

Learning Rich, Tractable Models of the Real World

MIT9904-09

Progress Report: July 1, 1999—December 31, 1999

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 1999

Work on this project really began in September 1999, when it was joined by two students. Our first task was to develop a propositional probabilistic rule-based representation of how the state of a complex world changes from one time to the next, depending on actions taken by the agent. Our representation has a well-defined probabilistic semantics, and for many domains is much more compact than the corresponding dynamic Bayesian network representation. In addition, it directly encodes the basic frame assumptions that most things don't change from one time step to the next. Although we eventually want to move to restricted first-order representations, getting the foundations correct for the propositional case was an important first step.

Once the representation was developed, we began to invent an algorithm for learning propositional probabilistic rules from data. Our algorithm is inspired by Drescher's scheme mechanism, but is in a much sounder probabilistic basis. We have implemented a prototype version of this algorithm, but have not yet experimented with it. We expect that it will have to be modified considerably as we use it and uncover its weaknesses. The rule formalism and the basic algorithm, along with some ideas about how to generalize this to the first-order case are documented in a research note titled "Learning Probabilistic Rules from Experience."

In addition to this, we researched the available software for building simulations of complex physical systems. We chose, and acquired a copy of, MathEngine. It will be able to simulate the actions of a robot arm in a domain consisting of multiple objects; the output of the simulation is rendered visual images, which will be input to the visual-processing stage of our system.

Finally, the PI lead a reading group, consisting of graduate students and researchers drawn from multiple groups within the MIT AI Lab, called "Reifying Robots". We read and discussed a wide variety of papers, all bearing on the question of how object-based representations can be acquired and used by robotic systems. This helped us understand the basic scientific questions better and also pointed out to us what little work there has been in this area.

Research Plan for the Next Six Months

In the next six months, we expect to move forward experimentally and theoretically.

Our first steps will be to experiment with the learning algorithm we have developed. We expect to modify it considerably as we gain experience with it. We will compare its performance with existing techniques for learning dynamic Bayesian networks and write a conference paper summarizing our results. We will then apply this algorithm to a simple propositional version of the simulated robotic test domain we are building.

We will use MathEngine to build at least two simulated domains of different complexities. We will apply and extend existing computer vision methods to perform segmentation and object recognition on the images generated from the simulation, in order to generate input appropriate to the learning algorithm.

We will extend our probabilistic rule format and learning algorithm to a case of limited first-order representation, in which we predict the behavior of particular objects as a function of the robot's actions. In addition, we will investigate the parameterization of the robot's actions with object descriptions, and the use of indexical-functional representation.