VU Einführung in die Künstliche Intelligenz

SS 2023

Hans Tompits

Institut für Logic and Computation
Forschungsbereich Wissensbasierte Systeme

www.kr.tuwien.ac.at

Human Judgment and Irrationality

- Decision theory is a *normative theory*, i.e., it describes how a rational agent *should* act.
- A *descriptive theory*, on the other hand, describes how agents (e.g., humans) *really do* act.
- Evidence suggests that these two kinds of theories do not coincide $\Longrightarrow$ humans appear to be "predictably irrational".

Allais Paradox

- Assume that there is a choice between lotteries A and B and then between C and D, which have the following prizes:
  - A: 80% chance of winning \$4000
  - B: 100% chance of winning \$3000
  - C: 20% chance of winning \$4000
  - D: 25% chance of winning \$3000
- Most people prefer B over A (taking the sure thing), and C over D (taking the higher EMV).
- However, the normative analysis yields a different result:
  - Assume, without loss of generality, a utility function with $U(\$0) = 0$.
  - Then, $B \succ A$ implies $U(\$3000) > 0.8 \cdot U(\$4000)$, and $C \succ D$ implies $0.2 \cdot U(\$4000) > 0.25 \cdot U(\$3000)$.
  - From the latter we obtain
$$U(\$3000) < \frac{0.2}{0.25} U(\$4000) = 0.8 \cdot U(\$4000).$$
  - ➥ There is no utility function consistent with theses choices!

Allais Paradox (ctd.)

- One possible explanation for the apparent irrational preferences is the *certainty effect*, i.e., people are strongly attracted to gains that are certain.
- Why is that?

Allais Paradox (ctd.)

- Possible answers:
  1. People may choose to reduce their computational burden: by choosing the certain outcomes, there is no need to estimate probabilities.
  2. People may mistrust the legitimacy of the stated probabilities (in particular, if stated by people with a vested interest in the outcomes).
  3. People may account their emotional state as well as their financial state.
     - People know they would experience *regret* if they gave up a certain reward (B) for an 80% chance of a higher reward and then lost.
     - I.e., in choosing A, there is a 20% chance of getting no money and *feeling like a complete idiot*, which is worse than just getting no money.

➥ Choosing B over A and C over D may not be irrational: just willing to give up $200 EMV to avoid a 20% chance of feeling like an idiot.

Ellsberg Paradox

- Prizes have an equal value, but probabilities are underconstrained.
- Payoff depends on the color of a ball chosen from an urn.
- You are told that the urn contains 1/3 red balls, and 2/3 either black or yellow balls, but you do not know how many black and how many yellow.
- Then, you are asked to choose between *A* and *B*, and then between *C* and *D*:
  - A: $100 for a red ball
  - B: $100 for a black ball

  - C: $100 for a red or a yellow ball
  - D: $100 for a black or yellow ball
- If you think there are more red than black balls, you should prefer *A* over *B* and *C* over *D*, and the opposite otherwise.
- *But* most people prefer *A* over *B* and *D* over *C*!
- ➥ People have *ambiguity aversion*.

Ellsberg Paradox (ctd.)

Ambiguity aversion (ctd.):

- A: \$100 for a red ball
- B: \$100 for a black ball

- C: \$100 for a red or a yellow ball
- D: \$100 for a black or yellow ball

➤ A gives you a $1/3$ chance of winning, while B could be anywhere between 0 and $2/3$.

➤ Likewise, D gives you a $2/3$ chance, while C could be anywhere between $1/3$ and $3/3$.

➥ Most people *elect the known probability* rather than the unknown one.

Decision Networks

- *Decision networks* (or *influence diagrams*) are a general framework for supporting rational decisions.
- They contain information about an agent's current state, its possible actions, the state that will result from the agent's action, and the utility of that state.
- Example of a decision network for the *airport siting problem:*

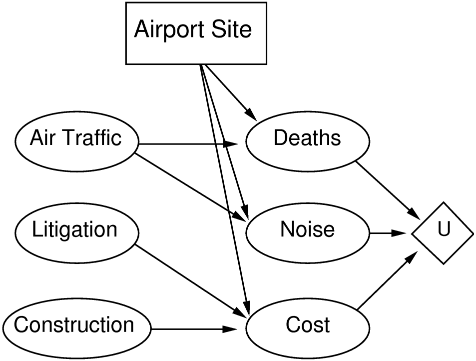

Decision Networks (ctd.)

