

Review of *Knowledge Engineering and Management*

Eleni Stroulia

This book describes the COMMONKADS methodology, the product of a family of knowledge and analysis design support (KADS) projects spanning the past two decades. The last of the book's authors has been involved in this effort since the beginning of 1983. Thus, the book is particularly interesting to those who have been following their work.

The original motivation for this work was to develop a tool kit of models and methods for expert system developers. The approach was based on two key observations: The first observation was that expert systems had to simulate, to some extent, human behavior to be understandable to their users and, in cases where they replaced experts, to integrate smoothly with the existing process in which they were introduced. The major bottleneck in the effort to make expert systems like the experts they were meant to simulate was *knowledge engineering*, that is, extracting domain-specific knowledge from the "expert's head" and representing it in a formal representation scheme that the system could interpret. The challenge then became to facilitate the knowledge-engineering process by designing a common language between the knowledge engineer and the domain expert. This challenge gave rise to the issue of the vocabulary of this common language. It could not be the domain knowledge itself because domains vary a lot. The second key observation that several research groups (Steels 1990; McDermott 1988; Wielinga and Breuker 1986; Clancey 1985; Chandrasekaran 1983) made independently yet almost simultaneously was that tasks were common across domains. Indeed, instances of the diagnosis task can be found in

medicine, car troubleshooting, software debugging, and so on. Identifying generic, domain-independent tasks, formalizing task representation, elucidating the role of the task in eliciting domain-specific knowledge, and standardizing the design and development of expert systems then became the major research problems of the field.

The COMMONKADS methodology, as it has evolved through the years, consists of three main components: First, it proposes a modeling framework for representing and analyzing tasks, the domain knowledge used to perform them, the agents who perform them, and the communication aspects of tasks performed by different agents. Second, it offers a set of task templates

Knowledge Engineering and Management: The COMMONKADS Methodology, Guus Schreiber, Hans Akkermans, Anjo Anjewierden, Robert de Hoog, Nigel Shadbolt, Walter Van de Velde, and Bob Wielinga, The MIT Press, 2000, 455 pp., ISBN 0-262-19300-0.

that can be used as initial approximations of specific tasks in a particular domain of interest. Third, it prescribes a set of activities for developing software systems in the overall context of an organization, its current processes, and its needs. There are three knowledge-modeling activities: (1) knowledge identification, (2) knowledge specification, and (3) knowledge refinement. During the first activity, the knowledge analyst is expected to become sufficiently familiar with the

domain that he/she can identify the potentially reusable task- and domain-related templates. Knowledge specification, includes the task decomposition and the specification of the domain information types and knowledge bases. The task decomposition can be guided by selecting to reuse some of the previously identified task templates. Finally, during knowledge refinement, the models are validated through simulation on paper or with prototyping, and the knowledge bases are refined. Depending on how familiar the analyst is with the domain, these activities might have to be performed repeatedly, and subsequent activities might provide feedback for corrections or extensions to the products of earlier ones.

The book is intended for practitioners in knowledge management. The terms *intellectual capital* and *knowledge assets* have become commonplace. Increasingly more organizations recognize that there is value in the knowledge implicitly embedded in their current business processes, and the question is how to tap into this hidden resource and how to manage it so its return can be maximized. Practitioners in the field will find that the COMMONKADS methodology can provide an initial answer to this question. It deals in a comprehensive manner with the problem of identifying knowledge assets, whether they lie in the artifacts that the current processes manipulate, the agents performing them, or the communication between them. It also offers support for restructuring these processes using the standard task templates. In this way, the book can guide the process of making implied knowledge explicit and embodying it in well-architected new or reengineered systems. The problem of evaluating the impact of such changes and making further decisions about the evolution of the process is always more complex, but analysis and rationalization of the current status is a big step in the right direction.

More generally, practitioners of traditional software engineering might also find this book of interest, especially because the iterative COMMONKADS process resembles the spiral model for the software life cycle; in

Joint Position Announcement:

The Learning Research and Development Center (LRDC) in conjunction with departments in the Faculty of Arts and Sciences and the School of Education is recruiting for five (5) tenure stream faculty positions beginning with the Fall Term 2001, subject to administrative approval. At least two of these are at open rank and three are at the Assistant Professorial level. Teaching responsibilities and tenure options will reside in the appropriate school or department such as Education, Psychology, and Computer Science. All appointments will be joint appointments between the appropriate department and LRDC.

