Notes for Meeting 17
Structural Learning

Learning in Cognitive Systems

We have discussed how cognitive systems represent knowledge and how they use it for inference, execution, and problem solving.

However, intelligent agents also have the ability to LEARN over time. We can define learning as:

- Given: Input from the environment for some class of tasks;
- Given: Existing content about this class of tasks (optional);
- Generate: New content that improves performance on these tasks.

Many standalone learning methods exist; we will focus on learning that is embedded in cognitive systems.

Varieties of Learning

People regularly learn many different types of long-term content:

- Memorizing factual information
- Storing episodic traces
- Learning stimulus-response pairs
- Acquiring and refining heuristics
- Creating new patterns and concepts
- Constructing new procedures

We will focus on the acquisition of procedural knowledge for use in achieving goals.

Types of Learned Procedural Structures

Research on procedural learning has examined a number of different types of knowledge:

- Action or operator models
- Search-control rules
- Macro-operators
- Hiearchical Task Networks
- Compiled procedures

Different structures are often associated with distinct classes of learning mechanisms.

Sources of Information

Humans can learn procedures from a variety of distinct sources:

- Direction instruction
- Worked-out solutions to problems
- Observations of others' behaviors
- Problem-space search
  - Successful solutions
  - Failed alternatives

All but the first require some form of generalization, at either storage or run time, from specific cases.

Empirical Approaches to Structural Learning

One paradigm for structural learning utilizes empirical methods to acquire procedural knowledge. This framework:

- Collects data about desirable or undesirable choices (e.g., states or operator instances during problem solving);
- Finds similarities among desirable choices or differences from undesirable ones;
- Creates rules or similar structures that specify when to select, reject, or prefer alternatives; and
- Uses these structures as heuristics to guide future execution or problem solving.

Some approaches to empirical learning (e.g., Langley, 1983) operate incrementally, like humans, but most rely on batch processing.

Analytical Approaches to Structural Learning

Another framework for structural learning uses analytical techniques to acquire procedural knowledge. This paradigm:

- Examines a single trace of a problem solution or failed attempt;
- Uses background knowledge to analyze reasons for success or failure;
- Constructs a deductive proof or explanation that identifies the elements on which this result depends;
- Compiles this proof into a rule or similar structure that omits intermediate steps and replaces constants with variables; and
- Uses such structures to guide future execution or problem solving.

Such "explanation-based" methods can learn very rapidly, but they depend heavily on accurate background knowledge.

Laird et al.'s (1984) Soar and Icarus use semi-analytic methods that have fewer guarantees but also require less analysis.

Learning Hierarchical Skills in Icarus

Recall that Icarus invokes a variant of means-ends analysis to solve novel problems.

The problem solver incorporates a simple learning mechanism that:

- Stores with each subproblem the initially satisfied goals and the order in which other goals are achieved;

- Whenever it solves a problem, creates a new skill that:

  - Uses the initially satisfied goals for conditions;

  - Orders subgoals based on the order in which they were achieved;

  - Links the skill's head to the goal that produced the problem.

Thus, whenever Icarus solves a novel problem, it creates a set of skills (HTN methods) that can solve similar ones in the future.

This suggests that hierarchical task networks are generalized traces of successful means-ends analysis.

```
               Open Issues in Structural Learning

Despite substantial progress in learning cognitive structures,
unsolved problems remain, including:

- Creating new conceptual predicates to alter representations;

- Improving the capacity for problem formulation;

- Combining analytical methods for rapid learning with empirical
  methods for knowledge revision; and

- Integrating methods for learning different information sources.

Taken together, advances on these topics would let us develop
more adaptive and robust cognitive systems.


                    Assignments for Meeting 18
                         Spatial Cognition

Read the articles:

- Forbus, K. D., Usher, J., & Chapman, V. (2004). Qualitative spatial
  reasoning about sketch maps. AI Magazine, 25, 61-72.

- Kuipers, B. (2008). An intellectual history of the spatial semantic
  hierarchy. In M. Jefferies & W. K. Yeap (Eds.), Robot and cognitive
  approaches to spatial mapping. Berling: Springer-Verlag. [optional]

- Gunzelmann, G., & Lyon, D. R. (2007). Mechanisms for human spatial
  competence. In T. Barkowsky, M. Knauff, G. Ligozat, & D. Montello
  (Eds.), Spatial Cognition V: Reasoning, action, interaction. Berlin:
  Springer-Verlag. [optional]

- Work on the sixth exercise (due 11:59 PM on 3/30/2011) and bring
  questions about it to class.
```