Decision network uses three types of nodes:

- *Chance nodes* (ovals): represent random variables.
  - E.g., the agent is uncertain about construction costs, the level of air traffic, the potential for litigation.
  - There are also the *Deaths*, *Noise*, and *Cost* variables, depending on the site chosen.
  - Chance nodes are associated with a conditional probability distribution that is indexed by the state of the parent nodes.
- *Decision nodes* (rectangles): represent points where a decision maker has a choice of actions; e.g., the choice of an airport site influences the cost, noise, etc.
- *Utility nodes* (diamonds): represent the agent's utility function.
  - It has as parents all variables describing the outcome that directly affect utility.

Evaluating Decision Networks

- Algorithm for evaluating decision networks:
  1. Set the evidence variables for the current state.
  2. For each possible value of the decision node:
     a) Set the decision node to that value.
     b) Calculate the posterior probabilities for the parent nodes of the utility node, using a standard probabilistic inference algorithm.
     c) Calculate the resulting utility for the action.
  3. Return the action with the highest utility.

☞ Decision networks are an extension of *Bayesian networks*, in which only chance nodes occur.

The Value of Information

- In the decision network analysis it is assumed that *all relevant information* is available before making a decision.
- In practice this is hardly ever the case:
  - ☞ *One of the most important parts of decision making is knowing what questions to ask.*
- *Information value theory* enables an agent to choose what information to acquire.
- Basic assumption:
  - the agent can acquire the value of any observable chance variables.
- These observation actions affect only the *belief state*, not the external physical state.
- The value of an observation derives from the *potential* to affect the agent's eventual physical action $\Longrightarrow$ this potential can be estimated directly from the decision model itself.

The Value of Information: Example

A simple example:

- An oil company plans to buy one of n indistinguishable blocks of ocean-drilling rights.
- One of the blocks contains oil worth C dollars, while all other are worthless.
- The price for each block is C/n Dollars.
- If the company is *risk neutral*, then it is indifferent between buying a block and not buying one.
- Now assume that the company can buy information (results of a survey) that says definitively whether block 3 contains oil or not.
- How much should the company be willing to pay for this information?

Example (ctd.)

To answer this question, we examine what the company would do if it had the information:

- With probability $1/n$, the survey will indicate oil in block 3.
  - In this case, the company will buy block 3 for C/n dollars and make a profit of $C - C/n = (n-1)C/n$ dollars.
- With probability $(n-1)/n$, the survey will show that block 3 contains no oil, hence the company will buy a different one.
  - Now, the probability of finding oil in one of the other blocks changes from $1/n$ to $1/(n-1)$, so the expected profit is $\frac{C}{(n-1)} - \frac{C}{n} = \frac{C}{n(n-1)}$ Dollars.
- Then, the resulting expected profit, given the survey information is

$$\frac{1}{n} \cdot \frac{(n-1)C}{n} + \frac{n-1}{n} \cdot \frac{C}{n(n-1)} = \frac{C}{n}.$$

➥ The company should be willing to pay up to C/n Dollars
$\Longrightarrow$ the information is worth as much as the block itself!

Remarks

- The value of information derives from the fact that *with* the information, one's course of action can be changed to suit the *actual* situation.
- One can discriminate according to the situation:
  - without the information, one has to do what is *best on average* over the possible situations.
- In general, the value of a given piece of information is defined to be the difference in expected value between the best actions before and after an information is obtained.

The Value of Perfect Information

- Assumption:
  - *Exact evidence* about the value of a random variable E_j can be obtained (i.e., we learn $E_j = e_j$).
  - ➥ We use the phrase *value of perfect information* (*VPI*).
- Given initial evidence $\mathbf{e}$, the value of the current best action α is defined by

$$EU(\alpha|\mathbf{e}) = \max_a EU(a|\mathbf{e}) = \max_a \sum_{s'} P(\text{Result}(a) = s'|a, \mathbf{e})U(s').$$

- The value of the new best action α_{e_j} after evidence $E_j = e_j$ is obtained is

$$EU(\alpha_{e_j}|\mathbf{e}, e_j) = \max_a \sum_{s'} P(\text{Result}(a) = s'|a, \mathbf{e}, e_j)U(s').$$

