

Intelligent Systems: Reasoning and Recognition

James L. Crowley

MoSIG M1

Winter Semester 2020

Lecture 14

26 march 2020

Symbolic Reasoning, Expert Systems and MYCIN

Outline:

Knowledge and Reasoning	2
Kinds of Knowledge	2
Introduction to Expert Systems	3
Expert System = Inference Engine + Domain Knowledge	3
The MYCIN Expert System.....	4
MYCIN: An Antibiotics Therapy Advisor.....	5
Facts	6
PARAMETERS	8
The MYCIN Confidence Factor	8
Backward Chaining Rules	9
Backward Rules Generate Goals	10
Independent Rules and the Combine Function.....	11
Co-routines: Findout and Monitor	12

Knowledge and Reasoning

What is knowledge? - Competence

Whatever enables the solution of problems.

Kinds of Knowledge

Cognitive Psychologists identify different categories of knowledge representation.

Declarative: A symbolic expression of competence.

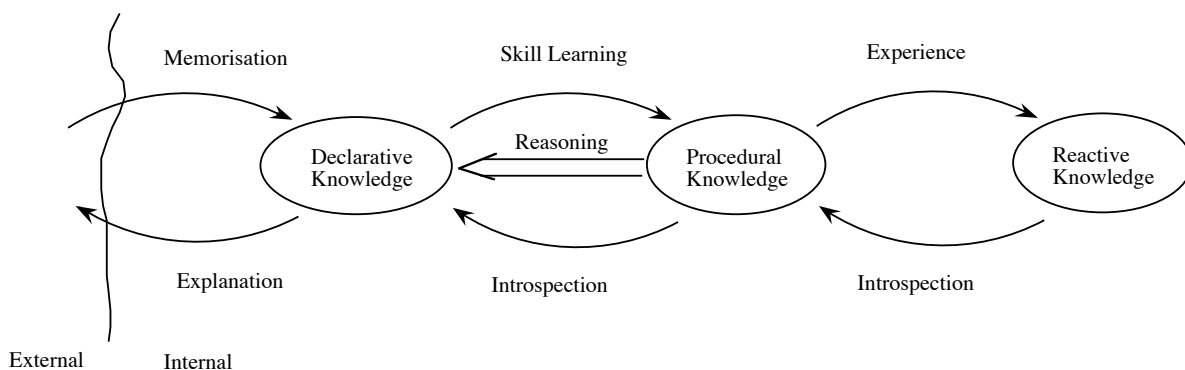
Declarative knowledge is abstract

Declarative knowledge is used to communicate and to reason.

Declarative knowledge must be interpreted to be used.

Procedural: A series of steps to solve a problem

Reactive: stimulus - response. Perception-Action.



Superficial knowledge provides reasoning without understanding. A common example of **superficial** reasoning is reasoning by symbol manipulation, without regard to the meaning of the symbols.

Deep knowledge requires the ability to predict and explain, and requires some form of model.

Expert Systems are a technology that attempts to solve problems purely with Declarative and Procedural representations.

Introduction to Expert Systems

Expert systems are a class of software that is useful for domains that are

Subjective,

Poorly formalized, and

require manipulating large numbers of poorly related facts.

Examples include diagnosis, counseling, debugging, game playing, design in complex spaces and problem solving.

Expert system provides an alternative to algorithmic programming.

In the 1980s and 1990s they provided a first viable technology for Artificial Intelligence.

Expert System = Inference Engine + Domain Knowledge

An expert system combines a domain independent "inference engine" with a database of domain knowledge. The domain knowledge is encoded symbolically as rules and facts.

Expert Systems are constructed by hand coding symbolic expressions of the expertise provided by a "domain expert". Such systems are constructed by iterative refinement through the collaboration of a programmer and a domain expert. The programmer "imitates" the system behavior, evoking corrections and advice by the expert. The programmer then encodes these as fact, rules and data structures. The result is a system that proposes solutions to problems using superficial reasoning. Such systems may appear to understand but are in fact superficial. They tend to reason without regard to meaning. The system can make non-sense statements if applied outside their domain.

