
Reviews of Books

The Architecture of Cognition. JOHN R. ANDERSON
Harvard University Press, 1983 345 + xi pp., \$25.00.

This book is basically the “final report” of John Anderson’s ACT (Adaptive Control of Thought) project, which involves nothing less than the creation of “a unitary theory of mind.” This is not the “final report” on ACT in the sense that Anderson will now stop working on it, but it is in the sense that Anderson considers the ACT theory to be in “final shape.” Thus, far from abandoning work on ACT, Anderson now feels that the theory is complete enough that he can text it in a variety of domains and subject it to exacting scrutiny by a wide variety of other investigators.

ACT is not a “unitary theory of mind” in the sense that it proposes that the key to mind is some unitary principle or mental representation (such as past proposals like stimulus-response pairs, propositions, schemes, etc.). ACT claims that a fairly complex set of representations and processes are necessary for the variety of abilities we characterize as mind. ACT is unitary in the sense that the same set of representations and processes are used for all aspects of mind. Thus, for example, according to ACT we do not have a separate linguistic capacity or mathematics capacity; both of these capacities (and the other capacities of mind) are manifestations of the same underlying processes and representations. Anderson provides a variety of evidence for this view during the course of the book, including showing that the kinds of generalization processes that operate while one is learning to prove geometry theorems are the same as the kinds that operate when one is learning to use a computer programming language and the same as the kinds that operated when one was first learning to speak a natural language.

ACT evolved from the earlier HAM (Human Associative Memory) theory that was the subject of a book Anderson wrote with Gordon Bower in 1973. The HAM theory involved a particular kind of semantic network that Anderson and Bower used to account for many memory phenomena. In a 1976 book, *Learning, Memory and Thought*, Anderson proposed an initial version of ACT that essentially added to the HAM network a processing component represented as a production system. As implied by the more grandiose name (Adaptive Control of Thought as opposed to Human Associative Memory), ACT aspires to cover much more territory than HAM did. Whereas HAM focused on simple laboratory studies of human memory, ACT tackles all of cognition — e.g., memory, language, learning, problem solving, etc.

The earlier and present versions of ACT are similar in that their basic architecture is composed of a declarative memory of facts, a production memory of procedures, and a working memory of active declarative facts. However,

they differ in that the original declarative memory representation was a semantic network with simple linguistics-like links (subject, predicate, relation, object) connecting undecomposable concept nodes, and the newer declarative representation has three kinds of complex cognitive units (temporal strings, spatial images, and abstract propositions) connected by a vast array of link types. Also, in the earlier theory, spread of activation through the declarative network was the major determiner of the time to perform a task, whereas in the newer theory the major determiner of performance time is the pattern matching processes used to match the conditions of the productions to the contents of working memory. Anderson convincingly motivates these revisions and others using both system design considerations and empirical results from psychology experiments.

Anderson is able to motivate and justify his theory by mustering a dazzling array of evidence of various kinds. For example, in discussing ACT’s pattern matching mechanisms, Anderson moves from letter and word perception to learning to prove geometry theorems, to planning one’s everyday errands. Even more surprising than these dramatic leaps from one domain to another is that the evidence cited from these diverse domains seems to fit together to support Anderson’s arguments.

An important contribution of this book, over and above the particular theory of cognition proposed, is the vast array of diverse results brought together in an organized fashion to constrain the architecture of cognition. I think any cognitive theorist would find these results important to keep in mind when devising his or her own theory even if that theory is quite different from Anderson’s. For example, I found Anderson’s set of evidence for the distinction between declarative and procedural information valuable, even though I did not conclude — as he did — that there was therefore a separate declarative and procedural memory. I personally find a cognitive architecture composed of cognitive units each containing related declarative and procedural information more compelling, but I also find Anderson’s arguments for the distinction persuasive. Thus, although I would propose a different cognitive architecture, I still find the evidence presented in the book invaluable.

Anderson has appropriately titled this book *The Architecture of Cognition* because he is indeed functioning as an architect here — and a modernistic one at that. In particular, Anderson gives us the spare superstructure of a cognitive edifice, but he leaves the exterior and interior design and finishing to others (and perhaps himself in the future). For example, he describes the general mechanisms needed for understanding a narrative text or learning to program a computer, but he does not attempt to account for the kinds of knowledge that the mechanisms need to use to understand a narrative (e.g., knowledge about underlying goals and plans)

or for the kinds of knowledge acquired when learning to program (e.g., the specific program plans learned). Thus, in Newell's apt terminology, this book is an account of the *system* level of cognition, but not the *knowledge* level.

Overall, I found this to be an interesting and valuable book for two reasons: first, it provides a thought-provoking proposal for the architecture of this system level of cognition and, second, it organizes a myriad of different kinds of results that provide general constraints as the architecture of cognition. I recommend this book as essential reading for researchers in all of the cognitive sciences.

John B. Black

Department of Psychology and
Department of Computer Science
Yale University

Expert-EaseTM

The Personal Tool For Building Expert SystemsSM

Now you can produce **commercially valuable** expert systems on an IBM® Personal Computer for a price far below the previous cost of such efforts.

Based on the work of Donald Michie and his colleagues, Expert-Ease enables non-technical users to build knowledge-based systems rapidly and easily. These systems can be used by non-experts with no technical training to generate decisions, troubleshooting strategies, diagnoses or other results which previously required the time and physical presence of an expert.

Expert-Ease enables you and your organization to produce **valuable knowledge-based systems for present use, while allowing you to assess the impact that expert systems will have on your organization in the future.**

Expert-Ease runs on the IBM PC or XT, and some "compatibles." It requires at least 128K of memory and one disk drive; two are advisable. It is available for \$2000.00 U.S. plus shipping and tax where applicable.

For additional information, or to purchase a copy of Expert-Ease, contact:

JEFFREY PERRONE & ASSOCIATES, DISTRIBUTORS.

3685 17th Street
San Francisco, California 94114

(415) 431-9562

or contact us on the Source electronic mail facility:

Source: BKP517

Expert Ease is a trademark of Expert Software International Ltd
"The Personal Tool For Building Expert Systems" is a service mark of
Jeffrey Perrone & Associates IBM is a trademark of IBM Corp