Cognitive Science 29 (2005) 307–311

# Introduction to the 2004 Rumelhart Prize Special Issue Honoring John R. Anderson

This special issue honors the research and mentorship contributions of Dr. John R. Anderson, the 2004 David E. Rumelhart Prize recipient. This prize was instituted in 2001, funded by the Robert J. Glushko and Pamela Samuelson Foundation. The prize is awarded annually to an individual or collaborative team making a significant contemporary contribution to the formal analysis of human cognition. Mathematical modeling of human cognitive processes, formal analysis of language and other products of human cognitive activity, and computational analyses of human cognition using symbolic or nonsymbolic frameworks all fall within the scope of the award. In 2004 the prize selection committee consisted of Alan Collins, Robert Glushko, Mark Liberman, Anthony Marley, and James McClelland (chair). Perhaps best known for his contributions to connectionist or neural network models, Dr. Rumelhart also exploited symbolic models of human cognition, formal linguistic methods, and the formal tools of mathematics. Reflecting this diversity, the first three winners of the David E. Rumelhart Prize are individuals whose work lies within three of these four approaches. Past recipients are Geoffrey Hinton, a connectionist modeler, Richard M. Shiffrin, a mathematical psychologist, and Aravind Joshi, a formal and computational linguist. Anderson is the leading proponent of the symbolic modeling framework, thereby completing coverage of the four approaches.

## Research Biography of John R. Anderson

John R. Anderson, Richard King Mellon Professor of Psychology and Computer Science at Carnegie Mellon University, is an exemplary recipient for a prize that is intended to honor “a significant contemporary contribution to the formal analysis of human cognition.” For the last 3 decades, Anderson has been engaged in a vigorous research program with the goal of developing a computational theory of mind. Anderson’s work is formulated within the symbol processing framework and has involved an integrated program of experimental work, mathematical analyses, computational modeling, and rigorous applications. His research has provided the field of cognitive psychology with comprehensive and integrated theories. Furthermore, it has had a real impact on educational practice in the classroom and on student achievement in learning mathematics.

Anderson’s contributions have arisen across a career that consists of five distinct phases. Phase 1 began when he entered graduate school at Stanford at a time when cognitive psychology was incorporating computational techniques from artificial intelligence. During this period and immediately after his graduation from Stanford, he developed a number of simulation models of various aspects of human cognition such as free recall (Anderson & Bower, 1972). The human association memory (HAM) theory that he and Gordon Bower developed (Anderson & Bower, 1973) immediately attracted the attention of everyone then working in the field. The book played a major role in establishing propositional semantic networks as the basis for representation in memory and spreading activation through the links in such networks as the basis for retrieval of information from memory. It also provided an initial example of a research style that has become increasingly used in cognitive science—the creation of a comprehensive computer simulation capable of performing a range of cognitive tasks and the comparison of this model with a series of experiments addressing the phenomena within that range.

Dissatisfied with the limited scope of his early theory, Anderson undertook the work that has been the major focus of his career to date, the development of the Adaptive Control of Thought (ACT) theory (Anderson, 1976). ACT extended the HAM theory by combining production systems with semantic nets and the mechanism of spreading activation. The second phase of Anderson's career is associated with the initial development of ACT. The theory reached a significant level of maturity with the publication in 1983 of *The Architecture of Cognition* (Anderson, 1983), which is the most cited of his research monographs (having received over 2,000 citations in the ensuing years). At the time of publication, the ACT* model described in this book was the most integrated model of cognition that had then been crafted and tested. It has had a major impact on the theoretical development of the field and on the movement toward comprehensive and unified theories, incorporating separation of procedural and declarative knowledge and a series of mechanisms for production rule learning that became the focus of much subsequent research on the acquisition of cognitive skills. In his *Unified Theories of Cognition,* Alan Newell had this to say: "ACT* is, in my opinion, the first unified theory of cognition. It has pride of place. … [It] provides a threshold of success which all other candidates … must exceed" (Newell, p. 29).

