Using Temporal Belief Networks for Planning

Jak Kirman
Department of Computer Science
Brown University, Box 1910, Providence, RI 02912
jak@cs.brown.edu (401) 863-7695

In our work on mobile robot control, we have adopted a decision-theoretic approach, encoding the problem in a belief network, a probabilistic model of those aspects of the world relevant to the goals of the robot. Belief networks provide a convenient way of modeling uncertain information and faulty sensors. In [Kirman *et al.*, 1991] we show how they can be used to deal with sensors that provide only probabilistic information. For example, a vision system used to identify a target may provide the probability that the target is present in the field of view[Dean and Kirman, 1992]. Our work has been primarily with temporal belief networks; an extension of belief networks that allow the representation of change over time.

Our research has focused on overcoming two difficulties that arise in the use of temporal belief networks. The first problem is in obtaining the conditional probabilities necessary to quantify the network. This may require large amounts of experimentation to gather statistics, and careful choices of the state spaces for the nodes. Some techniques to help cope with this problem are described in [Kirman *et al.*, 1991], and an application to a robotic navigation problem is described in [Basye *et al.*,]. The second problem is the computational expense of belief updating. For polytree networks, the updating is relatively inexpensive, roughly linear in the product of the sizes of the state spaces for the variables represented by the nodes in the network. However, if the state spaces for the nodes are not carefully chosen, this can result in lengthy computations. It is critical therefore to choose state spaces that capture the important distinctions between values but that are not unnecessarily large.

Some of our current research focuses on a computer-aided design environment to assist users in creating temporal belief networks and exploring the tradeoffs involving modeling accuracy and computational complexity. Belief networks present a concise graphical representation of probabilistic dependency, allowing the designer to visualize the relationships between the entities of interest in the problem; thus an interactive, visual graph-editing capability is helpful for visualizing the problem. Tools are provided to manipulate the network and to obtain performance and complexity measures for it.

We provide tools that perform topological manipulations, for example to generate the join tree of a network (an equivalent network that is a polytree), or to eliminate nodes that do not affect the measure of utility (barren nodes). We also provide tools to perform belief updating on the network in different ways, some using stochastic sampling and others using exact methods. A different category of tools provides ways of measuring aspects of interest. Some measures are of a theoretical nature, providing for example upper bounds on the time necessary to update beliefs in the network, or bounds on the amount of experimentation necessary to quantify the network. Other measures are empirical; in some cases it may not be possible to provide exact measures of the effect of a change to a network such as the removal of a dependency, but the change can be evaluated empirically using stochastic sampling techniques.

The major bottleneck in planning performance comes from the expense of belief updating. In order to help reduce the complexity of this operation, the environment provides guidance to the user in indicating where simplification would have a large impact on efficiency, where state spaces can be reduced without great loss of accuracy, and what modifications might be most beneficial.

References

[Basye *et al.*,] Basye, Kenneth; Dean, Thomas; Kirman, Jak; and Lejter, Moises. A decision-theoretic approach to planning, perception and control. *IEEE Expert*. Forthcoming.

[Dean and Kirman, 1992] Dean, Thomas and Kirman, Jak 1992. Representation issues in bayesian decision theory for planning and active perception. In *Proceedings of the DARPA Image Understanding Workshop*. DARPA.

[Kirman *et al.*, 1991] Kirman, Jak; Basye, Kenneth; and Dean, Thomas 1991. Sensor abstractions for control of navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 2812–2817.