Artificial Intelligence
Joint class 5314 / 4320

Expert Systems and Reasoning under Uncertainty

Main resources for this lecture:

- the Textbook of this class.
- *The essence of Artificial Intelligence*, by Alison Casey.

Websites of interest:

- Website of journal a journal on Decision and Reasoning under Uncertainty: http://www.ida.liu.se/ext/etai/dru/
- Research group: http://www.csse.monash.edu.au/research/ruug/index.shtml
- Article of interest: http://sern.ucalgary.ca/KSI/KAW/KAW99/papers/Stuckenschmidt1/PSMs_for_uncertainty.pdf
- etc. Shop around and help yourself on the web :-)

---

1 What is an expert system?

- a system which provides expert quality advice, diagnoses, and recommendations given real world problems.
- ES solve problems that would normally require a human expert (e.g., specialist doctor, minerologist).

1.1 Why use ES?

- because human expertise may be in short supply expensive, and experts hard to get hold of in a hurry.
- ES can be made "easily" available on demand.

1.2 General remarks

- building an ES may result unsuccessful at first attempt:
  - difficult to extract human expert's knowledge
  - most of expert knowledge is difficult to express exactly
  - obvious rules are often not mentioned by experts
- therefore, iterations are necessary between expert, tool, customer, before the ES is ready to be used.
- **consequence on the way ES are implemented:**
  - written in a way making it easy to modify and inspect
  - the system should be able to explain its reasoning (to expert, user, knowledge engineer) and answer questions about the solution process
  - updating the code should not involve rewritting a whole lot of code (just adding or deleting localized chunks of knowledge)

1.3 Approaches

Widely used: rule-based systems
But not only, as we'll see later.

1.4 Applications of ES

- medicine
- mathematics
- engineering
- geology: analysis of samples collected in oil exploration
- computer science: configuration of systems
- business
- law
- defence
- education
- etc.

In each domain: used to solve many different problems.
e.g., in electronics: designing circuits, diagnosing faults in circuits.

- designing circuits: very much related to problem-solving $\leftrightarrow$ designing a kitchen
- diagnosing faults in circuits: very much like diagnosing diseases

2 Designing an ES

2.1 Appropriateness of ES

Designing an ES may be time and money consuming... therefore it may be wise to first determine whether an ES is the answer to your problem.

- **is it worth to spend a lot of money and time on an ES to solve your problem?**
  - do not underestimate the time it needs: knowledge is difficult to acquire + many iterations necessary
  - are there other ways to solve your problem?
  - for highly specialized fields, where specialists are pricey ressource, it is a good idea
- **are ES techniques appropriate?**
  - physical tasks vs. advice
  - common-sense knowledge vs. highly technical and specialized knowledge
- **do you have experts at hand?**
  - you need to acquire knowledge from them, make sure that you have some available
- **do you have customers at hand?**
  - you need to get feed-back from them, make sure that you have some available

2.2 Acquiring/extracting knowledge

This is a central stage of ES design: extracting and representing knowledge. Knowledge engineers are in charge of this critical task, that involves interaction with both experts and users. Knowledge engineers need to get at least a little acquainted with the field of their ES.

Usual procedure: experts are given a series of example problems, and will explain aloud their reasoning. The knowledge engineer is supposed to abstract general rules from these explanations, and check them with the expert.

First prototype: the knowledge engineer needs to make (temporary) decisions about appropriate knowledge representation and inference methods (rules, logic, backward/forward chaining, etc.). The first prototype should only solve a small part of the overall problem. If the methods used seem to work well for the first prototype, it may be worth investing effort in completing it with rest of the knowledge. Otherwise, it will not have been to much time wasted.

