Solutions for Bayesian networks and decision graphs (second edition)

Finn V. Jensen and Thomas D. Nielsen

May 19, 2008

Solution for exercise 9.1

(i) $EU(Gd) = 14.7, EU(SB) = 15.5, EU(Dg) = 15.1$.

(ii) Let v_i denote the utility attached to the mark i. Then $EU(Gd) - EU(Sb) = 0.1v_0 - 0.1v_3$. Hence Gd can only be better than Sb if the mark 0 is given higher utility than the mark 3.

Solution for exercise 9.2

We say that two utility functions $U_r(X)$ and $U_s(X)$ are equivalent if for any two probability distributions $P_1(X)$ and $P_2(X)$ we have that $\mathrm{EU}_1^r \geq \mathrm{EU}_2^r$ if and only if $\mathrm{EU}_1^s \geq \mathrm{EU}_2^s$ (or similarly, $\mathrm{EU}_1^r - \mathrm{EU}_2^r \geq 0$ if and only if $\mathrm{EU}_1^s - \mathrm{EU}_2^s \geq 0$). By setting $U_r = aU_s + b$ we have $\mathrm{EU}_i^r = \sum_X P_i(X)U_r(X) = \sum_X P_i(X)(aU_s(X) + b) = a(\sum_X P_i(X)U_s(X)) + b = a\,\mathrm{EU}_1^s + b$, and therefore $\mathrm{EU}_1^r - \mathrm{EU}_2^r = (a\,\mathrm{EU}_1^s + b) - (a\,\mathrm{EU}_2^s + b) = a(\mathrm{EU}_1^s - \mathrm{EU}_2^s)$. Since $a > 0$ we have $\mathrm{EU}_1^r - \mathrm{EU}_2^r \geq 0$ if and only if $\mathrm{EU}_1^s - \mathrm{EU}_2^s \geq 0$.

Solution for exercise 9.3

See the Hugin network.

Solution for exercise 9.4

See the Hugin network.

Solution for exercise 9.5

See the Hugin network. Note that all the decisions can, in principle, be taken independently of each other, but since influence diagrams require a linear temporal ordering of the decisions we are forced to insert "artificial" informational arcs.

Solution for exercise 9.6

If $a_i = \arg\max_A \mathrm{EU}(A)$, then $\mathrm{EU}(a_i) \geq \mathrm{EU}(a_j)$ for all $j \neq i$. In particular, since $\mathrm{EU}(a_k) = \sum_H P(H)U(H, a_k) = P(h_k)$ we have $\mathrm{EU}(a_i) = P(h_i) \geq P(h_j) = \mathrm{EU}(a_j)$ for all $i \neq j$.

Solution for exercise 9.7

See Figure 1. To complete the decision tree you need to specify $4+4\cdot 3+36\cdot 3 = 124$ numbers.

Solution for exercise 9.8

The optimal strategy is to choose a_3. If the observation is the upper branch, then choose d_2, if lower branch, then choose e_2. The expected utility of the strategy is 1.905.

Solution for exercise 9.9

(i) See the Hugin network.

(ii) past(Action) $= \{MH0, MFC, MH1, OFC, MSC, MH2, OSC, MTC, MH, OTC\}$, past(MTC) $= \{MH0, MFC, MH1, OFC, MSC, MH2, OSC\}$, past(MSC) $= \{MH0, MFC, MH1, OFC\}$, and past(MFC) $= \{MH0\}$.

Solution for exercise 9.10

The partial temporal orderings are the same in both figures: $\{T_1\} \prec FV_1 \prec \{T_2\} \prec FV_2 \prec \{T_3\} \prec FV_3 \prec \{T_4\} \prec FV_4 \prec \{T_5\} \prec FV_5 \prec \{V_1, V_2, V_3, V_4, V_5\}$.

