

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Can ACT-R Realize "Newell's Dream"?

Permalink

<https://escholarship.org/uc/item/41p1r8g0>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 27(27)

ISSN

1069-7977

Author

Sainz, Javier S.

Publication Date

2005

Peer reviewed

Can ACT-R Realize “Newell’s Dream”?

Matthew F. Rutledge-Taylor (mrtaylo2@connect.carleton.ca)
Institute of Cognitive Science, Carleton University, 1125 Colonel By Drive
Ottawa, Ontario, K1S 5B6 Canada

Abstract

In “The Atomic Components of Thought”, John Anderson and Christian Lebiere claim that ACT-R (4.0) realizes “Newell’s Dream” of a unifying theory of cognition. In this paper it is suggested that each ACT-R model can account for only a finite set of cognitive processes, and cannot therefore be used to model an unbounded whole mind. It is suggested that this is due to an inherent context dependence of ACT-R models. This limitation runs counter to the intuitive criterion that a unifying theory of cognition ought to be able to provide an account of the mind as a system not bound to any particular context. It is suggested that thought in cognitive models ought to be conceived as temporary context specific operations based on persistent context independent knowledge. The basis for a new cognitive architecture, which differentiates thought from knowledge is proposed. This new architecture combines ACT-R with elements of Lawrence Barsalou’s situated simulation theory.

Keywords: ACT-R; cognitive architectures; knowledge representation; situated cognition; perceptual symbols systems.

Introduction

In 1972, at the Carnegie Symposium on Cognition, Allan Newell raised a concern about the course of research in psychology. He delivered a paper entitled “You can’t play 20 questions with nature and win”, in which he lamented the fact that there was very little that unified the wealth of knowledge that had been accumulated about individual human cognitive processes (1973a). In a paper published separately in the proceedings of the symposium, Newell suggested that production systems might serve as detailed models of the human control structure (1973b; 1990). Eighteen years later, Newell published a book entitled “Unified Theories of Cognition” in which he proposed that cognitive architectures hold the key to unifying psychology (1990).

There are many different cognitive architectures used to produce cognitive models of psychological phenomena, including Newell’s own SOAR architecture, which was first released in 1982 (Laird & Rosenbloom, 1996). However, the most popular architecture is ACT-R (Anderson, 1993; Anderson & Lebiere, 1998). This popularity is by no means accidental. The theory of cognition it implements has been well developed, and hence, has allowed a wide variety of researchers to produce theoretically grounded models of various cognitive phenomena. Additionally, and perhaps most significantly, ACT-R models typically fit the human experimental data they are designed to model, quite well.

This paper considers whether ACT-R, as it currently exists, realizes “Newell’s Dream” of a unifying theory of cognition, or not. It is suggested that given appropriately strict criteria ACT-R may be inadequate.

Newell’s Criteria

According to Newell: a theory is an explicit body of knowledge, from which answers to questions of a predictive, explanatory, or prescriptive type can be given; theories are approximate; theories cumulate; and, theories develop iteratively (1990, pp. 13-14). Newell defines a unified theory of cognition as “a single set of mechanisms for all of cognitive behavior” (1990, p. 15). He specifies these mechanisms as a prioritized list of areas of cognitive phenomena to be covered. They are, in order: problem solving, decision making, and routine action; memory, learning, and skill; perception, and motor behaviour; language; motivation, and emotion; and, imagining, dreaming, and, daydreaming. Thus, a complete unified theory of cognition, should account for all of these cognitive phenomena. However, given Newell’s views on theory development, an acceptable strategy would be to begin with a unified theory of the phenomena at the top of the list, and slowly augment the theory so as to accommodate successive items.

Background Theory

Ubiquitous in cognitive science is the view that cognitive systems can be analysed from a variety of perspectives. The tri-level hypothesis is that there are three basic levels of analysis (Dawson, 1998). Various researchers apply their own labels to these three levels. Newell divided them into the biological, cognitive, and rational (Newell, 1990); Zenon Pylyshyn makes use of the physical, syntactic, and semantic (Pylyshyn, 1999); and, Michael Dawson, the implementational, algorithmic, and computational levels (Dawson, 1998). Despite the difference in terms, there is arguably an equivalence between these hierarchies. The biological, physical, and implementational levels are, in the case of humans, the levels of description that (typically) appeal primarily to neural processes. The cognitive, syntactic, and algorithmic levels, describe human cognition in terms of operations on syntactic (or, otherwise, formal) structures. The rational, semantic, and computational levels are those at which the cognitive system is described in terms of its knowledge (i.e., goals, beliefs, and perceptions etc.).

