BOOK REVIEW

Barbara Gorayska and Jacob L. Mey (eds), *Cognition and Technology: Co-existence, Convergence and Co-Evolution*. Amsterdam: John Benjamins, 2004, vi + 369 pages. ISBN 90–272–3224–5 (Eur.) 1–58811–544–5 (US)

Reviewed by Iris van Rooij (Eindhoven University of Technology)

This edited book presents a collection of papers in *Cognitive Technology* (CT), a research discipline whose existence I had been unaware of before reading this book. The birth of CT, the editors tell us, can be traced to a lunch meeting in 1993, in old Hong Kong, where Barbara Gorayska and Jacob L. Mey conceived of the idea of "CT as a new discipline, combining findings from computer science, philosophy, psychology and pragmatics" (p. 2). Drawing on this interdisciplinary foundation, the goal of CT is to study the pragmatics and ethics of technologies that support, extend and emulate human cognition — i.e., *cognitive technologies*. The study of CT is contrasted with *Human Factors* (HF). While HF studies the cognitive make-up of users with the aim of designing technologies specifically adapted to human information-processing abilities and limitations, CT sets out to examine "the semantics, syntax, and pragmatics of information itself, and how its form of delivery might impact on the cognitive make-up of users" (p. 4).

Gorayska and Mey co-edited two earlier volumes in CT (Gorayska and Mey 1996; Marsh, Gorayska, and Mey 1999). The present volume, their third, reports on the convergence in the field of CT over the past decade. The editors are quick to point out "that the field of CT, despite all convergence, has evolved in ways that are not always easy to combine. (…) But in this field, as in all other areas of scientific research, the important thing to do is to let things develop, 'co-evolve' on their own, so as to obtain the maximum 'convergence' in the midst of 'coexisting', sometimes clashing views" (p. 1). A remark that is intended, I assume, to explain the title.

The book opens with an introductory chapter by the editors, in which they spell out the CT agenda. The list is long and includes many practical and ethical concerns, including the following:

– Should new technologies adapt to humans or should humans adapt to new technologies? If we are natural-born tool-users, or as Clark (2004) calls us

Pragmatics & Cognition 13:3 (2005), **647–655**.
ISSN 0929–0907 / E-ISSN 1569–9943

'natural-born cyborgs', would it not be more efficient (contrary to received view in HF) to let humans do the adapting?
– What responsibilities and dangers do new technologies bring? Do we really want or need technology, or "are our technological needs the result of someone else's wants?" (p. 13). Besides mind-extending effects, can cognitive technologies also have mind-numbing effects or other undesirable psychological consequences?

The list also includes more conceptual and philosophical questions, such as:

– If computers, robots and language can be viewed as cognitive technologies, can "the same be said about care-givers who help newborn babies acquire cognitive skills by simply interacting with them" (p. 12), and of the special-purpose cognitive processes that we develop during our lifetimes (e.g., the brain processes involved in reading, writing, imagery, and mental arithmetic)?
– Is there, in any real sense, a boundary between our minds and the technologies that augment them? Do cognitive technologies perhaps not only support and extend human cognitive abilities but also alter their very nature and in effect extend the material basis of our minds beyond 'the biological skin-bag' (Clark 2004)?

The remainder of the book is divided into three parts. Part I contains 6 papers addressing **Theoretical Issues**, Part II contains 5 papers presenting **Applications,** and Part III presents two papers as **Coda** to the entire collection. Several of the contributions are reprints of papers published in Vol. 1 of the *International Journal of Cognition and Technology*. The papers are by different authors and vary greatly in both length and focus. As is to be expected, they also vary in quality.

