Integrated Cognition: A Survey of Systems

Robert M. Rolfe, Brian A. Haugh
Institute for Defense Analyses
4850 Mark Center Drive
Alexandria, VA 22311
rolfe@ida.org, bhaugh@ida.org

Abstract

A survey of cognitive architectures and systems is presented based on a reference framework for integrated cognition (INCOG). Each cognitive architecture/system is described briefly and evaluated with respect to its support for cognitive capabilities along the six dimensions of the INCOG framework. Hypotheses are formulated about the potential contributions of existing systems to a cognitive system covering all the capabilities of the INCOG framework. Relationships between this framework and an emerging Standard Model of the Mind are discussed. Finally, some generalizations of the survey results are presented.

Overview

Many different cognitive architectures have been developed in the quest for artificial systems that approach the cognitive capabilities of humans. But, it is difficult to assess their progress towards supporting human-level cognition without some model of the essential elements of human cognition. In a companion paper in this symposium, a framework for integrated cognition (INCOG) is described, which has been designed to capture these cognitive "ingredients" in a structured framework which supports assessment of candidate cognitive architectures along six dimensions (see Figure 1).

In this paper, we present individual assessments of the coverage of INCOG framework capabilities for nine prominent cognitive architectures or systems and a hypothetical architecture developed at the time of this research. These assessments were performed in collaboration with the system architects, as cited below. Each architecture or system was assessed for its support of each capability in the INCOG framework. Support for a capability is indicated on framework diagrams by a keyed icon at the capabilities location along the axis of its dimension in the framework (see Figure 1). Degree of support by an architectures or system was not assessed, so this may vary widely between the different systems. But, distinctions are made between the core capabilities of an architecture or system, extensions in example applications, and potential uses, via the use of distinctive keyed icons in the figures.

These assessments, individually and collectively, provide a view of the state of cognitive architectures relative to one model of the requirements of human-level cognition at the time of assessment. We divide the examined cognitive models into two categories, established systems, and newer systems.[1] Every approach surveyed could contribute concepts, designs, or components to an integrated cognition architecture covering all elements of the INCOG framework.

Established Systems

Established systems have some maturity and have demonstrated some capability in a variety of integrated cognition dimensions, and have had an established user community.

SOAR

SOAR is a general purpose architecture designed as a unified theory of cognition by John Laird, Paul Rosenbloom, and Allen Newell (Rosenbloom, Laird, and Newell 1993). It is a production rule system based on the simple decision cycle - elaboration of state, choice of operators, selection of operator, and actions. Soar has a relatively large user base amongst existing cognitive architectures. It is supported by the University of Michigan and has been applied commercially by Soar Technology Inc. Input for assessment of Soar's support for the elements of the INCOG frame-work (Figure 1) was provided by John E. Laird of the University of Michigan and Robert Wray of Soar Technology.

Hypothesis: Soar is embeddable with extension and rule sets to implement many components of an integrated cognition system

[1]The set of cognitive architectures and systems assessments were identified at the time of the DARPA IPTO study in the year 2003. Some of the systems were newer at that time.

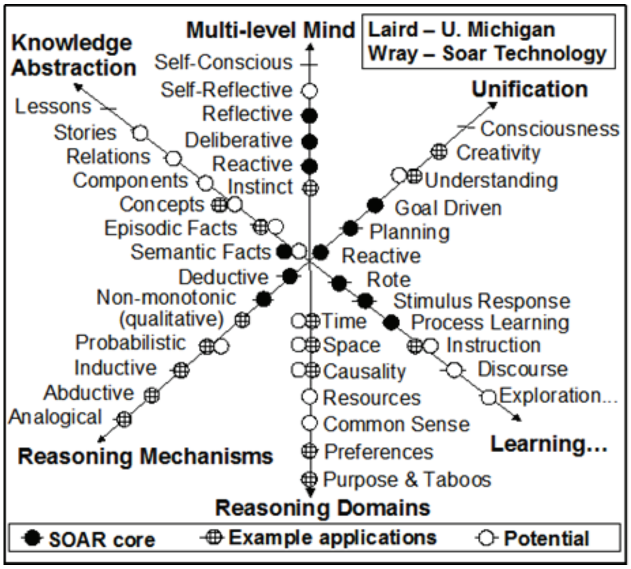

Figure 1 SOAR Ingredients

ACT-R

ACT-R (Anderson and Lebiere 1998) is a cognitive architecture using production rules developed at Carnegie Mellon University (CMU) by John Anderson and Christian Lebiere. It includes a detailed approach to integrating multiple modules that correspond to different cognitive functions. The fundamental controlling structure in cognition is reactive–where production rules respond to patterns of information in various cognitive buffers.

