INF5390 – Kunstig intelligens

Making Simple Decisions

Roar Fjellheim

Outline

- Uncertainty and utility
- Maximum expected utility
- Preference structures
- Decision networks
- Value of information
- Decision-theoretic expert systems
- Summary

AIMA Chapter 16: Making Simple Decisions

Agents and decision theory

- Agents need to make decisions in situations of *uncertainty* and *conflicting goals*
- Basic principle of decision theory: *Maximization* of *expected utility*
- *Decision-theoretic agents* are based decision theory, and need knowledge of *probability* and *utility*
- Here, we are concerned with “simple” (*one-shot*) decisions, can be extended to *sequential* decisions

Principle of Maximum Expected Utility (MEU)

- Let
  - $U(s)$ - Utility of state s
  - $RESULT(a)$ – Random variable whose values are possible outcome states of action a in current state
  - $P(RESULT(a) = s' \mid a, e)$ - Probability of outcome s', as a result of doing action a in current state, and given agent's available evidence e of the world
- Then the *expected utility EU* of *a*, given *e* is

$$EU(a \mid e) = \sum_{s'} P(RESULT(a) = s' \mid a, e) U(s')$$

- MEU: Agent should select *a* that maximizes *EU*

Problems with applying MEU

- Often difficult to formulate problem completely, and required computation can be prohibitive
- Knowing state of the world requires perception, learning, representation and inference
- Computing $P(RESULT(a) \mid a, e)$ requires complete causal model and NP-complete belief net updating
- Computing utility $U(s')$ may require search or planning since agent needs to know how to get to a state before its utility can be assessed

Preference and utility

- MEU appears to be a rational basis for decision making, but is not the only possible
  - Why *maximize average* utility, instead of e.g. *minimize losses*?
  - Can *preferences* between states really be compared by comparing two numbers?
  - Etc.
- We can state *constraints on preference structures* for a rational agent, and show that MEU is compatible with the constraints

Lotteries and preferences

- Lottery
  - Scenario with different outcomes with different probabilities
  - The agent have preferences regarding the outcomes
- Example *L* = [*p*, *A*; 1-*p*, *B*]
  - Lottery *L* with two outcomes, *A* with probability *p*, *B* with probability 1-*p*
- Preferences
  - $A > B$ *A* is preferred over *B*
  - $A \approx B$ Agent is indifferent between *A* and *B*
  - $A \geq B$ Prefers *A* over *B* or is indifferent
- Constraints on preferences include orderability, transitivity, etc.

Utility follows from preferences

- The constraints on preferences are the *axioms of utility*, from which utility principles follow
- *Utility principle*
  - If the agent's preferences obey axioms of utility, there exists a real-valued *utility function U* such that

$$U(A) > U(B) \Leftrightarrow A > B$$

$$U(A) = U(B) \Leftrightarrow A \approx B$$

- *MEU principle*
  - Utility of a lottery can be derived from outcome utilities

$$U([p_1, S_1; \ldots; p_n, S_n] = \sum_i p_i U(S_i)$$

Utility of money

- Utility theory comes from economics, and money is a common basis for utility functions