The technologies for "Expert Systems" are based on symbol manipulation, without regard for the meaning of the symbols. Thus Expert Systems are based on symbolic encoding of declarative and procedural knowledge from experts. The reasoning is syntactic. Thus it is shallow. The system cannot judge when it has been asked to solve a problem outside its domain.

While Expert Systems technologies provide useful solutions for many applications, there is a fundamental problem: Hand-crafting a knowledge base is generally a very expensive and difficult process.

The MYCIN Expert System

In the 1970's, at Stanford University, Edward Feigenbaum directed a research group named the "Heuristic Programming Project". Their central thesis was

Intelligence = Large quantity of domain knowledge and a little bit of reasoning.

This led to an investigation into Knowledge Representation Techniques.

From 1970 to 1973 they sought to build a system that could interpret data from Mass Spectrograms.

A Mass Spectrograph is a device that uses electrical or magnetic fields to determine the masses of atoms or molecules in a sample. A beam of ions is passed through the electrical or magnetic field. The field deflects the ions at different angles depending on their masses, thereby breaking the beam into separate, identifiable bands.

The result of their project was a system named DENDRAL.

DENDRAL was an un-maintainable "hack". However, by 1973 the group had learned to express declarative knowledge as "rules". It was decided to start over, building a "rule based" system for "anti-biotic Therapy".

Penicillin was discovered in 1929 and came into widespread use as an antibiotic in the 1940's. During the 1950's and 1960's a variety of new antibiotics were discovered. Each had unique properties and uses. By the 1970s, most medical doctors required consultation with a specialized expert to prescribe antibiotics.

The Stanford University Medical School was a world famous center for research in antibiotics. The Medical School asked the Computer Science School for help. Feigenbaum proposed to construct an "Artificial Expert" antibiotic therapy advisor.

The MYCIN system was created in the doctoral research of Edward Shortliffe, under that direction of Bruce Buchanan in the HPP team at Stanford.

The system developed from 1973 to 1978. It evolved into the first true "Expert System". As such it became the model for a new class of systems. It also revealed the importance of reasoning with uncertainty and the knowledge acquisition problem.

MYCIN: An Antibiotics Therapy Advisor.

By 1975, a large variety of Antibiotics were available. Each antibiotic was effective against a specific set of microbes, and triggered a specific set of side effects.

Patients were often allergic to certain families of antibiotics.

MYCIN was designed to be used by ordinary doctors who lacked the specialized training required to develop anti-biotic therapies.

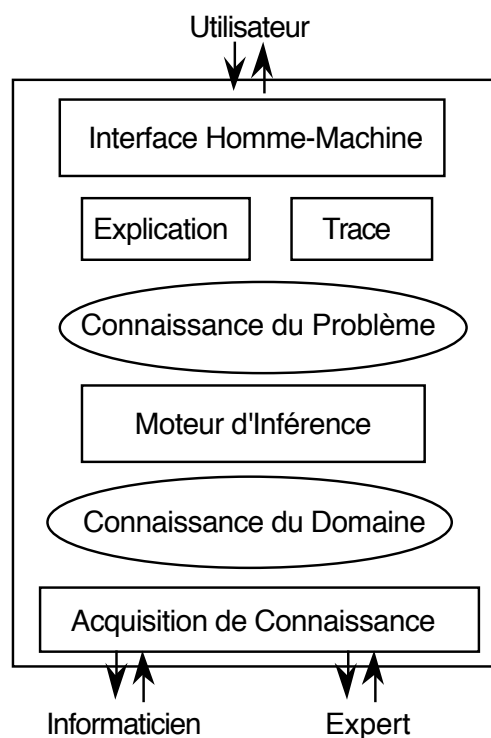
Specifications: The system was required to be:

- Easy to use
- Reliable
- Able to manipulate large numbers of unrelated facts.
- Able to use inexact and incomplete facts

Able to explain its advice.

MYCIN was composed of approximately 600 rules, manipulating a large base of structured facts. The rules provided procedural knowledge to

- 1) Request or infer the required information
- 2) Apply specialized knowledge to determine a therapy
- 3) Provide advice to doctors in the form of proposed therapy.
- 4) Respond to questions about its reasoning.