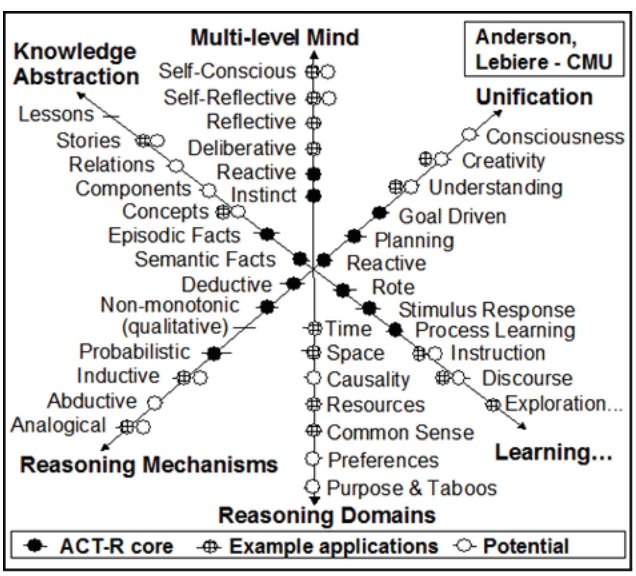

Figure 2. ACT-R ingredients

Successive versions of ACT-R have seen widespread applications to problems of cognitive and behavioral modeling. Input for assessment of ACT-R's support for the elements of the INCOG framework (Figure 2) was provided by John Anderson and Christian Lebiere of CMU.
Hypothesis: ACT-R based systems are embeddable for many components of an integrated cognition system

dMARS

The distributed Multi-Agent Reasoning System (dMARS) (dMARS) is a C++ implementation of an architecture based on the Belief, Desire, Intention (BDI) cognitive model (d'Inverno et al. 1998). It was developed by Michael Georgeff as a more powerful successor to the Procedural Reasoning System (PRS). dMARS has been applied to a very wide range of applications, including command and control of robotics and spacecraft; and situation awareness for the Australian Defense Forces. Input for assessment of dMAR's coverage of the INCOG framework (Figure 3) was provided by Michael Georgeff of Georgeff Inc.
Hypothesis: BDI concepts are useful for integrated cognitive systems.

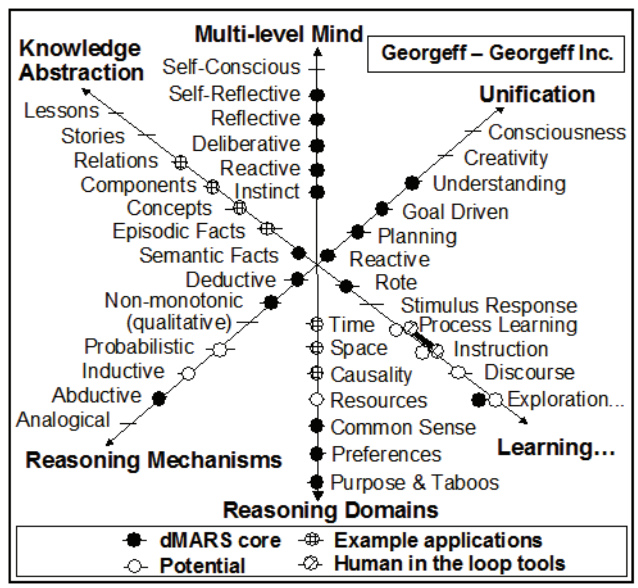

Figure 3. dMARS ingredients

ICARUS

ICARUS is an architecture for intelligent agents developed by Dan Shapiro and Pat Langley of the Center for the Study of Language and Information at Stanford University (Langley et al. 2002). ICARUS is distinguished by its incorporation of affective values into memory and behavior; the primacy of categorization over execution and of execution over problem solving; and the internal determination of tasks, intentions, and rewards. Input for assessment of

ICARUS's coverage of the INCOG framework (Figure 4) was provided by Pat Langley of Stanford University.
Hypothesis: ICARUS's is a general purpose architecture which is not yet broadly supported. Potential application to integrated cognition remain to be determined.

