COGNITIVE NEUROSCIENCE

Research Directions in Cognitive Science: European Perspectives Vol. 4

Edited by Guy A. Orban, Wolf Singer and Niels Ole Bernsen

LEA

Cognitive Neuroscience

Research Directions in Cognitive Science

European Perspectives Volume 4

Edited by

Guy A. Orban

Katholieke Universiteit te Leuven,
Laboratorium voor Neuro- en Psychofysiologie
Campus Gasthuisberg, B-3000 Leuven, Belgium

Wolf Singer

Max-Planck Institute for Brain Research,
Deutschordenstr. 46,
D-6000 Frankfurt/M 71, Germany

Niels Ole Bernsen

Cognitive Science Centre,
Roskilde University, DK 4000 Roskilde,
Denmark

LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS
Hove and London (UK) Hillsdale (USA)

Published on behalf of the Commission of the European Communities by:
Lawrence Erlbaum Associates Ltd., Publishers
27 Palmeira Mansions
Church Road
Hove, East Sussex, BN2 2FA

Publication No. EUR 11768 Vol.4 of the
Commission of the European Communities,
Directorate-General Telecommunications, Information
Industries and Innovation,
Scientific and Technical Communication Service,
Luxembourg.

British Library Cataloguing in Publication Data

Research directions in cognitive science.
1. Cognitive psychology.
I. Orban, Guy A. II. Singer, Wolf.
III. Bernsen, Niels O.
153

ISBN 0-86377-114-9
ISSN 0961-7493

Index compiled by Indexing Specialists, Hove, East Sussex BN3 2DJ
Printed and bound by BPCC Wheatons, Exeter

COGNITIVE NEUROSCIENCE
RESEARCH DIRECTIONS IN
COGNITIVE SCIENCE
EUROPEAN PERSPECTIVES
VOLUME 4

General Introduction
A European Perspective on Cognitive Neuroscience

Niels Ole Bernsen
Cognitive Science Centre
Roskilde University
DK 4000 Roskilde
Denmark

The present volume on Cognitive Neuroscience is one in a series of five presenting the findings of a joint European study in cognitive science 1987-88. The study was organised and funded as a collaborative network by the research unit FAST (Forecast and Assessment in Science and Technology) of the Commission of the European Communities and comprised about 35 scientists from the core disciplines of cognitive science. The research disciplines represented in the network were: cognitive psychology, logic and linguistics, cognitive neuroscience, human–computer interaction and artificial intelligence.

The aim of the network activity was to attempt a prospective mapping of research problems in cognitive science to be addressed over the next five to ten years. Prospective judgment of course has to be based on firm knowledge of the state of the art but a presentation of the state of the art of cognitive science as such was not our primary objective. This objective had already been addressed by a report presented to FAST in February 1986, *Cognitive Science in Europe* (ed. Michel Imbert et al.) and published by Springer-Verlag in October 1987.

As often happens in science nowadays, the dual aim of state-of-the-art presentation and prospective mapping realised through the FAST initiatives was linked to another, more practical, aim, namely that of making sure that cognitive science find its appropriate place in the European Community's long-term strategy for research and development in information technology. It is no doubt a pleasure to the contributors, and we hope to the European cognitive science community at large, that this has now happened to the extent that cognitive science has been included in the ESPRIT Basic Research Actions initiative which forms the upstream, basic research complement to the European Communities; ESPRIT programme in I.T.R. & D. In response to the first call for proposals for ESPRIT Basic Research Action (1988), some $65–70 million are currently being committed to basic research in cognitive science and artificial intelligence, computer science, and microelectronics. Moreover, now that European cognitive science is becoming increasingly visible partly through the FAST and ESPRIT initiatives, it seems reasonable to expect an increase in cognitive science funding at the national level.

As part of the more practical aim of funding procurement, the network agenda also included surveying and commenting upon the current institutional state and the state of collaboration in European cognitive science. Results and, we hope, some timely recommendations form part of a separate report of the network activity (Bernsen and the FAST Network, 1988).