Solution for exercise 9.11

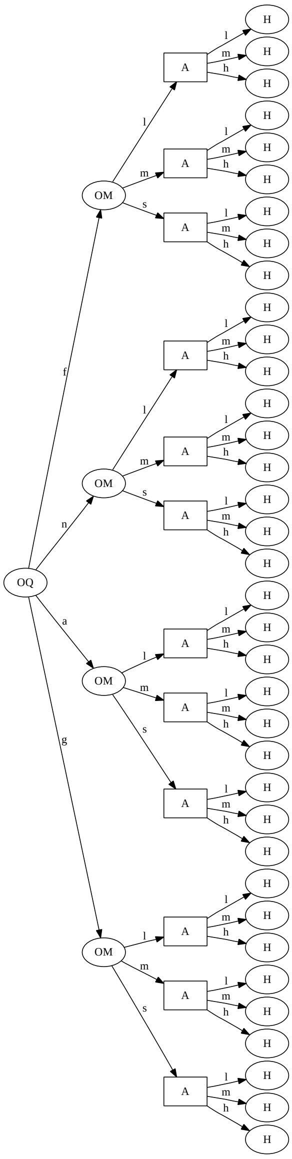

Figure 1: Decision tree representation of the mildew problem; due to the size of tree, the utility nodes are not shown.

(i) Take the probabilities from the Hugin network. The strategy is to test, and if the result is n, then do not drill - else drill. The expected utility of the strategy is $22,500.

(ii) See the Hugin network.

Solution for exercise 9.12

(i) The optimal strategy is to take test t_2. If there are no defects then you should buy the car. If there are one or two defects, then the car should be bought with an "antilemon guarantee".

(ii) See the Hugin network.

Solution for exercise 9.13

This is quite an easy exercise, since the solution can be found in Figure 9.5 in the book.

Solution for exercise 9.14

See the Hugin network and Figures 2 and 3.

Solution for exercise 9.15

See the solution to Exercise 9.14.

Solution for exercise 9.16

See the Hugin network.

Solution for exercise 9.17

See the Hugin network.