Pylyshyn asserts that a fundamental hypothesis in cognitive science is that this knowledge level is an autonomous (or, at least, partially autonomous) level of

description (1998, p. 4). That is, it is autonomous from the neural level of description. This conclusion depends on another belief that is ubiquitous to cognitive science: the computational theory of mind. The computational theory of mind is that the mind is a kind of computer with some functional equivalence to a universal Turing machine. It was Newell, along with his longtime collaborator, Herbert Simon, who asserted that physical symbol systems have the necessary and sufficient means for general intelligence (Newell & Simon, 1976). Pylyshyn (1999) suggests that it is the task of cognitive science to discover the details of the mechanisms that support mental computation; i.e., to determine what kind of computer the human mind is. Thus, it is the task of cognitive science to discover the cognitive architecture of the mind.

A cognitive (or, as it is sometimes referred, functional) architecture is a bridging of the biological and cognitive levels of description, whereby the computational primitives of human cognition are defined. Cognitive architectures are analogous to computer architectures in that both describe the physical constraints (e.g., memory capacity) on the algorithms that run in the architecture. Thus, the functional details of the cognitive architecture are constrained by the underlying physical system implementing the architecture. According to Pylyshyn (1998), facts about the knowledge level depend only on the functional details of architecture, and not on the details of how the architecture is implemented. This paper devotes attention to a particular cognitive architecture: ACT-R.

ACT-R

ACT (adaptive control of thought) theory was originally proposed by John Anderson in 1976. ACT was essentially a marriage of Anderson and Bower's existing model of declarative memory, HAM (1973), and a production system based on Newell's proposal (1973b). ACT evolved into ACT* (1983), then into ACT-R in 1993. ACT-R has undergone several significant revisions. In 1998, the release of ACT-R 4.0 coincided with the publication of the book "The Atomic Components of Thought" (Anderson & Lebiere, 1998). In the opinion of Anderson and Lebiere (1998), ACT-R 4.0 was the first version of ACT-R to legitimately realize "Newell's dream" of a unified theory of cognition.

ACT-R is a cognitive architecture implemented in LISP. ACT-R is also a theory of cognition. According to the designers of ACT-R, the atomic components of thought are chunks, which exist in declarative memory and production rules, which exist in procedural memory. A chunk is an independent pattern of information corresponding to a thing that one can be aware that one knows. Each chunk consists entirely of several slots (variables) with associated values. The slot values are either chunks themselves, or atomic features. The designers of ACT-R assert that chunks should have, on average, three or four slots, one of which must be an ISA slot. The ISA slot determines to what ontological type a particular chunk token corresponds (e.g., john841 is a person chunk). Respecting George Miller's "magic

number", chunks should have no more than seven (plus or minus two) slots (Miller, 1956).

Chunks originate primarily from two sources: perceived objects in the environment and records of solutions to past problems. Production rules specify how to retrieve and use chunks to solve problems. Each production rule consists of an if-then condition-action pair. The ISA value of chunks plays an integral role in determining whether a chunk matches conditions in the production. New production rules can be generated from chunks in memory via a process called production compilation.

In this paper, chunks will be written in the format [Chunk-Name: ISA Chunk-Type; slot-1-label slot-1-value; slot-2-label slot-2-value; etc.]. For example, an ACT-R chunk representing the knowledge that $3 + 4 = 7$ could be encoded [Fact3+4: ISA addition-fact; addend1 Three; addend2 Four; sum Seven]

Newell's criteria revisited

I suggest that whether ACT-R realizes "Newell's dream", or not, is unclear. As of version 4.0, ACT-R has been used to produce successful models of the cognitive phenomena in the first two groups of areas identified by Newell (i.e., problem solving, decision making, routine action, memory, learning, and skill). And, with the development of ACT-R 5.0 and a perceptual-motor system, ACT-R/PM (Byrne & Anderson, 2001), the areas in the third group (perception, and motor behaviour) have now been added to the list of ACT-R's successes. I take the adequacy of ACT-R models of the remaining areas identified by Newell to be, at a minimum, somewhat uncertain.

