

Can NLP Systems be a Cognitive Black Box? (Is Cognitive Science Relevant to AI Problems?)

Jerry T. Ball

Air Force Research Laboratory, Human Effectiveness Directorate
Warfighter Readiness Research Division
6030 S. Kent Street, Mesa, AZ 85212-6061
Jerry.Ball@mesa.afmc.af.mil

Abstract

This paper considers whether or not the internals of NLP systems can be a black box with respect to the modeling of how humans process language in answer to the question “*Is cognitive science relevant to AI problems?*”. Is it sufficient to model the input/output behavior using computational techniques which bear little resemblance to human language processing or is it necessary to model the internals of human language processing behavior in NLP systems? The basic conclusion is that it is important to look inside the black box of the human language processor and to model that behavior at a lower level of abstraction than input/output behavior. The development of functional NLP systems may actually be facilitated, not hindered, by adoption of cognitive constraints on how humans process language. The relevance of this position for the symposium is considered and some suggestions for moving forward are presented.

NLP Systems as a Black Box

Natural Language Processing (NLP) is a quintessential *AI Hard Problem*. For an NLP system to be successful, it must mimic human behavior at the level of input and output. Otherwise, successful communication with humans will not be achieved. Unlike many other AI systems, performing better than humans is not a desirable outcome (although performing as well as expert communicators is). The key question is whether or not human-like input and output can be achieved using computational mechanisms that bear little resemblance to what cognitive science and cognitive psychology tell us is going on inside the head of humans when they process language. *Can NLP systems be a cognitive black box (Is cognitive science relevant to AI problems)?*

To date, research in the development of functional NLP systems has largely adopted the black box approach—assuming that modeling the “*internals*” of human language processing is undesirable, and hopefully, unnecessary. It is undesirable because it would impose severe constraints on the development of functional NLP systems, and, besides, the basics of how humans process language have not been sufficiently worked out to support computational implementation. The NLP problem is too hard for us to

make progress if we accept the constraints on human language processing proposed by cognitive science researchers—especially given the many conflicting hypotheses they have put forward—and try to populate the black box with cognitively plausible systems.

But there is considerable evidence to suggest that AI researchers ignore the constraints of cognitive science and cognitive psychology at their own peril. The advent of the *parallel distributed processing* (PDP) strain of *connectionism* (Rumelhart and McClelland, 1986) was in part the product of cognitive science researchers focused on modeling the “internals” of human cognitive and perceptual behavior. Connectionist researchers highlighted many of the shortcomings of symbolic AI, arguing that they resulted from an inappropriate *cognitive architecture*, and proposing alternatives that delved inside the black box of cognition and perception, trying to specify what a cognitive system would be composed of at a level of abstraction somewhere above the neuronal level, but definitely inside the black box. The confrontation between connectionism and symbolic AI has largely subsided, and many of the connectionist claims have been shown to be overstated (cf. Pinker and Prince, 1988), but it seems clear that connectionist models do capture important elements of perception and cognition, particularly lower level phenomena. Many AI researchers are now working to integrate connectionist layers into their hybrid symbolic/subsymbolic systems (Sun and Alexandre, 1997) to capture the symbolic irregularities the connectionist systems revealed and which purely symbolic systems cannot easily model.

Within NLP, the problems of noisy input, lexical and grammatical ambiguity and non-literal use of language call out for adoption of techniques explored in connectionist and statistical systems, many of which come out of research in cognitive science. For example, *Latent Semantic Analysis* (LSA) (Landauer & Dumais, 1997), a statistical, data-reduction technique that is being used by psycholinguists to capture the latent (i.e. non-explicit) meaning of words as a multi-dimensional vector, offers hope for solving previously intractable problems in meaning representation—especially the problem of determining similarity of meaning without assuming

discrete word senses. If the LSA approach is successful, then an avalanche of NLP research on word sense disambiguation that is based on the identification of discrete word senses will need to be revisited.

Psycholinguistically motivated symbolic resources like *WordNet*—a computational implementation of a large scale *mental lexicon*—are also being adopted by many AI researchers (Fellbaum, 1998). The primary use of *WordNet* within the AI community is as a tool for scaling up NLP systems, without necessarily adopting the psycholinguistic principles on which it is based and without making claims of cognitive plausibility in the resulting systems. Interestingly, George Miller, the cognitive scientist leading development of *WordNet*, laments the fact that *WordNet* is not being used more extensively by the psycholinguistic community. However, since psycholinguists are not usually concerned with the development of large-scale systems, and since *WordNet*, like many other computational implementations of cognitive theories, has had to make some admittedly non-psychological choices, it has not had a major impact on this community.