Facts

All facts in MYCIN are represented by a "quadruple":

(Context, Parameters, Value, CF)

C = Context (a group of facts related to an entity)

P = Parameter (facts) about a context

V = Value (Boolean, Numeric, Symbol, List)

CF = Confidence Factor $CF : [-1,1]$

Contexts structure reasoning and provide control.

For the MYCIN antibiotic therapy advisor, 10 contexts were required.

PERSON: Data about the patient, including age, sex and weight

OPERS: Past Medical procedures

CURCULS: Medical cultures taken from the patient

CURDRUGS: Current drug therapies for the patient

CURORGS: Known microbial infections in the patient

OPDRUGS: Drugs used during recent medical procedures

PRIORCULS: Past medical cultures

PRIORDRUGS: Past medical therapies

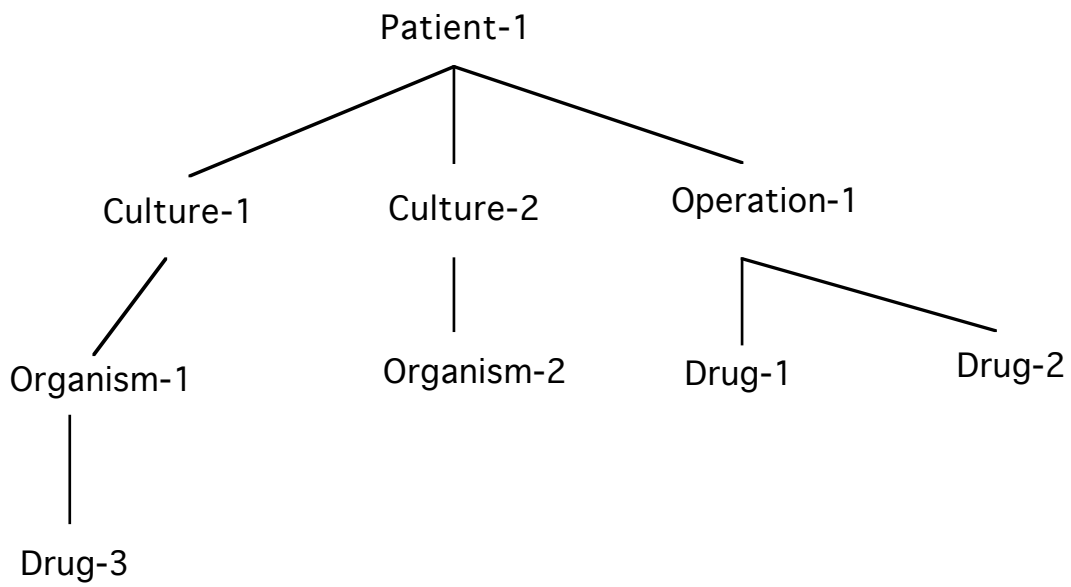
PRIORORGS: Past infections.

Contexts are organized in a tree.

MYCIN Context Tree was structured to respond to 4 questions.

- 1) What symptoms does the patient show
- 2) What microbes infect the patient
- 3) What antibiotics are effective against the microbes
- 4) What is the most appropriate antibiotic.

The context tree was called the "Dynamic Tree", and was composed of instances of each context.



The context tree served to focus reasoning.

PARAMETERS

Each context was composed of a number of parameters. Each parameter was described by a data structure.

Attributes of Parameters included:

Expect : {Y/N, NUMB, ONE_OF, ANY_OF}

PROMPT : A sentence to ask for the value of the parameter.

LABDATA (Y/N) : Whether the parameter should be requested from the doctor or inferred automatically by the system.

(LABDATA was later renamed "ASK-FIRST")

LOOKAHEAD: A list of rules for inferring the value of a Parameter.

TRANS: An English language explanation of the parameter and the meaning of its values.

Extensive pre-coding of English sentences allowed the system to appear capable of intelligent dialog. The MYCIN system could almost pass the Turing Test!

PARAMETER Categories:

SingleValued: Parameter could take a single value. If multiple values are provided the most likely must be determined.

MultiValued: The parameter could have multiple values. A fact was created for each value.

