

MIT9904-09

Learning Rich, Tractable Models of the Real World

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 simulated robotic system with an arm and 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 in First Year

Work on this project really began in September 1999, when it was joined by one graduate student and one undergraduate. 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 schema mechanism, but rests on 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. By the end of this project year, we expect to have a simulated environment up and running, and to have started on implementation of visual primitives operating on that domain.

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.

Progress in the first year has been slow, as the PI started this project too late to hire an appropriate post-doc. We have started a post-doc search and have three excellent candidates. In addition, we are actively recruiting new graduate students who would come with a strong background in learning probabilistic models. We are confident that we can start the next year with a more mature and effective research team. One effect of not being able to hire an appropriate post-doc and as many students as desired this year is that there is a projected budget surplus from this year of \$110,000. We have, accordingly, reduced our budgetary requests for year two by this amount.

Research Plan for the Second Year

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 extend our probabilistic rule format and learning algorithm to a very limited case of limited first-order representation, in which we predict the behavior of particular objects as a function of the robot's actions. Based on the results of this work, we will experiment with more and less rich representation languages, attempting to find a "sweet spot" in which the representation language gives us computational leverage in planning and learning without adding too much extra complexity.

In addition, we will investigate the parameterization of the robot's actions with object descriptions, and the use of indexical-functional representation. This would allow us to have, for example, a single action of "pick-up-the-block-I-am-looking-at" rather than a parameterized "pickup" action. Parameterized actions are practically and formally complicated, because they require a grounded way of designating the necessary objects. The indexical-functional (or deictic) approach will require addition of perceptual actions, and will actually require the agent to carry out longer chains of actions in order to achieve its goals. However, in very rich perceptual environments, such perceptual actions provide a way of directly controlling attention, which we expect will actually reduce the complexity of the problem overall.

We will invent sample-based planning and behavior-learning algorithms that work with the first-order representations. These will be inspired by work of Kearns et al. on sample-based planning for processes using very simple representations. We will also incorporate recent work by the PI and another one of her students on making sampling more efficient by taking into account the fact that we are simply trying to find the best action (actually, an action within epsilon of the best), rather than getting good estimates for the values of all actions. This insight allows us to make much more efficient sampling algorithms. The extension of these methods to the first-order case will require considerable insight and innovation, but we are prepared to do it.

By the end of the second year, the simulated robot system will be learning and using first-order rules to carry out behaviors in its environment that gain it reward. At this point, we will carry out an experimental study, comparing our representationally more complex learning representation to a more traditional one in the same domain. We feel confident that the change in representational strategy, although seemingly increasing the complexity of the learning problem, will actually decrease complexity and will therefore result in a system that is more effective and efficient than it would have been had we used traditional representation techniques.