The question remains: in what way is it unclear whether ACT-R realizes Newell's dream. The problem is with respect to what counts as meeting Newell's criteria. Newell asserts that a unified theory of cognition should have a single story to tell about the various areas he identifies (Newell, 1990, p. 15), which is consistent with all previously existing psychological knowledge of each area of study (Newell, 1990, p. 16). Interestingly, from a modeling perspective, unifying theories can take two distinct forms, with two different criteria. I will define here my weak and strong criteria for a universal theory of cognition qua modeling toolkit. The weak criterion is that a unified theory of cognition should provide a common set of tools and principles such that for any psychological phenomenon a model that accounts for the phenomenon can be created. This criterion does not imply that any pair of cognitive phenomena can be reconciled within a single model existing within such a unified theory of cognition. Hence, the strong criterion is that a unified theory of cognition should provide a common set of tools and principles such that a single model that accounts for every psychological phenomenon can be created. Simply, the strong criterion is met if and only if (in principle) a single model of the whole mind can be created. Additionally, it should be noted that I use the term "model" in the sense of a tokened instance of a type defined within the parameter space of a particular cognitive architecture. ACT-R has enjoyed a great deal of success in

meeting the weak criterion¹. However, I do not know of any attempts to satisfy the strong criterion, nor, have I come across any explicit reference to a distinction between the weak and strong criteria, in the ACT-R literature.

I suggest that in order to realize the spirit of Newell's dream, the strong criterion must be met. Cognitive science is about understanding the mind as a unified single entity, not as a finite collection of phenomena, or abilities. Hence, to be able to build a model of an entire mind should be our goal.

Computational primitives: The atomic components of thought

A computational primitive is a non-decomposable atom of thought. Although, there is in principle a distinction between an atom of thought and an atom of knowledge, there is no such distinction in ACT-R. Chunks are the atoms of declarative knowledge, and productions are the atoms of procedural knowledge. In ACT-R, declarative memory is an associative network of persistent chunks, each of which has a particular activation level. When the cognitive system engages in a cognitive process, the very same chunks that exist in memory are the ones recruited to take part in the process. Thus, according to ACT-R, the grain-size of one's basic level of knowledge of the world is identical to the grain-size of knowledge as it is used in active thought.

ACT-R is designed to model individual experimental tasks. Anderson and Lebiere write: "every ACT-R model corresponds to a subject performing in some particular experiment" (1998, p. 15); and, "most ACT-R models assume a system that starts out with substantial *relevant* knowledge, as is the case for the typical undergraduate subject" (1998, p. 16, my emphasis). If what is meant by relevant, is "task specific", then there may be issues with respect to the nature of persistent knowledge in declarative memory? A principle that contributed to the development of knowledge representation in ACT-R is that for every known entity in the world, there should be only one chunk in memory. In reference to an early version of ACT where nodes can be interpreted as chunks, Anderson wrote: "these memory systems assume only one node per individual, reflecting the spatio-temporal continuity of that individual" (1977, p. 430). Additionally, "there are no special 'subtoken' nodes connected to the individual to represent different aspects of the individual" (Anderson, 1977, p. 430). These constraints are presumably motivated by both theoretical and practical considerations. From a theoretical perspective, as suggested, Anderson seems committed to a sort of isomorphism between ontologically distinct entities in the world and representations of them in the mind. From a practical perspective, it is much easier to retain consistency in a knowledge base when there is no redundancy in representation.

The one-entity-one-chunk constraint seems to eliminate the possibility of distinct task specific chunks. Therefore, if ACT-R models are to scale up to modeling the entire mind,

chunks will have to encode information in a neutral manner that can be employed during a variety of tasks. Thus, chunks should not admit of context specific features, nor, should they admit of vague values. The reason for this latter condition is that it can be argued that vague predicates have no meaning outside of particular contexts. In addition, even within a given context a vague predicate can fail to have determinate truth-values. Although the relevance of a predicate's truth conditions to its meaning is a matter of some debate, I will assume that the two notions are, at a minimum, related.