Binary: A Boolean value; A single value that can be Yes or No

Reasoning between alternative single valued parameters required some form of evidential reasoning. For this, the MYCIN team invented a "Confidence Factor".

The MYCIN Confidence Factor

ALL facts in MYCIN are labeled with a confidence Factor $CF \in [-1, 1]$

ALL rules in MYCIN are labeled with a FORCE: $CF \in [-1, 1]$.

Backward Chaining Rules

(Context, Parameters, Value, CF)

Facts: Quadruple (C, P, V, CF)

C = Context (a group of facts related to an entity)

P = Parameter (facts)

V = Value (Boolean, Numeric, Symbol, List)

CF = Confidence Factor $CF : [-1,1]$

Backward chaining rules are interpreted as an "AND-OR tree"

$$(C, P_2, V_2, CF_2) \xrightarrow{CF_R} (C, P_1, V_1, CF_1)$$

$$CF_1 = CF_R \cdot CF_2$$

$$(C, P_3, V_3, CF_3) \text{ AND } (C, P_4, V_4, CF_4) \xrightarrow{CF_R} (C, P_1, V_1, CF_1)$$

$$CF_1 = CF_R \cdot \min\{CF_3, CF_4\}$$

Mycin rules can be disjunctive (use OR)

$$(C, P_2, V_2, CF_2) \text{ OR } ((C, P_3, V_3, CF_3) \text{ AND } (C, P_4, V_4, CF_4)) \xrightarrow{CF_R} (C, P_1, V_1, CF_1)$$

$$CF_1 = CF_R \cdot \max\{CF_2, \min\{CF_3, CF_4\}\}$$

Backward Rules Generate Goals


Backward Chaining with Rules generates Goals: $(C, P, ?V, ?)$

Our two rules are interpreted as:

IF Goal = $(C, P_1, ?V_1, ?)$ Then Ask $(C, P_2, ?V_2, ?)$

IF Goal = $(C, P_1, ?V_1, ?)$ Then Ask $(C, P_3, ?V_3, ?)$ AND $(C, P_4, ?V_4, ?)$

Graphical Representation:



Goals are then expanded recursively with additional rules.

Some rules can recursively open a new context. This context must be completed before the previous context.

Independent Rules and the Combine Function.

NOTE that disjunction is different from mutually independent Rules.

With two separate Rules, R1 and R2.

$$\begin{aligned} (C, P_2, V_2, CF_2) &\xrightarrow{CF_{R1}} (C, P_1, V_1, CF_1) \\ (C, P_3, V_3, CF_3) \text{ AND } (C, P_4, V_4, CF_4) &\xrightarrow{CF_{R2}} (C, P_1, V_1, CF_1) \end{aligned}$$

$$CF_1 = \text{Combine}(CF_{R1} \cdot CF_2, CF_{R2} \cdot \min\{CF_3, CF_4\})$$

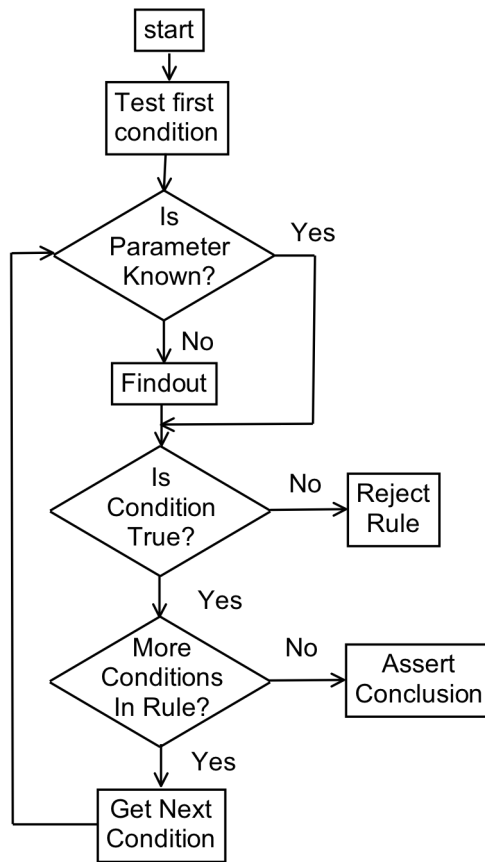
The confidence from independent rules is interpreted with the function Combine