Part I starts with a provocative paper by Andy Clark, presenting a "brief and impressionistic sketch" (p. 30) of the thesis that non-biological cognitive technologies not only extend our ability to cognize, but that they are, in a very real sense, part of what constitutes human cognition (the interested reader may wish to consult his new popular book *Natural Born Cyborgs* (2004) for more details). While reading Clark's paper one gets a sense of *panpsychism*, the deep philosophical implications of which seem to be ignored by the other contributions to this volume. If cognitive tools do not extend minds, but are constitutive of them, in what sense can we still speak of 'tools' and consider ourselves as 'tool-users'? It is clear that Clark conceives of such tools as "only tools in the thin and ultimately paradoxical sense in which my own unconsciously operating neural structures (…) are tools. I do not really 'use' my brain. There is no user quite so ephemeral" (p. 26). It is unclear, however, to what extent other researchers in CT are comfortable with Clark's view (Is damage to my personal computer equivalent to a lobotomy?) and its implications for the status of CT as a research discipline (Is CT a subdiscipline of Cognitive Psychology?).

The primary concern of Marcelo Dascal "is to show how several aspects of language and language use can fruitfully be conceptualized as cognitive technologies" (p. 40), but his paper can also be read as an introduction to CT in general. It contains one of the clearest and most explicit definitions of 'cognitive technology' that I could find in the book (pp. 36–37; but see also El Ashegh and Lindsay, pp. 175–176) and it presents a typology of cognitive technologies, distinguishing (1) 'strong' versus 'weak', (2) 'integral' versus 'partial', (3) 'complete' versus 'incomplete', and (4) 'constitutive' and 'non-constitutive' cognitive technologies (pp. 40–43). Subsequently, Dascal considers three ways in which language affects cognition: (1) as part of the *cognitive environment* in which much of our thinking takes place; (2) as a *cognitive resource*, for example, when used to gather, organize, store and retrieve information; and (3) as a *cognitive tool*, for example, when used to define concepts and create notational systems. Of particular interest for cognitive scientists, like myself, is his discussion of negative effects arising from the uncritical use of language as cognitive technology, leading to 'cognitive mistakes', 'category errors', and the like.

Compared to the concise and well-written contributions of Clark and Dascal, the unwieldy paper by Lindsay and Gorayska is somewhat of a disappointment. It starts by saying that Artificial Intelligence (AI) has overlooked the importance of understanding the human ability to judge what is relevant. Did AI not recognize this problem long ago, leading McCarthy and Hayes (1969) to formulate the 'frame problem'? Oddly enough, Lindsay and Gorayska make no reference to this classic work or any other, related work in AI (e.g., Haselager, 1997; Pylyshyn, 1989; cf. Lueg in Part II). Then, "by according to relevance the central role that it should have in explaining cognition", Lindsay and Gorayska promise "to clear up a considerable number of issues and problems that presently seem mysterious in connection with problem-solving, ethics, symbol-connection hybridism, and the motivation-action nexus" (p. 63). Their treatment does little, however, to clear things up for me. Consider, for example, their (informal) proposal for a connectionist learning mechanism for inferring relevance (pp. 68–78). How does this mechanism circumvent the problems, such as computational costliness (cf. Chandrasekharan, pp. 153–172) and context dependency (cf. Lueg, pp. 225–239), known to plague the problem of determining relevance?

The fourth contribution, by Rolf Pfeifer, is again a short and informative paper. It discusses how robots can be used as tools for the study of embodied cognition (*a la* Rodney Brooks). Since the study of cognition is itself a form of cognition, robots are cognitive tools for cognitive scientists (in a similar vein, AI programs, diagrams, formal languages and mathematical models are all cognitive tools for cognitive scientists). To illustrate the synthetic methodology of robotics (i.e., the ideology of "understanding by building" (p. 110) that robotics shares with traditional AI), Pfeifer discusses three case studies — 'the Swiss robots', 'the insect eye' and 'the dynamic walker' — as well as work in the field of developmental robotics. Even though the ideas presented in this paper are not new (nor does Pfeifer claim they are), the paper certainly contributes as a nice, short and informal introduction to robotics for the CT researcher.