This raises the question, why is it the case that vague predicates have been employed in atomic representations within ACT-R without difficulty, so far. Although, vagueness is different from context dependence (e.g., "Tim is above average in height" is context dependent, but not vague), what a vague predicate asserts of an entity depends on context (e.g., "Tim is tall" makes a different claim about Tim's height in the context of a set of basketball players as compared to a set of racehorse jockeys). Therefore, the truth-values of vague predications depend on more than just the entity of which the predicate is applied. It depends on, in the very least, an implied comparison set. Such an implied comparison set is furnished by the particular context in which the given ACT-R model exists. However, chunks that contribute to a general knowledge base ought not to be context dependent, if there is to be no redundancy in representation of knowledge (i.e., the casting of some fact about the world as multiple chunks, one for every possibly relevant context would entail some redundancy).

The chunk capacity problem

There is an issue related to which slot values ought to be associated with a particular chunk. It is obvious that my most basic knowledge of a whole entity includes more than a few features. Even of some stranger I observe walking down the street, I can perceive: Their apparent gender, ethnicity, mood, their basic physical features, what they are wearing, of whom their appearance reminds me, and so on. There are clearly more than seven plus or minus two features I may need to encode in my person chunk type². The problem is greater for people of whom I actually have a plethora of knowledge.

ACT-R provides two mechanisms for potentially avoiding this problem. The first is inheritance. When defining a chunk type, a superordinate chunk from which the new chunk inherits slots, may be specified. For example, my chunk token representing my cat Cougar could be of the species chunk type domestic cat, which inherits from the family type Felidae, which inherits from the order type Carnivora, which inherits from the class type Mammalia, and so on. By this system, each chunk type in the hierarchy would specify up to seven slot values representing things people know about examples of those types. Thus, every chunk can have up to 7n slots, where n is the chunk's rank in the hierarchy. This does not necessarily violate Miller's magic number constraint. For example, productions could

¹ Visit <http://act-r.psy.cmu.edu/publications/> for list of publications on ACT-R, a majority of which present ACT-R models of cognitive phenomena.

² I will rhetorically refer to myself as an ACT-R model of a human throughout this paper.

be limited to specifying a maximum of seven slot values per chunk for the purposes of goal matching and memory retrieval etcetera. Thus, only the seven most relevant aspects of a chunk would be featured in a particular cognitive process. This appears to be a clever way to accommodate context in ACT-R models.

The problem with this solution is the nature of inheritance in ACT-R. Each chunk may inherit from only one superordinate chunk. Thus, every chunk type can participate in only one conceptual hierarchy. The effect of this is that the definition of each chunk type is limited to only one epistemic stance. For example, if my “Cougar” chunk inherits from my domestic cat chunk type, which would place it in a biological hierarchy, it cannot also inherit from my pet chunk type, which would place it in a different (perhaps functional) hierarchy. Pace the standard ACT-R doctrine, it seems reasonable that an atomic component of knowledge ought to be able to participate in a variety of conceptual hierarchies, one for each unique epistemic perspective that can be taken of the thing represented by that atomic component of knowledge. The utility of this idea evidenced by the variety of computer programming languages that permit multiple-inheritance in object type definition such as C++ and (the increasingly popular) python programming language³.

The second mechanism provided by ACT-R for avoiding the problem of trying to accommodate all knowledge of a thing into a single chunk, is to create chunks that encode relational knowledge. Any logical expression can be encoded in a chunk. For example, “John loves Mary”, might be encoded by the chunk [John-loves-Mary: ISA relation; relation love; agent John; theme marry]

