Bayesian networks for logical reasoning

Jon Williamson*

Automated Reasoning Workshop 2001

Abstract

By identifying and pursuing analogies between causal and logical influence I show how the Bayesian network formalism can be applied to reasoning about logical deductions.

1 Preliminaries

A *Bayesian network* consists of a directed acyclic graph $\mathcal{G}$ over variables $C_1, \ldots, C_N$ together with a set of specifying probability values $\mathcal{S} = \{p(c_i|d_i) : d_i$ is a state $\pm c_{j_1} \wedge \ldots \wedge \pm c_{j_k}$ of the parents $D_i = \{C_{j_1}, \ldots, C_{j_k}\}$ of C_i in G, $i = 1, \ldots, N\}$. Now, under an independence assumption, namely that given its parents D_i, each node C_i in G is probabilistically independent of any set S of other nodes not containing its descendants, $p(c_i|d_i \wedge s) = p(c_i|d_i)$, a Bayesian network suffices to determine a probability function p over the variables. In particular, if the C_i are *causal* variables, and the graph $\mathcal{G}$ represents the causal relations amongst them, with an arrow from C_i to C_j if C_i is a direct cause of C_j, then it is thought that $\mathcal{G}$ will necessarily be acyclic, and that the independence assumption will be a valid assumption. If its graph is sparse then a Bayesian network may offer an efficient representation of a probability function: for example the space required to store a network is linear in N when there is a bound on the number of parents a node may have, and the time taken to infer probability values from the network is linear in N if the graph is singly-connected. It is generally presumed that causal graphs are simple enough for Bayesian networks to offer an efficient representation.

The components $\mathcal{G}, \mathcal{S}$, of a Bayesian network may often be thought of as the background knowledge of some rational agent, X say. The theory of objective Bayesianism holds that X's degrees of belief may be represented by a probability function, and that there is a *most rational* probability function μ that X should adopt. μ is the probability function that (i) is consistent with X's background knowledge $K = (\mathcal{G}, \mathcal{S})$, and (ii) maximises the entropy function $-\sum p(s) \log p(s)$ where the sum is taken over all states s of the form $\pm c_1 \wedge \ldots \wedge \pm c_N$.

An asymmetric relation R is an *influence relation* if the following condition holds:

*Department of Philosophy, King's College London WC2R 2LS UK. Email: jon.williamson@kcl.ac.uk.

- **irrelevance:** if $K_1 = (\mathcal{G}_1, \mathcal{S}_1)$ consists of the components of a Bayesian network on nodes $C_1, \ldots, C_n$, and $K_2 = (\mathcal{G}_2, \mathcal{S}_2)$ is the Bayesian network formed by adding an extra node C_{n+1} that is not an ancestor of any of $C_1, \ldots, C_n$, then $\mu_{K_2}^{C_1,\ldots,C_n} = \mu_{K_1}^{C_1,\ldots,C_n}$, i.e. the restriction of the new most rational belief function to the old domain is the same as the old most rational belief function. The new knowledge is irrelevant to beliefs on the old domain.

Causality, for example, is an influence relation: coming to learn that two variables have a common effect should not change one's beliefs in the original variables. It turns out that if $\mathcal{G}$ is knowledge of an influence relation, then the probability function that maximises entropy subject to the constraints imposed by $(\mathcal{G}, \mathcal{S})$ satisfies the Bayesian network independence assumption. The most rational belief function is just the probability function determined by the Bayesian network on $(\mathcal{G}, \mathcal{S})$.

2 Logical networks

A logical proof corresponds to a directed acyclic graph in which each node is a line in the proof and the parents of a node are the lines that are used in the logical step that yields that node. A *logical network* is the Bayesian network formed by such a graph $\mathcal{G}$ together with the appropriate probability specifiers $\mathcal{S}$. Logical implication is an influence relation and so, if the components $\mathcal{G}$ and $\mathcal{S}$ are thought of as background knowledge, one can use the Bayesian network to represent the most rational belief function that incorporates this knowledge, and to calculate desired probabilities. If there is an arrow from C_i to C_j in $\mathcal{G}$, then C_i is called a *logical influence* of C_j.

A causal network offers an efficient representation of a probability function in that it contains little redundant information. This is because (i) the independence assumption holds; (ii) each node raises the probabilities of its children; (iii) the graph is likely to be sparse. These features all carry over to logical networks.

Thus far a logical network is constructed from a finished proof. But we can also construct a logical network to reason about a partial proof. A logical network need not include every step in the proof, may leave out some of the logical influences on which the proof depends, and may include logical influences which are not part of the correct proof.

So logical networks provide a formalism for reasoning about proofs: they allow one to calculate probabilities and conditional probabilities of sentences involved in a proof. One can plan a proof by choosing a probable conjecture, finding plausible hypotheses and lemmas that render the conjecture most probable, and at each stage applying a rule of inference that is most likely to yield a desired sentence.

See `http://www.kcl.ac.uk/philosophy.ai/` for the full paper.[1]

[1]Thanks to David Chart, David Corfield and Donald Gillies for useful comments, and to the Arts and Humanities Research Board for funding this research.