple class of coupling functions.

One of the most widely studied computational applications of neural synchrony is as a solution to the superposition catastrophe problem [13]. In Chap. 3, Burwick proposes an intriguing alternative to typical phase-locking synchrony implementations. One manifestation of the superposition problem is that when a network learns to recognize patterns with substantial overlap at the input level (e.g., many pixels in common), a fine-tuned layer of inhibitory units is required to avoid the spread of activation, and synchronization, to all of the units when any given pattern is presented.

Instead of relying on inhibition and units oscillating at constant frequency (at least in steady-state), the author shows that an acceleration term added to the dynamics of the units has the effect of creating a coherent synchronization only for the units that encode the winning pattern; the other units, while active, fail to synchronize, and instead perform a precessional, non-stationary motion in the space of phases. The acceleration term depends only on the lateral connections between the units, and is responsible for the winning ensemble "leaving behind" the units that do not sufficiently recognize the pattern.

A drawback of oscillatory models is that it is difficult to observe patterns of activation in individual neurons consistent with the theory; oscillations seem to truly be a collective phenomenon. As a consequence, for synchronization to take place,

oscillatory networks require time in order to transmit phase information back and forth across the ensemble, limiting their computational capacity as well as their modeling scope. In Chap. 4, Achler proposes a mechanism to address the binding of such distributed representations, which is achieved by feedback inhibition, rather than oscillations. The recurring inhibitory feedback suppresses the activity of "distractors" that do not belong to the object-encoding ensemble. While it still requires time for the recurrence to play out, this mechanism may potentially result in faster convergence times, and furthermore provide an alternative physiologically plausible model for binding.

While optimization approaches are common in signal processing, computer vision and, under various guises, in theories of efficient coding for early sensory processing, little progress has been made towards generalizing optimization principles to the realm of oscillatory networks. In Chap. 5, Rao and Cecchi develop an approach towards understanding sensory coding in an oscillatory network based on maximizing the sparseness of the spatio-temporal representation of inputs. This leads to a network dynamics by which higher-level representations of object features are synchronized with the lower-level object primitives. Furthermore, a superposition of input objects produces higher-level representations that preserve the distinction between the objects via the phases of the oscillations. This behavior leads to a quantitative characterization of the network behavior in terms of its efficacy in classifying and disambiguating superposed objects. These quantitative measures of network behavior are a function of the network topology, and depend on the fanout of feed-forward, feedback and lateral connections. Rao and Cecchi show that these quantitative measures of network behavior are maximized when the network topology is qualitatively similar to the topology of brain networks.

1.2 Engineering Development and Applications

Recollection and prediction constitute two essential attributes of behavior in higher-order organisms such as vertebrate animals. In Chap. 6, Chung et al. demonstrate that these attributes can arise through neural network controllers embedded in a dynamic environment. Both recollection and prediction require a temporal comparison in that the state of the organism or environment in the past or future is compared against the current state. In this sense, temporal dynamics are essential in order to represent interactions between the organism and environment. The authors use two tasks, that of ball-catching and pole-balancing to illustrate their framework for representing recollection and prediction. They use classical feed-forward and recurrent networks, and a hierarchy of retro- and predictive configurations based on the access to external markers and to internal dynamical regularities. Their work provides a possible framework to understand the representation of temporal order in the brain.

Overcoming the computational bottleneck is a major challenge facing researchers interested in exploring temporal phenomena in neural networks. This problem arises because simulations of these networks need to be carried out iteratively over small time increments in order to preserve numerical stability and the accuracy of the

simulations. This implies that several iterations may be necessary to produce oscillations and phenomena such as synchronization. Hence, hardware implementations of the underlying models are particularly attractive, as real-time behavior can be achieved.

In Chap. 7, Saito et al. present hardware-based neural models that reproduce the biologically observed characteristics of neural responses. They offer several improvements over existing hardware implementations, such as eliminating the need for inductors, thereby enabling a realization using CMOS IC chips. Their system has been implemented and tested to demonstrate the generation of oscillatory patterns that govern locomotion in micro-robots.

In Chap. 8, Rast et al. describe an alternative method to address direct hardware implementations of neural network systems. They have successfully designed a neural chip multiprocessor called SpiNNaker. This platform allows users to specify network models using a high-level hardware description language. The authors have implemented spiking neuron models with support for asynchronous interactions between the neurons. Their goal is to be able to simulate networks with a billion neurons, which approaches the size of a mammalian brain. This research constitutes a promising exploration of the challenges in designing a highly scalable architecture for performing simulations of large neural networks.

The ultimate goal of much of the research in neural systems modeling is to be able to produce behavior that mimics the function of biological organisms, particularly in the way real brains operate. Several researchers work on component models that explain the functioning of specific pieces of the entire puzzle, such as the encoding and processing of visual sensory information, or motor control. However, there are relatively few efforts to integrate such component models together. In Chap. 9, Li and Nara present a roving robot that combines sensory processing with motor behavior in order to solve a maze navigation task. Such a task is ill-posed, and requires appropriate control rules to be able to navigate around obstacles in order to reach a desired destination. Li and Nara use the mechanism of chaotic dynamics to implement an effective navigation scheme. The desired destination is signaled by means of auditory signals. When confronted with an obstacle, their system is able to generate alternate paths to reach the destination.

