Towards a New Generation of Cognitive Architectures

Paul S. Rosenbloom

May 12, 2010

A *cognitive architecture* provides a hypothesis about the fixed mechanisms underlying intelligence and how they combine to yield appropriate environmental behavior in the presence of knowledge. Such an architecture will generally at least include memories, an ability to make decisions, learning mechanisms, and external interfaces. However, it may also embody additional mechanisms in support of more advanced capabilities such as reasoning, problem solving and planning, reflection, collaboration and emotion. A cognitive architecture may be intended as a model of human cognition and/or as the basis for artificially intelligent (AI) systems.

The ideal architecture provides an effective combination of *functionality* and *uniformity*. Functionality concerns the range of intelligent behaviors that are realizable, up to the full scope of human behavior when the goal is either a unified theory of human cognition or the achievement of human-level artificial intelligence. Uniformity focuses on providing such behaviors through the interactions among a relatively small number of general mechanisms, rather than hypothesizing a new mechanism for each additional behavior. Together uniformity and functionality determine the utility and elegance of the architecture as both a theory and a system. Yet they are too often in conflict in existing architectures as the desire for new functionality constantly strains what can be realized with existing combination of mechanisms. Reconciling functionality and uniformity demands a continual search for new mechanisms capable of increased generality and integrability.

For fifteen years (1983-1998) I co-led the *Soar project*, a multidisciplinary, multi-institutional effort with the goal of developing, understanding and applying a cognitive architecture that could effectively combine functionality and uniformity in service of both modeling human cognition and building effective AI systems. Soar was successful in yielding a wide range of intelligent behaviors across a broad scope of task domains from a small set of mechanisms, with the most striking results often due to originally unanticipated interactions among its mechanisms. However, as more and more functionality was desired, the strain eventually became too great for the existing uniform approach to sustain. The ultimate response on the part of the Soar project was to opt for functionality over uniformity, adding a large number of additional mechanisms to yield the Soar 9 architecture.

After a ten-year hiatus from work on cognitive architecture, I have recently returned to the fray with the goal of developing new architectures that dramatically extend the scope of the previous generations while simultaneously reconstructing them on a more uniform base. The intent is thus a new generation of architectures with significant improvements in both functionality and uniformity. Functionality is being approached via the development of *hybrid mixed architectures*. Cognitive architectures traditionally emphasize symbolic processing, although possibly with limited forms of numeric processing in perceptuomotor modules or via incorporation of forms of activation-based models, such as neural networks. A hybrid architecture tightly integrates symbolic and numeric/signal processing, enabling close coupling of perceptuomotor and cognitive behavior. A mixed architecture integrates symbolic and uncertain/probabilistic reasoning, enabling general reasoning under uncertainty. A hybrid mixed architecture supports all three – signals, probabilities and symbols – in an integrated manner.

Uniformity is being tackled by developing a common implementation of signal, probability, and symbol processing based on *graphical models*. Bayesian networks are the best-known form of graphical model, but other varieties include Markov networks (aka Markov random fields) and factor graphs. Factor graphs provide efficient computation with complex multivariate functions by decomposing them into a product of subfunctions that can be mapped onto undirected bipartite graphs containing nodes for variables and subfunctions. In conjunction with the summary-product algorithm, factor graphs yield state-of-the-art methods for many standard signal, probability and symbol processing problems, and thus raise the intriguing possibility of supporting highly uniform yet broadly functional hybrid mixed architectures. The long-term hope is that this approach will yield the next generation of general mechanisms in support of new architectures that far surpass today's best in terms of both functionality and uniformity. Progress to date is encouraging, but has not yet reached implementations of complete architectures.

This talk will include a brief review of the earlier work on Soar, along with the more recent work on Soar 9. This will be followed by a discussion of the new approach to constructing hybrid mixed architectures on top of uniform graphical models, plus results achieved so far with it and speculations on its future.

Biography

Paul S. Rosenbloom is a Professor of Computer Science at the University of Southern California (USC), Project Leader at USC's Institute for Creative Technologies (ICT) and Deputy Director of the Center for Rapid Automated Fabrication Technoloiges (CRAFT). He was at USC's Information Sciences Institute (ISI) for twenty years, most recently as its Deputy Director. Prior to coming to USC in 1987, he was an Assistant Professor of Computer Science and Psychology at Stanford University from 1984 to1987, and a Research Computer Scientist at Carnegie Mellon University from 1983 to 1984. He received a B.S. degree in Mathematical Sciences (with distinction) from Stanford University in 1976 and M.S. and Ph.D. degrees in Computer Science from Carnegie Mellon University in 1978 and 1983, respectively. Prof. Rosenbloom was elected a Fellow of the Association for the Advancement of Artificial Intelligence (AAAI) in 1994, and has served as Chair of the Association for Computing Machinery (ACM) Special Interest Group on Artificial Intelligence (SIGART), Councillor of the AAAI, Conference Chair for the AAAI, and Program Co-Chair of the National Conference on Artificial Intelligence (AAAI-92).

From 1983 until 1998, Prof. Rosenbloom was a co-PI of the Soar Project, a multi-disciplinary, multi-site attempt at developing, understanding, and applying a cognitive architecture capable of supporting general intelligence. Research on Soar spanned areas such as machine learning, problem solving and planning, production systems, intelligent agents, virtual humans, multi-agent systems, knowledge-based systems, neural networks, and cognitive modeling. The most significant applications were intelligent automated pilots and commanders for synthetic battlespaces, as deployed in Synthetic Theater of War '97 (STOW-97). From 1998 until 2007 Prof. Rosenbloom's focus was on exploring and/or developing new directions for ISI, such as blending entertainment and computing technologies and expertise for military training; virtual organizations consisting of mixtures of people, agents, robots, and/or computational systems; responding to the unexpected; high performance computing, scalable distributed computing, and computational science; biomedical informatics; automated construction; and technology and the arts. Currently, Prof. Rosenbloom is working on a new approach to cognitive architecture based on graphical models and writing a book tentatively titled *What is Computing? The Architecture of the Fourth Great Scientific Domain*.