- But the value of E_j is currently unknown, so to determine the value of discovering E_j, given current information $\mathbf{e}$, we average over all possible values e_{j_k} that might be discovered for E_j:

$$VPI_{\mathbf{e}}(E_j) = \left(\sum_k P(E_j = e_{j_k}|\mathbf{e})EU(\alpha_{e_{j_k}}|\mathbf{e}, E_j = e_{j_k})\right) - EU(\alpha|\mathbf{e}).$$

Some Properties of the VPI

- The expected value of information is *nonnegative*:

$$VPI_{\mathbf{e}}(E_j) \geq 0, \text{ for all } \mathbf{e} \text{ and all } E_j.$$

- VPI is *nonadditive*:

$$\text{in general, } VPI_{\mathbf{e}}(E_j, E_k) \neq VPI_{\mathbf{e}}(E_j) + VPI_{\mathbf{e}}(E_k).$$

- VPI is *order independent*:

$$VPI_{\mathbf{e}}(E_j, E_k) = VPI_{\mathbf{e}}(E_k, E_j).$$

Decision-theoretic Expert Systems

- *Decision analysis* (evolved in the 1950s and 1960s) studies the application of decision theory to actual decision problems.
- Traditionally, there are two roles in decision analysis:
  - the *decision maker*, stating preferences between outcomes; and
  - the *decision analyst*, who enumerates possible actions and outcomes, and elicits preferences to determine the best course of action.
- Early expert system research concentrated on answering questions rather than on making decisions.
- The addition of *decision networks* allows expert systems to recommend optimal decisions, reflecting preferences as well as available evidence.

Decision-theoretic Expert Systems (ctd.)

The process of creating a decision-theoretic expert system, e.g., for selecting a medical treatment for congenital heart disease (aortic coarctation) in children:

1. create a causal model (e.g., determine symptoms, treatments, disorders, outcomes, etc.);
2. simplify to a qualitative decision model;
3. assign probabilities (e.g., from patient databases, literature studies, experts subjective assessments, etc.);
4. assign utilities (e.g., create a scale from best to worst outcome and give each a numeric value);
5. verify and refine the model, evaluate the system against correct input-output-pairs, a so called *gold standard*;
6. perform sensitivity analysis, i.e., check whether the best decision is sensitive to small changes in the assigned probabilities and utilities.

Influence Diagram Example

Influence diagram for aortic coarctation:

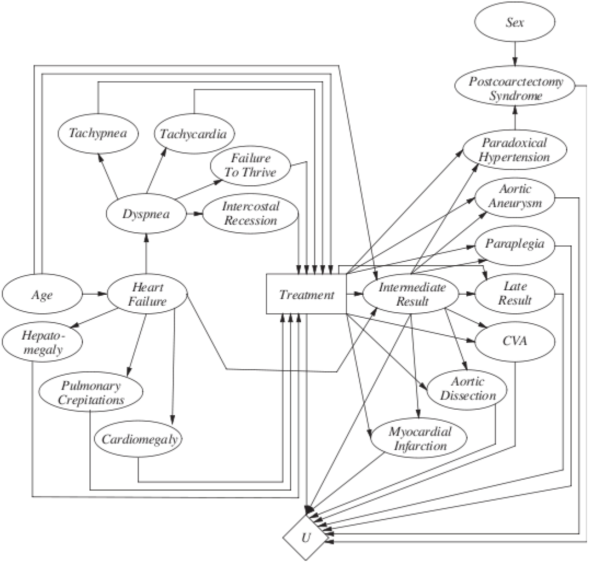

Summary

- *Decision theory* puts probability theory and utility theory together to describe what an agent *should do*.
- A *rational agent* makes decisions by considering all possible actions and choosing the one that leads to the best expected outcome.
- An agent whose preferences are consistent with a set of simple axioms possesses a *utility function*; furthermore, it selects actions as if maximising expected utility.
- The *value of information* is defined as expected improvement in utility compared with making a decision without the information.
- *Expert systems* that incorporate utility information are more powerful than pure inference systems:
  - they are able to make decisions and use the value of information to decide whether to acquire it, and
  - they can calculate their sensitivity to small changes in probability and utility assessments.