

This paper appeared in the *Proceedings of the Canadian Workshop on Machine Learning*, 1992.

Memory and Concepts in Reactive Learning*

SUSAN L. EPSTEIN

Department of Computer Science

Hunter College and The Graduate Center of The City University of New York

695 Park Avenue, New York, NY 10021

The extent to which concepts, memory, and planning are necessary to the simulation of intelligent behavior is a fundamental philosophical issue in AI [Kirsh, research community has taken the position that *reactive systems*, consisting of properly-organized, multiple low-level debate with fully operational systems has thus far been primarily on mobile robots that do simple Hsu and Simmons, 1991; Maes and Brooks, 1990; Mahadevan and Connell,

Hoyle is a reactive system in a cerebral, rather than a physical, domain. The program learns to perform well and quickly, often outpacing its human creators at two-person, perfect information a surprising amount of intelligent behavior can be treated as if it were situation-determined, that often planning is unnecessary, and that the memory required to support this learning is minimal. A series of ablation experiments, however, have demonstrated that concepts and memory are crucial to this reactive program's ability to learn and perform.

The thesis underlying this work is that reactive, hierarchical systems can minimize deliberation, but that both memory and explicitly represented concepts are necessary if a program is to learn to perform intelligently. *Concept* is defined here as some recognized set of regularities detected in some observed world. *Regularity* means repeated occurrence and/or consistency of use. A concept therefore includes not only the necessary and sufficient descriptions called definitions, but also defaults, associations, and expectations. Thus a concept is generalized domain knowledge, a description of what has been encountered. Although specific examples may be remembered, a concept is not a set of instances but a summary of experience.

Four other reactive game playing programs, Dooze, N-N/Tree, Morph, and Henri, employ Their pattern generalizations, however, are tailored to a single set of board-specific algorithms,

* This work was supported in part by NSF 9001936 and PSC-CUNY 668287.

and their memory requirements grow dramatically with the number of positions on the board. Hoyle outperforms these programs because it explicitly represents and exploits its concepts.

Hoyle's concepts organize the way it remembers experience, focus its attention on what is important to learn, force it to apply its experience, and permit it to discard experience that is judged unlikely to be useful. As a result, it is able to learn with far smaller memory requirements than the other four reactive game-playing programs, and to apply its compact useful knowledge much more flexibly. Although Hoyle is reactive, the full version of the program incorporates and remembers concepts: knowledge about the regularities that people learn, prefer, and exploit when playing games, and how people use those regularities. When the program is partially disabled and the results observed, it is clear that memory and the application of those concepts provide the program with its power.

Hoyle's ability to learn with only 15 relatively simple low-level agents suggests that more high-level behavior is available through low-level reactive processes than one might initially suspect. As the games Hoyle learns become more difficult, new concepts, like material and mobility, have been necessary to support performance. For Hoyle, learning high-level behavior efficiently with a limited memory requires concepts. A recent improvement enabled the program to learn which of the eight symmetries of the two-dimensional plane are applicable to a game, and to apply those to minimize both memory requirements and search time. Hoyle with symmetry learns faster and requires less memory. Low-level sensory data can offer an immediate improvement in high-level processing.

For the time being, several tasks have been relegated to the human system designer: the frame-based description of the game definition and useful knowledge, the correct identification of the culturally-determined *meta-principles* (commonsense regularities like efficiency, safety, and propriety, that are applicable to many different kinds of experience), the instantiation of the meta-principles to construct low-level agents, the hierarchical assignment of agents to tiers based upon knowledge about relations among meta-principles, the specification of which agents access which concepts, and the description of how the agents apply that knowledge. Current work is directed to training Hoyle on an increasingly difficult sequence of 14 games as well or better than its expert human competition, despite the program's restriction to two-ply search and its inability to plan.

In a domain that is not situation-determined, Hoyle is a successful, reactive, hierarchical system that retains only a small fraction of what it experiences. The program pays an interesting price for its reactivity, however. It must rely on concepts to learn to perform intelligently. Hoyle offers evidence that learning cerebral tasks demands more explicit concepts than Brooks would

like, and far fewer than Kirsh would assume. Hoyle may not resolve the controversy over whether representation is necessary for intelligence, but it should certainly influence our attitude on the significance of low-level agents in high-level tasks.

Hoyle's results demonstrate for at least one broad cerebral task, game playing, that a reactive system without memory is impractical, and that reliance only on extensive, detailed memory is brittle and often impossible. Concepts can structure resource-efficient memory, provide flexibility, and regularize knowledge to support performance. Hoyle is a hybrid; it combines reactive strategy with memory for learning and concepts for power. Will a reactive program ever, then, have to search and plan and believe? Hoyle's answer is not yet, perhaps not explicitly, and far less than we ever expected.

References

Many Civilizations. London: Oxford University Press.

Brooks, R. A. 1991. Intelligence without Representation. *Artificial Intelligence* 47 (1-3): 139-160.

Cobb, H. G. and Grefenstette, J. J. 1991. Learning the Persistence of Actions in Reactive Control Rules. In *Proceedings of the Eighth International Machine Learning Workshop*, 293-297. Morgan Kaufmann.

Connell, J. 1990. *Minimalist Mobile Robotics: A Colony-style Architecture for an Artificial Creature*. New York: Academic Press.

Epstein, S. L. 1992. Prior Knowledge Strengthens Learning to Control Search in Weak Theory Domains. *International Journal of Intelligent Systems*. To appear.

Esfahany, K. 1992. A Pattern Classifier that Learns to Play Games. In preparation.

Flax, M. G., Gelfand, J. J., Lane, S. H. and Handelman, D. A. 1990. Integrating Neural Network and Tree Search Approaches to Produce an Auto-Supervised System that Learns to Play Games. In *Proceedings of the Aerospace Applications of Artificial Intelligence Conference*.

Hsu, G.-T. and Simmons, R. 1991. Learning Football Evaluation for a Walking Robot. In *Proceedings of the Eighth International Machine Learning Workshop*, 303-307. Morgan Kaufmann.

Kirsh, D. 1991. Today, the Earwig, Tomorrow Man? *Artificial Intelligence* 47 (1-3): 161-184.

Levinson, R. and Snyder, R. 1991. Adaptive Pattern-Oriented Chess. In *Proceedings of the Eighth International Machine Learning Workshop*, 85-89. Morgan Kaufmann.

Maes, P. and Brooks, R. A. 1990. Learning to Coordinate Behaviors. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 796-802. AAAI Press.

- Mahadevan, S. and Connell, J. 1991. Scaling Reinforcement Learning to Robotics by Exploiting the Subsumption Architecture. In *Proceedings of the Eighth International Machine Learning Workshop*, 328-332. Morgan Kaufmann.
- Painter, J. 1992. Master's thesis, in preparation. Department of Computer Science, Hunter College of the City University of New York.
- Wierzbicka, A. 1985. *Lexicography and Conceptual Analysis*. Ann Arbor, MI: Karoma Publishers.
- Zaslavsky, C. 1982. *Tic Tac Toe and Other Three-in-a-Row Games, from Ancient Egypt to the*