Typical curve of utility vs. money

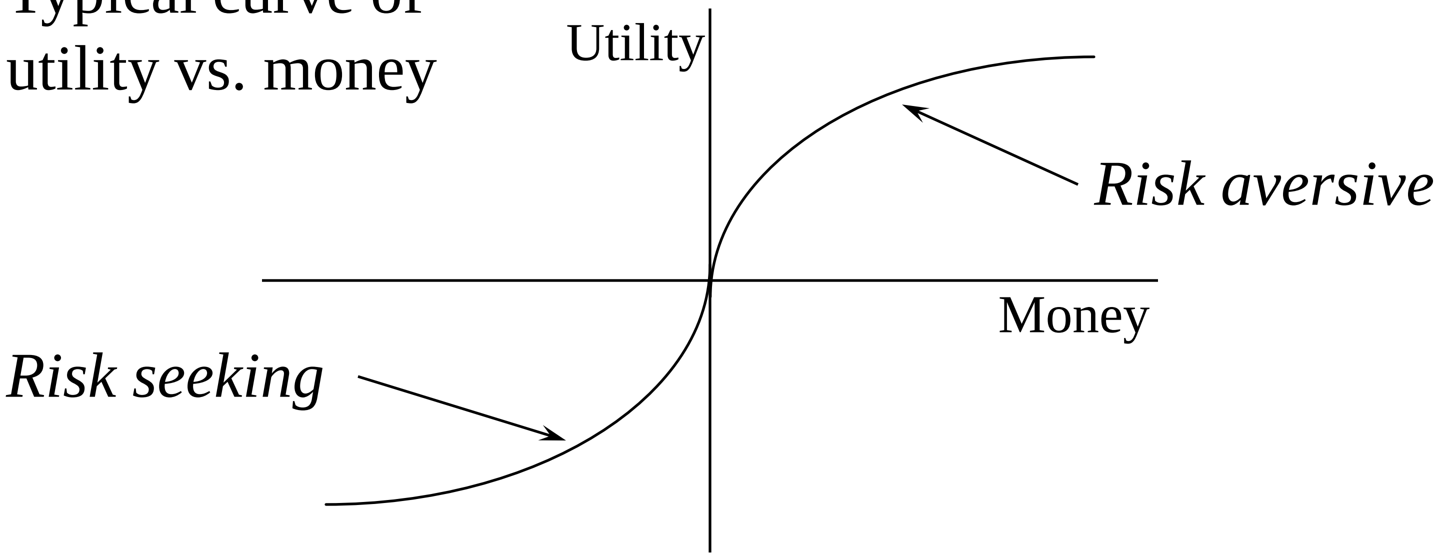

Human decision making

- Decision theory is *normative*, but not *descriptive*: People violate axioms of utility in practice
- Example
  - A: 80% chance of \$4000 C: 20% chance of \$4000
B: 100% chance of \$3000 D: 25% chance of \$3000
  - Most people choose B over A, and C over D. Since only the scale is different, there does not seem to be a utility function that is consistent with the choices
- Possible descriptive theory
  - People are risk-aversive with high-probability events (A-B)
  - People take more risks with unlikely payoffs (C-D)

Decision networks

- *Decision networks* (also called *influence diagrams*) are a general mechanism for making rational decisions
- Decision networks combine belief networks with nodes for actions and utilities, and can represent
  - √ Information about agent's *current state*
  - √ Agent's *possible actions*
  - √ States that will *follow* from actions
  - √ *Utilities* of these states
- Therefore, decision networks provide a substrate for implementing rational, utility-based agents

