

Learning Rich, Tractable Models of the Real World

MIT9904-09

Progress Report: January 1, 2000—June 30, 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 June 2000

Work on this project was largely suspended in the Spring of 2000, due to not having appropriate personnel available. The funds have been carried forward to the next period, allowing us to hire an exciting new post-doctoral researcher, who will work full-time on this project.

Our major contribution during this period was to refine the propositional probabilistic rule-based representation of knowledge about uncertain dynamic systems. We have worked out the semantics of such a rule-based representation in detail, generating both a sampling algorithm and a method for converting rule sets into dynamic Bayesian networks (DBN's). We will not, in practice, convert our rule sets to DBN's, but the conversion algorithm gives a precise semantics to the rule sets.

The sampling algorithm will be of importance in the context of planning. It will allow us to efficiently assess the effects of taking an action in some situation by drawing samples of possible next states according to their

probability distribution. We are currently implementing a very basic sample-based planning algorithm, which will allow us to test the sampling algorithm in complex domains.

Research Plan for the Next Six Months

In the next six months, we expect to move forward experimentally and theoretically. A new post-doctoral researcher, Tim Oates, will join the project in September and work on it full time. Tim's thesis work addresses the question of how a robot can learn linguistic terms from raw sensory experience. His expertise and interest fit directly with this project. In addition, two new PhD students, with other funding sources, will work on topics closely related to this project.

Our work plan for the next six months includes the following tasks:

- Complete implementation of learning algorithm for propositional probabilistic rules
- Test learning algorithm in simple simulated domain
- Compare to standard Bayesian network learning algorithm
- Extend rule formalism to restricted first-order case
- Invent learning algorithm for restricted first-order case
- Build complex simulated robotic hand-eye domain using MathEngine
- Implement vision algorithms for feature-detection in robotic domain
- Invent method for hypothesizing the existence of objects based on visual features of the environment
- Apply first-order rule learning algorithm to visually-detected objects
- Investigate the use of action descriptions that are parameterized by object representations
- Investigate the utility of deictic representations in learning and planning
- Explore connections between learning language and learning the world dynamics

In addition, we will again have a reading group that explores the literature relevant to our project in vision, learning, planning, and language.