Towards Uniform Implementation of Architectural Diversity

Paul S. Rosenbloom
Department of Computer Science & Institute for Creative Technologies
University of Southern California
13274 Fiji Way, Marina del Rey, CA 90292
rosenbloom@usc.edu

Abstract

Multi-representational architectures exploit diversity to yield the breadth of capabilities required for intelligent behavior in the world, but in so doing can sacrifice too much of the complementary benefits of architectural uniformity. The proposal here is to couple the benefits of diversity *and* uniformity through establishment of a *uniform graph-based implementation level* for diverse architectures.

Architectures for human level intelligence (HLI) seek to provide a coherent integration of capabilities sufficient for intelligent behavior in the world, whether as part of a detailed model of human cognition or a system more loosely tied to the specifics of human behavior. They require the integration of a wide range of cognitive capabilities for, among other things: memory and reasoning, problem solving and planning, reactivity, learning, reflection, interaction (including perception and motor control, use of language, etc), and the social aspects of cognition (such as emotion, collaboration, etc.).

A key issue with such architectures is what can be called the *diversity dilemma*: they need to be both diverse *and* simple. Diversity of capability is required to support intelligent behavior in a complex and uncertain world. Simplicity is critical for architectural elegance, integrability, extensibility, and maintainability. In Rosenbloom (2009a), a resolution to this dilemma was proposed based on analogy to Deering's (1998) *Internet Hourglass* and Domingos's (In press) call for an *interface layer* for AI. The idea is to seek a simple, uniform, *mesoscale level* in the cognitive hierarchy (Newell 1990) that can support increasing diversity above (for supporting intelligent behaviors) and below (for grounding in biological and computational technologies).

To be more specific, the proposal is to search beneath the architecture for a uniform *implementation level* based on *graphical models* (Jordan 2004). Traditionally, architectural implementation has been considered mere "implementation details", of pragmatic importance for efficiency and robustness, but of little theoretical interest. The one notable exception has been when a symbolic architecture is implemented via neural networks, as in Neuro-Soar (Cho *et al*. 1991) and SAL (Jilk *et al*. 2008).

Neural networks are graphical models of a sort, but they are far from the only such models. Bayesian networks [Pearl, 1988] are directed graphical models over random variables that have revolutionized probabilistic reasoning. Markov networks are undirected analogues of Bayesian networks. Factor graphs (Kschischang *et al*. 2001) are undirected, like Markov networks, but represent general multivariate functions, and add factor nodes in the graph rather than using separate clique potentials.

One of the most intriguing aspects of graphical models is their ability to uniformly process symbols, probabilities and signals via variants of the same graph structure and inference algorithm (the *summary-product algorithm*). This approach is not only uniform, but it subsumes state-of-the-art algorithms spanning these areas, such as arc consistency and production match algorithms (symbol processing), loopy belief propagation (probability processing), and Kalman filters and the forward-backward algorithm in HMMs (signal processing).

Although the roots of many HLI architectures are in symbol processing, both probabilities and signals are required to cope with the real world. When signal processing is implemented at all in existing architectures, it tends to be relegated to external modules having limited interaction with cognition. Neural networks do directly capture aspects of both probability and signal processing, but in turn have difficulty with symbol processing unless hybridized with an explicit symbolic component, such as in Clarion (Sun 2006) or SAL. Employing a more inclusive graphical implementation level should enable the uniform blending of all three representations into more elegant, integrable and extensible architectures. When conjoined with the mapping of neural networks onto graphical models (Jordan and Sejnowski 2001), it may also help bridge the gap between symbolic and neural processing.

To explore the notion of a graphical implementation level, a reimplementation and extension of the Soar architecture (Rosenbloom *et al*. 1993) has been initiated. Soar is particularly useful as a starting point because it: (1) is one of the most well developed and broadly applied HLI architectures; (2) has been explored as both a unified theory of human cognition and as an architecture for intelligent agents; and (3) exists in both uniform (versions 1-8) and diverse forms (version 9, which adds multiple forms of representation, memory and learning (Laird 2008)), enabling a strategy of starting reimplementation with the initial uniformity while seeking opportunities for a more uniform integration of the later diversity.