Founded in 1965, LRDC is a research institution at the University of Pittsburgh whose mission is the study of learning and instruction broadly conceived. Our interests vary widely with respect to the content and method of research: from brain imaging to reading instruction, from mathematics learning to museum learning, from computer tutoring to models of learning. Our research community also varies with respect to the purposes of their research: from theoretical development to the improvement of practice. LRDC exemplifies the idea that a strong multidisciplinary research base can provide a foundation for effective practices in a variety of teaching and learning settings and that the intricacies and complexities of real world situations provide important stimuli for the development of significant theory and research.

We are recruiting individuals whose research addresses issues of learning and/or instruction in three broad areas:

- Language and literacy, including reading, writing, and second language learning,
- Learning, motivation, and instruction in school subject matter areas, including, among others, mathematics, science, and history.
- Artificial intelligence and technology, including computer tutoring, natural language processing, and collaboration.

Successful candidates will have a demonstrated background of excellence in research and will be able to show ways in which their research relates to broader and more complex social arenas. LRDC is an institution with representatives from many disciplinary traditions, including psychology, education, computer science, and linguistics; successful candidates will have an interest in participating in the multi-disciplinary research goals of LRDC.

To apply send a letter of application, current resume, and representative papers. Also, please have three letters of recommendation sent (individuals currently holding faculty positions may instead provide the names of three individuals prepared to write letters of recommendation). We expect to begin reviewing applications and inviting candidates in for visits in early fall. To ensure full consideration, application should be received by November 1, although we will continue to review applications through February 15. The University of Pittsburgh is an Affirmative Action, Equal Opportunity Employer. Women and members of minority groups that are under-represented in academia are especially encouraged to apply.

Send materials to: Dr. Gaea Leinhardt, LRDC Joint Search, 3939 O'Hara Street, Pittsburgh, PA 15260

Note: These positions are also listed in departmental advertisements announcing positions recruited jointly with LRDC. A reply to either announcement will receive equal attention.

fact, chapter 15 on project management elaborates on the parallels between them. The COMMONKADS activities closely correspond with requirements analysis, requirements specification, and design. The other bridge between the book and the object-oriented software-engineering audience is the adoption of a subset of UML (unified modeling language) for representing CommonKADS models because the COMMONKADS methodology does not have its own notation.

The book could also be used as a text for a course in information systems. It is well organized, and the authors identify road maps for audiences with different main interests, such as knowledge analyst, knowledge manager, developer, and project manager. Every chapter starts with a summary of the key points and concludes with a section of pointers to related work. The book has quite a few case studies to put the proposed framework in perspective. There are short case studies dispersed in several of the chapters, each discussed in terms of

the points introduced by the chapter they belong in. In addition, chapter 10 contains a case study illustrating the use of the COMMONKADS analysis models and methods. Finally, there is an associated web address that contains exercises, demo systems, tools, and course slides.¹

The book represents a whole area of research that has matured, relevant not only to today's development issues but also highly applicable. Whether as a textbook or as a project-planning and management method, the book will shed useful insights on the process of building knowledge-intensive systems.

Note

1. www.commonkads.uva.nl.

References

- Chandrasekaran, B. 1983. Toward a Taxonomy of Problem-Solving Types. *AI Magazine* 4(1): 9–17.
- Clancey, W. J. 1985. Heuristic Classification. *Artificial Intelligence* 27:289–350.
- McDermott, J. 1988. Preliminary Steps

toward a Taxonomy of Problem-Solving Methods. In *Automating Knowledge Acquisition for Expert Systems*, ed. S. Marcus, 225–266. Dordrecht, The Netherlands: Kluwer Academic.

Steels, L. 1990. Components of Expertise. *AI Magazine* 11(2): 28–48.

Wielinga, B. J., and Breuker, J. A. 1986. Models of Expertise. Paper presented at the Seventh European Conference on Artificial Intelligence, 20–25 July, Brighton, United Kingdom.



Eleni Stroulia is a professor in the Department of Computer Science at the University of Edmonton, Alberta, Canada. Her e-mail address is stroulia@cs.ualberta.ca.