As these examples attempt to show, AI researchers have historically paid attention—if somewhat reluctantly—to the major trends in cognitive science and will continue to do so. But AI researchers must adapt the products of cognitive science to their own research agenda—the development of intelligent, functional computer systems. AI researchers are unlikely to spend years of research exploring specific cognitive phenomena. Such research does not lead to the development of functional systems. However, given the complexity of the systems they are building, AI researchers should seek awareness of advances in cognitive science that point the way towards computational implementation of more humanly plausible systems.

An awareness of cognitive science research is especially important for the development of functional NLP systems. The search space for possible solutions to the development of NLP systems is huge (perhaps infinite). To date most systems have been searching the part of that space consistent with our basic understanding of computation and symbol manipulation. But if philosophers like Searle (1980) and cognitive scientists like Harnad (1990, 1992) are right, ungrounded symbol systems will only prove suitable for solving a limited range of problems. Ultimately, our symbol systems will need to be grounded if they are to display the full range of human behavior and intelligence. Philosophers like Prinz (2002) and psychologists like Barsalou (1999) and Zwaan (2004) are exploring the implications of the perceptual grounding of symbols, and their research could well have important implications for NLP systems. Their research may open up additional subspaces in the search space for solutions to NLP problems that are ripe with interesting possibilities to be explored. To the extent that research in cognitive science is able to focus the search for solutions on fruitful paths and to prune the search tree by eliminating non-productive branches, it could actually facilitate the development of functional systems. This search argument hinges on the assumption that cognitively implausible systems are unlikely to be able to mimic human

input/output behavior in a domain as complex and human centric as language processing. It is the assumption that *NLP systems should not be cognitive black boxes (That cognitive science is relevant to AI problems)*. That we are more likely to be successful in developing NLP systems by modeling the human cognitive behavior inside the box than by applying computational mechanisms that only attempt to mimic input/output behavior.

Cognitive and Computational Constraints

Having argued for the adoption of cognitive constraints in the development of large-scale functional NLP systems, it must be admitted that there are few successes to date, and very few researchers who are even engaged in this line of research. An obvious way to apply cognitive constraints is to develop NLP systems within a cognitive architecture like ACT-R (Anderson et al., 2004: Anderson and Lebiere, 1998) or Soar (Rosebloom et al., 1993). NL-Soar (Lehman et al., 1995) is one of the very few NLP systems developed in a cognitive architecture. NL-Soar was used in the TacAir-Soar (Laird et al., 1998) project to provide natural language capabilities to synthetic agents which participated as enemy aircraft in a Tactical Air simulation. NL-Soar and TacAir-Soar were among the first successful uses of a cognitive architecture to build functional agents with language capabilities. However, during the course of the TacAir-Soar project, cognitive plausibility was de-emphasized in the interest of developing a functional system within the time constraints of the project. Even within a cognitive architecture it is possible to build cognitively implausible systems.

The AutoTutor system (Graesser et al., 2001) is another example of an NLP system influenced by cognitive science research. AutoTutor is a intelligent tutor that helps students learn how to solve physics problems. Although AutoTutor is not implemented in a cognitive architecture, it is based on extensive psycholinguistic research in discourse processing (Graesser et al., 2003) and it makes use of LSA to assess the meaning of student responses that cannot be fully processed by the higher level language understanding component. A key feature of AutoTutor is the integration of a *talking head*. Talking heads and *avatars* introduce a host of additional requirements for modeling human-like behavior in AI systems.

In my own research (Ball, 2004), I am using the ACT-R cognitive architecture to support the development of a functional language comprehension system (originally implemented in Prolog). A prototype system that constructs linguistic representations from written input currently exists and an ambitious research program leading to the development of language-enabled synthetic entities capable of interacting with humans and other agents in simulation environments is in the initial stages. As an initial attempt to scale up the prototype system, the use of Cyc to add commonsense knowledge was explored (Ball et al., 2004). More recently, an interface to *WordNet* has been adapted for use with ACT-R. A key challenge will be to integrate

WordNet and/or Cyc into the system in a cognitively plausible manner and not just as an external database of lexical and commonsense knowledge. That means transparently integrating these resources with the spreading activation mechanism of ACT-R and with any other cognitive mechanisms with which they interact. While this may seem like an unnecessary constraint on the development of a functional system, integrating WordNet with ACT-R's spreading activation mechanism is key to solving the word sense disambiguation problem (at least within ACT-R). An NLP system developed without access to ACT-R's spreading activation mechanism would need some alternative mechanism for solving this problem (e.g. LSA). Thus, although ACT-R's spreading activation mechanism constrains the development of the system, it provides a needed capability which is motivated by a mountain of psychological evidence, and this capability is seen as a benefit of the architecture and not a limitation.