$$\text{Combine}(CF_1, CF_2) = \begin{cases} CF_1 + CF_2 - CF_1 \cdot CF_2 & \text{if } CF_1 \geq 0 \text{ AND } CF_2 \geq 0 \\ -\text{Combine}(-CF_1, -CF_2) & \text{if } CF_1 < 0 \text{ AND } CF_2 < 0 \\ \frac{CF_1 + CF_2}{1 - \min(|CF_1|, |CF_2|)} & \text{if } CF_1 \cdot CF_2 < 0 \end{cases}$$

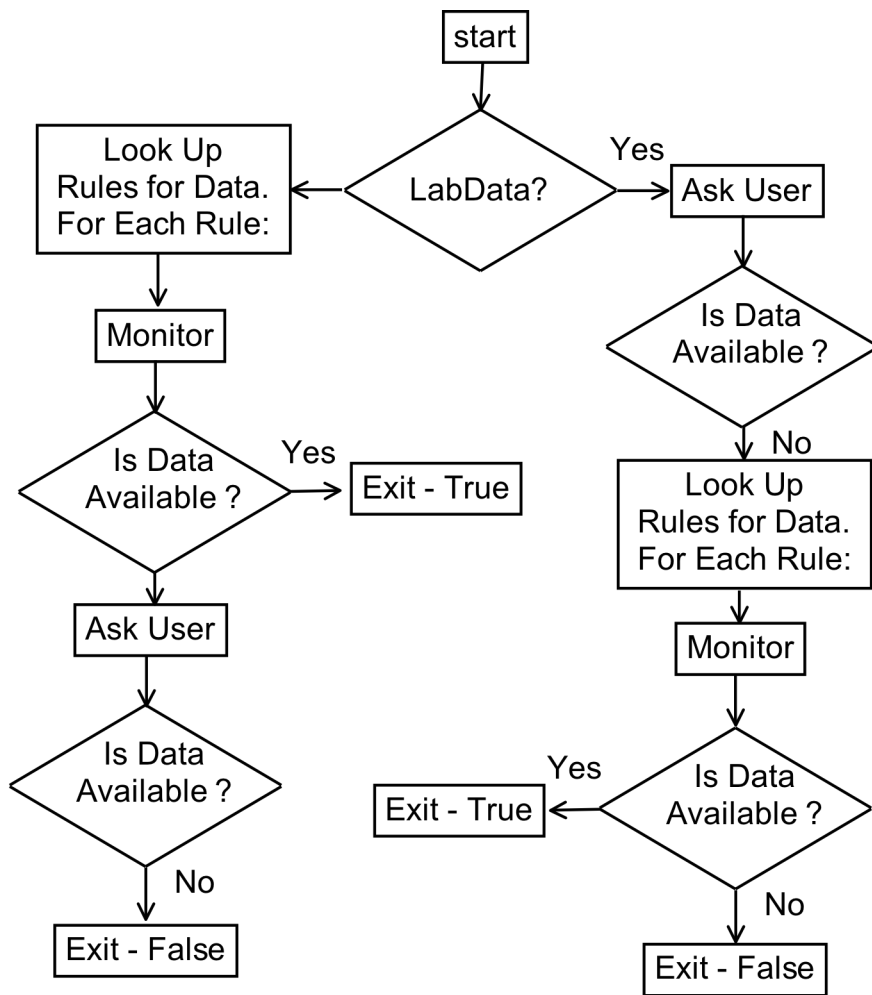
Co-routines: Findout and Monitor

Mycin rules are interpreted using two recursive Coroutines: Findout and Monitor

MONITOR :



FINDOUT:



At any instant the user may ask

WHY? The system provide an interpretation of the trace of reasoning

HOW: The system provides the source for a fact.

Coupled with the extensive use of preprogrammed sentences, this made the system appear to be intelligent. However the reasoning was shallow.

The Mycin system demonstrated the importance of reasoning with uncertain facts. This triggered many groups to investigate and propose alternatives to reasoning with uncertainty.

Ultimately (after 10 - 15 years) most of the community settled on Bayesian Reasoning.