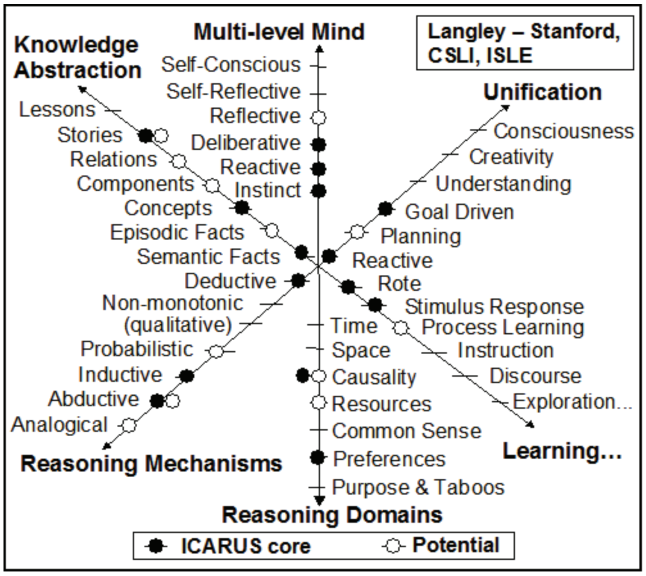

Figure 4. ICARUS – ingredients

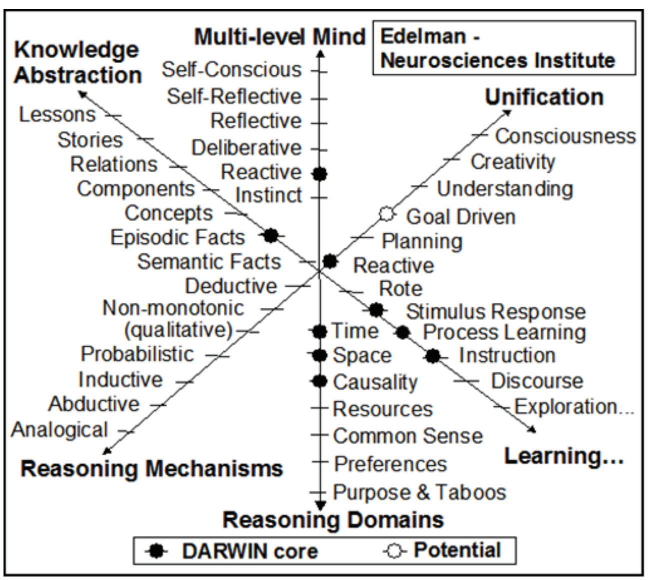

Figure 5. DARWIN Ingredients

DARWIN

DARWIN refers to a series of implementations of large scale (over 50,000 cells and 600,000 synapses) synthetic models of neural structures supporting the evolution of pattern recognition and sensorimotor coordination in a synthetic environment (Reeke Jr., Sporns, and Edelman 1990). It has been developed by Reeke, Sporns, and Edelman, of the Neurosciences Institute and Rockefeller University based on Edelman's Neural Darwinism (Edelman 1987).

Input for assessment of DARWIN's coverage of the INCOG framework (Figure 5) was provided by Gerald Edelman of the Neurosciences Institute. Hypothesis: Explores real time sensor steering and target tracking

UMPRS

UMPRS (the University of Michigan implementation of PRS) is a general purpose implementation of PRS (Lee, et al. 1994). It does not provide (i.e., "impose") specific capabilities or representations on agent programmers, rather providing a frameworks for their implementation. Hence, its core capabilities cover relatively few of the INCOG framework ingredients, although UMPRS applications have covered many more. Unification in UMPRS is focused on goals and planning and not reactive tasks. Input for assessment of UMPRS's coverage of the INCOG framework (Figure 6) was provided by Marcus J. Huber of Intelligent Reasoning Systems.

Hypothesis: UMPRS provides useful concepts for fully capable integrated cognition systems.

