Bayesian networks

Uffe Kjærulff Finn V. Jensen

Department of Computer Science
Aalborg University
Fredrik Bajers Vej 7E
DK-9220 Aalborg O
Denmark

{uk,fvj}@cs.auc.dk

18 November 1996

1 What is a Bayesian network?

A Bayesian network is a model of a problem domain, representing (causal) relations among domain variables, and is used for calculating probability distributions of the unobserved variables given the observed variables. The relations between the variables are given as conditional probabilities. Thus a Bayesian network consists of two parts: (1) a qualitative (or structural) part, and (2) a quantitative (or numerical) part. The qualitative part is given independently of the quantitative part, and encodes the causal relations among the domain variables in the form of an acyclic, directed graph (DAG). Furthermore, this DAG comprises a compact description of the assumed independence relations among the domain variables.

A sample Bayesian network is shown in Figure 1. This network is an oversimplified model providing decision support in the problem domain of insemination of cows. The network displays the causal relations among the domain variables. Pregnancy (P) causes hormonal changes (i.e., affects the hormonal state (H)) and shadings on scanning pictures, and hormonal changes cause changes in the outcomes of both blood test (B) and urine test (U). Now, the outcome of a blood test gives an indication of whether or not the insemination was successful. So, we want to calculate the effect in the opposite direction of the arrows, which can be done very efficiently using inference algorithms based on local computations.

Read in the opposite direction, the relations indicate that changes in hormonal state and shadings on scanning pictures are both symptoms of pregnancy, as well as the outcomes of both blood test and urine test are indicators of hormonal state. Thus, indirectly, pregnancy can be detected via both tests. So why aren't there causal links from P to B and U? This is where the independence relations embedded in the model come into play.

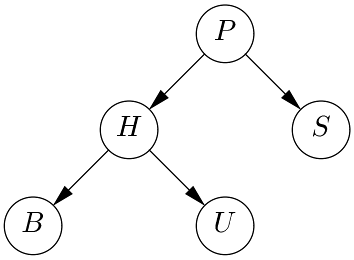

Figure 1: A model for test of pregnancy (P). Both blood test (B) and urine test (U) give indication of the hormonal state (H). Both the hormonal state and scanning (S) give indication of pregnancy.

The lack of arrows from P to B and U tells us that the model constructor has made the (reasonable) assumption that given that the hormonal state is known, information about the outcomes of blood test and/or urine test do not affect our degree of belief in the cow being pregnant or not. Similarly, knowing with certainty the hormonal state of the cow, the model constructor assumes that information about the outcome of a blood test does not provide any information about the possible outcomes of a urine test. Hence, he has put no link between B and U. On the other hand, when the hormonal state is not known with certainty, the outcome of a blood test does provide information about the possible outcomes of a urine test. The rules employed to read off the independence assumptions embedded in a Bayesian network are simple and easy to use.

Since the state of H will never be observed directly and since there is no immediate interest in the actual degree of belief associated with the various states of H, why bother to include it in the model? That is, why not just let B and U be immediate causal descendents of P, avoiding H? That alternative model would, however, not be equivalent to the model in Figure 1, which entails dependence of B and U even though the state of P is known. The alternative model, on the other hand, implies independence of B and U when the state of P is known, which is wrong! The reason why it is wrong lies in the fact that the causal relation between P and H is stochastic (i.e., pregnancy does not necessarily cause (significant) changes in the hormonal state, which is required for the outcomes of blood test and urine tests to be independent). In other words, the alternative model leads to over-estimation of the probability of P when both B and U are observed. So, intermediate variables like H may play an important role in Bayesian networks.

It is important to note (as illustrated above) that the causal relations are generally non-deterministic (i.e., stochastic). These stochastic relations are modelled through conditional probability distributions. For example, the causal relation between P and S is modelled via a conditional probability distribution over the states of S for each state of P (see Table 1), stating, for example, that pregnancy causes shadings on scanning pictures with probability 0.9.

To each variable, X, in a Bayesian network is associated a table of conditional probability distributions, one distribution for each possible combination of states, $y, \ldots, z$,

	$P = y$	$P = n$
$S = y$	0.9	0.01
$S = n$	0.1	0.99

Table 1: The conditional probability distributions over the states of S given the states of P.

of the parents, $Y, \ldots, Z$, of X (variables with links pointing towards X):

$$p(X \mid y, \ldots, z) = p(x_1, \ldots, x_n \mid y, \ldots, z),$$

which reads "the probability distribution of X given $Y = y, \ldots, Z = z$", where $x_1, \ldots, x_n$ are the possible states of X. Note that the tables for orphan variables specify marginal (i.e., unconditional) probability distributions. The product of all probability tables of a Bayesian network equals the joint probability distribution over all variable of the network. That is, the joint distribution factorizes according to the structure of the network, which is a consequence of the independence relations encoded in the network structure. This factorization of the joint probability distribution is also the corner stone of the efficient inference algorithms developed for Bayesian networks.