As another example of a cognitive constraint imposed by a cognitive architecture, ACT-R lacks a backward reasoning mechanism. Productions in ACT-R are only selected and executed via forward chaining. Again, there is a large amount of psychological evidence supporting this architectural constraint. Procedural knowledge is directional, and that direction is forward. On the other hand, backward chaining is a frequently used computational technique in AI systems (especially in reasoning systems). Many AI researchers would be reluctant to excise this technique from their computational toolbox. But if human language processing ability does not make use of backward chaining, then including such a capability in an NLP system distances the system from human language processing behavior and makes the mapping to human input/output behavior more problematic. Besides limiting production execution to forward chaining, ACT-R provides no mechanism for backtracking when production execution leads to a dead end. Consider the problem of processing *garden path sentences* like

The horse raced past the barn fell (Townsend and Bever, 2001)

It might be assumed that an algorithmic backtracking mechanism is needed to process such sentences—when a problem is encountered the processor backtracks and tries different alternatives (i.e. competing productions) until a solution is found. A large amount of psycholinguistic research demonstrates the difficulty humans have in processing such sentences, but there is little evidence that humans use anything like a backtracking algorithm to process such sentences. They do try to process garden path sentences multiple times, but not in an algorithmic backtracking manner. In my own experience of presenting this sentence to non-linguists, they are unable to make complete sense of it without elaborate explanation. Although most linguists consider this sentence to be perfectly grammatical—an instance of a reduced relative clause—and an NLP system with a backtracking mechanism and rules for recognizing reduced relative

clauses would eventually arrive at an appropriate parse, most non-linguists just cannot make full sense of it. They may reason that either “the horse” or “the barn” fell, but they have great difficulty in treating “the horse” as the patient of the verb “raced” rather than the agent—which is what a reduced relative clause reading leads to, as the unreduced and expanded form below demonstrates:

The horse (patient) that was raced past the barn by the jockey (agent) fell

An NLP system that is ignorant of such human limitations may generate garden path sentences which humans will be unable to interpret.

As another example, humans have great difficulty in processing *center embedded constructions* like:

The mouse the cat the dog chased bit ate the cheese

These constructions are easily handled using a stack mechanism capable of storing the noun phrases until the verbs are encountered and unwinding the stack to assign verbs to the appropriate noun phrases. I take the difficulty of these sentences for humans as an indication that they do not have a stack mechanism to support language processing (McElree et al., 2003), and more generally, that humans have great difficulty dealing with genuinely recursive structures. (Earlier versions of ACT-R contained a goal stack which was removed due to mounting evidence that humans do not have perfect memory for previous goal states.) An NLP system that makes use of a stack for language generation might well generate sentences that humans would be unable to process.

On the other hand, there are *right embedded* variants of the sentence above that humans can process:

The dog chased the cat that bit the mouse that ate the cheese

The fact that humans can process such sentences suggests that they are not truly recursive. It is well known that tail recursion in a higher level programming language like Lisp can be converted into iteration by the compiler, thereby reducing the memory demands of the system and improving performance. Given the extremely limited capacity of humans to retain separate chunks of information in awareness—current estimates indicate a capacity of around 4 chunks (Cowan, 2001)—the ability to make sense of such sentences suggests the chunking of the prior input (e.g. “the dog chased the cat”) before processing the subsequent input (e.g. “the cat that bit the mouse”), effectively implementing an iterative processing mechanism to process what most linguists consider to be a recursive structure. If I am right about this, then computational techniques developed in computer science and AI may inform our understanding of how humans process language (*Can artificial intelligence contribute to our understanding of human cognition?*).