1.3 Biological Experimentation

As already stated, one of the main goals of the inclusion of time in neural network is to create a theoretical framework to understand the behavior of the nervous system. In Chap. 10, Vasconcelos et al. ask the experimental counterpart to the theories of temporal processing: is there a codification of external objects that specifically relies on timing? They utilize a multi-electrode technique pioneered by Miguel Nicolelis and John Chapin that allows the simultaneous recording of the electrophysiological activity of up to hundreds of neurons, distributed across several cortical and sub-cortical areas of the rat's brain. With this setup, they show that different objects with which the rats interact display a unique activity pattern, when the activity

is represented by the average firing rate within time bins with sizes ranging from 40 to 1,000 msec. Interestingly, they find that the optimal temporal resolution and the relative contribution of the different brain areas, in terms of object coding, are object-dependent, suggesting that the intrinsic features of the objects, as well as the behavioral interaction patterns the animals engage in with the, determine the coding features.

Moreover, they also find that exposing the animals to the objects and allowing the interact with them increases the correlation between the firing patterns with those generated in the subsequent periods of slow-wave and REM sleep, a finding that parallels those obtained by Bruce McNaughton in the hippocampus and Dan Margoliash in the songbird system [3, 14].

In Chap. 11, Sturzbecher and de Araujo expound on the possibilities and limitations of combining simultaneous recordings of functional magnetic resonance imaging (fMRI) and electro-encephalograms (EEG), from the perspective of increasing the spatial and temporal resolution of functional data. Focusing on the practical application of identifying the brain region from which seizures originate in epileptic patients (the epileptogenic zone), they explain that the problem of colocating EEG and fMRI sources can be addressed by solving the surrogate problem of mapping the traces created by the very frequent interictal epileptic discharges (IED), typical of the disease even in the absence of seizures. The inference of fMRI/EEG source location is normally based on the general linear model (GLM), which assumes a fixed temporal profile for the electro-physiology to hemodynamic response function, and measures the amplitude of the fMRI response to the IED's using linear correlation. The authors show that by using Kullback–Leibler divergence to measure these responses, which does not require any assumption about the transfer response, it is possible to significantly increase the accuracy of source localization. Their finding illustrates a concrete, practical application of an explicit inclusion of time in the analysis of real neural data.

References

1. Berdyyeva TK, Reynolds JH (2009) The dawning of primate optogenetics. Neuron 62(2):159–160
2. Cardin JA et al (2009) Driving fast-spiking cells induces gamma rhythm and controls sensory responses. Nature 459(7247):663–667
3. Dave AS, Margoliash D (2000) Song replay during sleep and computational rules for sensori-motor vocal learning. Science 290:812–816
4. Fries P et al (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. Science 291(5508):1560–1563
5. Gray CM et al (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. Nature 338:334–337
6. Haider B, McCormick DA (2009) Rapid neocortical dynamics: cellular and network mechanisms. Neuron 62(2):171–189
7. Han X et al (2009) Millisecond-timescale optical control of neural dynamics in the nonhuman primate brain. Neuron 62(2):191–198
8. Jefferys JGR, Traub RD, Whittington MA (1996) Neuronal networks for induced '40 Hz' rhythms. Trends Neurosci 19:202–208

9. Joliot M, Ribary U, Llinás R (1994) Human oscillatory brain activity near 40 Hz coexists with cognitive temporal binding. Proc Natl Acad Sci USA 91(24):11748–11751
10. Kuramoto Y (1984) Chemical oscillations, waves, and turbulence. Springer, Berlin
11. Liu H et al (2009) Timing, timing, timing: fast decoding of object information from intracranial field potentials in human visual cortex. Neuron 62(2):281–290
12. Rodriguez E et al (1999) Perception's shadow: long-distance synchronization of human brain activity. Nature 397:430–433
13. von der Malsburg C (1999) The what and why of binding: the modeler's perspective. Neuron 24:95–104
14. Wilson MA, McNaughton BL (1994) Reactivation of hippocampal ensemble memories during sleep. Science 265:676–679
15. Winfree AT (2001) The geometry of biological time, 2nd edn. Springer, Berlin

Chapter 2
Adaptation and Contraction Theory for the Synchronization of Complex Neural Networks

Pietro DeLellis, Mario di Bernardo, and Giovanni Russo

Abstract In this chapter, we will present two different approaches to solve the problem of synchronizing networks of interacting dynamical systems. The former will be based on making the coupling between agents in the network adaptive and evolving so that synchronization can emerge asymptotically. The latter will be using recent results from contraction theory to give conditions on the node dynamics and the network topology that result into the desired synchronized motion. The theoretical results will be illustrated by means of some representative examples, including networks of neural oscillators.

2.1 Introduction

Synchronization and coordination of motion are key features of many biological and living systems. For example, circadian rhythms regulate the functions of cells in our bodies which are entrained to the day/night cycle via the sophisticated action of various gene regulatory networks, see, for example, [14]. More notably, synchronization has been proposed as a powerful mechanism to explain some of the patterns observed in the brain [1, 3, 19, 27]. For instance, as noted in [18], partial synchrony in cortical networks can be used to explain several brain oscillatory patterns such as the alpha and gamma EEG rhythms. Also, synchronization of brain waves has

P. DeLellis · M. di Bernardo · G. Russo
Department of Systems and Computer Engineering, University of Naples Federico II, Naples, Italy

P. DeLellis
e-mail: pietro.delellis@unina.it

M. di Bernardo
e-mail: mario.dibernardo@unina.it

G. Russo
e-mail: giovanni.russo2@unina.it

M. di Bernardo (✉)
Department of Engineering Mathematics, University of Bristol, Bristol, UK
e-mail: m.dibernardo@bristol.ac.uk