The overall view of cognitive science taken in the papers published in the present volumes is fairly comprehensive though not exhaustive. Choice of topics for presentation and discussion has been made with a view to potential long-term relevance to information technology. Authors have been encouraged to take a personal view of their respective fields rather than a more comprehensive, and perhaps less exciting, encyclopaedic view. Each contribution has been written in order to make it comprehensible to cognitive scientists from other disciplines.

Since the general characteristics of current cognitive science are not, as such, addressed in the individual introductions or papers, a brief sketch may be in place here.

Sometimes a new theory can provide a unifying perspective to a number of hitherto disparate scientific endeavours and thus motivate a potentially drastic regrouping among the sciences. This is the case in cognitive science, where a new theory of the most general type which I shall call a *research programme* currently has this effect. The new research programme offering a unifying perspective to large parts of the sciences of logic, linguistics, psychology and neuroscience came from computer science and artificial intelligence and had been gaining ground steadily since the 1950s. It consists of the general idea that intelligent agents should be looked upon as information processing systems, that is, as systems receiving, manipulating, storing, retrieving, transmitting and executing information. Some of the general questions to ask concerning intelligent agents according to this research programme are: what information do such systems have? how is it represented? how is the information processed? and how are the processes implemented? The theoretical language of cognitive science is that of computation and information processing.

The objectives of cognitive science are to define, build and test information processing models of the various sub-systems (and of *their* sub-systems) making up intelligent agency, whether human (biological, natural) or artificial, and eventually to make them fit together into general cognitive theories and systems. The knowledge obtained can then be applied in various ways. Examples of cognitive sub-systems is vision, speech, natural language, sensory-motor control, memory, learning and reasoning. Each of these sub-systems are highly complex and may be further broken down into a number of functional components. Today's limited-capacity autonomous robots and knowledge based systems having a (written or spoken) natural language interface are examples of technologically implemented steps towards more general systems.

In somewhat more detail, the research programme of cognitive science may be characterised as follows:

1. Intelligence or cognition is physically implemented. However, a central level of analysis is the description of cognitive systems as systems for the manipulation of representations. Representations may consist of discrete symbols or may be of other types, such as distributed representations.

2. Widely different types of physical implementation are capable, in principle, of manipulating the same representations in the same ways: chips made of silicon or galliumarsenide, optical devices, mechanical or hydraulic devised, organic-biological systems.

3. Artificial, i.e. non-biological intelligence and hyperintelligence is therefore possible, at least in principle. Cognitive science is an investigation of both biological and artificial intelligence.

4. The level of description at which cognitive systems are described as manipulating representations cannot be reduced to: (a) The physical implementation of the system ; (b) the behaviour of the system ; (c) the conscious experiences of the system, if any.

5. Cognitive science is mechanistic. Intelligence or cognition, including semantics or meaning, and consciousness, is regarded as being produced by, in a wide sense, mechanical operations.

6. Acceptance of some version of functionalism. Functionalism states that cognition is constituted by the information processing functions which are physically implemented in the system.

7. Historically as well as in scientific substance and methodology, cognitive science is closely related to the computer and its information processing potential as studied by computer science and artificial

intelligence. The paradigms of cognitive science still derive from contemporary computer systems, whether serial or parallel, classical or connectionist. Use of computer simulations is essential to cognitive science except for the few areas, like problems involving social or organisational aspects, where specific computer modelling is not yet feasible.

8. Cognitive science is multidisciplinary. Methodologically, cognitive science aims at increased collaboration and cross-fertilisation between disciplines, the central assumption being that this is the most promising way of accelerating the achievement of the research programme, and hence also of realising the application potential of cognitive science. If anything has become apparent over the last twenty years, it is that cognition or intelligence are extremely complex phenomena whose investigation requires the full exploitation of a wide range of methodological tools. The basic idea behind the interdisciplinarity of cognitive science, then, is that each discipline employs its own *particular* methods in order to add *constraints* to the construction of *common* models and theories of the cognitive functions and their interrelationships. These models and theories are expressed in the common language of cognitive science, that is, in the language of the research programme and of one or other of the paradigms (discussed later). In addition, the basic idea behind interdisciplinarity assumes that each discipline could significantly contribute to the development of models and theories. So each discipline should, in order to belong to cognitive science, be concerned with both knowledge and processing (or competence and performance) in cognition, abstract programme and implementation, peripheral and central processes, sub-system integration, and the understanding of intelligent performance in complex, real-life tasks.