The paper by Kerstin Dautenhahn on 'narratives' seems out of place in this volume on CT. This is not because narratives cannot be usefully conceived of as cognitive technology (so much is clear after reading Dascal), but because Dautenhahn does not explicitly take this CT perspective on narrative. Instead her prime interest seems to be in the cognitive ethology of narrative. After detailing hypotheses about the origins, form and function of narratives, the connection with CT is finally made in the concluding section. The conclusion, however, seems dissonant with the very idea of CT. From Dunbar's (1993) observations of the relationship between primate brain size and social group size, Dautenhahn infers that humans cannot maintain social networks of more than 150 people unless they somehow evolve larger neo-cortices. But what about the potential of non-biological mind-extending and mind-altering technologies as envisioned by CT researchers like Gorayska, Mey and Clark?

The last paper in the **Theoretical Issues** section, by Sanjay Chandrasekharan, explains how the World Wide Web (WWW) not only serves as a 'knowledge repository' but also functions as an 'action-enabling space', for example, for buying flowers, sending cards, booking rooms or tickets, making money transactions, etc. This second conception of the WWW leads Chandrasekharan to propose that the *Semantic Web* effort (see http://infomesh.net/2001/swintro/) "should focus more on the possible actions humans and artificial agents can execute on the Web" (p. 153). Inspired by distributed cognition (*a la* Kirsch and Hutchins), Chandrasekharan advocates an Active Design approach (which he contrasts with the Passive Design approach of both Good-Old Fashioned and Brooksian AI) and describes how the Semantic Web effort may realize and benefit from an affordance-model of formal ontologies. I believe the paper will be of great interest for anyone working in Web design, but also cognitive scientists may find ideas of interest here.

The first paper in Part II is by El Asheg and Lindsay. Following Meenam and Lindsay (2002), they propose what seems to be the mirror thesis of Clark. While for Clark the mind extends outside the brain and body into technological artifacts, El Asheg and Lindsay take the word 'technology' to include our brains and bodies whenever they are used to reach goals. This they refer to as 'natural technology' (Is this a form of 'pantechnologism'?). Although El Asheg and Lindsay exclude skills like eating and walking, the word 'natural technology' includes much of what cognitive psychologists consider their domain (and possibly all of it). They go on to study one such natural technology: the body-image generator (BIG) module. The BIG module is a hypothetical "internal cognitive mechanism that underlies an individual's view of his/her physical appearance and provides the images towards which evaluations of the physical self are directed" (p. 181). El Asheg and Lindsay report on an empirical study meant to test two hypotheses: (1) that the BIG module exists, and (2) that it is a natural technology. The study is quite interesting, but with respect to these hypotheses the results are far from convincing. First of all, the results fail to show the alleged "double dissociation" (p. 216) taken to support (1), for the results are *correlational* and the authors seem unaware of the controversy surrounding the use of double dissociation methodology for inferring the existence of mental modules (e.g., Dunn and Kirsner, 2003). Second, it is unclear how their (or any conceivable) study instantiates an empirical test of (2).

Christopher Lueg considers the problem of designing 'context-aware artifacts'. Imagine, for example, a mobile phone that senses the appropriate level of intrusiveness for the current situation (e.g., loud ringing, buzzing, vibrating, or no signal at all), or a room that senses the social interaction between its users and accommodates to their needs (e.g., change the light intensity, put on music, adjust the temperature). Leug draws a pertinent analogy between this design problem and the general AI problem. Even though "researchers in context-aware artifacts and ubiquitous computing do not consider their work as AI work", Lueg convincingly argues that they too "run into problems, such as the frame problem or the problem of reliably predicting human behavior, that have been haunting AI researchers for decades" (p. 237). Knowledge of such fundamental difficulties not only aids context-aware artifact design (e.g., by preventing attempts at doing the impossible and encouraging more pragmatic and modest design goals); it also means that research in context-aware artifacts, like robotics, "may help contribute to gaining a better understanding of the complexity of human behavior and human social life" (p. 236).

