Reasoning with Uncertainty

Exercise 1

a. A reasoning system employs the certainty-factor calculus for reasoning with uncertain information. Consider the following five uncertain implications (rules):

$$\{R_1 : a \vee b \rightarrow f_{0.5},$$
$$R_2 : b \rightarrow e_{0.5},$$
$$R_3 : c \rightarrow e_{1.0},$$
$$R_4 : c \wedge d \rightarrow f_{0.8},$$
$$R_5 : e \rightarrow f_{0.9}\}$$

We wish to derive how certain we are about proposition f. Furthermore, it is known that $a_{0.7}$, $b_{0.8}$, $c_{0.5}$ and $d_{0.7}$ are true with the associated uncertainty, which are shorthand notations for $\mathrm{CF(a, e')} = 0.7$, etc. Determine $\mathrm{CF}(f, e')$.

b. The certainty-factor calculus employs the following combination functions to compute the conjunction and disjunction, respectively, of uncertaint propositions:

$$\mathrm{CF}(e_1 \wedge e_2, e') = \min\{\mathrm{CF}(e_1, e'), \mathrm{CF}(e_2, e')\}$$

$$\mathrm{CF}(e_1 \vee e_2, e') = \max\{\mathrm{CF}(e_1, e'), \mathrm{CF}(e_2, e')\}$$

Show by means of a counter-example that these functions are incorrect from the point of view of probability theory.

c. (This exercise is similar to exercise a). A reasoning system employs the certainty-factor calculus for reasoning with uncertain information. Consider the following four uncertain implications (rules):

$$\{R_1\colon\ a \vee b \vee c \rightarrow f_{1.0},$$
$$R_2\colon\ c \wedge d \rightarrow f_{0.5},$$
$$R_3\colon\ f \rightarrow g_{0.2},$$
$$R_4\colon\ e \rightarrow f_{0.6}\}$$

The user supplies the following certainty factors 0.8, 0.4, 0.7, 0.6 and 1.0 for the propositions a, b, c, d and e. Compute the certainty factor $\mathrm{CF}(g, e')$ for proposition g using these facts and rules.

d. Show that the combination function for co-concluding rules $f_{\mathbf{co}}(x, y) = x + y(1 - x)$ is not idempotent. (Idempotency of an operator o means that if you apply the operator again, you get the same result, i.e., $o(z) = o(o(z))$.) Do you think that idempotency would be a good property of $f_{\mathbf{co}}$? Provide motivation for your solution.

Exercise 2

Let $\mathcal{B} = (G, P)$ be Bayesian network with acyclic directed graph $G = (V(G), A(G))$ and associated joint probability distribution P, as shown in Figure 1.

a. Which probabilistic information must be locally available for vertex V_2 to compute locally, i.e., in V_2, the (marginal) probability distribution $P(V_2)$?

b. Suppose that $V_1 = \mathit{true}$ has been observed. Determine now $P^{\mathcal{E}}(V_3) = P^{V_1=\mathit{true}}(V_3) = P(V_3 \mid V_1 = \mathit{true})$.

c. Now, suppose that in addition to $V_1 = \mathit{true}$ also the value *false* has been observed for variable V_3. Determine now the probability distributions $P^{\mathcal{E}}(V_i) = P(V_i \mid \mathcal{E})$, for $i = 1, 2, 3$ and $\mathcal{E} = \{V_1 = \mathit{true}, V_3 = \mathit{false}\}$.

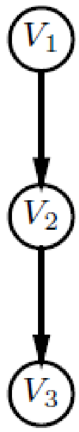

$$P(v_1) = 0.4, P(\neg v_1) = 0.6$$

$$P(v_2|v_1) = 0.3, P(\neg v_2|v_1) = 0.7$$
$$P(v_2|\neg v_1) = 0.9, P(\neg v_2|\neg v_1) = 0.1$$

$$P(v_3|v_2) = 0.7, P(\neg v_3|v_2) = 0.3$$
$$P(v_3|\neg v_2) = 0.1, P(\neg v_3|\neg v_2) = 0.9$$

Figure 1: Bayesian network; the notation $\neg v$ stands for $V = \mathit{false}$ and v for $V = \mathit{true}$.

Exercise 3

Major parts of the certainty-factor calculus can be translated to Bayesian networks. This offers a lot of insight into the nature of this calculus.

a. Give the translation of the combination function f_{prop} from the certainty-factor calculus, and give the conditions that must be satisfied for the resulting probability distribution P.

b. Show how we might redefine f_{prop} such that is comes closer to probability theory in general by dropping one of the conditions. How would an uncertain implication look like in this case?

c. The combination function f_{co} can be interpreted in probability theory as a noisy-OR Bayesian network. Show that the combination function would look like if the noisy-OR definition is replaced by the noisy-AND. (In the noisy-AND a logical AND is used to determine the probability distributions $P(E \mid C_1, \ldots, C_n)$.)

Exercise 4

Consider the following sentences:

> Bob teaches the course `AI101`. John and Mary will likely (with probability 0.8) follow `AI101`. The person who follows a course is a student of a person that teaches a course. For each student of Bob, Bob becomes happy with probability 0.5.

a. Given this knowledge, what is the intuitive chance that Bob will be happy?

b. Model the knowledge above to Horn clauses and find all (SLD) derivations that prove that Bob is happy.

c. Use these derivations to answer the following questions:

- What is the probability that Bob will be happy? Does this coincide with the answer that you gave in (a.)?
- If Bob is happy, what is the probability that John follows the course?
- If Bob does not have a student, what is the probability that Bob is happy? Do you agree with this conclusion given the knowledge above?

d. Estimate how many students are needed like Mary and John to make sure that the probability of Bob being happy is at least 95%.

Exercise 5

A physician observes a symptom S which is indicative of a disease D. The prevalence of the disease is 30% of the admitted population. The symptom will be observed 70% or 10% of the time, depending on whether the disease is present or absent, respectively. The physician can give treatment T . The utility of this treatment is given by $u1(d, t) = 100$, $u1(d, \neg t) = 100$, $u1(\neg d, t) = 10$, $u1(\neg d, \neg t) = 0$. Furthermore, there is a cost to treatment: $u2(t) = 20$, $u2(t) = 0$. Assume the utilities are additive. Determine the optimal policy for T and compute the expected utility given this policy.

Exercise 6*

In a television show, the contestant must choose between three closed doors. Behind one door a prize awaits (a car). Behind both other doors are goats. The contestant chooses a door, e.g. number one. The host, who knows where the car is, opens another door, e.g. number three, and a goat appears. The contestant is now given the opportunity to choose between the two remaining doors (one and two). What is the better choice from this point of view? To stay with the door he originally chose, or to switch to the other closed door? Model this problem using a couple of well-chosen random variables, and show by probabilistic reasoning what is the best choice (you only need probabilities, although one can always frame such a problem in a decision network as well if you like, but this would make things more complex and one would need to decide how a car compares to a goat).