Brain and Cognition **47,** 570–573 (2001)
doi:10.1006/brcg.2001.1459, available online at http://www.idealibrary.com on IDEAL®

REVIEW

Hybrid Modeling of Cognition: Review of *The Atomic Components of Thought*. John R. Anderson and Christian J. Lebiere. Mahwah, NJ: Erlbaum, 1998

Reviewed by Philip Pavlik, Jr.

Northern Michigan University

Published online November 14, 2001

The Atomic Components of Thought (Anderson & Lebiere, 1998) is a report on an ambitious work in progress. As with most attempts to construct a unified theory of cognition this project harkens to the call of Allen Newell (1990) when he said: ''Psychology has arrived at the possibility of unified theories of cognition—theories that gain their power by positing a single system of mechanisms that operate together to produce the full range of human cognition.''

This is certainly the goal of Anderson and Lebiere and the ACT-R (Adaptive Components of Thought—Rational) team at Carnegie Mellon. In pursuing this goal, Anderson and his colleagues have constructed one of the most fully articulated and complex hybrid production rule architectures extant. They describe it as a hybrid architecture because it blends aspects of production rule modeling (conditional action rules interacting with a data structure) with a certain amount of subsymbolic processing (such as a memory activation function) similar to connectionist modeling. The latest version of this theory, ACT-R 4.0, allows finer grained modeling of thought and has clearer differentiation between symbolic and subsymbolic processes than previous ACT theories. ACT-R 4.0 is the subject of *The Atomic Components of Thought*.

In the introduction to this book, Anderson and Lebiere present a ''no-magic doctrine'' consisting of six principles that they claim have driven their work with ACT-R and its realization in this book. These six principles can be taken as a benchmark from which they believe their work with the ACT-R theory should be measured. These principles do provide a good basis from which to evaluate their work.

The first of these six principles insists that ACT-R remains *experimentally grounded*. Because psychology and cognitive science are attempts to describe reality, this is a necessity, and the authors do well in this arena with some small caveats. The book's examples all carefully compare the performance of the ACT-R model with the performance of actual subjects, for example. More importantly though, this book, in Chapters 5 and 6, carefully explains the new work with ACT-RPM, an expansion of their model that includes a visual interface, auditory processing, speech production, and motor processing. These new features bring ACT-R a long way toward being more firmly grounded in the experimental data and are utilized in some of the models described later in the work. One could hope that at some point all

Address correspondence and reprint requests to Philip Pavlik, Jr., Psychology Department, Baker Hall, Carnegie Mellon University, Pittsburgh, PA 15213. Fax: (412) 268-2798. E-mail: ppavlikjr@hotmail.com.

0278-2626/01 $35.00

modeling in ACT-R would require the integration of perceptual input modeling and motor outputs to be considered a serious effort.

In their second principle, they lay claim to modeling that achieves *detailed and precise accounting of data*. They show how their theory makes predictions of more than just a basic aspect or two of a subject's behavior, such as number correct or latency. For instance, in the case of modeling experiments where human subjects make choices, ACT-R is able to model latencies, orders of subject actions, number correct, and all other responses measurable by the program used to analyze the human subjects. Additionally, through its partial matching equation ACT-R can accurately model the types of errors that subjects commit. The detail with which ACT-R modeled the behaviors analyzed was impressive.

One of the most notable sections of the book addresses the third principle of their ''no-magic doctrine.'' This third principle is that the model should be capable of *learning*. While most ACT-R models start out with significant background knowledge, similar to what we would hope the undergraduate research subject possesses, ACT-R also has mechanisms for learning. In a meticulous demonstration of these mechanisms, this book documents how ACT-R has been used to model the math experience of the typical American child from kindergarten to college and the corresponding progression of math skills.

The fourth principle stakes a claim for ACT-R to be able to model *complex cognitive phenomena*. In demonstrating that ACT-R can achieve this, the authors attempted to model analogy use in physics problem solving and the scientific discovery process in psychology. In both attempts they took a simplified version of the process being investigated and designed rather complex production rule models of the phenomena. In both cases, their attempts were largely successful and useful both in understanding the processes they examined and in showing that ACT-R can be used effectively in modeling complex cognition.