Solution for exercise 9.18

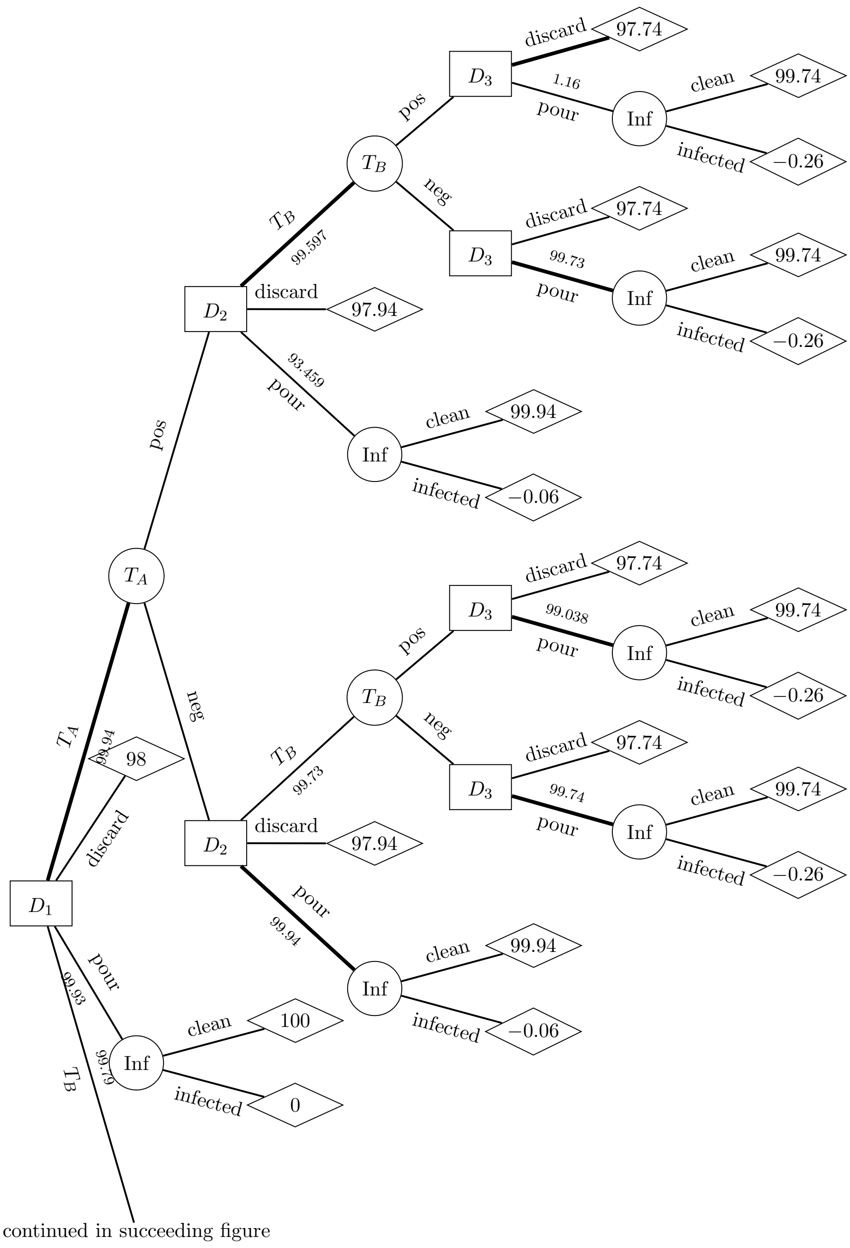

Figure 2: The graphical part of a decision tree for the milk problem from Exercise 9.14. The thick arcs represent the optimal strategy. Note that the decision tree have been modified according to the errata list.

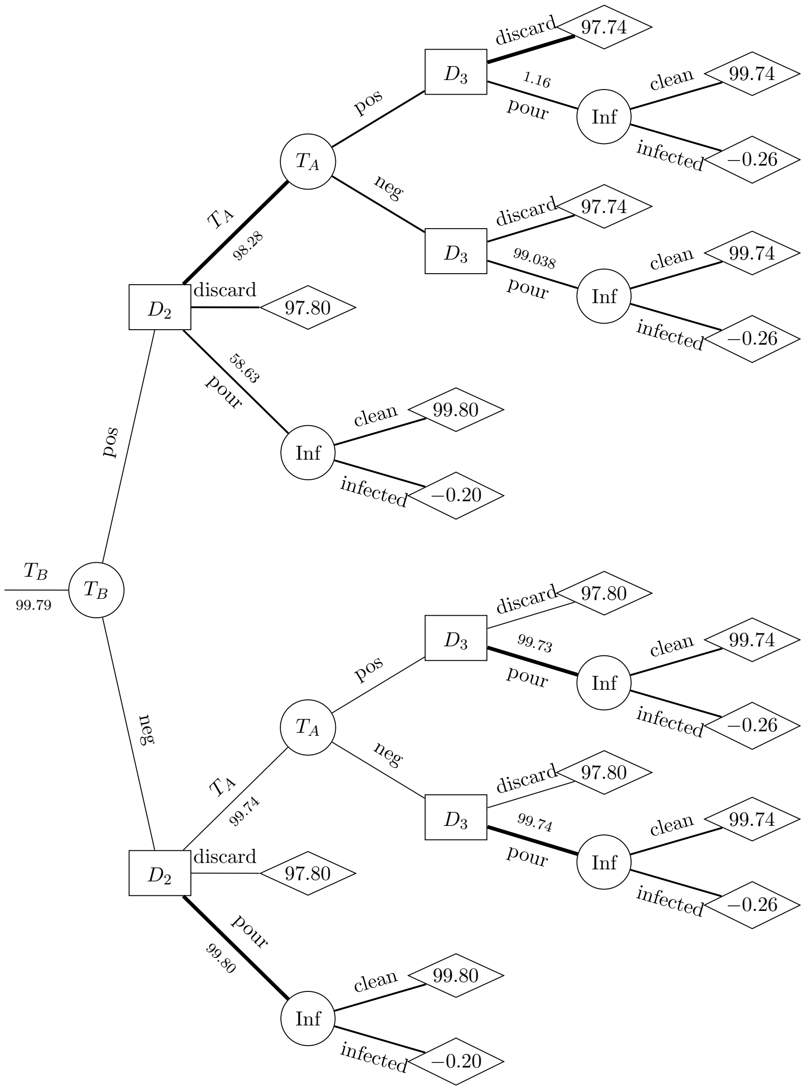

Figure 3: Continuation of diagram in Figure 2.