Initial experiments have proceeded from the bottom up, with reimplementations of Soar's *elaboration* and *decision* cycles (Rosenbloom 2009a; 2009b). The elaboration cycle uses parallel rule match and firing to retrieve information from long-term memory about the current situation.

*Proceedings of the AAAI Fall Symposium on Multi-Representational Architectures for Human-Level Intelligenc*e, 2009.

Reimplementation of the elaboration cycle was via factor graphs. Each rule became a graph, structured à la junction trees, while working memory (WM) became a 3D array, structured as an *octree*, that provided the messages passed in the graphs. The overall result was a new graph-based rule-match algorithm with worst-case cost exponential in the *treewidth* rather than in the number of conditions.

Soar's decision cycle executes elaboration cycles until quiescence and, based on retrieved preferences, then chooses what to do next or generates an impasse if it can't decide. The decision cycle was reimplemented via *Markov logic* – a combination of first-order logic and Markov networks – and its associated Alchemy language (Domingos *et al*. 2006). The result was an enhanced decision cycle, combining both symbolic and probabilistic reasoning, and enabling the addition of both a simple semantic memory and the kinds of *trellis structures* implicated in sequential signal processing (such as speech).

Many outstanding issues remain with these fragmentary initial reimplementations, but they start to reveal the potential of a uniform, graphical implementation level to support diverse HLI architectures. Current work is focused on completing a uniform implementation of a multi-representational decision cycle capable of integrated symbol, probability and signal processing to support combining procedural knowledge, declarative knowledge (such as Soar 9's semantic and episodic memories) and perceptual knowledge. Future work will include reimplementing the remainder of Soar, while enhancing it with further capabilities, such as perception, decision-theoretic planning, Markov decision processes, and theory of mind. Beyond Soar, it will also be essential to explore reimplementations of other leading architectures, and hybrids among them, as well as new multi-representational architectures that are more directly inspired by the uniform multipotency of a graphical implementation level.

Acknowledgements. This work was enabled by sabbatical support from the USC Viterbi School of Engineering and funding from the USC Institute for Creative Technologies.

References

Cho, B., Rosenbloom, P. S. and Dolan, C. P. 1991. Neuro-Soar: A neural-network architecture for goal-oriented behavior. *Proceedings of the 13th Annual Conference of the Cognitive Science Society* (pp. 673-677). Erlbaum.

Deering, S. (1988). *Watching the waist of the protocol hourglass*, Keynote address at ICNP '98.

Domingos, P. In press. What is missing in AI: The interface layer. In P. Cohen (Ed.), *Artificial Intelligence: The First Hundred Years*. Menlo Park, CA: AAAI Press.

Domingos, P., Kok, S., Poon, H., Richardson, M. and Singla, P. 2006. Unifying logical and statistical AI. *Proceedings of the Twenty-First National Conference on Artificial Intelligence* (pp. 2-7). AAAI Press.

Jilk, D. J., Lebiere, C., O'Reilly, R. C. and Anderson, J. R. 2008. SAL: An explicitly pluralistic cognitive architecture. *Journal of Experimental and Theoretical Artificial Intelligence*, *20*, 197-218.

Jordan, M. I. 2004. Graphical models. *Statistical Science (Special Issue on Bayesian Statistics)*, *19*, 140-155.

Jordan, M. I. and Sejnowski, T. J. 2001. *Graphical Models: Foundations of Neural Computation*. Cambridge, MA: MIT Press.

Kschischang, F. R., Frey, B. J. and Loeliger, H. 2001. Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory*, *47*, 498-519.

Laird, J. E. 2008. Extending the Soar cognitive architecture. In *Artificial General Intelligence 2008: Proceedings of the First AGI Conference*. IOS Press.

Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufman.

Rosenbloom, P. S. 2009a. Towards a new cognitive hourglass: Uniform implementation of cognitive architecture via factor graphs. *Proceedings of the 9th International Conference on Cognitive Modeling* (ICCM 2009).

Rosenbloom, P. S. 2009b. A graphical rethinking of the cognitive inner loop. *Proceedings of The IJCAI International Workshop on Graph Structures for Knowledge Representation and Reasoning.*

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

Sun, R. 2006. The CLARION cognitive architecture: Extending cognitive modeling to social simulation. In R. Sun (ed.), *Cognition and Multi-Agent Interaction*. Cambridge University Press, New York.