A Standard Model of the Mind

Christian Lebiere (cl@cmu.edu)
Psychology Department
Carnegie Mellon University

You Can't Play 20 Questions…

- Large and growing number of phenomena
- Binary oppositions never seem to get resolved
- Divide-and-conquer strategy not working
- Model control structure
- ***Put them all together***
  - Complete process models
  - Analyze complex task
  - 1 program for many tasks

YOU CAN'T PLAY 20 QUESTIONS WITH NATURE AND WIN: PROJECTIVE COMMENTS ON THE PAPERS OF THIS SYMPOSIUM

Allen Newell
Carnegie-Mellon University

I am a man who is half and half. Half of me is half distressed and half confused. Half of me is quite content and clear on where we are going.

My confused and distressed half has been roused by my assignment to comment on the papers of this symposium. It is curious that it should be so. We have just listened to a sample of the best work in current experimental psychology. For instance, the beautifully symmetric RT data of Cooper and Shepard (Chapter 3) make me positively envious. It is a pleasure to watch Dave Klahr (Chapter 1) clean up the subitizing data. The demonstrations of Bransford and Johnson (Chapter 8) produce a special sort of impact. And so it goes. Furthermore, independent of the particular papers presented here, the speakers constitute a large proportion of my all-time favorite experimenters--Chase, Clark, Posner, Shepard. Not only this, but almost all of the material shown here serves to further a view of man as a processor of information, agreeing with my current theoretical disposition. Half of me is ecstatic.

Still, I am distressed. I can illustrate it by the way I was going to start my comments, though I could not in fact bring myself to do so. I was going to draw a line on the blackboard and, picking one of the speakers of the day at random, note on the line the time at which he got his PhD and the current time (in mid-career). Then, taking his total production of papers like those in the present symposium, I was going to compute a rate of productivity of such excellent work. Moving, finally, to the date of my chosen target's retirement, I was going to compute the total

You Can't Play 20 Architectures…

- The branding problem in cognitive architectures
  - Credit assignment goes to architecture "brand", not collection of mechanisms, hindering understanding
  - Evolving architectures leave it unclear what set of incremental results can be claimed unless re-validated
  - Modeling efforts might still not be truly incremental to achieve "the ingredients of accumulation"
- 20 Architectures - 3 possible outcomes
  - The field remains mired in conflicting accounts
  - One architecture becomes the established account
  - A ***consensus*** emerges, implicitly or explicitly (AAAI FS IC13)

A Standard Model of the Mind

Laird, J. E., Lebiere, C. & Rosenbloom, P. S. (in press). A Standard Model of the Mind: Toward a Common Computational Framework across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. *AI Magazine.*

- AI, cognitive science, neuroscience and robotics all concerned with understanding forms of minds
- Develop ***Standard Model*** of human-like minds
  - Coherent baseline for shared cumulative progress
  - Focus research efforts in most crucial areas
  - Improve communication within and between areas
  - Serve as integrative framework across areas
- What it is not intended to be
  - Complete or definitive
  - A single reference implementation

ACT-R, Soar and Sigma

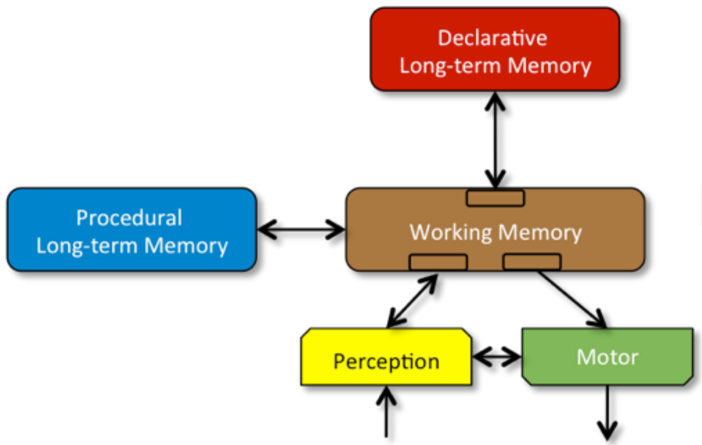

Figure 5. The structure of the standard model.

