What is an expert system?

Peter Jackson, Introduction to Expert Systems:

*An **expert system** is a computer program that represents and reasons with knowledge of some specialist subject with a view to solving problems or giving advice.*

How does an expert system differ from more conventional application programs and other kinds of AI programs?

- It *simulates human reasoning* about a problem domain
- It performs reasoning over *representations of human knowledge*
- It solves problems using *heuristic or approximate methods*
- It addresses *problems of realistic complexity* that have genuine scientific or commercial interest
- It must *exhibit high performance* in terms of speed and reliability
- It must be able to *explain and justify solutions or recommendations*

Examples of early expert systems

1965	DENDRAL	Stanford	analyze mass spectrometry data
1965	MACSYMA	MIT	symbolic mathematics problems
1972	MYCIN	Stanford	diagnosis of blood diseases
1972	Prospector	SRI	mineral exploration
1975	Cadeceus	UPitt	internal medicine
1978	Digitalis	MIT	digitalis therapy advice
1979	Puff	Stanford	obstructive airway diseases
1980	R1	CMU	computer configuration
1982	XCON	DEC	computer configuration
1983	KNOBS	Mitre	mission planning
1983	ACE	AT&T	diagnose faults in phone cables
1984	FAITH	JPL	spacecraft problem diagnosis
1986	ACES	Aerospace	satellite anomaly diagnosis
1987		AMEX	credit authorization
1992	MAX	NYNEX	phone network troubleshooting

Architecture of a typical expert system

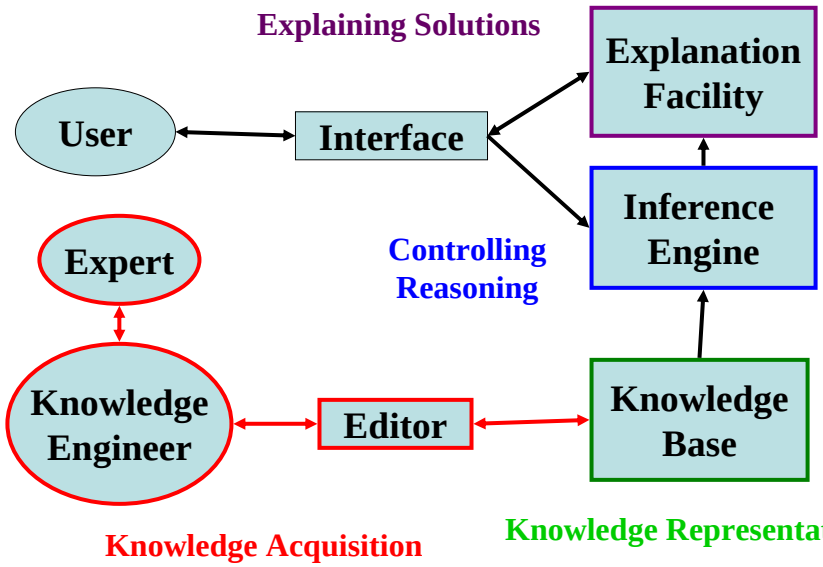

MYCIN diagnoses bacterial infections of the blood:

…
Have you obtained positive cultures indicating general type? *Yes*
What type of infection is it? *Primary bacteremia*
When did symptoms first appear? *May 5*
Let's call the most recent positive culture C1. From what site was C1 taken? *From the blood*
When? *May 9*
Let's call the first significant organism from this culture U1. Do you know the identity of U1? *No*
Is U1 a rod or a coccus or something else? *Rod*
What is the gram stain of U1? *Gram-negative*
Have there been positive cultures from sites that are not associated with the infections about which you are seeking advice? *No*
Did U1 grow aerobically? *Yes*
…
My recommendation will be based on the opinion that the identity of U1 may be

1. Pseudomonas-aeruginosa	2. Klebsiella-pnemoniae	3. E. coli
4. Bacteroides-fragilis	5. Enterobacter	6. Proteus-nonmirabilis

…
My first recommendation is as follows: to cover items 1,2,3,5 and 6, give gentamycin using a dose of 119 mg (1.7 mg/kg) q8h IV (or IM) for 10 days. Modify dose in renal failure. Also, to cover item 4, give clindamycin using a dose of 595 mg (8.5 mg/kg) q6h IV (or IM) for 14 days. If diarrhea or other gastrointestinal symptoms develop, check for pseudomembranous colitis.