Present-day cognitive science is an interdisciplinary endeavour rather than a new science in its own right and speaking about core disciplines suggests that insights from other sciences like mathematics, physics, biology, computer science, anthropology, and the philosophy of science, mind and language actually do contribute to the advance of cognitive science. Furthermore, numerous sub-disciplines exist linking the core disciplines together, such as computational linguistics, computational logic, psycholinguistics, neuropsychology, and so on.

9. Cognitive science is closely related to application, in particular, though not exclusively, to the application of information technology. Applications of cognitive science are of, at least, three types.

(a) Specifying information processing models of the various sub-systems making up intelligent agency is essential to the building of increasingly intelligent artifacts such as the coming generations of vision systems, speech systems, natural language interfaces, robots, and knowledge-based systems. Interactions between the various disciplines of cognitive science have in the past produced such important AI knowledge representation and reasoning techniques as semantic networks, production systems, and logic programming, as well as significant results in areas like vision, speech and natural language processing. Future interaction will have to face still other areas of research where humans continue to perform far better than current artificial systems, as in casual reasoning, reasoning about time, plans and intentions, learning, or fluent, skill-based behaviour.

(b) It has become clear that the actual design of information processing systems should go hand in hand with research on their interactions with human agents in real-life task situations. Through the computerisation of work in all sectors of society, information technology has become an important tool at the interface between humans and their work. Successful system design, whether of large control systems, computer networks, manufacturing systems, office systems, tutoring systems, speech and language systems, or expert systems, not only depends on the training of users but also on the system's inherent adaptability to users. If a system's design is not successful, users are not likely to want to use it, and if they use it, serious accidents may occur, as in nuclear power plants or large chemical installations. In this situation, cognitive science research is strongly needed in the interaction between I.T. tools, task domain and work context, the cognitive resources of users, and the new patterns of social interaction arising from the use of computers. Thus, the rapidly evolving field of human–computer interaction research could be included among the core disciplines making up cognitive science as being sufficiently distinct from, and somewhat orthogonal to, the others to merit a disciplinary label of its own.

(c) Although these two points describing the application potential of cognitive science for information technology have been central to the present network activity, it should be noted that they do not exhaust the application potential of cognitive science. The human information processing system can be damaged or inoperative in various ways and from various causes, with neurological and psychological disorders or loss of certain mental abilities as a result. Studying the system and its behaviour in information processing terms promises better ways for diagnosing, repairing, and retraining the system as well as better ways of supplying the system with efficient prostheses.

The aforementioned points 1–9 are by no means uncontroversial among cognitive scientists. And needless to say, interdisciplinary collaboration among traditionally separate disciplines is not uncontroversial either. What is interesting, however, is that points 1–9 currently do seem to represent international convegence towards a common conception of cognitive science.

A research programme is in itself nothing very important. What matters are the research paradigm(s) demonstrating the practical viability of the programme. A research paradigm consists of one or more successful, specific applications of the research programme to particular problems falling within its scope. These applications, in casu models of specific cognitive functions, are seen by the scientific community as evidence that the principles on which they are based might be generalised to account for a much larger class of cognitive phenomena, and possibly to all of cognition. The central scientific task, then, is to implement and test this assumption. Cognitive science currently appears to have to consider two different research paradigms. The relationship between them is not clear at this point and is subject to strong, ongoing debate (e.g. Fodor and Pylyshyn, 1988; Smolensky, 1988).