Decision network for airport location

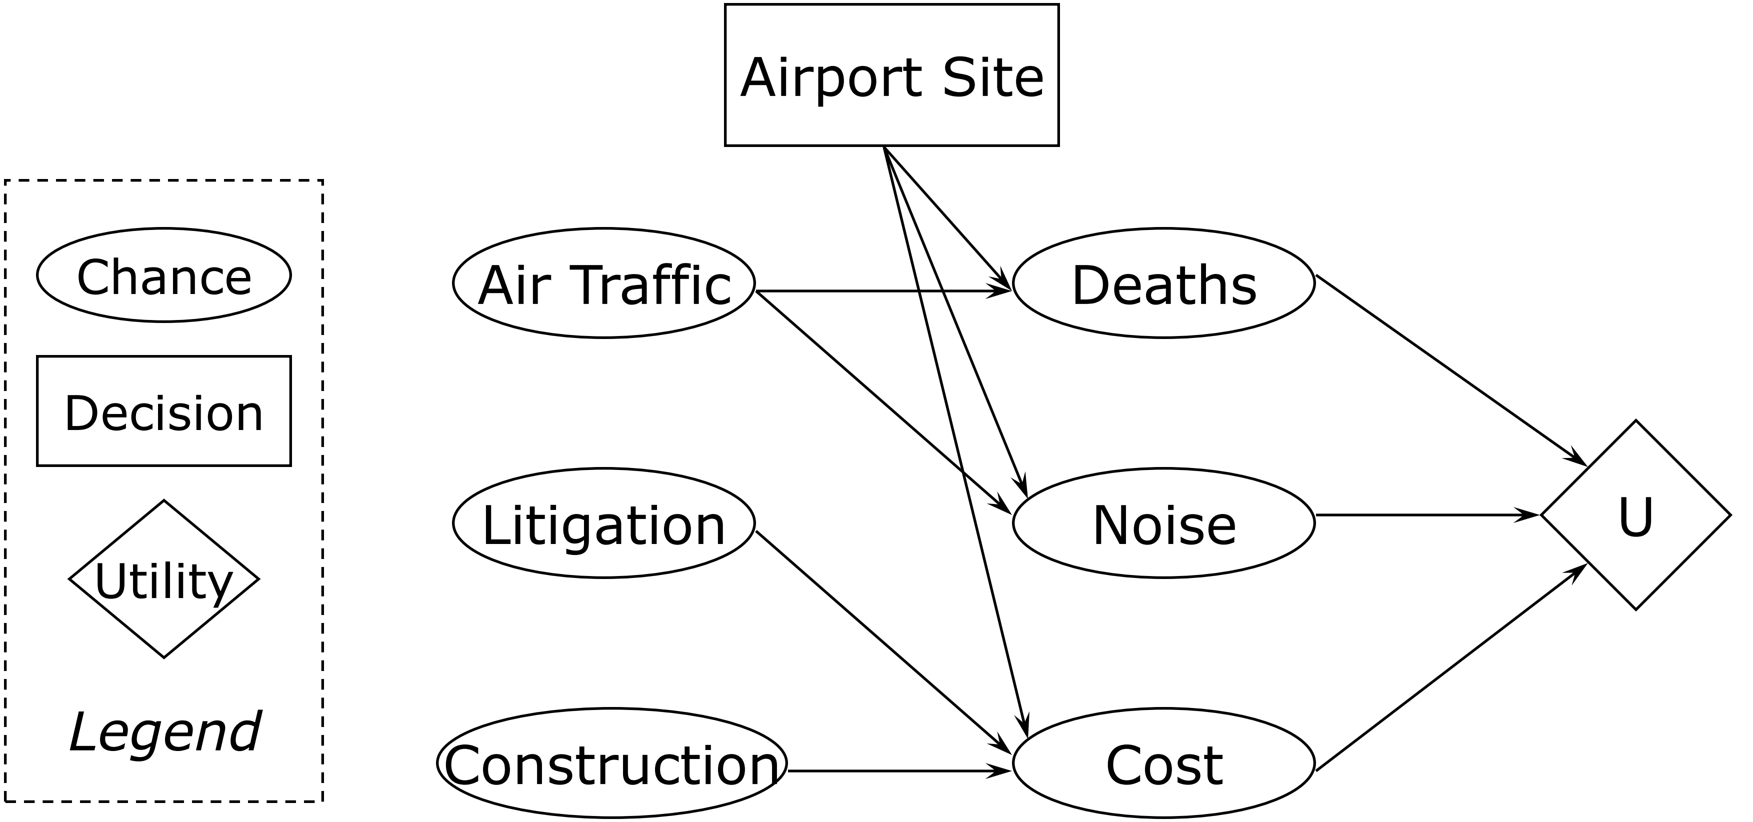

Node types in decision networks

- *Chance* nodes (ovals)
  - Represent random variables (as in belief networks), with associated conditional probability table (CPT) indexed by states of parent nodes (decisions or other chance nodes)
- *Decision* nodes (rectangles)
  - Represent points where the decision maker has choice of actions to make
- *Utility* nodes (diamonds)
  - Represent the agent's utility function, with parents all nodes that directly influence utility

Evaluating decision networks

- Set the evidence variables (chance nodes with known values) for the current state
- For each possible value of the decision node
  - √ Set decision node to that value (from now on, it behaves like a chance node that has been set as an evidence variable)
  - √ Calculate posterior probabilities for parent nodes of the utility node, using standard probabilistic inference methods
  - √ Calculate resulting utility for the action
- Return the action with the highest utility

