EXPERT SYSTEMS

Overview

- The last lecture looked at rules as a technique for knowledge representation.
- Now, we will look at one of the main applications of rules as a knowledge representation system.
- This is as the basis of rule-based expert systems.
- Rule-based expert systems were the big AI success story in the early 80s.
- At the time they were over-hyped, and since have fallen from favour.
- However, they are still used, and are worth learning about because of their historical importance.

What is an expert system?

- A *expert system* is a computing system that is capable of expressing and reasoning about some domain of knowledge.
- Typical domains are:
  - internal medicine (INTERNIST)
  - geology (PROSPECTOR)
  - chemical analysis (DENDRAL)
- The purpose of the expert system is to be able to solve problems or offer advice in that domain.

Expert systems can be distinguished from other kinds of AI program because:

- They deal with subjects of considerable complexity—subjects that normally require a good deal of human expertise.
  - provide expert-level solutions to complex problems.
- They must be fast and reliable;
- They must be capable of explaining and justifying solutions
  - be understandable.
- They must be sufficiently flexible that new information may be easily accomodated.

- Architecture of an expert system:

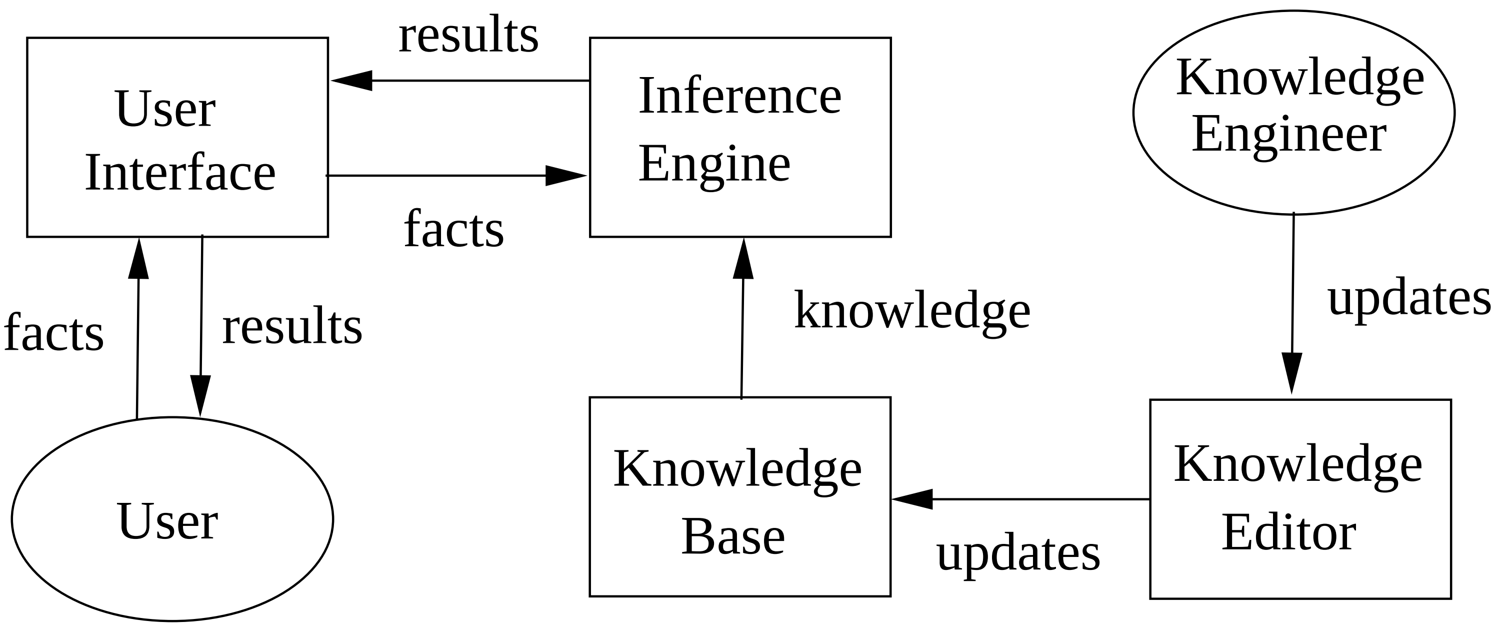

- The *knowledge base* holds the expertise that the system can deploy.
- The knowledge-base is constructed by the *knowledge engineer* in consultation with the domain expert.
- For many expert systems knowledge representation is through the use of rules (could also be frames, semantic nets etc.)
- In use, some facts are added to the working memory, which represent *observations* about the domain.
- The inference engine allows permits new inferences to be made from the knowledge in the knowledge base.
- These new facts represent *conclusions* about the state of the domain given the observations.
- For a rule-based expert system, the inference engine would be a mechanism for carrying out forward and/or backward chaining.