Computer science provided the prevailing *computer metaphor of the mind* which has driven an extensive amount of research in cognitive science and cognitive psychology. While most cognitive scientists and

psychologists no longer accept the cognitive validity of this metaphor, much research is still devoted to showing where and how the computer metaphor fails as an explanation of human cognitive behavior. Cognitive scientists and psychologists ignore advances in AI and computer science at their own peril. The advent of *situated cognition* (Clancy, 1997) within AI and the focus on development of robots capable of navigating and communicating in the real world (Mavridis and Roy, 2005) are likely to provide results that will inform and guide cognitive science research. Recent successes in vision research (Umbugh, 1998) and speech recognition (Huckvale, 1997) and synthesis need to be accommodated by the cognitive science community. As computational cognitive models are scaled up to larger and more complex problems and domains, the successes and failures of AI research will become more and more relevant. Computational cognitive scientists will need to adapt the results of AI research to suit their own research agenda—the development of cognitively plausible computational cognitive models (*Can artificial intelligence contribute to our understanding of human cognition?*).

Speech Recognition: at a Crossroads?

Speech recognition systems have achieved phenomenal success in the last decade. Dragon Dictate can transcribe the spoken input of a user on which the system has been sufficiently trained with impressive accuracy. Although already impressive, the quality of speech recognition systems continues to improve. According to X. Huang (in Barker, 2002) word error rates have been reduced an average of 10% each year over the preceding decade and improvements in performance are likely to continue.

Speech recognition systems are typically based on the use of Hidden Markov Modeling and statistical and search techniques which are not cognitively motivated and it is acknowledged that engineering and cognitive science approaches to speech research have significantly diverged (Huckvale, 1998). Do speech recognition systems constitute a counter example to the basic premise of this paper? Despite the successes cited above, the performance of speech recognition systems on large vocabulary, speaker independent, continuous speech in noisy environments falls well short of human performance. While it still remains to be seen if the use of powerful computational and search techniques which are not cognitively motivated will ultimately succeed in matching or exceeding human performance in speech recognition, there is considerable consensus in the cognitive science community that many of the simplifying assumptions of existing systems limit their potential to match human performance. For example, most systems are based on the assumption that speech can be represented as a sequence of phones (a basic unit of speech). However, there is ample evidence of coarticulatory effects between neighboring phones which call this assumption into question. There is also considerable evidence that the syllable and not the phone is the most perceptually salient unit of speech. In addition to providing

an acoustic model based on a string of phones, most speech recognition systems also provide a language model based on either word co-occurrence (bigram or trigram frequency) or some sort of finite state grammar. If a grammar is used, it is typically fully expanded into the language model to support statistically integration with the acoustic model. From the perspective of higher level language processing, the language model of most speech recognizers is overly simplistic and is incapable of representing the full range of structures which occur in unconstrained language. A large amount of research in speech recognition is currently focused on the development of techniques for integrating higher level linguistic knowledge into speech recognition systems. For example, Microsoft has been trying for several years to integrate its speech recognition system with its separately developed NLP system (Barker, 2002). Unfortunately, there is a basic conflict in the way the two systems were developed. The NLP system processes input from right to left starting at the end of a sentence! What seemed like a reasonable engineering decision has precluded the integration of the NLP system with the separately developed speech recognition system. Had the NLP system developers taken cognitive plausibility more seriously, they might have avoided this major disconnect.

But how relevant is cognitively based speech research for the development of speech recognition systems, and vice versa. Most cognitively based models of speech recognition make simplifying assumptions which limit their relevance. In particular, most such models begin with abstract representations of the speech input which circumvent many of the problems faced by speech recognition systems. As Scharenborg et al. (2005) note, a cognitive model which assumes that the input is a sequence of phonemes (a linguistic unit which is an abstraction over phones), sweeps many of the problems of speech recognition under the rug. Scharenborg et al. go on to suggest that cognitively based speech recognition systems should take advantage of research in automatic speech recognition by beginning with the raw acoustic input. They present a model of human spoken-word recognition which does just this. Like Huckvale (1998), they argue for increased communication between researchers who take an engineering or AI approach to speech research, and those who take a human-centered cognitive approach. They claim that this can be achieved by a focus on computational-level (Marr, 1982) descriptions of speech models which are relevant to both AI and cognitive science research in speech, rather than focusing on algorithmic or implementation-level descriptions which are not typically relevant across research areas. Although I agree in spirit with their argument, I am unsure to what extent algorithmic issues can be avoided in descriptions of speech recognition systems. For example, their cognitive model employs competition between competing hypotheses, whereas most speech recognition systems search for the most likely hypothesis without competition. This is a critical algorithmic difference and it may turn out that competition is an unnecessary element of their cognitive model (just as they claim that feedback is an unnecessary element of other

cognitive models since feedback from lexical to pre-lexical levels cannot actually improve pre-lexical perception). Indeed, most speech recognition systems use neither competition nor feedback.