Anderson then began a major program to test whether ACT* and its skill acquisition mechanisms actually provided an integrated and accurate account of learning. He started to apply the theory to development of intelligent tutoring systems; this defines the third phase of his research. This work grew from an initial emphasis on teaching the programming language LISP to a broader focus on high school mathematics (Anderson, Corbett, Koedinger, & Pelletier, 1995), responding to perceptions of a national crisis in mathematics education. These systems have been shown to enable students to reach target achievement levels in a third of the usual time and to improve student performance by a letter grade in real classrooms. Anderson guided this research to the point where a full high school curriculum was developed that was used in urban schools. Subsequently, a separate corporation was created to place the tutor in hundreds of schools, influencing tens of thousands of students. The tutor curriculum was recently recognized by the Department of Education as one of five "exemplary curricula" nationwide. Although Anderson does not participate in that company, he continues research to develop better tools for tracking individual student cognition, and this research continues to be informed by the ACT theory. His tutoring systems have established that it is possible to impact education with rigorous simulation of human cognition.

In the late 1980s, Anderson began work on what was to define the fourth phase of his research, which was an attempt to understand how the basic mechanisms of a cognitive architecture were adapted to the statistical structure of the environment. Anderson (1990) called this a rational analysis of cognition and applied it to the domains of human memory, categorization, causal inference, and problem solving. He used Bayesian statistics to derive optimal solutions to the problems posed by the environment and showed that human cognition approximated these solutions. Such optimization analysis and use of Bayesian techniques have become increasingly prevalent in cognitive science.

Subsequent to the rational analysis effort, Anderson has turned his full attention back to the ACT theory, defining the fifth and current phase of his career. With Christian Lebiere, he has developed the Adaptive Control of Thought–Rational (ACT–R) theory, which incorporates the insights from his work on rational analysis (Anderson & Lebiere, 1998). Reflecting the developments in computer technology and the techniques learned in the applications of ACT*, the ACT–R system was made available for general use. A growing and very active community of well over 100 researchers is now using it to model a wide range of issues in human cognition, including dual tasking, memory, language, scientific discovery, and game playing. It has become increasingly used to model dynamic tasks such as air traffic control, where it promises to have training implications equivalent to the mathematics tutors. Through the independent work of many researchers, the field of cognitive science is now seeing a single unified system applied to an unrivaled range of tasks. Much of Anderson's own work on the ACT–R has involved relating the theory to data from functional brain imaging (Anderson, Qin, Sohn, Stenger, & Carter, 2003).

In addition to his enormous volume of original work, Anderson has found the time to produce and revise two textbooks, one on cognitive psychology (Anderson, 2005) and the other on learning and memory (Anderson, 2000). The cognitive psychology textbook, now in its sixth edition, helped define the course of study that is modern introductory cognitive psychology. His more recent learning and memory textbook, now in its second edition, is widely regarded as reflecting the new synthesis that is occurring in that field in animal learning, cognitive psychology, and cognitive neuroscience.

Anderson has previously served as president of the Cognitive Science Society and has received a number of awards in recognition of his contributions. In 1978 he received the American Psychological Association's Early Career Award, in 1981 he was elected to membership in the Society of Experimental Psychologists, in 1994 he received APA's Distinguished Scientific Contribution Award, and in 1999 he was elected to both the National Academy of Sciences and the American Academy of Arts and Science. Currently, as a member of the National Academy, he is working toward bringing more rigorous science standards to educational research.

## The Special Issue on Theoretical Advances and Applications of Unified Computational Models

This special issue of *Cognitive Science* features work by John Anderson's students, colleagues, and collaborators. The diversity of these articles attests to the fertility and generality of the ACT–R framework. It is a rare framework that can be employed for cognitive processes from driving simulations (Salvucci) and Web navigation (Pirolli) to sentence parsing (Lewis) and algebra (Anderson). It might be thought that a framework could only accommodate this diverse range of tasks by being a rather unconstrained general-purpose language. These contributions disprove this contention. ACT–R is powerfully grounded by rational considerations, and specific Bayesian formalisms lie at its core. Interestingly, ACT–R is not only constrained from above by considerations of rational use of probabilistic information, but is also constrained from below by recent neuroscience evidence (Anderson et al., 2003). The contributions by Lewis, Lovett, Pirolli, Salvucci, and Taatgen (this issue) show that a variety of detailed situations can be implemented in ACT–R, whereas Anderson's article (this issue) suggests how ACT–R, in turn, may be implemented in the brain.

ACT–R's combination of neural, computational and mathematical considerations impose strong constraints on the timing and control of cognitive processes, and impressively these constraints have been repeatedly shown to apply to human performance (for examples, see Lewis' constraints on sentence parsing, and Lovett's constraints on learning utility values of information sources). These articles attest not only to the broad applicability of ACT–R framework, but also to its genuine predictiveness. In short, the framework has proved to be a felicitous combination of a general and appropriately constrained language for expressing theories.