MYCIN

- One of the most important expert systems developed was MYCIN
- This is a system which diagnoses and treats bacterial infections of the blood.
- The name comes from the fact that most of the drugs used in the treatment of bacterial infections are called:
  - "Something"mycin
- MYCIN is intended to be used by a doctor, to provide advice when treating a patient.
- The idea is that MYCIN can extend the expertise of the doctor in some specific area.

- Rules in MYCIN are of the form:

```
IF
1. The gram stain of the organism is gramneg, and
2. The morphology of the organism is rod, and
3. The aerobicity of the organism is anaerobic
THEN
there is suggestive evidence that the identity
of the organism is bacteroides.
```

- Another example:

```
IF
1. The identity of the organism is
       not known with certainty, and
2. The gram stain of the organism is gramneg, and
3. The morphology of the organism is rod, and
4. The aerobicity of the organism is aerobic
THEN
there is strongly suggestive evidence that the
identity of the organism is enterobactericeae.
```

- The antecdent is allowed to be a mixture of AND and OR conditions.

- An example of a rule with OR conditions is:

```
IF
1. The therapy under consideration is:
cephalothin, or
clindamycin, or
erythromycin, or
lincomycin, or
vancomycin
   and
2. Meningitis is a diagnosis for the patient
THEN
It is definite that the therapy under consideration
is not a potential therapy.
```

- Note that we have rules about treatment as well as about diagnosis.

- We can also have OR in the consequent of the rule:

```
IF
The identity of the organism is bacteroides
THEN
I recommend therapy chosen from among the
following drugs:
1. clindamycin
2. chloramphenicol
3. erythromycin
4. tetracycline
5. carbenecillin
```

How MYCIN works

- MYCIN has a four stage task:
  - decide which organisms, if any, are causing significant disease.
  - determine the likely identity of the significant organisms.
  - decide which drugs are potentially useful.
  - select the best drug, or set of drugs.
- The control strategy for doing this is coded as meta-knowledge.

- The relevant rule is:

```
IF
1. There is an organism which requires therapy, and
2. Consideration has been given to possible
   other organisms which require therapy
THEN
1. Compile a list of possible therapies, and
2. Determine the best therapy.
ELSE
Indicate that the patient does not require therapy.
```

- This is applied in a mixture of forward and backward chaining:

  How do I decide if there is an organism requiring therapy? Well, Rule 90 tells me that organisms associated with significant disease require therapy. But I don't know about any organisms, so I'll ask the user ... now I can apply RULE 90 to each of these ... but I need to know if the organism is significant. Now I have a set of rules which tell me whether this is the case. Let's see, RULE 38 says that an organism from a sterile site is significant. I don't have any rules for saying if the site was sterile, so I'll ask the user ...

- So MYCIN starts by trying to apply the control rule, and this generates *sub-goals*.
- The first of these is to determine if there is an organism which needs to be treated.
- This generates another sub-goal; whether the organism is significant.
- This provokes a question to the user.
- The answer allows other rules to be fired, and these lead to further questions.
- Eventually the IF part of the control rule is satisfied, and the THEN part compiles a list of drugs, and chooses from it.

A consultation with MYCIN

- As we saw in the last lecture, MYCIN chains back from its overall goal of deciding what organisms need treatment until it finds it lacks information, and then asks the user for it.
- Using MYCIN is thus an interactive process:
  1. MYCIN starts running.
  2. MYCIN asks a question.
  3. The user answers it.
  4. MYCIN asks another question.
  5. …

```
1) Patient's name: (first-last)
** FRED BRAUN
2) Sex
** M
3) Age
** 55
4) Are there any illnesses for Fred Braun which may
be related to the present illness,
and from which organisms have been grown
in the microbiology laboratory?
** Y
```

```
-------- CULTURE-1 ----------
5) From what site was the specimen for CULTURE-1 taken?
** BLOD
= BLOOD
6) Please give the date and time
when CULTURE-1 was obtained.
** JUN 20 1977
The frst organism isolated from the blood
culture of 20-JUN-77 (CULTURE-1) will be
referred to as:

-------- ORGANISM-1 ---------
```

```
7) Enter the laboratory-
reported identity of
ORGANISM-1
** UNKNOWN
8) The stain (gram or Ziehl-
Neelsen acid-fast) of
ORGANISM-1
** NEG
9) Is ORGANISM-1 rod or coccus
** ROD
10) What is the form of the
individual organisms?
** FUSIFORM
```