The principle problem with this solution is the proliferation of chunks in declarative memory. That is, there would be thousands, if not millions of relational chunks in memory, one for every relation that one had ever entertained. This is a problem, because not all of these relations are independent of one another. Specifically, changes to some relations should affect others. For example, if I learn that that John has died, every (or at least most) chunk(s) relevant to him must be updated. Unfortunately, ACT-R does not provide an easy way to modify all chunks relevant to a particular other chunk; in ACT-R, qualitative knowledge is only modified by production firings. To update every chunk of which John is a slot value, a separate retrieval attempt for every combination of chunk type and slot in which John could potentially occur must be made. Additionally, there would probably have to be a unique production for updating each kind of chunk. For example, chunks [John-wrote-the-book-10-ways-to-catch-a-sparrow: ISA authorship-relation; agent John; theme ‘10 way to catch a sparrow’], [John-is-married-to-Hilary: ISA marriage-relation; person1 John; person2 Hilary], and, [John-is-smelly: ISA an-odour-fact; individual John; scent-type smelly] should be updated differently, if at all. To make matters worse, this speculation about the means of updating chunks is not even possible! In the case of real minds, new knowledge of this sort is, more-or-less, seamlessly integrated into the knowledge base. In contrast,

ACT-R productions fire serially, and require a minimum of 50 ms each to fire. Thus, it would take an ACT-R model of the mind hours to complete firing the potentially thousands of productions required to update it’s knowledge base. ACT-R is not designed to operate over a large knowledge base, in real time.

Dreyfus, GOFAI, and ACT-R

In general, ACT-R falls victim to many of the criticisms made against traditional (good old fashioned) A.I. systems. Limiting current ACT-R models to individual experimental scenarios is akin to limiting Terry Winograd’s SHRDLU (Winograd, 1972) to interacting with a drastically circumscribed blocks micro-world. Hubert Dreyfus argued that SHRDLU failed as a model of understanding because understanding is a concept whose meaning depends on a vast network of other human concepts, which are absent from the block micro-world (Dreyfus, 1979). Dreyfus writes:

In our everyday life we are, indeed, involved in such various ‘sub-worlds’ as the world of the theatre, of business, or of mathematics, but each of these is a ‘mode’ of our shared everyday world. That is, sub-worlds are not related like isolable physical systems to larger systems they compose; rather they are local elaborations of a whole which they presuppose... Since,..., micro-worlds are not worlds, there is no way they can be combined and extended to the world of everyday (1979, p. 151).

This criticism is consistent with the theme of my weak and strong criterion distinction. No collection of ACT-R models, each performing in it is own micro-world, will collectively inform us about how a single human mind accomplishes each of these tasks in our ‘everyday’ world.

Dreyfus was an early advocate of the notion of situated cognition, which continues to be researched by the likes of Andy Clark (1997), among others. The similarity in structure of ACT-R chunks and Marvin Minsky’s frames makes Dreyfus’ attack on Minsky’s Frame theory (1974) relevant to the current topic:

No piece of equipment makes sense by itself... What makes an object a chair is its function, and what makes possible its role as equipment for sitting is its place in a total context... There is no argument why we should expect to find elementary context-free features characterizing a chair type, nor any suggestion as to what these features might be. They certainly cannot be legs, back, seat, and so on, since these are not context-free characteristics defined apart from chairs which then ‘cluster’ in a chair representation (Dreyfus, 1979, p. 163-4).

Dreyfus is questioning if it is even in principle possible provide adequate formal criteria for our everyday human concepts. This skepticism strikes directly at the basis of ACT-R’s knowledge representation scheme.

³See <http://www.python.org/>

A different kind of cognitive architecture

I have argued that a network of ACT-R chunks cannot form the basis for humans' basic level of knowledge of the world. However, it seems entirely plausible that when we actually engage in consciously accessible cognitive processes, ACT-R has the grain-size right, vagueness and all. What I suggest is that ACT-R is a good candidate for a model of high-level cognitive processing which sits on top of a low-level dynamical system responsible for encoding knowledge in a holistic manner.

A researcher who has taken a great interest in the prospect of modal knowledge representation is Lawrence Barsalou also known for his work on ad-hoc category representations (Barsalou, 1983). Barsalou suggests that perceptual symbol systems via his situated simulation theory underlie much of our conceptual knowledge (Barsalou, 1999; Barsalou, 2003; Barsalou, Simmons, Barbey, Wilson, 2003). Barsalou writes, "A concept is not a single abstracted representation for a category, but is instead a skill for constructing idiosyncratic representations tailored to the current needs of situated action" (2003, p.521). An example to which Barsalou theory speaks is the use of visual imagery.

