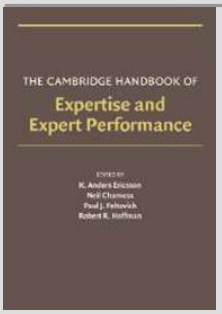

The Cambridge Handbook of Expertise and Expert Performance

Edited by K. Anders Ericsson, Neil Charness, Paul J. Feltovich, and Robert R. Hoffman

Cambridge University Press, 2006
$65.00 902 pages

Reviewed by Paul Harmon

There has been quite a bit of discussion on BPTrends, and elsewhere, about the changing nature of business processes. Various authors have proposed that process analysis is different today, because analysts are increasingly being asked to focus on tasks that involve knowledge workers or experts. Some have suggested that methodologies designed to analyze manufacturing procedures do not work as well when you are trying to analyze complex cognitive tasks and have proposed new methodologies with new modeling notations.

In my July 18th BPTrends Advisor on Alternative Approaches to Process Analysis, I suggested we could conceptualize the entire field of methods using two continua. On the horizontal axis, I suggested we discriminate between the scope of our analysis efforts. Architecture projects that aim at defining a value chain consider processes at such a level of abstraction that we don't consider the details of specific processes, no matter how complex. The analysis of mid-level processes and activities, on the other hand, can be very sensitive to human task complexity.

I qualified my remarks by describing this history of expert systems development in the Eighties and early Nineties. I noted that expert systems developers made an effort to capture expert knowledge in the Eighties. Most of these systems failed for a very simple reason. Experts are individuals who are constantly learning and changing their mental models and heuristics as knowledge evolves and experience with specific new problems is accumulated. It turns out that, even if you could build an expert system that captured the knowledge of an expert, as of June 2006, you would have to revise it in July. It is simply easier to maintain the human expert than to try to create and maintain a software system. Put a different way, there are jobs you try to train people to do, and there are jobs you hire people to do, trusting that their previous experience will equip them for the tasks they will face. You don't try to train creative people; you hire them. Which is just a way of admitting we don't know much about analyzing those tasks – we simply know what successful output looks like. Thus, we don't need to worry about analysis models or notation for mid to lower level processes and tasks performed by senior managers, new product designers or movie directors. We analyze the high level tasks to get benchmarks and success criteria, hire the best people we can, and then get out of their way.

I went on to suggest that the challenges clearly lie in the analysis of mid-level and specific tasks of medium complexity. There are, of course, lots of companies working on the analysis and modeling of mid-level processes that lie in this range. I know many people who have worked in this area, myself included, who have concluded that for most of these tasks, a workflow notation, especially one supplemented with swimlanes, notes, and business rules, works just fine. On those occasions when you encounter a task that involves people working together and exchanging emails, you usually don't analyze the task below the level at which a box indicating that the work is going on occurs. Keep in mind, the goal isn't to model everything, but only to create such models as are needed by business people or IT folks to accomplish their daily tasks. We analyze to see if we can determine how to perform the process better, or to specify what

needs to be automated. We usually don't need to try to analyze how a loan negotiator carries out each step of the negotiation.

I used Figure 1 to illustrate the business process analysis space, and used yellow to highlight the areas in which innovative modeling techniques might be appropriate. In a review last month of *Working Minds: A Practitioner's Guide to Cognitive Task Analysis*, I also suggested that we could divide the complexity of tasks along the same task complexity continuum shown in Figure 1. In that review, however, I added that the simple procedures were the domain of ordinary work, that the tasks of mid-level complexity was the domain of knowledge workers and that the tasks requiring creativity and true expertise was the domain of experts. Following that line of reasoning, I might re-label the blue workflow area as the area of conventional task analysis, and the yellow, innovative modeling area as the area of cognitive task analysis.

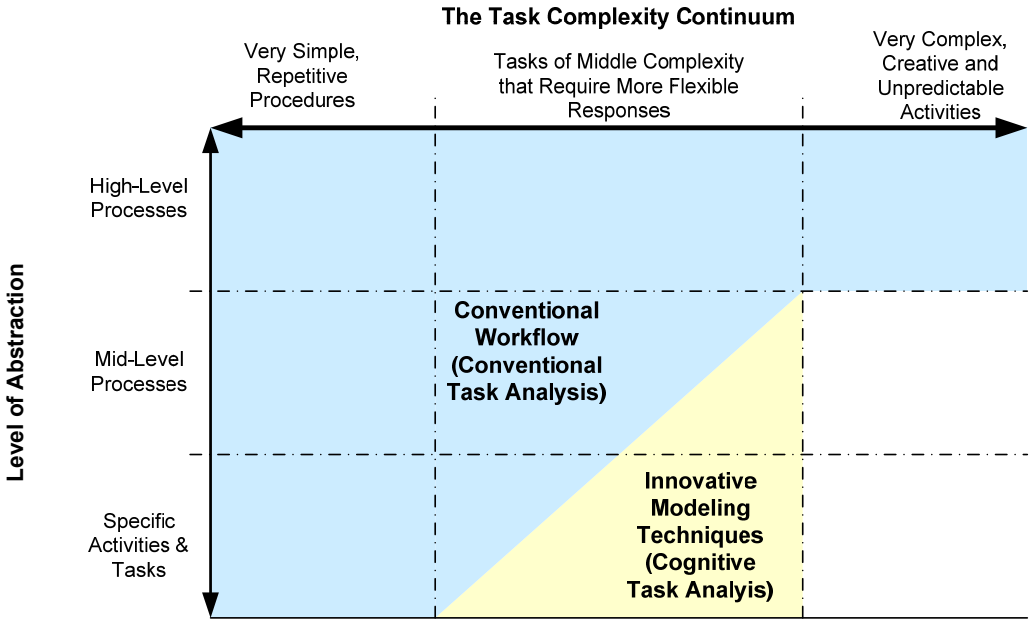

Figure 1. The business process analysis space.

