

Itiel E. Dror and Stevan Harnad (eds.)

Cognition Distributed

BENJAMINS CURRENT TOPICS

16

Cognition Distributed

How cognitive technology extends our minds

Edited by

Itiel E. Dror

Stevan Harnad

University of Southampton

John Benjamins Publishing Company

Amsterdam / Philadelphia



™ The paper used in this publication meets the minimum requirements of American National Standard for Information Sciences – Permanence of Paper for Printed Library Materials, ANSI Z39.48-1984.

Library of Congress Cataloging-in-Publication Data

Cognition distributed: how cognitive technology extends our minds / edited by Itiel E. Dror, Stevan Harnad.

p. cm. (Benjamins Current Topics, ISSN 1874-0081 ; v. 16)

Includes bibliographical references and index.

1. Cognition. 2. Cognitive science. 3. Artificial intelligence. I. Dror, Itiel E. II. Harnad, Stevan R. III. Series.

BF311.C5548734 2007

153--dc22

2008044247

ISBN 978 90 272 2246 6 (Hb; alk. paper)

© 2008 – John Benjamins B.V.

No part of this book may be reproduced in any form, by print, photoprint, microfilm, or any other means, without written permission from the publisher.

John Benjamins Publishing Co. · P.O. Box 36224 · 1020 ME Amsterdam · The Netherlands
John Benjamins North America · P.O. Box 27519 · Philadelphia PA 19118-0519 · USA

Cognition Distributed
How cognitive technology extends our minds

Benjamins Current Topics

Special issues of established journals tend to circulate within the orbit of the subscribers of those journals. For the Benjamins Current Topics series a number of special issues have been selected containing salient topics of research with the aim to widen the readership and to give this interesting material an additional lease of life in book format.

Volume 16

Cognition Distributed. How cognitive technology extends our minds

Edited by Itiel E. Dror and Stevan Harnad

Some of these materials were previously published in *Pragmatics & Cognition* 14:2 (2006)

*We dedicate this to Tim Berners-Lee who freely bequeathed the World Wide Web
to humanity as its Cognitive Commons*

About the Authors

Stevan Harnad, born in Hungary, did his undergraduate work at McGill and his doctorate at Princeton and is currently Canada Research Chair in Cognitive Science at the University of Quebec, Montreal and adjunct Professor at Southampton University, UK. His research is on categorisation, communication and cognition. Founder and Editor of *Behavioral and Brain Sciences*, *Psychology*, and *CogPrints Archive*, he is past President of the Society for Philosophy and Psychology, Corresponding Member of the Hungarian Academy of Science, and author and contributor to over 150 publications: <http://www.ecs.soton.ac.uk/~harnad/>

Itiel Dror is a Senior Lecturer in cognitive sciences at the University of Southampton, UK. He holds a number of graduate degrees, including a Ph.D. in cognitive psychology from Harvard (USA). He specializes in the fields of human cognition and behaviour, training and skill acquisition, technology and cognition, expertise, and biometric identification. Dr. Dror has worked in universities around the world and conducted research and consultancy to numerous organizations, including the UK Passport Services, the USA Air Force, the Japanese Advance Science Project, the European Aerospace Research & Development Agency, the Israeli Aerospace Industry, the BBC, and many commercial companies. He has published widely in both basic science and applied domains. For more information: <http://users.ecs.soton.ac.uk/id/>

Pierre Poirier is research-professor in the philosophy of cognitive science, Université du Québec à Montréal. Three questions currently interest him. What does mobile robotics teach us about human cognition? How can evolutionary considerations be introduced in cognitive science? What is the best epistemological framework for comprising all systems, from genes to brains to institutions, which properly can be described as possessing/producing knowledge. Living in Montréal, he is happy to have never found it necessary to get a driver's license.

Guillaume Chicoisne is research and planning officer in the Cognitive Science Institute, Montréal. As a proponent of multidisciplinary approaches, his research interests are manifold, but three themes emerge from his past and current projects: virtual communities, multi-agent systems, and man-computer interaction. Nevertheless, his favourite way of working involves mainly pen, paper and a whiteboard.

John Sutton is currently head of the Philosophy Department at Macquarie University in Sydney. He is interested in the interdisciplinary study of memory, and in testing the methodological implications of distributed cognition and the 'extended mind' hypothesis by seeing if these frameworks have anything to offer memory researchers in a range of other disciplines. His other interests include Descartes, animal spirits, habit and skill, dreaming, the history and philosophy of science, and the foundations of cognitive science.