Relationship to This Symposium

To some extent the topic of this symposium is a false dichotomy. AI and Cognitive Science are not distinct fields of research. On one view, AI is one of the Cognitive Sciences. On another view, AI—as the general study of intelligence—encompasses both human and artificial intelligence (Boden, 1990). The history of AI and Cognitive Science is full of cross-disciplinary research, collaboration and contrast. A recent example is the *Modular Construction of Human-Like Intelligence* conference sponsored by AAAI. Individual researchers in AI and Cognitive Science must select their own approach to research—be it the development of functional systems, small scale models or empirical research. At the ends of the spectrum—the development of efficient algorithms and formal theories of computation within AI, versus the focus on empirical study of narrowly circumscribed cognitive phenomena in Cognitive Psychology—the fields do appear distinct. But the middle ground encompassing the development of systems that mimic human behavior and research aimed at explaining human behavior—what might be called “*Cognitive AI*”—provides a fruitful ground for interaction and exchange of ideas.

Of course, there are real challenges to doing research in this middle ground. To the extent that AI research is dominated by research on the development of more efficient algorithms and formal theories (which from my perspective appeared to be the case at AAAI 2004), it will be difficult to get research on the development of cognitively plausible AI systems accepted—especially if those systems are not of the scale expected within the AI community. On the other hand, to the extent that sacrifices in cognitive plausibility are required to develop functional computational cognitive models, and to the extent that such models are not directly supported by empirical studies, it will be difficult to publish in the cognitive modeling and cognitive science community, which expects models to be closely tied to specific empirical results. Although the major venues for publication of AI and Cognitive Science research pose problems, smaller venues like this symposium and the human-like intelligence conference mentioned above do provide opportunities for publication and exchange of ideas.

Perhaps this symposium can be seen as an attempt to raise awareness of the value of research at the intersection of AI and Cognitive Science within the AI community. A similar symposium at the Cognitive Science conference (or in some other Cognitive Science venue) might also be productive. The use of cognitive architectures for the development of large-scale AI systems could provide a basis for grounding the discussion. It should be emphasized that these cognitive architectures have only been available

to support such research for the last decade or so and it has not yet been demonstrated that they scale (TacAir-Soar may be an exception). However, the DARPA Biologically Inspired Cognitive Architectures (BICA) program, which just started and has substantial funding, is directly aimed at extending architectures like Soar and ACT-R to insure they scale and to handle the full range of cognitive and perceptual phenomena needed for the development of functional systems capable of interacting in real world and simulated environments. Enticing AI researchers to explore the potential of these architectures for building functional systems and enticing Cognitive Modelers to scale up their models and focus less on modeling specific empirical results may prove to be a difficult sale, but, if successful, it will reinvigorate the connections between AI and Cognitive Science research.

The key question asked in this paper is whether or not NLP systems can be a cognitive black box. More generally, can AI systems be developed using computational techniques lacking in cognitive plausibility. For an important subset of AI systems—those systems that are necessarily closely tied to human behavior—the answer appears to be “no”. It does not appear to be possible to match the nuances of human behavior in terms of input/output behavior without considering the internal processing mechanisms involved. Master chess programs may outperform most humans, but they fail to pass the Turing test in terms of their input/output behavior. So long as the interest is in developing good chess programs, this is not an issue, but if the goal is to develop a chess program that mimics human behavior, then the processing mechanisms, cognitive constraints and limitations that humans display become relevant. NLP systems are in the subset of AI systems that need to model human behavior closely enough to necessitate consideration of internal processing mechanisms. This does not mean that some abstraction away from the wetware of human cognition and perception is not possible, but that the appropriate level (or levels) of abstraction is not the level of input/output behavior. *What the appropriate level of abstraction is for any particular AI system, is (or should be) a key topic of this symposium.* For those AI systems which need to consider what goes on inside the cognitive black box, the products of research in cognitive science will be relevant and important.