Value of information

- The agent will normally not have all required information available before making a decision
- Important to know which information to seek, by performing tests that may be expensive and/or hazardous
- The importance of tests depend on
  - √ Will different outcomes make significant difference to the optimal action
  - √ What is the probability of different outcomes
- *Information value theory* helps agents decide which information to seek, by using *sensing actions*

Motivating example

- Oil company to buy one of n indistinguishable blocks, exactly one block contains oil worth C, price for each block is C/n
- A seismologist offers to investigate block 3, determining if it has oil or not. How much is this information worth?
  - With probability $1/n$, block 3 has oil. Then the company will buy block 3 for C/n, and make profit $C-C/n = (n-1)C/n$
  - With probability $(n-1)/n$, block 3 is empty. The company will buy another block. Probability of oil there is $1/(n-1)$, with profit $C/(n-1)-C/n = C/n(n-1)$
- Expected profit given the survey information

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

- The information is as much worth as the block itself!

Considerations for information gathering

- Information has value if it is likely to cause a change of plan, and if the new plan will be significantly better than the old
- An information-gathering agent should
  - Ask questions in a reasonable sequence
  - Avoid asking irrelevant questions
  - Take into account importance of information vs. cost
  - Stop asking questions when appropriate
- Requirements met by using *VPI(E) - Value of Perfect Information* of evidence *E*. Properties:
  - Always *non-negative*
  - Depends on current state and is *non-additive*
  - *Order-independent* (simplifies sensing actions)

An information gathering agent

function INFORMATION-GATHERING-AGENT(*percept*) **returns** an *action*

persistent: *D*, a decision network

integrate *percept* into *D*

j <= the value that maximizes $VPI(E_j) / Cost(E_j)$

if $VPI(E_j) > Cost(E_j)$

  then return $REQUEST(E_j)$

  else return the best action from *D**

*non-information seeking action

Comments on information-gathering agent

- Information-gathering agent is *myopic*, i.e. it just considers one evidence variable at a time
- It may hastily select an action where a better decision would be based on two or more information gathering actions
  - √ "Greedy" search heuristic - often works well in practice
- A perfectly rational agent would consider all possible sequences of sensing action that terminate in an external action
  - √ May disregard permutations due to order-independence

Decision analysis vs. expert systems

- Decision analysis (application of decision theory)
  - √ Focus on *making decisions*
  - √ Defines possible actions and outcomes with preferences
  - √ Roles
    - *Decision maker* states preferences
    - *Decision analyst* specifies problem
- Expert systems ("classical" rule-based systems)
  - √ Focus on *answering questions*
  - √ Defines heuristic associations between evidence & answers
  - √ Roles
    - *Domain expert* provides heuristic knowledge
    - *Knowledge engineer* elicits & encodes knowledge in rules

Decision-theoretic expert systems

- *Decision-theoretic expert systems*
  - √ Inclusion of decision networks in expert system frameworks
- Advantages
  - √ Make expert preferences explicit
  - √ Automate action selection in addition to inference
  - √ Avoid confusing likelihood with importance
    - Common pitfall in expert systems: Conclusions are ranked in terms of likelihood, disregarding rare, but dangerous conclusion
  - √ Availability of utility information helps in knowledge engineering process

Knowledge engineering for decision-theoretic expert systems

- Create causal model
- Simplify to qualitative decision model
- Assign probabilities
- Assign utilities
- Verify and refine model
- Perform sensitivity analysis

Summary

- *Probability theory* describes what an agent should believe based on evidence, and *utility theory* describes what an agent wants
- *Decision theory* combines the two to describe what an agent should do
- Decision theory can be used to build *a rational agent*, that considers all possible actions and chooses the one with the best expected outcome
- Under certain reasonable assumptions, outcomes can be scored by a real-valued *utility function*
- Rational agent acts to *maximize expected utility*

Summary (cont.)

- *Decision networks* can be used to express and solve decision problems,
- They extend belief networks with *decision* and *utility* nodes in addition to *chance* nodes
- *Value of information* is expected improvement in utility compared to deciding without information
- *Decision-theoretic expert systems* combine decision networks and inference
- They can make decisions, choose to get more information, and perform sensitivity analysis