David Kirsh is Professor in the Cognitive Science Department at the University of California San Diego, where he has been since its inception in 1989. His D.Phil is from Oxford University in Philosophy and he held a post doc and later was research faculty at MIT in the Artificial Intelligence Lab. His interests are broad but recently concentrate on understanding principles for designing interactive environments.

Arthur Glenberg is professor of Psychology and Educational Psychology at the University of Wisconsin. He has held visiting scientist positions at the University of Colorado, the Institute for Interdisciplinary Studies (ZiF), the Max Planck Institute for Cognitive Science and Neuroscience, and the Department of Neuroscience in Parma, Italy. His research is developing a theory of language comprehension based on the embodied cognition framework.

Angelo Cangelosi is Reader in Artificial Intelligence and Cognition at the University of Plymouth. His research interests include the modeling of language evolution in cognitive agents and robots, experimental and computational study of spatial language, the symbol grounding problem, and the combination of neural nets and genetic algorithms for the evolution of neural architectures. In Plymouth he leads the "Adaptive Behaviour and Cognition" research group.

Luc Steels is professor in AI at the Free University of Brussels (VUB) and director of the Sony Computer Science Laboratory in Paris. He has contributed to many areas of Artificial Intelligence research ranging from knowledge-based systems to robotics. His most recent research topic is concerned with models for the origins and evolution of language using computer simulations and robotic experiments.

Todd Gureckis is a postdoctoral fellow at Indiana University on the NIH-NIMH Cognitive Modeling Training Grant. He received his Ph.D. in psychology in 2005 and his B.S. in Electrical Engineering in 2001 both from the University of Texas at Austin.

Robert Goldstone is a Chancellor's Professor of Psychological and Brain Sciences and the Director of the Cognitive Science program at Indiana University. He was awarded a 2000 APA Distinguished Scientific Award for Early Career Contribution in Cognition and Learning, and a 2004 Troland research award from the National Academy of Sciences.

Daniel L. Schwartz is Associate Professor of the Learning Sciences in the School of Education at Stanford University. He received his Ph.D. in Human Cognition from Columbia University. His work combines cognitive science and classroom research to understand processes of learning and problem solving and to help design and evaluate novel forms of instruction using advanced technologies.

Taylor Martin is Assistant Professor of Science and Mathematics Education at the University of Texas at Austin. She received a B. A. in Linguistics from Dartmouth College, an M.S. in Psychology from Vanderbilt University, and a Ph.D. in Education from Stanford University. Her primary research interest is how people learn content in complex domains from active participation. She examines how hands-on activities impact mathematics learning and investigates the development of adaptive expertise.

Jiajie Zhang is a Professor and the Associate Dean for Research in the School of Health Information Sciences at the University of Texas at Houston. His research interests include human-computer interaction, cognitive science and engineering, distributed cognition, and health informatics.

Vimla L. Patel is a Professor of medical informatics and medical psychology at Columbia University. Her research interests include medical cognition, cognitive evaluation of human-computer interaction, collaborative healthcare teams, and naturalistic decision making in critical-care environments.

Jeffrey M. Stibel is currently President and CEO, Web.com, Inc. (NASDAQ: WWW). He previously served as General Manager of United Online, Inc. (NASDAQ: UNTD), owners of NetZero, Juno, and Classmates.com. Stibel also serves on the Board of Directors for Brown University's Entrepreneurship Program, Thinmail, The Search Agency, Sea Vista Development, Axon Labs, and Worldwide MediaWorks. Stibel studied business, economics and brain science at MIT and Brown University, where he was a Brain and Behavior Fellow.

Chris Baber has been a lecturer at the University of Birmingham since 1990 and was promoted to Reader in 2004. He is the author of some 50 journal papers and over 200 conference papers or book chapters. His research concerns human interaction with technology, in particular through the use of everyday skills and abilities. This work includes the study of speech-based interaction with computers and wearable computing, and has been funded by the UK Ministry of Defence, EPSRC and the European Union. He is a member of the EPSRC College, and has been on the Scientific Committee of the IEEE International Symposium of Wearable Computers.

Paul Smith has been a Crime Scene Investigator with Leicestershire Constabulary since 1999. He received an MSc in Forensic Archaeology from the University of Bournemouth in 2004. A Registered Forensic Practitioner (RFP) and a qualified Crime Scene Manager and Coordinator, he is currently seconded to the University of Birmingham to assist in developing methods to support shared analysis for the forensic investigator. He is also working towards a PhD in Wearable Technological Applications for Crime Scene Investigation.