In the physics problem-solving example, the authors presented subjects and the ACT-R model with a series of simple problems involving calculations like flux, velocity, and volume. The subjects and the ACT-R model were presented on a computer with a solved example that could be selectively viewed by using the mouse to reveal specific areas of the display as they were needed in order to find and map procedures from the sample to a test problem. The model and the subjects were tested in two conditions, one in which the variables matched with the quantity names (such as resistance or mass), and one in which there was no matching. Two ACT-R models, with many overlapping productions were developed to deal with these two cases, and even though they were constrained by this overlap, their performance was very close to subject data in latencies and numbers correct. (This effort also highlights a growing controversy in the ACT-R community over the utility and necessity of using both declarative and production rule representation in modeling. Two versions of this model were devised, one relying on declarative representations and another using purely production rules, and both modeled subject data equally well.)

In the scientific discovery example, an ACT-R production rule model was compared with task experts (psychologists), domain experts (cognitive psychologists), and both high- and moderate-ability undergraduates. Subjects and model were tested on their scientific discovery using a rather complex simulation of a memory experiment lab that allowed varying of a number of parameters in order to run experiments with the purpose of testing two original hypotheses. Not only did this model elucidate the scientific discovery process at a very fine-grained mechanistic level, but also various aspects of the ACT-R implementation reflected the sophistication of this modeling architecture. Specifically, this model interleaved specific and general production rules, coded for temporal context, and was fully integrated with the working memory model implicit in the ACT-R architecture. Further, this model utilized the visual interface in such a way that at some points the model was making decisions based on working memory activations from past experiments and sometimes through a review of data presented visually.

The fifth principle applies to the use of *parameters*. Anderson and Lebiere are quite cognizant of the concerns over the number and justification of parameters in ACT-R modeling. While this issue is a real concern, it is something that is being appropriately dealt with by the authors. Many parameters vary little from model to model in the ACT-R architecture, and there is a constant effort to reduce the number of variable parameters. Eventually we could hope for modeling that did not require any parameters to vary; however, until that point arrives ACT-R is effective in minimizing parameterization of models.

Maintaining a modeling environment that is *neurally plausible* is the final principle to which *The Atomic Components of Thought* is dedicated. Unfortunately, only seven pages deal with this principle, describing ACT-RN, a connectionist implementation from 1993 of an earlier version of the ACT-R theory described in this book. While these seven pages went fairly far to demonstrate that a connectionist implementation of ACT-R might be possible, they did not explain in any complete and systematic fashion how the latest theory would be mapped into a connectionist model. The authors provide only a slight explanation of how ACT might correspond with neural architecture. This sixth principle should have been dealt with more completely in a fashion that was more referential to the current theory.

One tentative criticism is the *de novo* nature of the modeling in the ACT-R architecture. When designing a production rule system to model some subject performance, the general practice is to analyze the task, extrapolate a likely schematic structure of the production rules involved, and then implement the model in ACT-R. While this is effective for modeling and learning about tasks, it fails to account for the developmental course of the assemblage of production rules as would occur in a maturing human. Just as it is important to retain neural plausibility, it also would seem important to maintain this sort of developmental plausibility.

A second tentative problem in the implementation of modeling in this architecture is the reliance on explicit goal structures for the task at hand. Obviously, the subjects that the models are being compared to have much more complex motivational bases of their performance than the goal structure ACT-R posits. With this in mind, it seems that phenomena such as anxiety, boredom, and emotional states must be integrated into goal structures in order to fully model performances of subjects.

Because ACT-R is itself a theory in the development these concerns may well be addressed in a future revision. At no point do Anderson and Lebiere exclude certain approaches in the future or indicate that their theory is static. While they remain well fixed in certain basic aspects of their model, such as the serial nature of production rule execution, they seem to base their positions firmly in relation to the data, and appear ready to change their model if their position becomes untenable.

This book provides a well-reasoned and coherent theoretical and research overview of the most complete unified theory for modeling human information processing. It appears to be directed toward an audience of graduate students and above, but may be approachable by the advanced undergraduate. Researchers in cognition and perception will be most interested in this book; however, other fields may benefit from an understanding of this maturing modeling language.

REFERENCES

Anderson, J., & Lebiere, C. (1998). *The atomic components of thought.* Mahwah, NJ: Lawrence Erlbaum Assoc.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard Univ. Press.