The articles reveal current lively areas of research for ACT–R specifically, and unified cognitive models more generally. One area of recent activity, featured prominently in the articles by Lovett, Salvucci, and Taatgen, is the need for mechanistic models of control processes. This focus is a sign of the maturity of cognitive science. In early computational models of cognition, it was considered sufficiently impressive to model performance on one task. Several of the articles in this special issue raise the bar for cognitive models, arguing that we need to know how a system behaves when confronted with several ongoing tasks. When the need to accomplish multiple tasks is acknowledged, mechanisms are needed to shift between tasks, foreground backgrounded processes, prioritize tasks, integrate and/or modularize processing, and monitor success and update plans accordingly. The authors tackle all of these hard issues with rigorous aplomb and take sizable steps toward grounding cognitive control in productions (Lovett, Taatgen, Salvucci) and brain structures (see Anderson's discussion of the role of the anterior cingulate in attentional control).

Another crucial aspect of human cognition that has been historically underrepresented is the surprising degree to which people can program themselves to strategically process their world. Failures of strategic programmability are prominent in cognitive science. Phenomena such as Stroop interference (see Lovett's article, this issue) are intriguing because they suggest that we are not in complete control of our cognitive processes. When instructed to name the color of a word's ink, participants are slower if the word is the name of a conflicting color than if it is a neutral word. Similarly, social psychologists have studied situations where people automatically tend to associate faces from other races with negative words at the same time that they associate faces from their own race with positive words, despite their stated desire not to engage in racial stereotypes and their explicit denial of the validity of these stereotypes. Demonstrations of these failures to have the thoughts we want to have are important because we generally believe we are in control of our own mental processes. However, it is equally important to reflect on the fact that these failures of control are only striking precisely because we know that we often *can* program ourselves with impressive levels of efficiency and reliability. Parking behavior is radically affected by presenting facts to oneself such as "The parking meters are not checked on Sundays" and "Snowplows come through this street at 3:00 a.m. when it snows." A single sentence of instructions to pay attention to color and not shape (Lovett), road conditions rather than incoming cell phone calls (Salvucci), or to use a specific method for solving algebraic problems (Anderson) can radically affect performance. Most computational models of cognition have ignored that people can be programmed by language to immediately alter their cognition. Several of these articles in this special issue (Anderson, Lovett, Salvucci, & Taatgen) embrace the unmet challenge of constructing cognitive devices that can be dynamically programmed by task needs, instructions, and changing strategies.

A final theme common to several of the manuscripts (most notably in the articles by Anderson, Pirolli, and Taatgen) is the fertility of applying rational and Bayesian principles to human cognition. The apparent wisdom of this approach appears on the surface to be cast in doubt by the large number of empirical demonstrations of suboptimal and irrational human decision making, including the research for which Daniel Kahneman was awarded a 2002 Nobel Prize in economics. Failures of rational decision making notwithstanding, enormous progress in many fields has been made with Bayesian modeling. This point is emphasized by the articles in this issue—they demonstrate surprising depth to rational analyses of memory, inference, search, and information access. Humans may not be built in the best way possible, but comparing our performance to formally devised gold standards of rationality has revealed both informative deviations from optimality and surprising cases where cognition apparently conforms to normative accounts.

## Selected Bibliography

Anderson, J. R. (1976). *Language, memory, and thought.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Anderson, J. R. (2000). *Learning and memory* (2nd ed.). New York: Wiley.

Anderson, J. R. (2005). *Cognitive psychology and its implications* (6th ed.). New York: Worth.

Anderson, J. R., & Bower, G. H. (1972). Recognition and retrieval processes in free recall. *Psychological Review, 79,* 97–123.

Anderson, J. R., & Bower, G. H. (1973). *Human associative memory.* Washington, DC: Winston.

Anderson, J. R., Corbett, A. T., Koedinger, K., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *Journal of the Learning Sciences, 4,* 167–207.

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

Anderson, J. R., Qin, Y., Sohn, M-H., Stenger, V. A., & Carter, C. S. (2003). An information-processing model of the BOLD response in symbol manipulation tasks. *Psychonomic Bulletin & Review, 10,* 241–261.

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