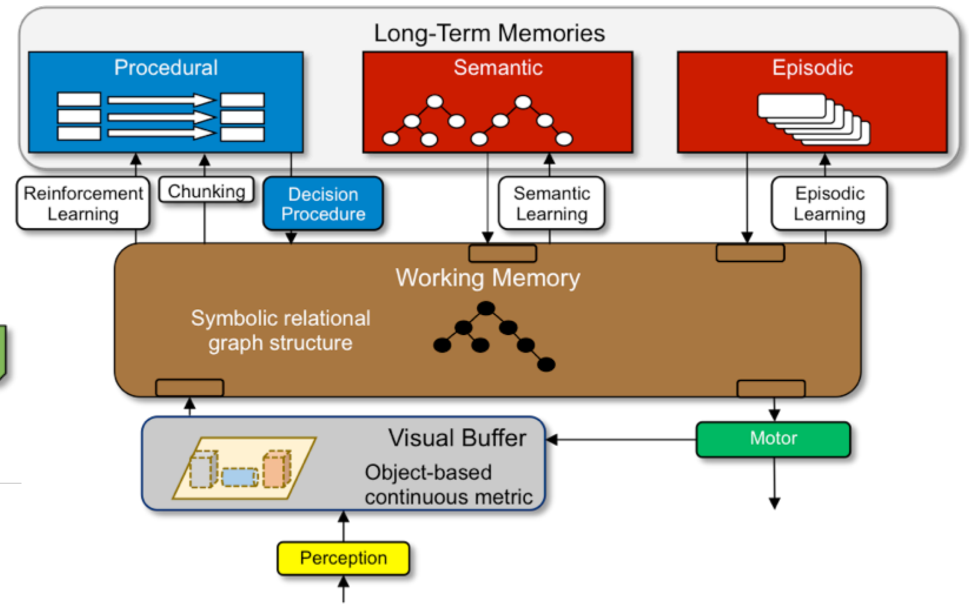

Figure 3. Soar cognitive architecture.

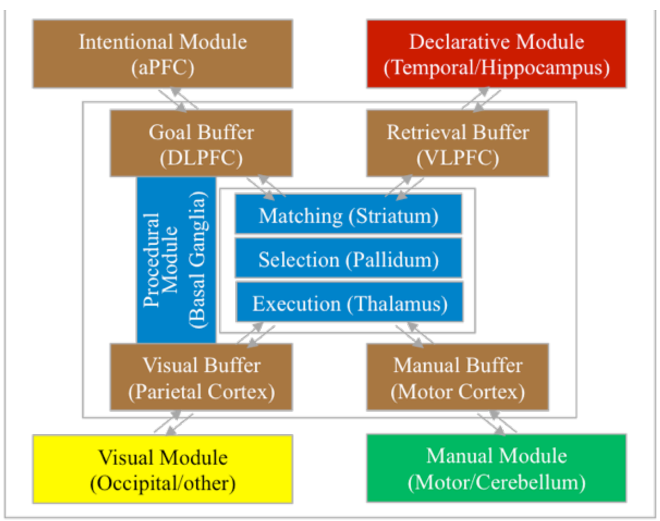

Figure 2. ACT-R cognitive architecture.

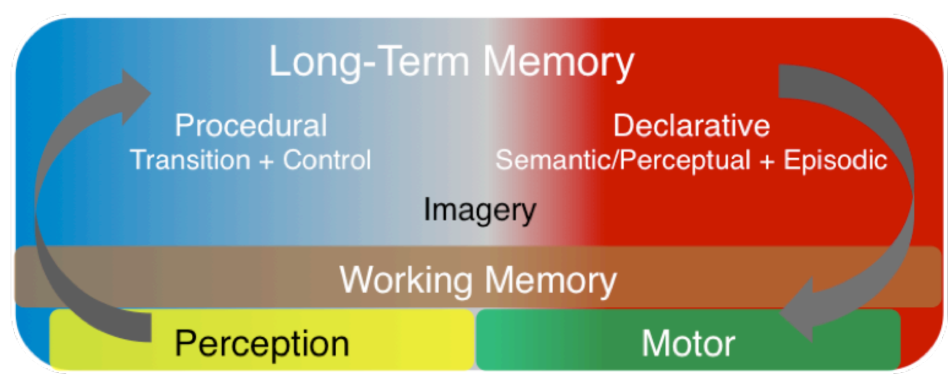

Figure 4. Sigma cognitive architecture.

A. Structure and Processing

1. The purpose of architectural processing is to support bounded rationality, not optimality
2. Processing is based on a small number of task-independent modules
3. There is significant parallelism in architectural processing
    a. Processing is parallel across modules
        i. ACT-R & Soar: asynchronous; Sigma: synchronous
    b. Processing is parallel within modules
        i. ACT-R: rule match, Sigma: graph solution, Soar: rule firings
4. Behavior is driven by sequential action selection via a cognitive cycle that runs at ~50 ms per cycle in human cognition
5. Complex behavior arises from a sequence of independent cognitive cycles that operate in their local context, without a separate architectural module for global optimization (or planning).