According to the *Classical AI* paradigm, an intelligent system's input and output consist of physical signals and movements, whereas a large central part of the information processing linking input and output consists of automatic computation over language-like, discrete, and combinatorial symbolic codes, as in conventional serial or more recent parallel computers (Fodor, 1976; Newell, 1980; Pylyshyn, 1984). According to the *Neural Network Computation* paradigm or the *Connectionist* paradigm, which was been strongly revived in the 1980s, computation over discrete combinatorial symbols exists to a lesser extent, or does not exist at all, in intelligent biological systems. Instead, the complex cognitive abilities of higher organisms are based on the information processing abilities arising from the collective behaviour of large populations of highly interconnected and very simple processing elements, such as nerve cells or simple artificial processing elements. Consequently, it is maintained, cognitive scientists should develop and implement their theories of intelligent information processing in ways that resemble much more closely the way in which the brain actually operates (McClelland and Rumelhart, 1986).

Today, both paradigms can claim a number of successes in terms of concrete models jointly covering most areas of cognitive science.

This ongoing debate over research paradigms is a very real one because, at the present time at least, the two paradigms clearly do generate different systems providing different functional primitives (i.e. different elementary information processing capabilities) and possibly different behaviours.

Thus, the virtues of connectionist systems include their ability to rapidly acquire and apply large amounts of knowledge in noisy situations not governed by rigid laws but by context-sensitive regularities having many exceptions. And the paradigms generate importantly different directions for research and different technologies, and tend to attract different core disciplines of cognitive science. Thus, many researchers in classical AI tend to be sceptical about the potential of artificial neural network systems, despite the success of similar systems in nature; others do not doubt the importance for AI of "massive parallelism", but argue that connectionist systems do not really represent an alternative research paradigm, only a specific way of physically *implementing* classical *cognitive architectures*. Logicians and most linguists tend to disregard connectionist systems, whereas many cognitive psychologists and virtually all cognitive neuroscientists, who never really adopted the classical AI paradigm anyway, tend to embrace the neural network computation paradigm as the first firm basis for realistic models and general theories of cognition. The formal language of the classical AI paradim is that of symbolic logic and algebra, whereas the formal language of the connectionist paradigm is that of dynamic systems theory belonging to mathematical physics. Also, researchers studying the peripherals of cognitive systems like speech, low-level vision, or movement, all of which involve considerable signal processing, appear to be more strongly attracted by neural network computation than those studying central processes. Not least, this latter point has led many cognitive scientists to believe that the two paradigms really are basically *different cognitive architectures*, but that they are compatible in the sense of being apt to model different types of cognitive function or different parts of cognitive functions, such as voluntary, introspectively accessible, attentive, and controlled processes versus skilled, automatic, pre-attentive, probably massively parallel processes. Research in the next five to ten years will no doubt result in important attempts to integrate these two approaches or paradigms of cognitive science.

Two interrelated themes dominate the network findings and cut across the distinction between scientific substance and methodology. These themes can be viewed as constituting some central tendencies of current research covering all the core disciplines of cognitive science. Since the themes or trends are based on prospective analyses of most areas of cognitive science, from research on vision and speech to research on natural language, logic, and reasoning, they appear to form a stable pattern. These trends should be encouraged by an appropriate research policy. The themes are integration in theory, computer models, and actual working implementations, and real-world validity of theories, models and applications.

The theme of integration covers the following aspects:

Integration of different cognitive functions;
Integration of cognitive sub-functions into cognitive functions;
Integration of models into more general theories;
Integration of partial models into full models;
Integration and convergence of approaches, methods and results from different disciplines;
A more theory-driven approach in traditionally experimental disciplines like cognitive psychology and neuroscience.

Integration clearly means a trend towards the construction and testing of general theories and towards increased interdisciplinarity. Moreover, as mentioned earlier, the possibility of integrating the two current research paradigms of cognitive science is currently the subject of lively debates.