References

- Anderson, J. R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review* 111, 1036-1060.
- Anderson, J. R., & Lebiere, C. (1998). *The Atomic Components of Thought*. Mahway, NJ: LEA.
- Ball, J. (2004). A Cognitively Plausible Model of Language Comprehension. In *Proceedings of the Thirteenth*

- Conference on Behavior Representation in Modeling and Simulation*. ISBN: 1-930638-35-3.
- Ball, J., Rodgers, S., & Gluck, K. (2004). Integrating ACT-R and Cyc in a Large-Scale Model of Language Comprehension for Use in Intelligent Agents. In *Papers from the AAAI Workshop*, Technical Report WS-04-07, 19-25. Menlo Park, CA: AAAI Press.
- Barker, D. (2002). Microsoft Research Spawns a New Era in Speech Technology. *PC AI*, 16:6, 18-27.
- Barsalou, LW (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577-609.
- Boden, M. (ed.) (1990). *The Philosophy of Artificial Intelligence*. New York: Oxford University Press.
- Clancy, W. (1997). *Situated Cognition*. Cambridge, UK: Cambridge University Press.
- Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behavioral and Brain Sciences*: 24(1), 87-114.
- Fellbaum, C. (ed.) (1998). *WordNet, An Electronic Lexical Database*. Cambridge, MA: The MIT Press
- Graesser, A., Gernsbacher, M., & Goldman, S. (2003) (eds.) *Handbook of discourse processes*. Mahwah, NJ: LEA
- Graesser, A., VanLehn, K., Rose, C., Jordan, P., & Harter, D. (2001). Intelligent tutoring systems with conversational dialogue. *AI Magazine*, 22, 39-51.
- Harnad, S. (1992) Connecting Object to Symbol in Modeling Cognition. In: A. Clark and R. Lutz (Eds) *Connectionism in Context*. Springer Verlag, 75-90.
- Harnad, S. (1990). The Symbol Grounding Problem. *Physica D*, 42:335-346.
- Huckvale, M. (1997). 10 things engineers have discovered about speech recognition. Presented at *NATO ASI Workshop on Speech Pattern Processing*.
- Huckvale, M. (1998). Opportunities for Re-Convergence of Engineering and Cognitive Science Accounts of Spoken Word Recognition. In *Proceedings of the Institute of Acoustics Conference on Speech and Hearing*.
- Laird, J., Jones, R. & Nielsen, P. (1998). Lessons learned from TacAir-Soar in STOW-97. In *Proceedings of the Seventh Conference on Computer Generated Forces and Behavioral Representation*. Orlando, FL
- Landauer, T. & Dumais, S. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, 104, 211-240.
- Lehman, J., VanDyke, J., & Rubinoff, R. (1995). Natural language processing for IFORs: comprehension and generation in the air combat domain. In *Proceedings of the Fifth Conference on Computer Generated Forces and Behavioral Representation*
- McElree, B., Foraker, S. & Dyer, L. (2003). Memory structures that subserve sentence comprehension. *Journal of Memory and Language*, 48, 67-91.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Mavridis, N. & Roy, D. (2005). Grounded Situation Models for Robots: Bridging Language, Perception and Action. In *Proceedings of the AAAI-05 workshop on modular construction of human like intelligence* 32-39.
- Pinker, S. & Prince, A. (1988) On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, 28, 73-193.
- Prinz, J. (2002). *Furnishing the Mind*. Cambridge, MA: The MIT Press.
- Rosenbloom, P., Newell, A., & Laird, J. (1993). *The Soar Papers: Research on Integrated Intelligence*. Cambridge, MA: The MIT Press.
- Rumelhart, D. & McClelland, J. (1986). *Parallel Distributed Processing* Vols 1 and 2. Cambridge, MA: The MIT Press.
- Scharenborg, O., Norris, D., ten Bosch, L & McQueen, J. (2005). How Should a Speech Recognizer Work? *Cognitive Science*, 29, 867-918.
- Sun, R. & Alexandre, F., (eds.) (1997). *Connectionist-Symbolic Integration*. Mahway: NJ: LEA.
- Searle, J. R. (1980) Minds, brains and programs. *Behavioral and Brain Sciences* 3: 417-424.
- Townsend, D. & Bever, T. (2001). *Sentence Comprehension*. Cambridge, MA: The MIT Press.
- Umbaugh, S. (1998). *Computer Vision and Image Processing*. Prentice Hall.
- Zwaan, R. (2004). The immersed experiencer: toward an embodied theory of language comprehension. In: B.H. Ross (Ed.), *The Psychology of Learning and Motivation*, Vol. 44 35-62. New York: Academic Press.