Learning Rich, Tractable Models of the Real World

Tim Oates & Leslie Pack Kaelbling

Artificial Intelligence Laboratory
Massachusetts Institue Of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: We tend to think of the world as being made up of objects. There are chairs and apples and clouds and meetings. Certainly, part of the basis for this view is that there are clumps of coherent physical material that tend to be well-described in the aggregate. Even without engaging in the philosophical debate about whether objects really exist, it is hard to imagine a truly intelligent agent that does not conceive of the world in terms of objects and their properties and relations to other objects.

It is crucial for an agent living in our world to be able to take advantage of knowledge of the form:

> If object A is on object B, then if I move object B, object A will probably move too.

Such statements offer an ability to compactly express generalized information that cannot be approached without the description of the world in terms of objects.

We propose to represent the dynamics of the world with models that

- allow strong generalization through representation in terms of objects,
- represent the uncertainty in the world through sampling, and
- never require a complete description of the state of the world.

In this project, we emphasize the object-based representation of world models; this research will occur in parallel with work addressing the other issues.

Motivation: The everyday world of a household or a city street is exceedingly complex and dynamic, from a robot's perspective. Although model-free reinforcement learning enables robots to acquire low-level perceptuo-motor skills, it does not help with flexible behavior in complex high-level domains. 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.

Previous Work: In the last 10 years, there has been a great deal of progress in technical methods for making robust, adaptive systems for uncertain environments. These techniques include neural networks and other function approximation strategies, probabilistic reasoning and Bayesian networks, Markov decision processes and reinforcement learning. These methods have allowed us to build much more robust and effective intelligent systems than before. However, they all have severe representational limitations.

Approach: Our approach is to learn a world model in the form of a restricted set of first-order rules, quantified with probabilities. An example rule might be:

> For all objects x and y, if in(x,y), then after taking the action dump-out(y), not in(x,y) and on-table(x) with probability 0.9.

How can we learn a model like this, and what does it mean?

The first step is to learn a model that describes the world in terms of state variables rather than objects. A very nice, but theoretically ungrounded system for doing this was described by Drescher [1]. We would begin by altering Drescher's method to put it on a formal foundation. This would result in learning rules of the form *If my hand is in front of me, then after I turn my head forward, I will see my hand with probability 0.9.*

The main technical issues involve the semantics of the rules and their associated probabilities. It is important that, given a particular initial state and action, that the rules (implicitly) describe a well-defined probability distribution over possible outcomes. If multiple rules match the current situation, then issues arise as to how to combine their probabilities. A default assumption of independence is probably warranted, with the learning algorithm noticing important correlations and building specific rules to describe them. A more important problem arises when there are multiple rules with the same outcome: the "noisy-or" model can be used to combine probabilities from multiple possible causes of an event. Another important issue at this stage is to see whether more classical Bayesian-network learning would work as well or better than a Drescher-like learning method. If so, we will adopt it as a background instead.

The next step is to extend the model learning to deal with objects. Traditionally, AI systems have employed fairly general logics for representing objects, their properties and relations. However, these general logics typically have intractable inference properties. Our plan is to learn rules of very restricted forms. First, we assume that the perceptual system will deliver, essentially, an existentially quantified formula, such as

> There exist a, b, c, d, and e, such that a is a table, b is a cup, c is a marble, d is a box, and e is a box-top; b is on a, c is in b, d is empty, and e is on d.

An existential input, coupled with universally quantified rules, will yield an existential output; that is, a description of the next state of the world in the same form as the input state was described.

Thus, although it looks like we're using first-order predicate calculus, we're using it in a very restricted way that will maintain tractability. Furthermore, we expect to use planning techniques based on sampling from the distribution of the outcome state, rather than computing the entire probability distribution, which considerably simplifies the inference problem.

Our strategy for learning such rules will combine ideas from Drescher for guiding the search for good rule structures with ideas from the Bayesian network community [2] for deriving correct probability assignments for the rules. The rule-learning algorithm will not be particularly complex, but it will require a large degree of parallelism to search for appropriate rules. It will be possible to test its feasibility on a single computer; but in the second year, we expect to use the loosely-coupled parallelism of a network of workstations.

Difficulty: This problem is very hard but also ready to be addressed. It is relatively easy to see how to start, but there are a number of important issues that will have to be addressed along the way:

- How should we deal with negation?
- How could we synthesize appropriate predicates and relations?
- When is indexical representation useful?

Another very important question that will be addressed is that of the necessity of object-based representations. Some researchers would argue that this representational approach is too "old-fashioned" and that there is no need for direct representation of objects. One study that we will do early on is to solve a problem using non-relational, neural-network type representations and solve it again with the object-based representations. We feel sure that our minimalist approach to object representation will make systems much better able to learn and function in their environments than traditional full-blown logic systems or systems without any representation of objects at all.

Impact: This work may serve to build a bridge between traditional high-level AI knowledge-representation work and real perception-action systems.

Research Support: This work is supported in part by a grant from the Nippon Telegraph and Telephone Corporation and in part by the Office of Naval Research.

References:

[1] Gary L. Drescher. *Made-up Minds: A Constructivist Approach to Artificial Intelligence*. The MIT Press, Cambridge, Massachusetts, 1991.

[2] Daphne Koller and Avi Pfeffer. Learning probabilities for noisy first-order rules. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence*, Nagoya, Japan, 1997.