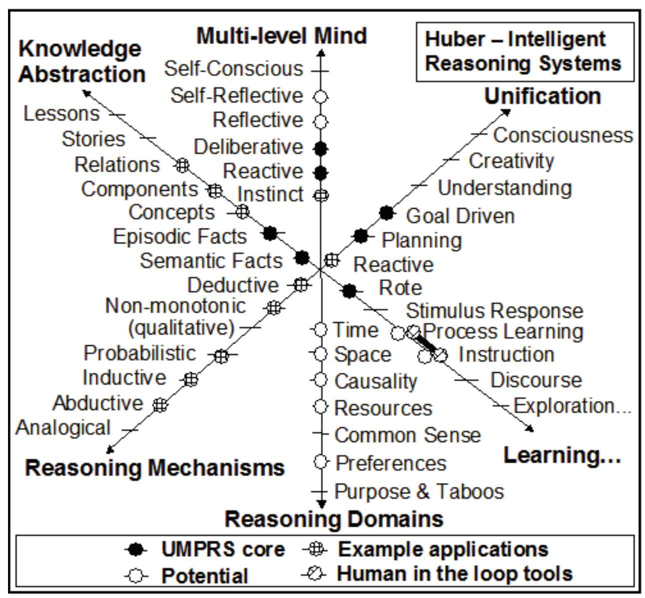

Figure 6. UMPRS ingredients

Newer Systems

Newer systems have newer integration strategies and mechanisms relative to the established systems previously described. Further, the capabilities may not be yet implemented, and generally lacked a large user community at the time of this assessment.

Shruti/Smirti

Shruti and Smirti are related architectures developed by Lokendra Shastri of UC, Berkeley (Shastri 1999). They demonstrate how simple, neuron like, elements can encode a large body of relational causal knowledge and provide a basis for reactive, rapid inference. Input for assessment of Shruti/Smirti's coverage of the INCOG frame-work (Figure 7) was provided by Lokendra Shastri of UC, Berkeley.

Hypothesis: Shruti provides a key cognitive real-time component that provides reactive text understanding and may provide a general model for composition for integrated cognition systems.

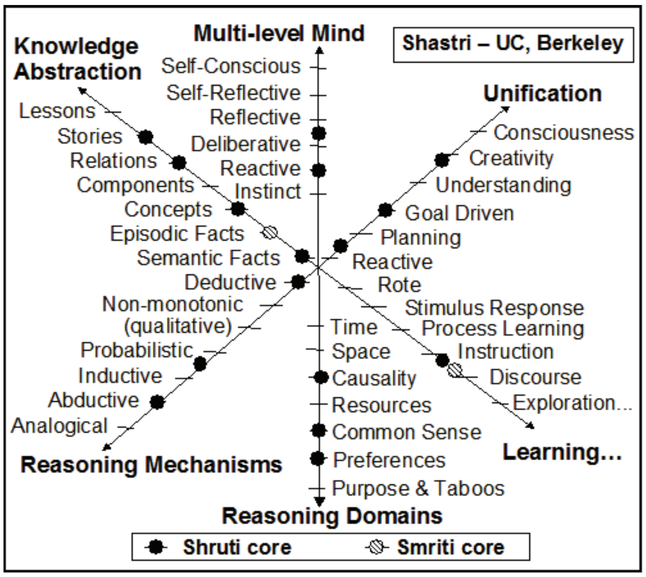

Figure 7. Shruti and Smirti ingredients of neuro-logically inspired cognitive systems

SAGE

SAGE (Self-Aware Adaptive Generalization Engine) is a cognitive-based architecture that is adaptive, self-reliant, and can reason by analogy (like people do) in order to discover meaningful relationships between seemingly dissimilar data. It blends connectionist/neural networks with symbolic systems. Its self-supervised learning uses self-reflective algorithms that allows the system to acquire new knowledge, learn from its past, and avoid extensive human intervention by guiding its own performance. Applications includes roles as a network security watch dog, decision aid for intelligence, or strategic agent for military simulation". Input for assessment of SAGE's coverage of the INCOG framework (Figure 8) was provided by Chris Furmanski and John Hummel of HRL/UCLA.

Hypothesis: A very general integrated model that provides composition concepts and components for integrated cognition systems.