2 How to create a Bayesian network?

There are two main questions connected to the construction of a Bayesian network:

- Where does the structure of the network come from?
- Where do the numbers (i.e., probabilities) come from?

There are two main sources of information, namely *domain experts* and *data*. Thus, possible answers to the above questions are:

1. From domain experts.
2. From data.
3. From domain experts and data.

Most often, however, the structure is provided by domain experts alone, but tools exist for data-based construction of the network structure. The probability tables, on the other hand, are often generated from data using a variety of statistical methods. It should be emphasized, however, that the probabilities need not be "objective" (i.e., based on statistics), but may be based on subjective assessments by domain experts. Combining expert knowledge and statistical methods often provides the best practical solution to constructing robust models.

3 Modelling decisions in Bayesian networks

Bayesian networks are models providing decision support. Therefore, features for explicit modelling of decision actions is an obvious extension to Bayesian network models. Such an extension can be achieved through addition of decision and utility nodes, resulting in models which are often called influence diagrams.

Say we wish to extend our insemination network with a decision node, D, representing the insemination decision. Suppose that six weeks after the insemination there are the two actions w (wait another six weeks) and r (repeat the insemination). Assume that the decision to perform an insemination is based on the probability of pregnancy, the utility of pregnancy, and the cost of performing an insemination. Thus, for each combination of states of P (y and n) and D (w and r), we must specify an associated utility, U. Let the utility of pregnancy be 200 units, the cost of keeping the cow non-pregnant for six weeks be 30 units, and the cost of insemination be 10 units. So, the utility of $D = w$ given $P = y$ is 200 and if $P = n$ we cannot hope for more than $200p - 2 \cdot 30 - 10$, where $p = 0.87$ is the success rate of inseminations. The utilities of $D = r$ is $200p - 30 - 10$ regardless of the state of P (see Table 2).

	w	r
y	200	134
n	104	134

Table 2: Utility table $U(P, D)$ for insemination.

Since U depends on P and D, the original Bayesian network model is extended by adding "causal" links from P and D to U (see Figure 2).

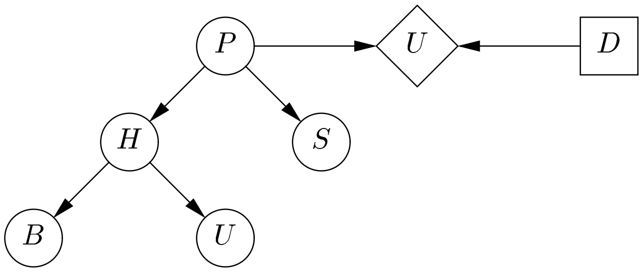

Figure 2: The insemination model extended with an action node and a utility node.

Given this influence diagram, we may now calculate the expected utilities given actions w and r, i.e., $EU(w)$ and $EU(r)$. Six weeks after the insemination we have $p(P) = (0.87, 0.13)$. Thus,

$$EU(w) = U(y, w)p(y) + U(n, w)p(n) = 200 \cdot 0.87 + 104 \cdot 0.13 = 188$$

and

$$EU(r) = U(y, r)p(y) + U(n, r)p(n) = 134 \cdot 0.87 + 134 \cdot 0.13 = 134.$$

So, as expected, in this situation the optimal action is w. Now, negative blood and urine tests may result in a revised belief about P, which in turn may result in $EU(r) > EU(w)$ meaning that the model recommends a repeated insemination.

Influence diagrams often includes several decision nodes representing a sequence of dependent decision actions.

4 When is Bayesian networks the appropriate tool?

Hopefully, the above discussion has shed some light on the question of when Bayesian network models are appropriate solutions to building decision support systems. That is, what characterizes a problem domain for which Bayesian networks are suited as domain models? Similarly, when are other approaches more appropriate? In summary, the ideal domain for the Bayesian network approach has the following characteristics:

- There is a well established knowledge of the *causal structure* of the domain.
- Parts of the causal relations are of *stochastic* nature (i.e., the same cause may in different cases have different impact).
- The causal structure is *chained* (i.e., a cause has an impact which again has another impact, etc.).
- The events in the domain are *well-defined.*
- The set of all possible cases forms a *closed world* of manageable size.
- The task is *therapeutic* (i.e., by observing symptoms and risk factors you estimate the probabilities of the possible causes, and from this you decide for a treatment).

Domains for which you should use other approaches has the following characteristics:

- The events are ill-defined (like "heavy", "strong", "pain", etc.). Then fuzzy logic may be a good choice.
- The causal structure is unknown and not of any interest (e.g., classification of handwritten letters). Then neural networks may be a good choice.
- The causal structure is deterministic. Then propositional logic may be adequate.

Other tutorial texts on Bayesian networks can be found on the homepage of Hugin Ltd.: `http://www.hugin.dk`. It also contains references to further information.