The paper by Satinder P. Gill reports on an empirical study comparing two drawing surfaces: the white board and one of its electronic counterparts, the SMARTBoard. Unlike the whiteboard, the SMARTBoard allows only one person at a time to draw on it. Gill is interested in the effects this has on the behavior of two people (she studied landscape architects) working at the surfaces. Her transcription data illustrate how participants get frustrated when confronted with the inability to coordinate their actions at the SMARTBoard. To gain control over the surface the participants engage in disruptive communicative acts, a pattern of behavior that contrasts with the coordinated movements of participants working at the whiteboard. Gill connects her work with Polanyi's (1966) notion of 'tacit knowledge'.

Bowman, Hinkley, Barnes and Lindsay return to the topic of natural technology. Their candidate is the mechanism underlying autism. While current theories of autism assume the condition arises from an innate cognitive deficit, the authors believe the source is an emotional hypersensitivity. In their view the cognitive deficits characteristic of autism are natural technologies designed to prevent emotionally laden information from entering the system. Bowman *et al.* perform two experiments to show that children actively avoid looking at people's eyes so as to shut out their emotional content. Here I comment on the second experiment. In this experiment autistic children and a control group were presented with pictures of faces. On each face a heart shaped target was superimposed on either the mouth, the cheek, or an eye. The child's task was to identify the target's location as fast as possible. No effect of location was found for the control group. The autistic children were overall slower, but relatively faster to detect the target on the mouth than on either the cheek or eye (the latter two showing equal response times). The authors see support in these results for their hypothesis. But what was the purpose of including both the mouth *and* cheek condition, if not to show that the aversion was *specific* to the emotionally laden eyes? As they stand, the results show a 'mouth bias' in autistic children. The reason awaits further experimentation.

The **Applications** section closes with Jirotka and Luff's exposition of a semi-formal language for modeling sequential activities, called *Communicating Sequential Processes* (CSP). CSP is characterized by its closeness to natural language and its ability to model parallel and non-deterministic processes. The authors explain how their work has a three-fold connection with CT: (1) The modeling language CSP can serve as a cognitive resource for designers of new technologies (cf. Dascal); (2) CSP supports the descriptive study of sequential 'users' activities in the work place (cf. Gill); and (3) CSP modeling practice may contribute to fundamental research into the nature of human social conduct (cf. Lueg; Pfeifer). Jirotka and Luff present an illustration of CSP in the context of a 'financial dealing room'. Although the illustration is insightful and suggestive of a general utility of CSP as a cognitive tool, it does not provide a full-scale specification of CSP. Researchers interested in applying CSP in other domains will probably have to consult cited sources for more details.

The **Coda** (Part III) starts with a paper by Syed M. Ali. He considers the Schizophrenia Problem:

> As a partner, the computer tends to resemble a schizophrenic suffering from severe 'intrapsychic ataxia' — the psychiatric term for a radical separation of cognition from emotion. (…) Interacting in accordance with the requirements of its programs, the computer, like the schizophrenic forces us to empathize one-sidedly with it and communicate with it on its own terms. And the suspicion arises that the better we can do this the more like it we become (p. 333, quoted by Ali from Janney, 1997, p. 1).

According to Ali a solution to the Schizophrenia Problem presupposes a solution to Chalmers' (1996) Hard Problem of consciousness (Ali then goes on to argue that neither Dreyfus nor Heidegger can solve the latter problem). I do not see why this should be so. If we can conceive of beings behaviorally equivalent to humans but without conscious experience (i.e., Chalmers' zombies), then why cannot computers remain "in the dark" and yet be emotionally adequate partners for humans? On the flip side, even if we were to solve Chalmers' Hard Problem, how would this help us design emotionally adept computers?