James Cross graduated from the University of Birmingham with an MEng in Electrical Engineering, and went on to obtain a PhD in 2003. His main research interests lie in the development of wearable computing, imaging systems and hardware design.

John Hunter was appointed Professor of Archaeology and Ancient History at the University of Birmingham in 1996. Apart from following an extensive scheme of research excavation and survey in Scotland, he began to develop forensic archaeology in 1988, co-wrote the first textbook on the subject and has routinely lectured to police and forensic groups ever since. He founded the Forensic Search Advisory Group, is a Director of the Centre for International Forensic Assistance (CIFA), a lead assessor for the Council for the Registration of Forensic Practitioners (CRFP), and member of CRFP's Incident Investigation panel. He works actively on police and war crime missions throughout the world and has recently returned from the Balkans as a consultant to the International Commission on Missing Persons (ICMP).

Richard McMaster received a BSc (Hons) in Psychology and then an MSc in Occupational Psychology from the University of Sheffield in 1999 and 2002. He joined the University of Birmingham in 2004 as a Research Associate and is also undertaking a part-time PhD in Distributed Cognition in the Emergency Services.

Christophe Heintz is working at the Institut Jean Nicod on the relations between cognitive and social factors of knowledge production. His research is focused on methodological questions in science studies, on the study of mathematical and anthropological knowledge, and on the role of the Internet in knowledge production. He has been trained in mathematics and philosophy at the Universities of Paris and Cambridge, and is currently completing his Ph.D. at the Ecole de Hautes Etudes en Sciences Sociales (Paris).

Dimitri Kanevsky is a Research Staff Member at IBM Thomas J. Watson Research Center. He holds 90 patents and has a “Master Inventor” status. He is a finalist of the 2005 Word Technology Award for IT-Software. His conversational biometrics based security patent was recognized by MIT, Technology Review, as one of five most influential patents for 2003.

Sara Basson currently works in IBM Research, where she is developing accessibility opportunities for all users of technology, irrespective of their abilities or disabilities. She holds an M.B.A. from NYU, and a Ph.D. in Speech and Hearing Sciences from The Graduate Center of the City University of New York. She serves on the Board of Directors of AVIOS.

Alexander Faisman works at the IBM Thomas J. Watson Research Center. He developed Vi-aScribe software, which is the hub of the Liberated Learning project, a consortium of 10 universities worldwide. Alexander received a Masters of Science in Mathematics in 1985 from Tashkent University. He is also the author of many scientific papers presented in international conferences as well as in publications, and he holds numerous granted and pending patents.

Leonid Rachevsky is a software systems analyst in the “Human Language Technologies” group at the IBM T.J. Watson Research Center. He obtained his MSc in Mathematics at Kazan State University and his PhD in Technical Science (Applied Mathematics) at the Kazan Institute of Chemical Engineering. He has worked extensively as a software engineer and senior software analyst in Israel, Canada, and the United States.

Sarah Conrod is Lead Researcher of the Baddeck Showcase and Executive Director of the Alexander Graham Bell Institute at Cape Breton University. Her work centers on developing and implementing projects that blend technology and accessibility. She also sits on the Board of Directors for the Cape Breton Branch of the Canadian National Institute for the Blind. Sarah holds a Bachelor of Science, a Certificate in Public Administration, and a Masters of Education in IT.

Alex Zlatsin is a senior software engineer at IBM T.J. Watson Research Center. His latest projects include the development of the managed client systems for AIX and Linux and Grid computing, in particular the deployment of a Speech Recognition Engine as the voice grid application. He holds over 30 patents and is an IBM master inventor. He holds the MS degree in Applied Mathematics from Moscow University of Railway Transport.

Barbara Dragsted has a PhD in translation from the Faculty of Languages, Communication and Culture at the Copenhagen Business School, where she is currently employed as an assistant professor teaching translation and English grammar. She is involved in different research projects under the Center for Research and Innovation in Translation and Translation Technology (CRITT). Her research interests include cognitive processes in translation, translation technology, terminology and knowledge sharing, and the study of translation workflows in multilingual communication

Table of contents

About the Authors	ix
-------------------	----

Introduction

Offloading cognition onto cognitive technology <i>Itiel E. Dror and Stevan Harnad</i>	1
--	---

