Learning Rich, Tractable Models of the Real World
MIT9904-09

Progress Report: July 1, 2000—December 31, 2000

Leslie Pack Kaelbling

Project Overview

The everyday world of a household or a city street is exceedingly complex and dynamic, from a robot's perspective. In order for robots to operate effectively in such domains, they have to learn models of how the world works and use them to predict the effects of their actions. In traditional AI, such models were represented in first-order logic and related languages; they had no representation of the inherent uncertainty in the world and were not connected up to real perceptual systems. More recent AI techniques allow model-learning directly from perceptual data, but they are representationally impoverished, lacking the ability to refer to objects as such, or to make relational generalizations of the form: "If object A is on object B, then if I move object B, object A will probably move too."

We are engaged in building a robotic system with an arm and camera (currently, in simulation) that will learn relational models of the environment from perceptual data. The models will capture the inherent uncertainty of the environment, and will support planning via sampling and simulation.

Progress Through December 2000

In September, 2000, we added a postdoctoral researcher, Tim Oates, and two research assistants, Natalia Hernandez and Sarah Finney to the project. During these three months, we have explored a wide range of issues related to relational model learning. In particular, we have studied the use of deictic expressions, both in propositional and in relational representations. In addition, we have designed an architecture for goal-directed model-learning and have implemented most of the algorithmic components. We summarized our work, with a detailed discussion of a number of design issues in the project, in a report titled “Learning in Worlds with Objects”, which is available on the project web site and will be presented at the AAAI Stanford Spring Symposium in March, 2001.
A list of concrete achievements is as follows:

- Implemented Utree algorithm for learning in partially observable domains
- Implemented neuro-dynamic programming algorithm as a basis for comparison for Utree
- Implemented Benson’s algorithm for generating deictic expressions in propositional domains
- Tested Oates’ method for rule induction

- Designed an architecture for combining Utree with rule induction and automatic generation of deictic expressions
- Implemented simple symbolic simulated domain
- Wrote an extensive report on the subject, informed by interactions with our partner, Shigeru Katagiri

Research Plan for the Next Six Months

We are currently integrating the implementations described above. We expected to have empirical results very soon that demonstrate the utility of deictic representations, even in fairly simple domains. Concretely, our plan is to:

- Test Utree and neuro-dynamic programming with deictic and non-deictic representations in simple domain
- Prepare and submit a paper on these results to the International Conference on Uncertainty in Artificial Intelligence
- Add rule-learning and automatic deictic expression generation algorithms to allow the system to learn world model
- Use learned world model to plan to solve new tasks more efficiently
- Invent algorithms for applying these ideas to the first-order case
- Study the role of probability in these representations, with particular emphasis on making an efficient near-deterministic approximation, but generalizing robustly when necessary.
- Design vision algorithms that learn to do object segmentation based first on optical flow, with a goal of making them more robust by allowing the agent to manipulate the objects