

Modeling Problem-Solving Strategies as the Deliberate Retrieval of Actions and Goals

Christopher R. Waterson (WATERSON@EECS.UMICH.EDU)

Randolph M. Jones (RJONES@EECS.UMICH.EDU)

Artificial Intelligence Laboratory

University of Michigan

1101 Beal Avenue

Ann Arbor, MI 48109-2110

Past research with learning problem solvers has typically addressed action selection and goal selection as distinct mechanisms. Usually, significant effort is devoted to learning when to select particular actions from a known set of operators during reasoning (e.g., Minton, 1988; Mitchell, Utgoff, & Banerji, 1983). The learned action-selection strategies usually depend on descriptive features of intermediate problem states and the initial goals of the problem. If the reasoning system is capable of creating subgoals during reasoning, there is often a fixed mechanism for determining what the subgoals should be. For example, in means-ends analysis (Fikes, & Nilsson, 1971; Minton, 1988; Newell & Simon, 1972), new subgoals are fixed as the unmatched preconditions of selected operators.

Our research focuses on the advantages of viewing both action selection and subgoal creation as two instances of the same process: the deliberate retrieval of internal features from memory. When a reasoning system is presented with a problem, the problem is represented as a set of features describing an initial state and some goal conditions. These features serve as input to a memory model, which returns a new set of features retrieved from an experiential memory. The retrieved features can be viewed as symbols indicating which action should next be taken, or they can be symbols that simply remain in memory to feed back in to the next cycle of retrieval. The former type of symbols correspond to action selections, and the latter have the effect of subgoals.

This view is advantageous because it provides a purely functional purpose to subgoals: they serve strictly to focus the retrieval of new subgoals and actions. However, another strength of this uniform view is that both action selection and subgoal creation can arise from the same experiential learning mechanism, because they have similar representations and functions. A reasoning system that can learn both about action selection and subgoal creation can exhibit seemingly rigid reasoning strategies like means-ends analysis, but they have the flexibility to mix goal-driven and opportunistic reasoning when their experience indicates it is appropriate.

Jones and VanLehn (1994) investigated the application of a uniform learning model to action selection and subgoal retrieval, and successfully used such a system to model strategy acquisition in children learning to add. However, this system only took a first step at uniformly representing the retrieval

of subgoals and actions. Our current research focuses on developing a truly integrated model of retrieval and learning for problem solving, within the Soar (Newell, 1990) architecture.

This research implements retrieval as the result of an associative learning mechanism. Features describing the current state of a problem combine with symbols representing current goals and subgoals, to recall potential actions to use for the next reasoning step. If an action cannot immediately execute, symbols representing the action combine with problem features to recall new internal features (subgoals) that lead in turn to the retrieval of new actions or subgoals. Retrieval of both actions and subgoals adapts with experience and feedback.

The current model uses a symbolic feature association mechanism, called SCA (Miller & Laird, 1997), that has been developed within the constraints of the Soar architecture. SCA is a successful model of human category learning, and therefore provides an interesting candidate for use as the retrieval mechanism governing human problem solving.

References

- Fikes, R. E., & Nilsson, N. J. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2, 189–208.
- Jones, R. M., & VanLehn, K. (1994). Acquisition of children's addition strategies: A model of impasse-free, knowledge-level learning. *Machine Learning*, 6, 11–36.
- Miller, C. S., & Laird, J. E. (1997). Accounting for graded performance within a discrete search framework. *Cognitive Science*, 20, 499–537.
- Minton, S. (1988). *Learning effective search control knowledge: An explanation-based approach*. Boston: Kluwer Academic.
- Mitchell, T. M., Utgoff, P. E., & Banerji, R. (1983). Learning by experimentation: Acquiring and refining problem-solving heuristics. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*. Los Altos, CA: Morgan Kaufmann.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.