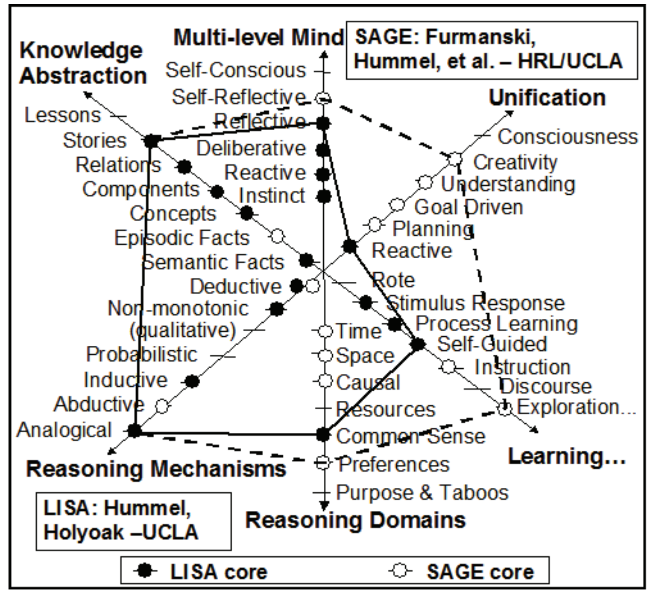

Figure 8. LISA and SAGE ingredients of neurologically inspired cognitive systems

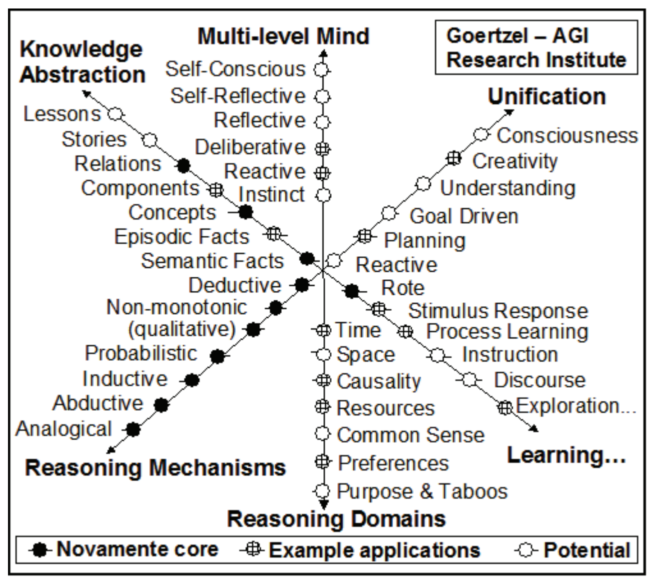

Figure 9. Novamente ingredients

Novamente

Novamente is a system organized with distributed atoms of knowledge that may be employed in an unlimited number of contexts, atoms have truth value and attention value, mind agents operating on these atoms, learn how to learn, an attempt at the holy grail, an artificial mind. Input for assessment of Novamente's support for the elements of the INCOG framework (Figure 9) was provided by Goertzel of AGI Research Institute.

Hypothesis: A very general integrated model that provides composition concepts and components for integrated cognition systems.

INCOG Strawman Architecture

An INCOG Architecture was developed for the DARPA IPTO as a strawman to start of a set of programs in the year 2003. The architecture was further evolved with a hypothesized set of software functional packages that together would cover the full scope of the capabilities of the INCOG framework. This INCOG Architecture is described in some detail in (Rolfe and Haugh 2004).

Figure 10 provides a top-level view of the INCOG strawman architecture. This figure–based on Dr. Ronald Brachman's proposed cognitive architecture (Brachman 2002)–has been adapted here to better capture distinctions made in the INCOG framework. In particular, this INCOG architecture distinguishes more finely between different levels of the multi-level mind, separating Brachman's Reactive Processes into Programmed Instinct and Learned Reactions; and adding several other levels above the Reflective. This architecture also highlights discourse as a key area of human-level cognition by separately identifying it and its relations to other input and output processing.