...

```
It is important to cover for the
following probable infections and
associated organisms:

INFECTION-1 is BACTEREMIA

<ITEM-1> E.COLI
<ITEM-2> KLEBSIELLA
<ITEM-3> ENTEROBACTER
<ITEM-4> KLEBSIELLA-PNEUMONIA
```

Explanation in MYCIN

- An important feature of any expert system is that it can explain how it comes up with conclusions.
- This allows users to understand what the system is doing, and so see that it is doing sensible things.
- One of the advantages of rules is that they make it possible to generate good explanations.
- Essentially the explanation is just a presentation of the *proof tree*.
- We can illustrate this with an example explanation dialogue from MYCIN.

```
34) Was penicillinase
    added to CULTURE-1?
** Why?

[3.0] This will aid in
 determining whether
 ORGANISM-1 is a contaminant.
 It has already been
 established that:
 [3.1] the site of
 CULTURE-1 is blood, and
 [3.2] the gram stain of
 ORGANISM-1 is grampos
```

```
Therefore if
 [3.3] penicllinase
       was added to this
       blood culture
then
 there is weakly suggestive
 evidence that ORGANISM-1
 is a contaminant [RULE 39]

** Why?
```

```
[4.0] ... in order to determine
 whether there is significant
 disease associated with
 ORGANISM-1.
 It has already been
 established that:
  [4.1] this blood culture
     was taken from a sterile
     source
```

```
Therefore if
  [4.2] ORGANISM-1 is not a
    contaminant
 then
  there is strongly suggestive
  evidence that there is
  significant disease associated
  with ORGANISM-1
```

- When asked "WHY", MYCIN:
  - uses the current rule.
- If asked "WHY" again, MYCIN
  - uses the rule which caused the current rule to fire
- Remember that MYCIN is chaining backwards from the overall goal of showing significant disease.
- This explanation can be continued by asking more "WHY"s.

Other Kinds of Expert System

- Many expert systems are rule-based like MYCIN.
- However, there is no reason why expert systems cannot be based on other forms of knowledge representation:
  - Frames
  - Semantic networks
  - Bayesian networks
- What makes a program an “expert system” is not its use of rules, but its expert level performance.

How Expert Systems are Used

- As we saw, MYCIN was intended as an advisor.
- The idea behind many expert systems was similar:
  - Capture expertise known to a few
  - Make it available to many.
- During the "expert systems boom" many ESs were built which tried to *replace* experts.
- Such systems often ran into difficulties:
  - how to construct them?
  - how to field them?

Problems with Expert Systems

- Problems with construction:
  - Knowledge aquisition bottleneck.
  - Machine learning.
- Problems with representation:
  - What does "significant evidence" mean?
  - Handling uncertainty.
  - Bayesian networks.

- Problems with acceptance
  - Operational issues.
  - Legal issues.
- Problems with domain
  - Brittleness.
  - Common sense knowledge and CYC.

Why Expert Systems *Aren't* Agents

- Remember that agents have a direct connection with their environment:

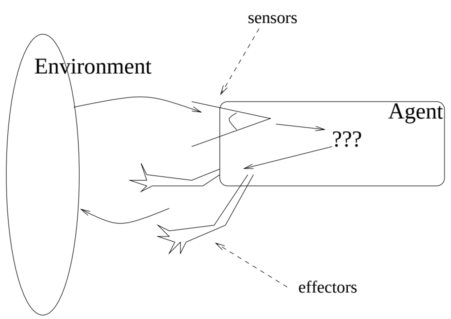

- Expert systems typically have a human mediator:

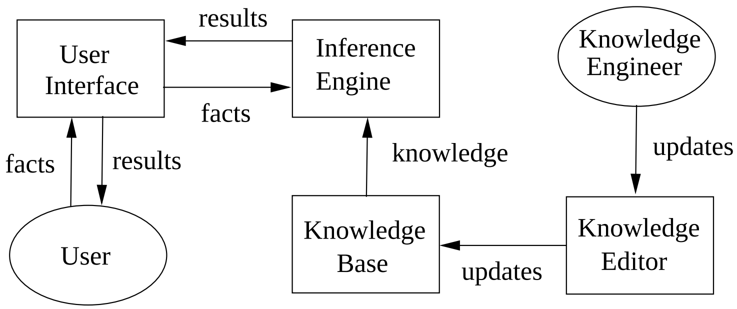

Summary

- This lecture introduced the idea of expert systems through the example of the MYCIN system.
- The lecture described how rules and meta-rules can be used to represent knowledge and control inference.
- The lecture also described how explanation may be performed, and discussed some of the problems that building and deploying expert systems may run into.