In 1978, Steven Kosslyn published research on visual imagery (Kosslyn, et al., 1978). Briefly, participants studied a map with various landmarks; with their eyes closed, they were asked to imagine the map and focus on a particular landmark; a different landmark was named, and the participant was instructed to hit a button once they could 'see' the second landmark. It was found that the time taken to 'see' the second landmark was proportional to the distance between the landmarks on the map. Thus, Kosslyn concluded that people were using mental images, via the 'mind's eye' to perform the task. These findings are consistent with the idea that knowledge about visual information is stored experientially in the visual cortex. According to Barsalou's theory, when the subject is asked a question about the map, they simulate seeing the map and respond in a manner similar to how they would if they were viewing the original map. They do not rely on explicit symbolic long-term memory.

Pylyshyn responded to the results of Kosslyn's classic 'mind's eye' experiment, by suggesting that spatial information can be encoded symbolically (Pylyshyn, 1999). Pylyshyn argued that the correlation between time and distance in the results could merely be an artifact of the 'imagining' process and not due to operations on a spatial mental representation of the map in the participants' minds. This argument is supported by the fact that in a response study, Pylyshyn was able to show that people could answer questions about aspects of a map instantaneously, if the subjects were not asked to imagine traversing the map (Pylyshyn, 1999). This suggests that it is at least possible that a declarative symbolic representational system underlies one's ability to engage in mental imagery.

Pylyshyn's counter-example suggests only that people have the capacity to encode visual information in a declarative manner, and not that people systematically do so. The use of examples of visual imagery may be misleading. People are largely very good at interpreting visual

information declaratively. However, a much smaller proportion of the general population can do so with auditory information. For example, consider the case of pitch comparison. It is not obvious that people who do not have any musical training have very detailed declarative knowledge of the music to songs they know. When asked to determine whether the note associated with "dash" in the song jingle bells is higher or lower than the note associated with "sleigh", most people have no alternative but to "simulate" singing the song to themselves and make an online comparison between the two notes via auditory imagery (Zatorre, 2004). Barsalou would argue that, fundamentally, our knowledge of the song jingle bells resides in our auditory sensory system, and not in an amodal memory system such as ACT-R's declarative memory system.

Barsalou's theory that all conceptual knowledge is modal may be too extreme. However, is it possible that a symbolic representation underlies my ability to sing a song to myself, but is not flexible enough for me to directly access information about the notes associated with each word in the song? What seems more likely is that people can know jingle-bells in at least two different ways. The non-musically trained know it only via the experience of hearing and singing it. The musically trained can have explicit knowledge of the musical notation associated with the song. Interestingly, it is well known that Beethoven was deaf during the time period that he composed some of his best known works. Clearly, Beethoven had never heard his later compositions. However, he may very well have been able to simulate the experience of hearing his music by directly translating his explicit knowledge of the sheet music associated with a given work to a phenomenal (internal) experience of music, via the mind's ear.

I suggest that Barsalou's situated simulation theory can be reconciled with ACT-R. Chunks would be abandoned as a form of atomic knowledge representation (as opposed to atomic thought representation as discussed above). Modular perceptual symbol systems would take on the burden of encoding knowledge at the lowest level relevant to cognitive psychological research. The perceptual symbol systems would be augmented with a sort of amodal knowledge system which underlies people's declarative knowledge base. ACT-R would sit on top of these perceptual symbol systems. In a given context the perceptual symbol systems and the amodal system would generate, on the fly, a set of temporary task specific chunks and productions which would exist only while the system is attending to a particular cognitive task. The temporary chunks generated would only serve as atomic components of (current) thought, and no longer as atomic components of (persistent) knowledge.

This proposed unification of Barsalou's situated simulation theory with ACT-R is an example a hybrid cognitive architecture which I take to be necessary for building models of the whole mind. There is a sort of autonomy of high-level cognitive processes from the underlying physics. However, the autonomy of a transient ACT-R model of a task is due to the fact that it exists only as an abstraction from the underlying knowledge base of the system as a whole.