Articles

A framework for thinking about distributed cognition <i>Pierre Poirier and Guillaume Chicoisne</i>	25
Distributed cognition: Domains and dimensions <i>John Sutton</i>	45
Distributed cognition: A methodological note <i>David Kirsh</i>	57
Radical changes in cognitive process due to technology: A jaundiced view <i>Arthur M. Glenberg</i>	71
The grounding and sharing of symbols <i>Angelo Cangelosi</i>	83
Collaborative tagging as distributed cognition <i>Luc Steels</i>	93
Thinking in groups <i>Todd M. Gureckis and Robert L. Goldstone</i>	99
Distributed learning and mutual adaptation <i>Daniel L. Schwartz and Taylor Martin</i>	117
Distributed cognition, representation, and affordance <i>Jiajie Zhang and Vimla L. Patel</i>	137

Categorization and technology innovation <i>Jeffrey M. Stibel</i>	145
Crime scene investigation as distributed cognition <i>Chris Baber, Paul Smith, James Cross, John Hunter, and Richard McMaster</i>	159
Web Search engines and distributed assessment systems <i>Christophe Heintz</i>	185
Speech transformation solutions <i>Dimitri Kanevsky, Sara Basson, Alexander Faisman, Leonid Rachevsky, Alex Zlatsin, and Sarah Conrod</i>	207
Computer-aided translation as a distributed cognitive task <i>Barbara Dragsted</i>	237
Index	257

Offloading cognition onto cognitive technology

Itiel E. Dror & Stevan Harnad

University of Southampton / Université du Québec à Montréal

“Cognizing” (e.g., thinking, understanding, and knowing) is a *mental state*. Systems without mental states, such as cognitive technology, can sometimes contribute to human cognition, but that does not make them *cognizers*. Cognizers can offload some of their cognitive functions onto cognitive technology, thereby extending their performance capacity beyond the limits of their own brain power. Language itself is a form of cognitive technology that allows cognizers to offload some of their cognitive functions onto the brains of other cognizers. Language also extends cognizers’ individual and joint performance powers, distributing the load through interactive and collaborative cognition. Reading, writing, print, telecommunications and computing further extend cognizers’ capacities. And now the web, with its network of cognizers, digital databases and software agents, all accessible anytime, anywhere, has become our “Cognitive Commons,” in which distributed cognizers and cognitive technology can interoperate globally with a speed, scope and degree of interactivity inconceivable through local individual cognition alone. And as with language, the cognitive tool par excellence, such technological changes are not merely instrumental and quantitative: They can have profound effects on how we think and encode information, on how we communicate with one another, on our mental states, and on our very nature.

Introductory overview

With the development and wide use of cognitive technologies (Dror, 2007; Dascal & Dror, 2005), questions arise as to their effects on their human users and society, as well as on their own scope and limits: Can cognitive technologies (i) increase cognitive capacities and thus enhance human efficiency? (ii) affect how individuals and society go about achieving their goals? (iii) highlight and transform how we view ourselves and our goals? (iv) modify how we cognize and thus change our mental states and nature? (v) give rise to new forms of cognition (such

as distributed cognition) and mental states that are either distributed across or even embodied in cognitive technology?

These issues are examined as follows:

- (1) The notion of an “extended mind” – with mental states (i.e., felt states) “distributed” beyond the narrow bounds of the individual brain – is not only as improbable as the notion that the US government can have a distributed migraine headache, but arbitrary.
- (2) “Cognition” – if it is simply defined as the ability to do the kinds of things that cognizers like us can do, plus the underlying functional mechanisms for doing them – can be arbitrarily defined to be as wide or as narrow as we like.
- (3) Vagueness about the nature, locus and scope of cognizing leads to a dissociation of “cognitive states” from mental states. However, their co-occurrence had been our only basis for distinguishing cognitive performance capacity from other capacities and functionality (animate or inanimate, narrow or wide).
- (4) If cognitive states are indeed not mental states, it follows that “cognitive technology” is not just something *used* by cognizers, but a functional part of the cognitive states themselves, because the boundary between user and tool disappears, and cognitive states become merely instances of functional states in general.
- (5) We then do not need the terms “cognitive” and “distributed cognition” at all, and can just talk about relatively complex and wide or narrow functional states, leaving it a coincidence and mystery (at least at this stage) that every single case of what we used to call “cognitive” also happened to be mental.
- (6) A way to resolve this is to accept that only mental states are cognitive states, that cognition is only narrow, and that the only place cognition is “distributed” is within a single cognizer’s brain.
- (7) The only kind of “technology” that might really turn out to be *intrinsically cognitive*, rather than just being a tool used by cognizers, would be a robot that could pass the Turing Test (TT) – because such a TT-scale robot would almost certainly have mental states, and hence it would be a cognizer in its own right.
- (8) Whatever distributed activity was going on within the functional mechanism generating such a TT robot’s performance capacity would then indeed be a case of distributed cognition (exactly as the distributed activity within our own brains is distributed cognition) – even if not all the components of its generating mechanism were located inside the robot’s head.
- (9) The “cognitive technology” used by such a TT robot, however, would still not be part of its distributed cognitive (hence mental) state, just as it is not a part of ours.

