Expert System

One Definition of Expert System

- A computing system capable of representing and reasoning about some knowledge rich domain, which usually requires a human expert, with a view toward solving problems and/or giving advice.

  - the level of performance makes it “expert” - Some also require it to be capable of explaining its reasoning.

  - Does not have a psychological model of how the expert thinks, but a model of the expert’s model of the domain.

Categories of Expert Systems

Category	Problem Addressed
Prediction	Inferring likely consequences of given situations
Diagnosis	Inferring system malfunctions from observations, a type of interpretation
Design	Configuring objects under constraints, such as med orders
Planning	Developing plans to achieve goals (care plans)
Monitoring	Comparing observations to plans, flagging exceptions
Debugging	Prescribing remedies for malfunctions (treatment)
Repair	Administer a prescribed remedy
Instruction	Diagnosing, debugging, and correcting student performance
Control	Interpreting, predicting, repairing, and monitoring system behavior

Knowledge in a Knowledge Base

- Knowledge specific to the domain + facts specific to the problem being solved
- A medical KB is defined in HANDBOOK of MEDICAL INFORMATICS as: -

"a systematically organized collection of medical knowledge that is accessible electronically and interpretable by the computer."

- They note "a medical KB usually:

  - includes a lexicon (vocabulary of allowed terms) and - specifies relationships between terms in the lexicon. "

- For example, in a diagnostic KB, terms might include:

  - patient findings (e.g., fever or pleural friction rub), - disease names (e.g., nephrolithiasis or lupus cerebritis) and - diagnostic procedure names (e.g., abdominal auscultation or chest computed tomography).

- Knowledge Representation is the key issue

  - Aim is usually to present the knowledge in as "declarative" a fashion as possible

Traditional Feature Comparisons: E/KBS versus ANN

E/KBS	Neural Networks
• Symbolic	• Numeric
• Logical	• Associative
• Mechanical	• Biological
• Serial	• Parallel
• Rule Based	• Example Based
• Needs “Rules”	• Finds “Rules”
• Much Programming	• Little Programming
• Requires Reprogramming	• Adaptive System
• Needs an Expert	• Needs a Database

But much of this too simple, KBS are not really “logical” and can use examples etc.