ACT-R 6

ACT-R is an evolving architecture. It has undergone significant revisions and changes over the years, and soon, the next version, ACT-R 6 will be released. ACT-R 6 will incorporate several changes to ACT-R 5. One change that is relevant to my concern with the current version of ACT-R, is increased modularity (Bothell, Byrne, Lebiere, Taatgen, 2004). In previous versions of ACT-R, all cognitive processes were implemented as productions operating on the contents of the goal buffer, which is a place holder for a single active declarative chunk. This necessarily serialized all high level cognition. ACT-R 6 will allow for multiple modules each of which will have its own declarative chunk buffer. This change should increase the amount of parallelism in ACT-R somewhat. However, more significantly, this change allows for a degree of autonomy of some cognitive processes (e.g., visual perception) from operations on the goal buffer. This decentralisation is a necessary step in adjusting the ACT-R paradigm towards the hybrid architecture that I have suggested.

This current trend in the evolution of ACT-R may naturally lead it towards solutions to the problems I have identified with the architecture, which may be similar to those that I have suggested. Whether future changes, less radical than what I have proposed will suffice, is yet to be seen.

References

- Anderson, J. R. (1977). Memory for information about individuals. *Memory and Cognition*, 5, 430-442.
- Anderson, J. R., & Lebiere, C. (1998) *The Atomic Components of Thought*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Barsalou, L. W. (1983). Ad hoc categories. *Memory & Perception*, 11, 211-227.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577-609.
- Barsalou, L. W. (2003). Situated simulation in the human conceptual system. *Language and Cognitive Processes*, 18, 513-562.
- Barsalou, L. W., Simmons, W. K., Barbey, A. K., and Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *TRENDS in Cognitive Science*, 7, 84-91.
- Bothell, D., Byrne, M. D., Lebiere, C., & Taatgen, N. A. (2004). ACT-R 6 Proposals. <<http://act-r.psy.cmu.edu/act-r6/ACT-R6proposal.pdf>>
- Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: The MIT Press.
- Dawson, M. R. W. (1998). *Understanding cognitive science*. Malden, MA: Blackwell.
- Dreyfus, H. L. (1979). From micro-worlds to knowledge representation: AI at an impasse. In J. Haugeland (Ed.), *Mind design II* (pp. 143-182). Cambridge, MA: The MIT Press.
- Kosslyn, S. M., Ball, T. M., & Reiser, B. J. (1978) Visual images preserve metric spatial information: Evidence from studies of image scanning. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 46-60.
- Laird, J. E., & Rosenbloom, P. S. (1996). The evolution of the soar cognitive architecture. In Steier, D. M., & Mitchell, T. M. (Eds.), *Mind Matters: A tribute to Allen Newell* (pp. 1-50). Mahwah, NJ: Erlbaum.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81-97.
- Minsky, M. (1974). A framework for representing knowledge. In J. Haugeland (Ed.), *Mind design II* (pp. 111-142). Cambridge, MA: The MIT Press.
- Newell, A. (1973a). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual Information Processing* (pp. 283-310). New York: Academic Press.
- Newell, A. (1973b). Production systems: Models of control structures. In W. G. Chase (Ed.), *Visual Information Processing* (pp. 463-526). New York: Academic Press.
- Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.
- Newell, A. & Simon, H. (1976). Computer science as empirical enquiry: Symbols and search. In J. Haugeland (Ed.), *Mind design II* (pp. 81-110). Cambridge, MA: The MIT Press.
- Pylyshyn, Z. W. (1998). Introduction: Cognitive architecture and the hope for a science of cognition. In Z. W. Pylyshyn (Ed.), *Constraining Cognitive Theories* (pp. 1-7). Stamford, CT: Ablex Publishing Corporation.
- Pylyshyn, Z. W. (1999). What's in your mind. In: E. Lepore & Z. Pylyshyn (Eds.), *Rutgers Lectures in Cognitive Science*. Blackwell
- Winograd, T. (1972). Understanding natural language. *Cognitive Psychology*, 1, 1-191.
- Zatorre, R. (2004, November 26). Processing of music and speech by the human auditory cortex: Neuroimaging evidence. [Carleton Cognitive Science Research Seminar.] Ottawa, ON.