Revisiting the Chunking of Goal Hierarchies: Learning Hierarchically Controlled Behavior with Extremely Limited Working Memory

Richard Lewis
University of Michigan

Hierarchical goal structures remain ubiquitous in computational models of complex tasks in AI, HCI and cognitive psychology. Recently, ACT-R has moved away from direct architectural support for goal hierarchies (Altmann & Trafton, 2002; Anderson & Douglass, 2001). This talk presents work in progress that addresses several critical questions about hierarchical goal processing that are motivated in part by such a move.

(1) How can a system with an extremely limited working memory and no architectural support for goal hierarchies support hierarchically controlled behavior?

(2) Can we recapture the benefits of a learning mechanism like Soar's chunking, which operates over architecturally distinguished goal/subgoal relations?

(3) Can we find direct empirical evidence for the traversal of relatively fixed goal/subgoal hierarchies in tasks that do not involve search, as well as evidence for the collapsing of such hierarchies via learning?

(4) What are the implications of abandoning architectural goals and subgoals for the general control structure of cognition?

Toward this end we present a theory of learning hierarchically controlled behavior and its realization in an ACT-R model of a simple HCI task (operating a simulated automated teller machine). ACT-R is given a declarative specification of the task in the form of a goal-subgoal hierarchy. A set of general productions interprets the declarative specification to perform the task, and over time the production composition mechanism (Taatgen & Lee, in press) gradually collapses the goal hierarchy.

The result bears striking similarity to Rosenbloom and Newell's original work on chunking goal hierarchies. One important and novel aspect of the model is the creation of new control symbols that mediate processing and permit efficient navigation of the hierarchy with an extremely small working memory load.

We also present new data from humans performing the task and analyze their reaction times as a function of goal hierarchy depth over multiple trials. Two groups of subjects read slightly different instructions that we hypothesized would induce two different goal hierarchies for the identical stream of button presses. This kind of design and analysis provides more direct evidence for both the hierarchy and its collapse than fitting overall learning curves to a power law. Finally, we draw out the implications of this work for theories of control structure, using existing cognitive architectures (Soar and Epic) as concrete comparison points.

This is collaborative work with Roger Remington, Alonso Vera, Bonnie John, and Mike Mattessa.