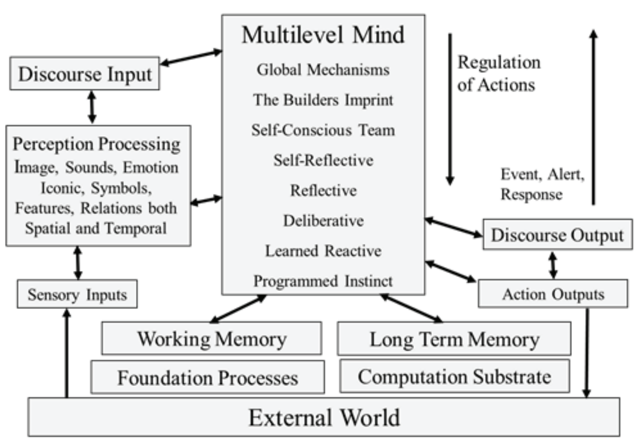

Figure 10 INCOG Example Architecture

This architecture diagram represents the world external to the cognitive agent by the External World box at the bottom. The rest of the boxes represent information and processes of the cognitive agent. Raw Sensory Inputs come into the agent from its sensors and are processed initially via Perception Processing, which hands off linguistic data to Discourse Input. The results of perception and discourse input processing are fed to various levels of the Multilevel Mind, as appropriate. The Multilevel Mind uses inputs from Working Memory and Long-Term Memory to place new inputs in context and to help determine its responses and other activities. The Multi-level Mind also stores information in Working Memory and Long-Term Memory as warranted. The processing and exchange of information throughout is enabled by a host of Foundation Processes, as well as the Computational Substrate. The results of the execution of cognitive processes may then find expression in discourse via Discourse Output or through other Action Outputs. The Discourse Outputs consist of the intended discourse, which is communicated to the appropriate physical effectors as Action Outputs in order to effect speech, writing, and any other forms of communication.

Hypothesis: A very general model designed to support all the levels of mind and other INCOG capabilities given suitable elaboration of all of its components.

Standard Model of the Mind

A Standard Model of the Mind has been recently proposed to provide a consensus high-level model of human-like minds (Laird, Lebiere, and Rosenbloom In press). This model is not intended as a full-blown cognitive architecture, but as an abstraction of a common core from multiple architectures. As such, it is difficult to assess the Standard Model of the Mind for its support of the capabilities of the INCOG framework. It does not include the details that would be required for assessing many of those capabilities. However, given its instantiation in current versions of prominent cognitive architectures, such as Soar and ACT-R, we can infer that the Standard Model of the Mind is capable of providing some level of support for most, if not all, of the INCOG framework capabilities.

Although existing architectures compatible with the Standard Model of the Mind have shown substantial support for the capabilities of the INCOG framework, it is unclear how extensive that support can be for some of those capabilities. In particular, this model appears to limit the structure of much of its memory to relations over symbols supplemented by quantitative metadata annotating them. Such relations and metadata are not adequate to represent the detailed structure of complex visual/spatial and auditory information. Humans do not reason by symbols alone, we compare and manipulate complex dynamic 3-dimensional spatial representations in conscious working memory and retain those memories long-term. These representations play key roles in human spatial reasoning, such as planning and pattern recognition. Hence they cannot be delegated to a perception module unless that module includes reasoning in a working memory that is the focus of attention, as well as interaction with a long-term memory. However, this model does not identify any such reasoning capabilities for its Perception module, which is described as just converting external signals into symbols, relations, and their metadata for input into working memory.

Thus, the Standard Model of the Mind appears to have limited support for spatial reasoning due to an apparent exclusion from some memory modules of the complex representations required for reasoning with spatial models. This limitation appears to limit its potential support for many INCOG framework capabilities, such as some inductive, abductive, and analogical reasoning that relies on such spatial models.