Processing Comparisons

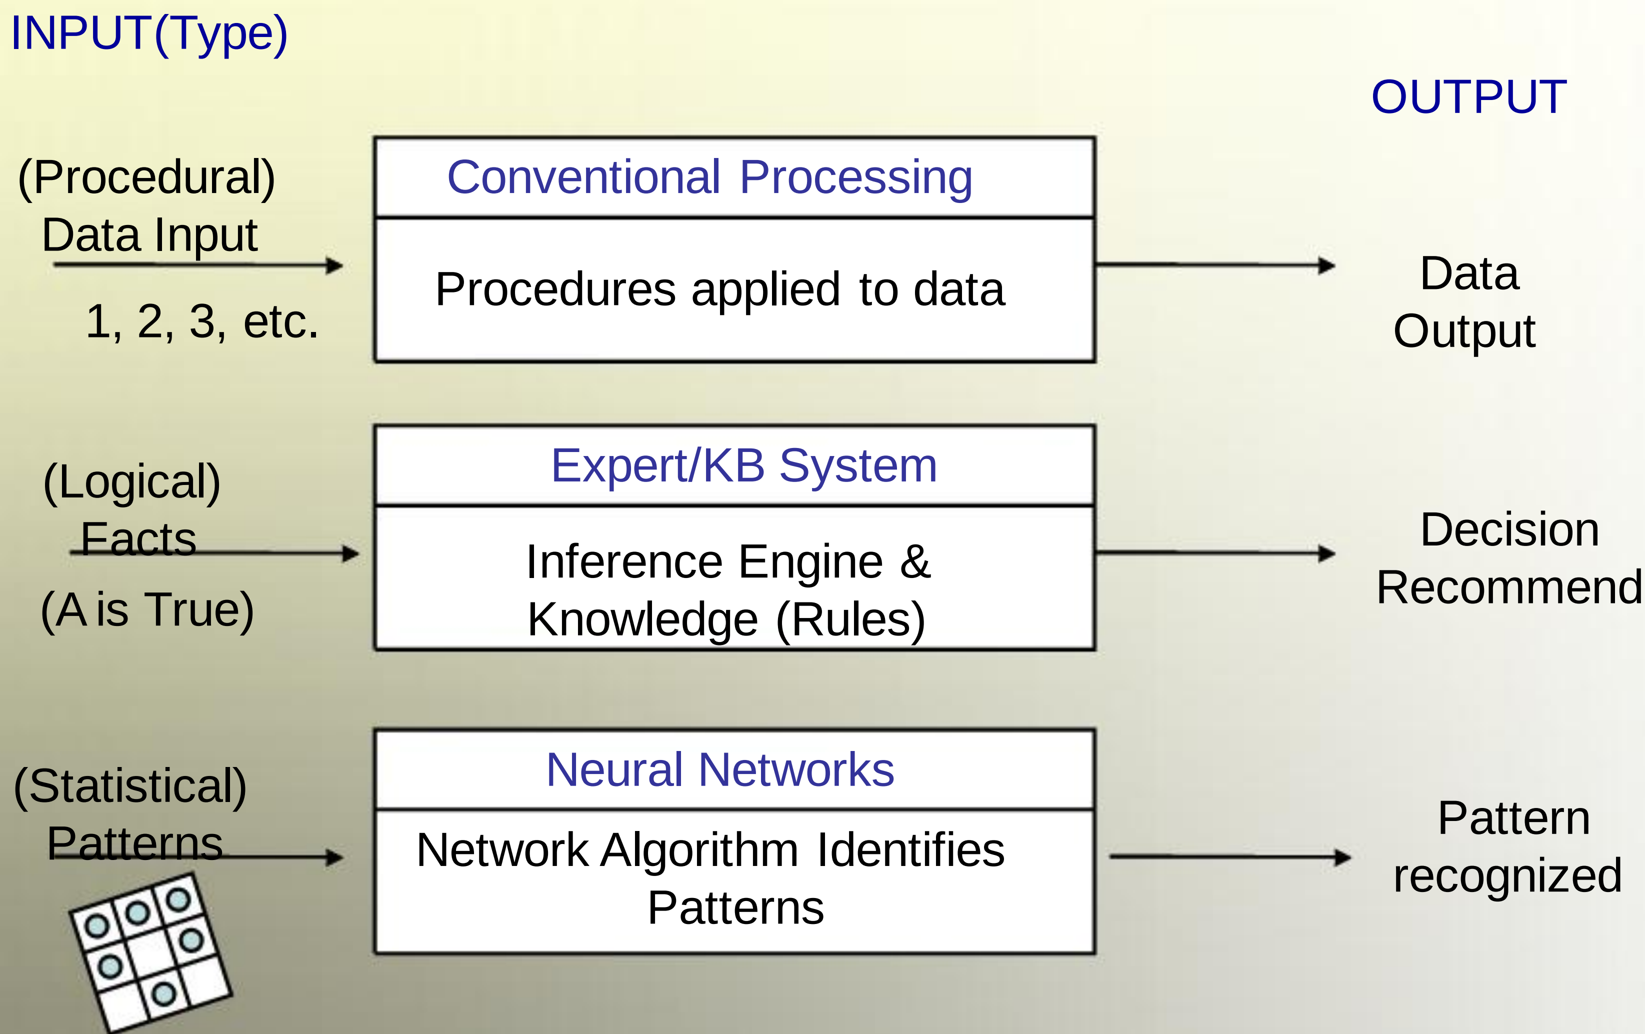

Medical Expert and KB systems

- are designed to give expert-level, problem-specific advice in the areas of :

  - medical data interpretation,
  - patient monitoring,
  - disease diagnosis, - treatment selection, - prognosis, and
  - patient management.

- Research in medical expert and knowledge-based systems and the development of such systems has been most significant to the broad realm of quality assurance and cost containment in medicine.

One Distinction Between an Expert System and a Knowledge-Based System

• To be classified as an ‘expert system’ the system must be able to explain the reasoning process. • This is often accomplished by displaying the rules that were applied to reach a conclusion.

Some Basic Concepts

- Knowledge representation deals with the formal modeling of expert knowledge in a computer program.
  - Important questions in this respect concern the given degree of structuralization of the medical domain under consideration, the necessity to include vagueness of medical terms and uncertainty of medical conclusions into the chosen formal representation, as well as the extent and completion of the respective knowledge domain.
- Reasoning mechanisms are inference methods which draw medical conclusions from given patient data by means of the stored medical knowledge.
  - Most important is the selection of the appropriate formal approach with respect to the given medical domain.
  - One differentiates methods to infer logical conclusions (e.g., propositional and predicate logic, three-valued logic, fuzzy logic, non-monotonic logic) and to combine medical evidence (e.g., Bayes theorem, certainty factors, Dempster-Shafer theory of evidence).

Assertional Knowledge

- It might be a detailed description of a complex domain like a disease, a linguistic structure, etc.
- This type of knowledge is used to describe a given clinical situation usually in an object structure. • This is done by associating the different elements or objects characterizing the context inside the same framework with the consideration of the relationships between these objects.
  - Example: an exhaustive description of a specific disease organized following: the set of its symptoms, its possible treatments, medicines, etc.

