Probabilistic Graphical Models for Diagnosis and Decision Making

Finn V. Jensen, Aalborg University, fvj@cs.aau.dk

The core of a probabilistic graphical model is a directed graph, where the links represent cause-effect relations (See Figure 1).

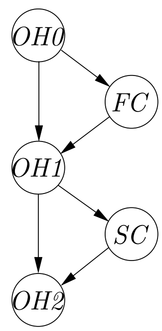

Figure 1: A structure for a poker game. The variables $OH0$, $OH1$, and $OH2$ represent my opponent's hands during play. FC and SC represent her change of cards.

When the strength of the causal links are represented by conditional probabilities, then the structure is called a *Bayesian network*. Bayesian networks are primarily used for *belief updating*: given case specific evidence, what is then the posterior probability distributions for some specific variables in the model. For example, what is $P(OH2 \,|\, FC = 3, SC = 1)$? The knowledge representation offers other features like adaptation to experience, analysis of conflicting evidence, and sensitivity analysis.

Bayesian networks can in various ways be extended to also represent decision scenarios. You introduce nodes representing decisions and nodes representing utilities (gains and losses). See Figure 2. A solution of a decision graph is a set of optimal policies, one for each decision node.

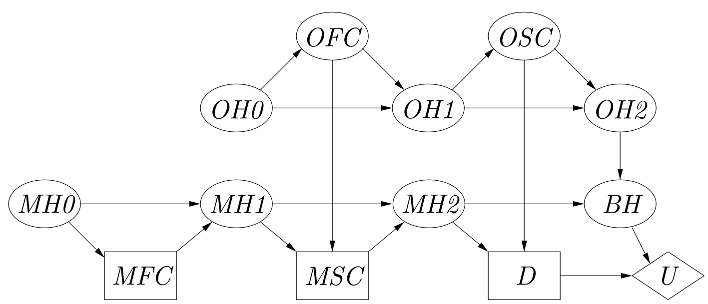

Figure 2: The poker model extended with decision nodes (MFC, MSC, D), and a utility node (U).

There are rather good algorithms for learning Bayesian networks from databases. With respect to decision graphs, learning of preferences from databases is still a major scientific challenge.