The last chapter by Will Fitzgerald is a reflective and inspiring paper — an appropriate coda indeed. Fitzgerald asks himself (and us) what AI could have been had Martin Luther King attended the first AI workshop in 1956. Would AI have defined itself as the descriptive and prescriptive study of intelligent (read: rational) thinking and action? Or would it have included other values and qualities in its conception of "human being", such as *justice*, *revenge*, *politics*, *dignity*, *violence*, *forgiveness* and *love*? "A really good AI model of *forgiveness*, for example, is, I suspect, no harder to create than a good AI model of *temporal reasoning*, and no easier as well" (p. 350; *emphasis in original*). Looking back now, as well as looking forward, can we perhaps have such a thing as "computational humanism" (p. 351)?

The present volume shows a great diversity of papers. While some authors discuss the conceptual and philosophical foundations of CT, others focus more on practical problems. Few contributors, however, explicitly address ethical issues on the CT agenda. According to the editors all papers were bound by a common concern: "How to make the most of technology without 'losing our soul'" (p. 2). In my view, only the paper by Fitzgerald really touches on this issue. Nonetheless, the volume gives a clear picture of the aims and scope of CT. I can recommend this book to anybody interested in Human/Technology Interaction and in CT in particular. For those with general (not necessarily CT-related) interests in Cognitive Science I can also recommend the papers by Clark, Dascal, Pfeifer, Chandrasekharan, Lueg, Jirotka and Luff, and Fitzgerald.

References

Chalmers, D. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.

Clark, A. 2004. *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Oxford: Oxford University Press.

Dunbar, R.I.M. 1993. "Coevolution of neocortical size, group size and language in humans". *Behavioral and Brain Sciences* 16: 681–735.

Dunn, J.C. and Kirsner, K. 2003. "What can we infer from double dissociations?" *Cortex* 39(1): 1–7.

Gorayska, B. and Mey, J.L. (eds) 1996. *Cognitive Technology: In Search of a Humane Interface*. Amsterdam: North Holland.

Haselager, W.F.G. 1997. *Cognitive Science and Folk Psychology: The Right Frame of Mind*. London: Sage.

Janney, R.W. 1997. "The prosthesis as partner: Pragmatics and the human-computer interface". In J.P. Marsh, C.L. Nehaniv, and B. Gorayska (eds), *Proceedings of the Second International Cognitive Technology Conference: Humanizing the Information Age*. Los Alamitos, CA: IEEE Computer Society Press, 1–6.

Marsh, J., Gorayska, B., and Mey, J.L. (eds). 1999. *Humane Interfaces: Questions of Methods and Practice in Cognitive Technology.* Amsterdam: North Holland.

McCarthy, J. and Hayes, P.J. 1969. "Some philosophical problems from the standpoint of artificial intelligence". In B. Meltzer and D. Michie (eds), *Machine Intelligence, Vol. 4*. Edinburgh: Edinburgh University Press, 463–502.

Meenan, S. and Lindsay, R.O. 2002. "Planning and the neurotechnology of social behavior". *International Journal of Cognition and Technology* 1(2): 233–274.

Polanyi, M. 1966. *The Tacit Dimension*. Garden City, NY: Doubleday.

Pylyshyn, Z. (ed). 1987. *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*. Norwood, NJ: Ablex.

Reviewer's address

Iris van Rooij
Eindhoven University of Technology
Human Technology Interaction
P.O. Box 513
5600 MB Eindhoven
The Netherlands

i.v.rooij@tm.tue.nl
http://www.ipo.tue.nl/homepages/ivrooij/

About the reviewer

Iris van Rooij (M.A. in experimental psychology, Radboud University Nijmegen, The Netherlands; Ph.D. in cognitive psychology, University of Victoria, Canada). Her dissertation investigated computational tractability as a guiding principle in cognitive theory formation. She currently holds a position as postdoctoral researcher in Human/Technology Interaction at the Eindhoven University of Technology in the Netherlands. She has published in psychology, cognitive science and computer science, and has special interests in the philosophy of mind and the philosophy of science.