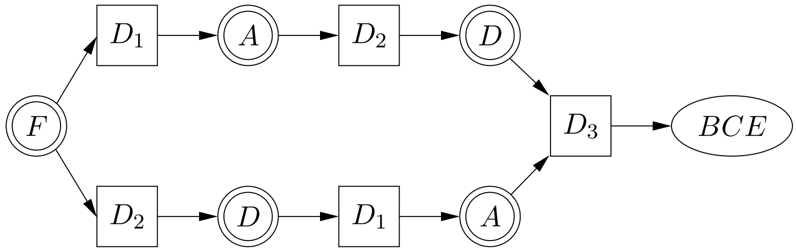

Figure 4: An S-DAG for the UID in Figure 9,28.

See the Hugin network.

Solution for exercise 9.19

See the Hugin network.

Solution for exercise 9.20

See Figure 9.17.

Solution for exercise 9.21

There are several types of asymmetry present in the dating problem. If Joe decides to ask Emily out (and she accepts), he will have the option of either going to a night club or staying at home (*structural asymmetry*). If Joe decides to go to a night club, then his overall night club experience will depend on the live music and whether he meets any friends their (*order asymmetry*); these two observations are modeled as two variables with no pre-specified observation order. At first glance, this example of order asymmetry may seem a bit artificial, since the observation order does not influence the optimal strategy. On the other hand, from a modeling perspective it may be advantageous to be able to specify such partial orderings. Note that an example of a partial order that *does* influence the optimal strategy is given in the diagnosis problem (Example 9.5). The dating problem does not contain an example of *functional asymmetry*, but an example of this can be found in the reactor problem (Example 9.3).

Solution for exercise 9.22

See Figure 4.

Solution for exercise 9.23

In order to represent the tic-tac-toe game as a Markov decision process, we need to represent the possible states of the world (the possible board-positions) and the possible actions of our player (let's call him A). To represent the board positions, one can use a factorized representation consisting of two variables S_A and S_B encoding the positions of the pieces for each of the players. The states of each variable correspond to the possible positions of the players' pieces on the board. Thus, each variable will have $\binom{9}{1} + \binom{9}{2} + \binom{9}{3} = 129$ states. For the decision variable we need to encode the possible movements of the pieces. Thus, the state space of the decision variable is defined by the number of ways we can move from one board position (9+1; the tenth "state" represents the pieces not yet in play) to another (9). This gives a total of 90 configurations. Illegal moves are represented by large negative values in the utility function, ensuring that they will not be part of the optimal strategy. Finally, we also need to represent the policy of the opponent B. This is encoded in the conditional probability distributions associated with the S_B nodes (see Figure 5).

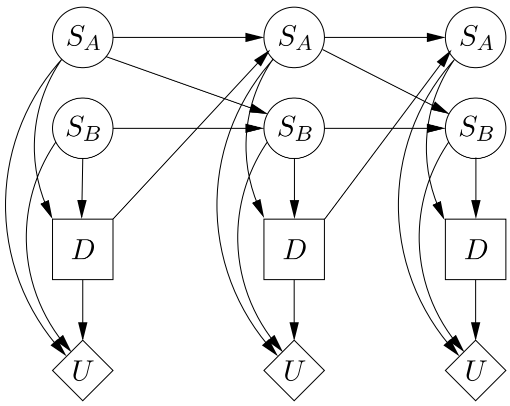

Figure 5: MDP representation of the tic-tac-toe problem.

Solution for exercise 9.24