Chapter 10 (part 3): *Using Uncertain Knowledge*

- *Making Decisions Under Uncertainty*

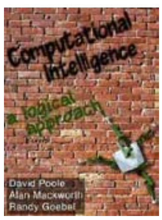

D. Poole, A. Mackworth, and R. Goebel, *Computational Intelligence: A Logical Approach*, Oxford University Press, January 1998

Making Decisions Under Uncertainty

- *What an agent should do depends on:*
  - *What the agent believes. You may be tempted to say "what is true in the world", but when agent does not know what is true in the world, it acts only on its beliefs. Sensing the world updates the agent's beliefs by conditioning on what is sensed.*
  - *The agent's goals. When an agent has to reason under uncertainty, it has to consider not only what will "most likely" happen but everything that "may" happen (role of expectations).*

Decision theory specifies how to trade off the desirability (reducing good outcomes but reducing disastrous outcomes) and probabilities of the possible outcomes for competing actions.

Decision Variables

- *Decision variables are like random variables (without pdf) that an agent gets to choose the value of.*
- *A possible world (outcome) specifies the value for each decision variable and each random variable.*
- *The probability of a proposition (or outcome) is undefined unless you condition on the values of all decision variables.*

Decision Tree for Delivery Robot

- *The robot can choose to wear pads to protect itself or not.*
- *The robot can choose go the short way past the stairs or a long way that reduces the chance of an accident.*
- *There is one random variable of whether there is an accident.*

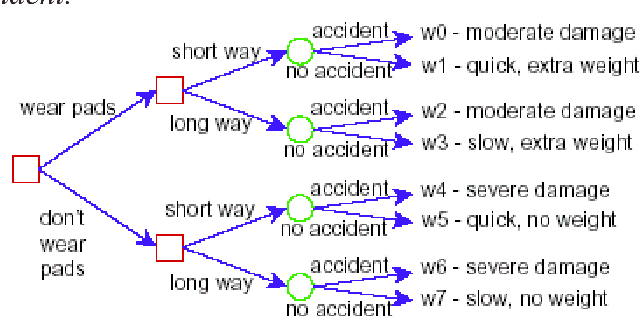

Expected Values

- *The expected value of a numerical random variable is its average value, weighting possible worlds by their probabilities.*
- *Suppose V is a numerical random variable and w is a possible world. Let $\rho(V, w)$ be the value x such that $w \models V = x$.*
- *The expected value of V is*

Probability of the outcome w

$$E(V) = \sum_{v \in \Omega} \rho(V, w) * \mu(w)$$

- *The conditional expected value of V given e is*

$$E(V \mid e) = \sum_{w \models e} \rho(V, w) * \mu_e(w)$$

Probability of outcome w given the evidence e

Utility

- *Utility (real random variable) is a measure of desirability of worlds to an agent (the world that results from the decision!). Utility = 1/cost*
- *Let U be a real-valued random variable such that $\rho(U, w)$ represents how good the world (or outcome) is to an agent.*
- *Simple goals can be specified by: worlds that satisfy the goal have utility 1: other worlds have utility 0.*
- *Often utilities are more complicated: for example, made up from the amount of damage to a robot, how much energy it has used up, what goals are achieved, and how much time it has taken.*

Single decisions

- *In a single decision, the agent chooses a value for each decision variable. Let the compound decision variable d be the tuple of all original decision variables. The agent can choose $d = d_i$ for any $d_i \in dom(d)$.*
- *The expected utility of decision $d = d_i$ is $E(U|d = d_i)$.*
- *An optimal single decision is the decision $d = d_{max}$ whose expected utility is maximal:*
- *Stock-Market Prediction: is it time to sell, to buy, or to hold shares in order to maximize the profit (utility) ?*

$$\mathbf{E(U \mid d = d_{max}) = \max_{d_i \in dom(d)} E(U \mid d = d_i)}$$

Example

- ***Delivery Robot:***
  - ***The single decision is the complex decision variable:<pads, which_way>***
    a possible world = <wear_pads,short,accident>

E(U | pads=wear_pads ∧ which_way = short) =

P(accident | wear_pads ∧ which_way=short)

. Utility(w_0) + (1-P(accident | wear_pads ∧ which_way = short)) . Utility(w_1).

- ***Utility(w_0), Utility(w_1) are known by the decision maker***
- ***Conditional probabilities are provided as well.***
- ***We can then compute the optimal decision.***

Sequential Decisions

- *An intelligent agent doesn't make a multi-step decision and carry it out without considering revising it based on future information.*
- *A more typical scenario is where the agent: observes, acts, observes, acts, …*
- *Subsequent actions can depend on what is observed.*
  What is observed depends on previous actions.
- *Often the sole reason for carrying out an action is to provide information for future actions.*
  For example: diagnostic tests, spying.

Decision Networks

- *A decision network is a graphical representation of a finite sequential decision problem. Decision networks extend belief networks to include decision variables and utility.*
- *A decision network is a DAG with three types of nodes:*
  - *Chance nodes (ellipses) as in belief networks.*
  - *Decision nodes (rectangles) labeled with decision variables. Arcs into decision nodes show the information that will be available when the decision is made.*
  - *A value node (diamond) represents the utility. Arcs into the value node show the variables the utility depends on.*

Decision Network for the Alarm Problem

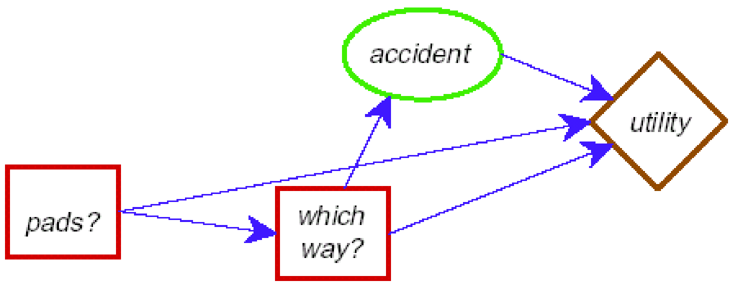

- *This shows explicitly which nodes affect whether there is an accident.*

Decision Network for the Alarm Problem