In discussing Working Minds, I also spent some time discussing the techniques that can be used to represent knowledge. In essence, I suggested that there were two elements that worked together. First there is a cognitive network (a frame or object network if you are more familiar with software implementations) which captures the concepts (vocabulary) and the relationships between concepts. Most expert systems eventually used models that described objects and attributes (Birds have wings and feathers and lay eggs. Mammals have fur. Fish have scales. Animals breathe air and are mortal. Examples of Animals are Birds, Mammals and Fish.) The arrows that connect the objects capture relationships, as, for example, that Birds are Animals, etc. In some cases the relationships are simple, but in many cases they imply inheritance. Thus, because Birds are Animals, they inherit the characteristics of animals. There are, of course, exceptions to this rule, and thus, as the knowledge to be captured becomes more extensive these cognitive networks become extremely complex.

The second feature of most knowledge systems is rules. A rule is a logical statement of the form: If X, then Y. In essence, if you know that a given object has certain attributes, you can reason that other things are true or false. Thus, if the animal has feathers, it is a bird, and it lays eggs.

Similarly, if the animal has fur, it's a mammal and very unlikely to lay eggs. The fact that there can be exceptions, like egg laying mammals (e.g. a platypus), suggests that really complex knowledge systems must deal with probabilistic rules, and then things really start to become fun.

The Cambridge Handbook of Expertise and Expert Performance is a collection of articles that define the current state of what we know about expertise and expert performance. It has been edited by four experts from Florida who represent a gamut from psychology to artificial intelligence, and one of the authors, Robert R. Hoffman, was also a co-author of the *Working Minds* book I reviewed last month. The Handbook is a massive book of some 900 pages and a relatively expensive book at $65. It's hardly bedside reading.

I went to the trouble of defining the business process space to suggest who might want to study this book. If you are concerned with ordinary business problems, you don't need to consider this book. On the other hand, if you are focused on improving business processes that rely on the performance of human knowledge workers or experts, this book can provide a wealth of useful information. The book's title, which emphasizes expertise, uses expertise more broadly than I do and often includes discussions of tasks that I would describe as knowledge work rather than genuine expertise.

The chapters are articles written by experts, and there are 42 of them. I won't list them all, but I can give you an overview by listing the sections into which the chapters are grouped:

I. Introduction and Perspective
II. Overview of Approaches to the Study of Expertise – Brief Historical Accounts of Theories and Methods
III. Methods for Studying the Structure of Expertise
IV. Methods for Studying the Acquisition and Maintenance of Expertise
V. Domains of Expertise
    Professional Domains
    Arts, Sports, & Motor Skills
    Games and Other Types of Expertise
    Generalizable Mechanisms Mediating Expertise and General Issues

The book provides a great and comprehensive introduction into the whole filed of knowledge and expertise. The chapters on what make experts provide a detailed review of the research. For example, human experts usually rely on something like 10,000 rules. They normally maintain concept networks that are organized into around seven hierarchal levels. Thus, some rules are used to analyze a problem from a more abstract perspective, some are used for more specific analysis, while still other rules are very concrete and are only used when specific types of problems are encountered. Experience and new information play vital rules in the maintenance of expertise, and thus, an expert, (or a software system) separated from conferences and specific problems, soon begins to lose his or her edge.

Some of the articles won't interest you if that isn't your particular domain of focus, but everyone who is working on cognitive analysis will probably find several useful articles. I especially liked the summary of Expert Systems, by Buchanan, Davis and Feigenbaum, an article on task analysis by Schraagen, an article on representing the knowledge of experts by Hoffman and Lintern, an article on protocol analysis by Ericsson and one on simulation by Ward, Williams and Hancock. I also enjoyed articles an expertise in software design by Sonnentag, Niessen and Volmer and one on Writing Expertise by Kellogg. And I found the article on Decision-Making Expertise by Yates and Tschirhart particularly useful.

I had hoped for more studies of human expertise in business domains. I would have liked to see articles on decision making in banks or at stock firms, or on management decision making at the departmental level, or the use of knowledge to optimize the flow through supply chains. The articles in this book provide insights into how human knowledge workers must approach problems in these domains but did not provide the specifics that I would have liked.

The problem, for the editors, of course, is that it's expensive to study human performance issues with care. It's relatively easy to do it in academic settings, and thus, we get all those interesting studies of how college students make decisions or play games. In business environments, however, it is usually easier to hire an expert and simply trust that he or she has what it takes to perform.

There are exceptions, as I mentioned in reviewing *Working Minds*. Some jobs need to be done very efficiently and require that the employer train the workers. Thus, the military routinely hires young men and women whom they must train for complex and esoteric tasks, like sonar interpretation, rapid ship deployment planning, and battlefield management. Thus, the US department of defense has funded a number of interesting studies on how to analyze what's involved in complex cognitive tasks and how best to train for those tasks. Increasingly, other industries that employ large numbers of knowledge workers are finding themselves faced with similar challenges and will need to learn about the analysis of knowledge and expertise. One thinks of all the new rapidly established industries designed to capitalize on new technologies that require technicians.

If you are a process analyst who is tasked with analyzing the performance of knowledge workers, you will probably find this book a very useful reference book. It may not tell you how to solve the specific problem you face, but it will certainly suggest what has been tried and what has been learned to date, and suggest the latest approaches that are being explored.

Paul Harmon is the Executive Editor of Business Process Trends (www.bptrends.com). He is a recognized BPM analyst and consultant and the author of several books on expert systems, and *Business Process Change*.