B. Memory and Content

1. Declarative and procedural long-term memories contain symbol structures and associated quantitative metadata
    a. ACT-R: chunks with activations and rules with utilities; Sigma: predicates and conditionals with functions; Soar: triples with activations and rules with utilities
2. Global communication is provided by a short-term working memory across all cognitive, perceptual, and motor modules
3. Global control is provided by procedural long-term memory
    a. Composed of rule-like conditions and actions
    b. Exerts control by altering contents of working memory
4. Factual knowledge is provided by declarative long-term memory
    a. ACT-R: single declarative memory; Sigma: unifies with procedural memory; Soar: semantic and episodic memories

C. Learning

1. All forms of long-term memory content, whether symbol structures or quantitative metadata, are learnable
2. Learning occurs online and incrementally, as a side effect of performance and is often based on an inversion of the flow of information from performance
3. Procedural learning involves at least reinforcement learning and procedural composition
    a. Reinforcement learning yields weights over action selection
    b. Procedural composition yields behavioral automatization
        i. ACT-R: rule composition; Sigma: under development; Soar: chunking
4. Declarative learning involves the acquisition of facts and tuning of their metadata
5. More complex forms of learning involve combinations of the fixed set of simpler forms of learning

D. Perception and Motor

1. Perception yields symbol structures with associated metadata in specific working memory buffers
    a. There can be many different such perception modules, each with input from a different modality and its own buffer
    b. Perceptual learning acquires new patterns and tunes existing ones
    c. An attentional bottleneck constrains the amount of information that becomes available in working memory
    d. Perception can be influenced by top-down information provided from working memory
2. Motor control converts symbolic relational structures in its buffers into external actions
    a. As with perception, there can be multiple such motor modules
    b. Motor learning acquires new action patterns and tunes existing ones

Table 1: Standard model architectural assumptions

Progress

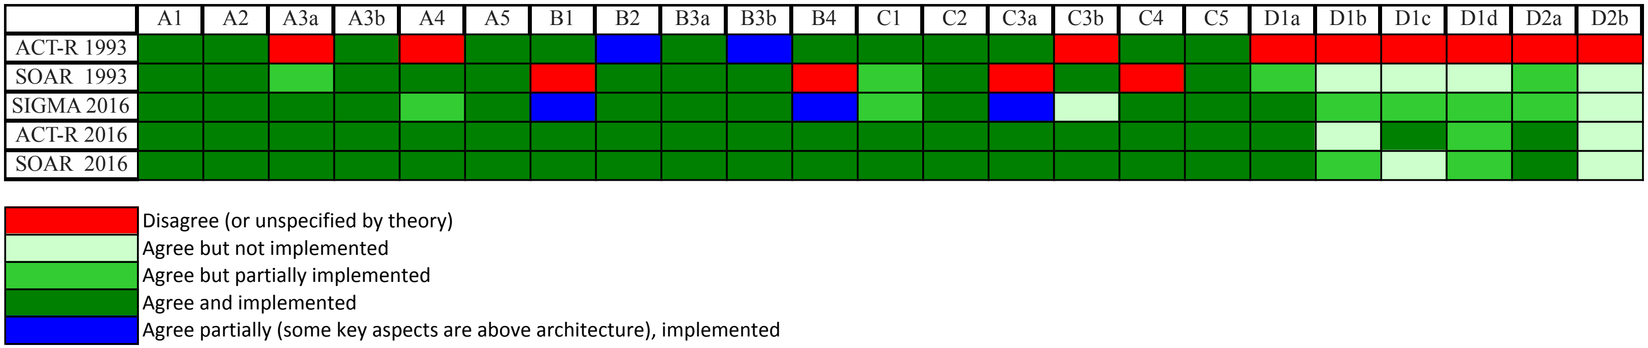

- Tactical approach to start with coherent subset
- Degree of progress is both substantial and limited
- Extend to increasingly broader class of architectures
- Extend to other levels above and below deliberate act
- Fundamentally interactive, ***social***, bottom up process
- AAAI 2017 Fall Symposium CFP http://sm.ict.usc.edu

Your Mind is Not a Toaster

- 2 levels of organization
  - Top level is assembly of a few interacting pieces
  - Bottom level is uniform substrate of many identical elements (atoms, chunks)
- A mind has intermediate levels of organization
  - Mind holds information
  - ***Content has structure***
  - Properties like reusability and composability