- (10) Nor would a group of such TT robots, interacting and collaborating, be a case of distributed cognition; it would merely be a case of collaborative cognition among individual (narrow) TT-robot cognizers, just as it is in the case of a group of collaborating human cognizers.
- (11) Cognitive technology does, however, extend the scope and power of cognition, exactly as sensory and motor technology extends the scope and power of the bodily senses and movement.
- (12) Just as we can *see* further with telescopes, move faster with cars, and do more with laser microsurgery than we can do with just our unaided hands and heads, so we can *think* faster and further, and do more, with language, books, calculators, computers, the web, algorithms, software agents, plus whatever is in the heads of other cognizers.
- (13) Both sensorimotor technology and cognitive technology extend our bodies' and brains' performance capacities as well as giving us the *feeling* of being able to do more than just our bodies and brains alone can do.
- (14) Sensorimotor and cognitive technology can thus generate a perceptual change, rather like virtual reality (VR), making us feel a difference in our body image and causal power (perhaps not unlike what the physical metamorphosis from caterpillar to butterfly might feel like, as one sensed one's newfound somatic capacity to fly).
- (15) This change in perceived body image is indeed a change in mental state; but although its distal inputs and outputs certainly extend wider than the body (as all sensory inputs and all motor outputs do), the functional mechanism of that altered mental state is still just proximal -- skin and in -- exactly as when it is induced by VR technology.
- (16) Hence, although sensorimotor and cognitive technology can undeniably extend our bodies' sensorimotor and cognitive performance powers in the outside world, only their sensorimotor *input and output contact points* with our bodies are part of our cognitive (= mental) state, not the parts that extend beyond.
- (17) Perhaps it could be otherwise too, as in the case of a hypothetical TT-robot whose generating mechanism is indeed partly located outside its body: Maybe parts of our brain could be removed and still functionally integrated with the rest wirelessly, through telemetry or some other action at a distance: But that would just be a widened, spatially distributed *body*.
- (18) The resultant distributed cognitive state would still not be the same thing as considering a telescope, car, library or calculator as parts of a distributed cognitive state (for either a human or a TT robot): Those would still just be parts of the sensorimotor I/O to and from the cognizer's body.

- (19) We are not aware of the generating mechanism underlying our cognitive capacity, however, only of its outcome: Hence retrieving a word from memory or retrieving a word via a Google search feels much the same to us.
- (20) Does the fact that cognizing is a conscious mental state, yet we are unconscious of its underlying functional mechanism, mean that the underlying functional mechanism could include Google, Wikipedia, software agents and other human cognizers' heads after all?
- (21) The worldwide web, a distributed network of cognizers, digital databases and software agents, has become our "Cognitive Commons," in which cognizers and cognitive technology can share cognizing anytime and anywhere, and interact globally with a speed, scope and degree of interactivity that yield distributed cognizing with performance powers inconceivable within the scope of individual cognition.
- (22) Such changes go beyond mere quantitative increase in efficiency and performance power. As we increase our use and reliance on cognitive technologies, they effect and modify how we cognize, how we do things and what we do. Just as motor technology extended our physical ability and modified our physical life, cognitive technology extends our cognitive ability and modifies our mental life.

Part I: What distributed cognition is not

Meaning: Narrow and wide. Philosophers, in wrestling with the problem of meaning ("Is meaning in the head or is it in the world?") have sometimes resorted to saying that there are *two* kinds of meaning, "narrow" and "wide," with the former located between the ears and the latter distributed across the entire universe – both the Newtonian universe of distant stars and the Platonic universe of the eternal truths of logic and mathematics. The wide meaning of "apple," for example, includes not only whatever it is that I may have in mind when I think of or say "apple," but also what apples really are, out there in the world.¹

That, however, is all metaphysics, and concerns the existence and "reality" of some elusive entity called "meaning." The mission of cognitive science is more modest: Humans and other organisms have certain functional capacities, including metabolism, reproduction, and locomotion. It is clear that each of these capacities is "narrow," even though it sometimes involves a local interaction between the organism and part of the world around it (be that other organisms or the inanimate world).