The real-world theme covers the trend towards explaining, simulating and actually building larger-scale, more general-purpose, real-time, closer-to-real-life systems of speech and grammar, natural language and communication, vision, perception and movement or action, and problem-solving. This trend also receives strong support from human–computer interaction research, which from the outset has to face human information processing in complex, real-life situations. Real-world research in cognitive science contrasts with, e.g. research in cognitive psychology on the performance of abstract and ecologically meaningless tasks in the laboratory or AI research on system performance in "micro-worlds". The real-world trend marks an important step beyond these classical approaches in cognitive science and implies the disappearance from the field of the sharp, traditional distinction between basic and applied research.

The two themes of integration and real-world validity are closely related because, in a large number of cases, explanation and synthesis of performance in complex, real-life situations require an integration of different cognitive functions and systems, and of different approaches. The themes are also closely related to the technological applicability of models, systems, and theories because integration and real-world validity is what is needed, both in order to extend the range of applicability of systems and in order to adapt them to users. In many cases, computer simulations may function both as theoretical test-beds and as software prototypes of potential machines for technological applications.

Numerous examples of the above trends can be gathered from the network papers. I shall leave it to the reader to find these examples and to judge whether they are sufficient to justify the conclusions stated above. If they are, then it can confidently be stated that contemporary cognitive science in Europe demonstrates the viability of the research programme, the productivity of the research paradigms, the convergence of disciplines towards common models, theories and problems to the extent allowed by the existence of two different paradigms, and the potential applicability of results. Not everything is idyllic, however; nor could or should this be so within an emerging science. Most of the basic questions still remain unanswered, with the prospect that cognitive science may look very differently in ten years time from now.

I would like to thank all network participants for their friendly collaboration during the past two years. We all learned, I think, that even today, large-scale, multidiciplinary European collaboration in science is not a matter of course, but something that requires a substantial effort. I am especially grateful to the leaders of the four "network institutions" and Special Editors of the first four volumes in the series: Alan Baddeley (Cambridge APU), Michel Imbert (University Paris VI), Jens Rasmussen (Riso National Laboratory) and Helmut Schnelle (Bochum University). Without their judgment, vigilance, patience and collaborative spirit the network would never have been set up, let alone have produced anything. Derek Sleeman (Aberdeen), Special Editor of Volume V, entered the collaboration at a later stage and has demonstrated impressive efficiency in catching up with the work that had already been done.

We are all deeply indebted to Dr Riccardo Petrella, Head of FAST, whose sensitivity to emerging trends in science and technology first brought cognitive science to the attention of EC scientific programmes and whose dynamism, non-hesitant support, and constant goodwill have made the network possible. I must personally thank Dr Petrella and the Danish Science and Engineering Research Council for making my one-year stay at FAST possible and the EC's ESPRIT programme for allowing me time to complete the work while assisting in setting up the ESPRIT Basic Research Actions.

REFERENCES

Bernsen, N.O. and the FAST Cognitive Science Network (1988). *Cognitive science: A European perspective.* (Report to the FAST Programme.) FAST, EC Commission, Brussels.

Fodor, J.A. and Pylyshyn, Z.W. (1988). Connectionism and cognitive architecture, *Cognition, 28,* (1-2), 3-71.

Fodor, J.A. (1976). *The Language of Thought.* Sussex: The Harvester Press.

Imbert, M., Bertelson, P., Kempson, R., Osherson, D., Schnelle, H., Streitz, N.A., Thomassen, A., Viviani, P. (Eds.) (1987). *Cognitive Science in Europe.* Springer-Verlag.

McClelland, J.L., Rumelhart, D.E. and the PDP Research Group (Eds.) (1986). *Parallel Distributed Processing* Vols. 1-2, Cambridge MA: Bradford Books, MIT Press.

Newell, A. (1980). Physical Symbol Systems, *Cognitive Science, 4,* 135-183.

Pylyshyn, Z.W. (1984). *Computation and Cognition. Toward a Foundation for Cognitive Science.* Bradford Books, Cambridge MA, MIT Press.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences Vol. 11*,(1), 1-74.

Contents