MYCIN is a *rule-based deduction system*

Rule:

if *condition 1*
condition 2
…
then ***p***
inference 1
inference 2
…

Reasoning proceeds

- Forward from assertions
- Backward from hypotheses

Sample rules:

(defrule 52
if (site culture is blood)
(gram organism is neg)
(morphology organism is rod)
(burn patient is serious)
then 0.4
(identity organism is pseudomonas)

(defrule 71
if (gram organism is pos)
(morphology organism is coccus)
(growth-conformation organism is clumps)
then 0.7
(identity organism is staphylococcus)

Designing rules to identify animals

One extended rule for each animal?

Rule 0:

if *?x* has hair
has hooves
has long legs
has long neck
has tawny color
has dark spots
chews cud
gives milk

then *?x* is a giraffe

Working Memory

Stretch has hair
Stretch has long legs
Stretch has long neck
Stretch has tawny color
Stretch has dark spots

Problem: only partial information

Solution: smaller rules yielding intermediate conclusions

Designing better rules to identify animals

R1: if *?x* has hair
then *?x* is a mammal

R2: if *?x* gives milk
then *?x* is a mammal

R3: if *?x* is a mammal
eats meat
then *?x* is a carnivore

R4: if *?x* is a mammal
has pointed teeth
has claws
has forward-pointing eyes
then *?x* is a carnivore

R5: if *?x* is a carnivore
has tawny color
has dark spots
then *?x* is a cheetah

R6: if *?x* is a carnivore
has tawny color
has black stripes
then *?x* is a tiger

Working Memory

Swifty has hair
Swifty eats meat
Swifty has tawny color
Swifty has dark spots
Swifty is a mammal
Swifty is a carnivore
Swifty is a cheetah

Forward-chaining reasoning

Loop through the rule base until no rule produces a new assertion, or goal is reached (e.g. animal is identified)

- For each rule:
  - Try to match each of the rule's antecedents by matching it to known facts in working memory
  - If all the rule's antecedents are supported, assert each consequent unless there is an identical assertion already

Rule Base

R1: …
…
Ri: if antecedent 1
antecedent 2
…
then consequent 1
consequent 2
…
…
Rn: …

Working Memory

assertion 1
assertion 2
…
assertion n
new assertions

Backward-chaining reasoning

R1: if *?x* has hair
then *?x* is a mammal

R2: if *?x* gives milk
then *?x* is a mammal

R3: if *?x* is a mammal
eats meat
then *?x* is a carnivore

R4: if *?x* is a mammal
has pointed teeth
has claws
has forward-pointing eyes
then *?x* is a carnivore

R5: if *?x* is a carnivore
has tawny color
has dark spots
then *?x* is a cheetah

R6: if *?x* is a carnivore
has tawny color
has black stripes
then *?x* is a tiger

Working Memory

Swifty has hair
Swifty eats meat
Swifty has tawny color
Swifty has dark spots
Swifty is a mammal
Swifty is a carnivore
Hypothesis: Swifty is a cheetah

Dealing with uncertainty in the MYCIN system

Sources of uncertainty:

(1) data that is provided as input to the reasoning process

(2) possible inferences that can be drawn from data

In MYCIN, the certainty of an assertion is quantified with a *certainty factor*, cf:

(false) $-1 \leq cf \leq +1$ (true)

• User can specify certainty of input data:
> Is U1 a rod or a coccus or something else? *(Rod 0.8)*

cf

• Rules specify the certainty of the consequents, *given the certainty of the antecedents*

(defrule 52
if (site culture is blood)
(gram organism is neg)
(morphology organism is rod)
(burn patient is serious)
then 0.4
(identity organism is pseudomonas)

Propagating uncertainty through rules

(defrule 52
  if (site culture is blood) 1.0
    (gram organism is neg) 0.8
    (morphology organism is rod) 0.7
    (burn patient is serious) 1.0
  then 0.4
    (identity organism is pseudomonas)

cf of antecedents =
minimum of individual cf's
= **0.7**

cf of final assertion =

cf of antecedents x certainty of consequent

cf of (identity organism is pseudomonas) =

0.7 x 0.4 = 0.28

Combining multiple sources of evidence

(defrule 52
  if (site culture is blood)
    (gram organism is neg)
    (morphology organism is rod)
    (burn patient is serious)
  then 0.4
    (identity organism is pseudomonas)

cf: A = 0.28

(defrule 75
  if (gram organism is neg)
    (morphology organism is rod)
    (compromised-host is yes)
  then 0.6
    (identity organism is pseudomonas)

cf: B = 0.36

final cf of assertion =

A + B – AB A,B > 0 **cf = 0.28 + 0.36 – (0.28 x 0.36) = 0.54**

A + B + AB A,B < 0

(A + B)/(1 - min(|A|, |B|)) A,B have opposite signs