2.3 General architecture

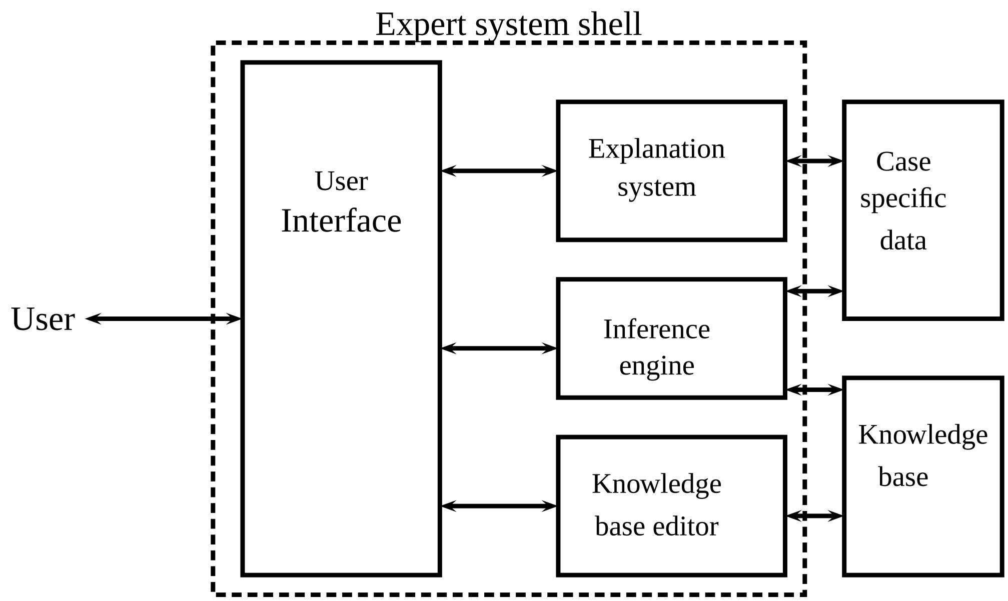

Rule-based systems and backward/forward chaining methods have been presented in class, so we jump now to Reasoning under uncertainty.

3 Reasoning under uncertainty

All knowledge representation and inference methods presented so far assumed that knowledge was certain. And regarding what we have done so far, if a logical agent cannot check whether a condition holds, it will never be able to make any conclusion. This also holds for checking if a conclusion holds or not (cf. $p \rightarrow q \equiv \neg p \vee q$).

When we have no certain information, we may have ideas of what is better or worse. And given this, the objective is to maximize the performance: act rationally (cf. lecture notes on intro to AI, and agents).

3.1 Handling uncertainty in knowledge

Purpose of this section: what is uncertainty?
In particular, diagnosis involves in general uncertain knowledge (e.g., medicine, auto repair).
Example:

$$\forall p, Symptom(p, Toothache) \rightarrow Disease(p, Cavity)$$

is not a true rule. We could instead consider:

$\forall p, Symptom(p, Toothache) \rightarrow (Disease(p, Cavity) \vee Disease(p, GumDisease) \vee Disease(p, Abscess) \vee \ldots$

But this rule is not true either because we cannot be exhaustif for:

- lack of will: we are lazy :-)
- lack of theoretical knowledge: we may not know all disease related to symptom p
- lack of practical knowledge: we may not know of all specific cases

What the agent can provide instead a **degree of belief** given relevant sentences → probability theory: assigns to each sentence a degree of belief (between 0 and 1).

Probability summarizes the uncertainty that comes from our laziness and ignorance :-)
0.8 means degree of belief of 80%, and NOT 80% true.

Degrees of truth are dealt with by **fuzzy logic**.

Remarks:

- "The patient has a cavity" is either true or false in logic: result of the observation of the world
- on the other hand, "the probability that the patient has a cavity is 0.8" is the agent's belief, not an observation of the world: such belief has been formed based on percepts or database of cases (which constitute percepts too).
- probability are constantly updated as new evidence comes: e.g., the example with the shuffled stack of cards (cf. textbook)

3.2 Making rational decisions under uncertainty

Involves:

- **preferences:** the agent must have preferences among the possible outcomes of the possible actions to take
- **utility[1] functions:** to help determine which action should be taken: represent and reason with preferences
  - represent preferences: give outcomes a degree of preference for instance
  - reason with preferences: combines the degrees to end up with the most preferred alternative

[1] Utility: used in the sense "the quality of being useful".

in utility theory, every state has a degree of usefulness (or utility), to an agent, and the agent will prefer states with higher utility. The utility function is agent-dependent. Consider a multi- player game. The utility function for one player will not be the same as for another one because both players don't have the same goal (or in this case, they have the same one – winning – which is not compatible).

Preferences (represented by utility values) are combined with probability theory in decision theory (theory of rational decisions).

Def. *An agent is rational if and only if it chooses the action that yields the highest expected utility, averaged over all the possible outcomes of the action.*
= **Maximum Expected Utility (MEU)** principle.

3.2.1 Design of a decision-maker agent

```
function DM-agent (percepts) return action
    static: belief-state= probabilistic beliefs about the current state of the world
            action= the agent's action

    update belief-state based on action and percepts
    calculate outcomes probabilities for actions
        given action descriptions and current belief-state
    select action with highest expected utility
        given probabilities of outcomes and utility information
    return action
```

3.3 Basic probability theory

Probability of x: degree of belief in x
1=can't believe it is not true; 0=necessarily false; 0.5=true and false are equally probable.