A Bayesian approach to object identification

Hanna Pasula
pasula@cs.berkeley.edu
Computer Science Division
University of California, Berkeley

There are many real world domains where an agent can observe the world state only partially and intermittently, using noisy sensors. Merely keeping track of the objects present in such a system is non-trivial. The problem may be complicated further if the system dynamics are not fully known or unpredictable, so that some on-line learning is necessary. I have been working on a principled approach to state estimation and prediction under these realistic conditions. So far, I have focused mostly on *object identification*, deciding if some newly observed object is the same as a previously observed one. The work has been applied to the surveillance of a large metropolitan freeway system. The vehicles are observed through scattered cameras, and both the viewing and traffic conditions are highly variable.

The general approach is based on a probabilistic formulation of the domain. The conditional independence and variable density distribution assumptions are made explicit, so the domain can be expressed as a *dynamic probabilistic network* (DPN). The right assumptions can simplify the problem considerably, although care must be taken to make them realistic. As an example, consider the traffic surveillance domain. There, we can assume that certain observations are dependent on the sensor where they are made, and the object generating them, but not directly dependent on what happens at other sensors. Color, for example, is constant for each vehicle, and the way it appears at each observation site depends on the characteristics of that site. The introduction of hidden variables representing such features of objects in the system enables us to remove many inter-sensor dependencies. As a result, we can use small, local probability models. We learn model parameters using the *Expectation Maximazation* (EM) algorithm.

Unfortunately, inference in large, complex problems will remain intractable even if the probabilistic network is simplified. The correspondence between observations and objects is unknown, and so inference in the system should involve summing over all the possibilities. Since the number of possible observation-object assignments is exponential in the number of objects, solving such a summation exactly is a $\#P$ problem. Our solution has been to depend on an approximation method based on *Monte Carlo Markov Chain* (MCMC) sampling. We define a Markov chain over the space of all observation-object assignments. Then, we define the probabilistic transitions in the chain so that the stationary distribution is equivalent to the probability distribution over these assignments, given all the observations. The chain is run to convergence, and then the samples it generates are used to approximate this distribution. Jerrum and Sinclair [1997] have shown that this method is a tractable and efficient approximation to an explicit summation. Constructing the chain is not difficult, and running it involves some specialized DPN inference, which can be made very efficient. One important application of this method is the approximation of the hidden variable values called for by the EM algorithm.

Interacting with the world in a real-time manner is a further problem. For one, we have to be able to keep up with changing parameter values. This can be done by calling EM at intervals, but, if every EM iteration depends on a full MCMC run, this may be slow. We are currently experimenting with on-line EM, where each call involves only a single iteration. Then, there is the question of how to run MCMC on a changing state space. This space will expand as new observations are made, and we need a principled way of cutting out old observations. At the same time, we may want to keep the old information around for as long as it may help us improve our models. A simple answer might be to simply incorporate the information from dropped observations into the parameter priors. It might, however, prove more useful to add a meta-level which would reason about issues such as the information content of old data, and the long-term changes of the models. This level could also be helpful when the approach is extended to predict, as well as estimate, the world state.

The system has been tested extensively on simulated data, as well as on small quantities of real data. It has been shown to outperform existing approaches, such as that of Huang and Russell [1998].

To conserve space, full citations are available upon request.