Conclusions

This survey of these cognitive architectures supports the following general observations:

1. Working implementations tend to require significant low level programming for ingredients not explicitly in their architectures, E.g. SOAR, ACT-R, dMARS, UMPRS.
2. Learning capabilities are limited to refinement within the scope of initial knowledge bases, few of these systems sought to understand or invoke learning strategies beyond process learning related to current knowledge.
3. Current systems tend to be weak with respect to self-reflection, and knowledge sharing and consequently would be difficult to employ in a heterogeneous integrated cognitive architecture without extensions.
4. Current systems have core capabilities that cluster near the center of the INCOG multi-dimensional framework.
5. Newer or proposed systems have extended coverage of the strawman multi-dimensional framework to near the periphery in some dimensions.
6. Coverage is sparse within some dimensions because of the large number of potentially interrelated reasoning domains not addressed in any of the architectures.
7. There are significant sets of advanced components (in established cognitive systems) that could be included in integrated architectures that together would provide significant new capability for cognitive systems.
8. The Standard Model of the Mind provides a structure capable of supporting some level of capability for most, if not all, of the INCOG framework ingredients.
9. The Standard Model of the Mind appears to have limited support for spatial reasoning and related reasoning capabilities due to an apparent exclusion from long-term memory of the complex representations required for detailed reasoning with dynamic 3-D spatial models.

Acknowledgement

We thank Dr. Ronald Brachman and Dr. Barbara Yoon of the DARPA Information Processing Technology Office (IPTO) for their support of this work. We thank Dr. John Salasin for his vision in conceiving these investigations into Integrated Cognition (INCOG). We also thank Professor Pat Langley of the Computational Learning Laboratory at Stanford University, who motivated development of the INCOG framework diagram to communicate the scope and difficulty of constructing intelligent machine architectures. We also thank the many researchers who contributed input to assessments of the coverage of their architectures and systems of the ingredients of the INCOG framework.

References

Anderson, J. and Lebiere, C. 1998. The atomic components of thought. Mahwah, NJ: Erlbaum.

Brachman, R., 2002 "Developing Cognitive Systems. A Convergence of Thinking." Information Processing Technology Office (IPTO) Defense Advanced Research Projects Agency, briefing presented August 15, 2002.

d'Inverno, M., Kinny, D., Michael Luck, and Michael Wooldridge: 1998. A Formal Specification of dMARS.in Singh, Rao, and Wooldridge (Eds.): Intelligent Agents IV. 155-176. New York: Springer Verlag.

Edelman, G. 1987. Neural Darwinism: The Theory of Neuronal Group Selection. Basic Books.

Goertzel, B., Pennachin, C., Senna, A., Maia, T., and Lamacie, G. 2003. Novamente: An Integrative Architecture for Artificial General Intelligence. In Proceedings of the International Joint Conference on Artificial Intelligence-2003.

Huber, M. 1999. JAM: A BDI-theoretic Mobile Agent Architecture, 236-243. In Proceedings of the Third International Conference on Autonomous Agents and Multiagent Systems.

Huber, M. 2000. JAM: A BDI-theoretic Mobile Agent Architecture. AgentLink News, Issue 5: 2-5.

Hummel, J. and Holyoak, K. 2003. Relational Reasoning in a Neurally-plausible Cognitive Architecture: An Overview of the LISA Project. Los Angeles, CA: University of California.

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.

Langley, P., Shapiro, D., Aycinena, M., and Siliski, M. 2002. A Value-Driven Architecture for Intelligent Behavior. Stanford, CA: Stanford University.

Lee, J., Huber, M., Kenny, P., and Durfee, E. 1994. UMPRS: An Implementation of the Procedural Reasoning System for Multirobot Applications. In Proceedings of the AIAA/NASA Conference on Intelligent Robotics in Field, Factory, Service, and Space: 509-519. Houston, TX. American Institute of Aeronautics and Astronautics/National Aeronautics and Space Ad-ministration.

Reeke Jr., G.N, Sporns, O., and Edelman, G.E. 1990. Synthetic Neural Modeling: The "Darwin" Series of Recognition Automata, 1498-1530. In Proceedings of the IEEE, Vol 78, No.9.

Rolfe, R. M.; and Haugh, B.A. 2004. Integrated Cognition – A Proposed Definition of Ingredients, A Survey of Systems, and Example Architecture. Arlington, Virginia, IDA Paper P-3855 and DTIC, URL: www.dtic.mil/dtic/tr/fulltext/u2/a451750.pdf

Rosenbloom, P., Laird, J., and Newell, A. 1993. The Soar Papers: Readings on Integrated Intelligence. Cambridge, MA: MIT Press.

Shastri, L. 1999. Advances in Shruti - A neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony. Applied Intelligence 11(1): 79-108.