The Architecture of Complexity

- Many complex systems with large number of interacting parts
- ***Complexity as hierarchy***
- 2 watchmakers parable
- Emergence of hierarchy
- Dynamics of hierarchies
- Span of control: flat vs deep hierarchies
- Also: Newell levels, JRA 7 orders of magnitude

THE ARCHITECTURE OF COMPLEXITY

HERBERT A. SIMON*

Professor of Administration, Carnegie Institute of Technology

(Read April 26, 1962)

A NUMBER of proposals have been advanced in recent years for the development of "general systems theory" which, abstracting from properties peculiar to physical, biological, or social systems, would be applicable to all of them.[1] We might well feel that, while the goal is laudable, systems of such diverse kinds could hardly be expected to have any nontrivial properties in common. Metaphor and analogy can be helpful, or they can be misleading. All depends on whether the similarities the metaphor captures are significant or superficial.

It may not be entirely vain, however, to search for common properties among diverse kinds of complex systems. The ideas that go by the name of cybernetics constitute, if not a theory, at least a point of view that has been proving fruitful over a wide range of applications.[2] It has been useful to look at the behavior of adaptive systems in terms of the concepts of feedback and homeostasis, and to analyze adaptiveness in terms of the theory of selective information.[3] The ideas of feedback and information provide a frame of reference for viewing a wide range of situations, just as do the ideas of evolution, of relativism, of axiomatic method, and of operationalism.

In this paper I should like to report on some things we have been learning about particular kinds of complex systems encountered in the behavioral sciences. The developments I shall discuss arose in the context of specific phenomena, but the theoretical formulations themselves make little reference to details of structure. Instead they refer primarily to the complexity of the systems under view without specifying the exact content of that complexity. Because of their abstractness, the theories may have relevance—application would be too strong a term—to other kinds of complex systems that are observed in the social, biological, and physical sciences.

In recounting these developments, I shall avoid technical detail, which can generally be found elsewhere. I shall describe each theory in the particular context in which it arose. Then, I shall cite some examples of complex systems, from areas of science other than the initial application, to which the theoretical framework appears relevant. In doing so, I shall make reference to areas of knowledge where I am not expert—perhaps not even literate. I feel quite comfortable in doing so before the members of this society, representing as it does the whole span of the scientific and scholarly endeavor. Collectively you will have little difficulty, I am sure, in distinguishing instances based on idle fancy or sheer ignorance from instances that cast some light on the ways in which complexity exhibits itself wherever it is found in nature. I shall leave to you the final judgment of relevance in your respective fields.

I shall not undertake a formal definition of

* The ideas in this paper have been the topic of many conversations with my colleague, Allen Newell. George W. Corner suggested important improvements in biological content as well as editorial form. I am also indebted, for valuable comments on the manuscript, to Richard H. Meier, John R. Platt, and Warren Weaver. Some of the conjectures about the nearly decomposable structure of the nucleus-atom-molecule hierarchy were checked against the available quantitative data by Andrew Schoene and William Wise. My work in this area has been supported by a Ford Foundation grant for research in organizations and a Carnegie Corporation grant for research on cognitive processes. To all of the above, my warm thanks, and the usual absolution.

[1] See especially the yearbooks of the Society for General Systems Research. Prominent among the exponents of general systems theory are L. von Bertalanffy, K. Boulding, R. W. Gerard, and J. G. Miller. For a more skeptical view—perhaps too skeptical in the light of the present discussion—see H. A. Simon and A. Newell, Models: their uses and limitations, *in* L. D. White, ed., *The state of the social sciences,* 66–83, Chicago, Univ. of Chicago Press, 1956.

[2] N. Wiener, *Cybernetics,* New York, John Wiley & Sons, 1948. For an imaginative forerunner, see A. J. Lotka, *Elements of mathematical biology,* New York, Dover Publications, 1951, first published in 1924 as *Elements of physical biology.*

[3] C. Shannon and W. Weaver, *The mathematical theory of communication,* Urbana, Univ. of Illinois Press, 1949; W. R. Ashby, *Design for a brain,* New York, John Wiley & Sons, 1952.

PROCEEDINGS OF THE AMERICAN PHILOSOPHICAL SOCIETY, VOL. 106, NO. 6, DECEMBER, 1962

You Can't Play 20 Models…

Warwick, W., Walsh, M., Rodgers, S., & Lebiere, C. (2016). Integrating Heterogeneous Modeling Frameworks using the DREAMIT Workspace. In Proceedings for the First International Conference on Human Factors and Simulation..

- Tradeoff between levels of fidelity/complexity
- Compose component behavior models
- ***General, reusable models***
- Support and enforce modeling modularity
- Data flows for learning and validation
- Scalable modeling process