Alternative KB Approaches

- Rule-based approach

  - Events trigger firing of rules (condition/action pattern)
  - e.g. Arden Syntax and Medical Logic Modules (MLM)

- Case & Model-based approach

  - Create a model (template) of clinical guidelines
  - e.g. PRODIGY, EON, PROforma, GLIF

But AI is a broad field - a tree representation

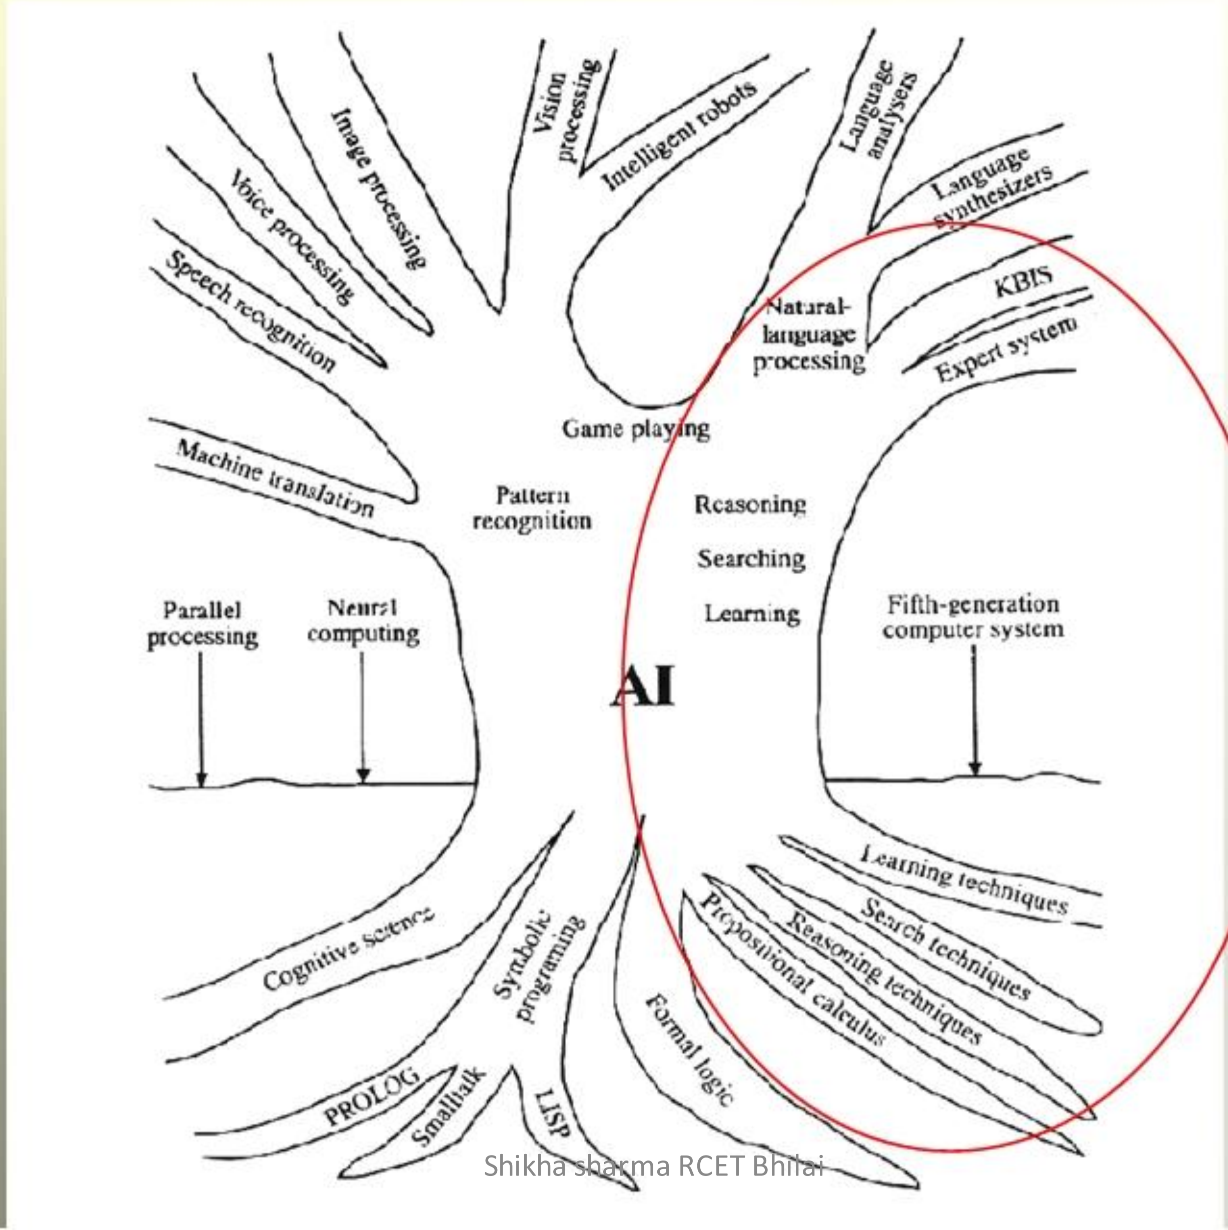

Knowledge-base may really include many things

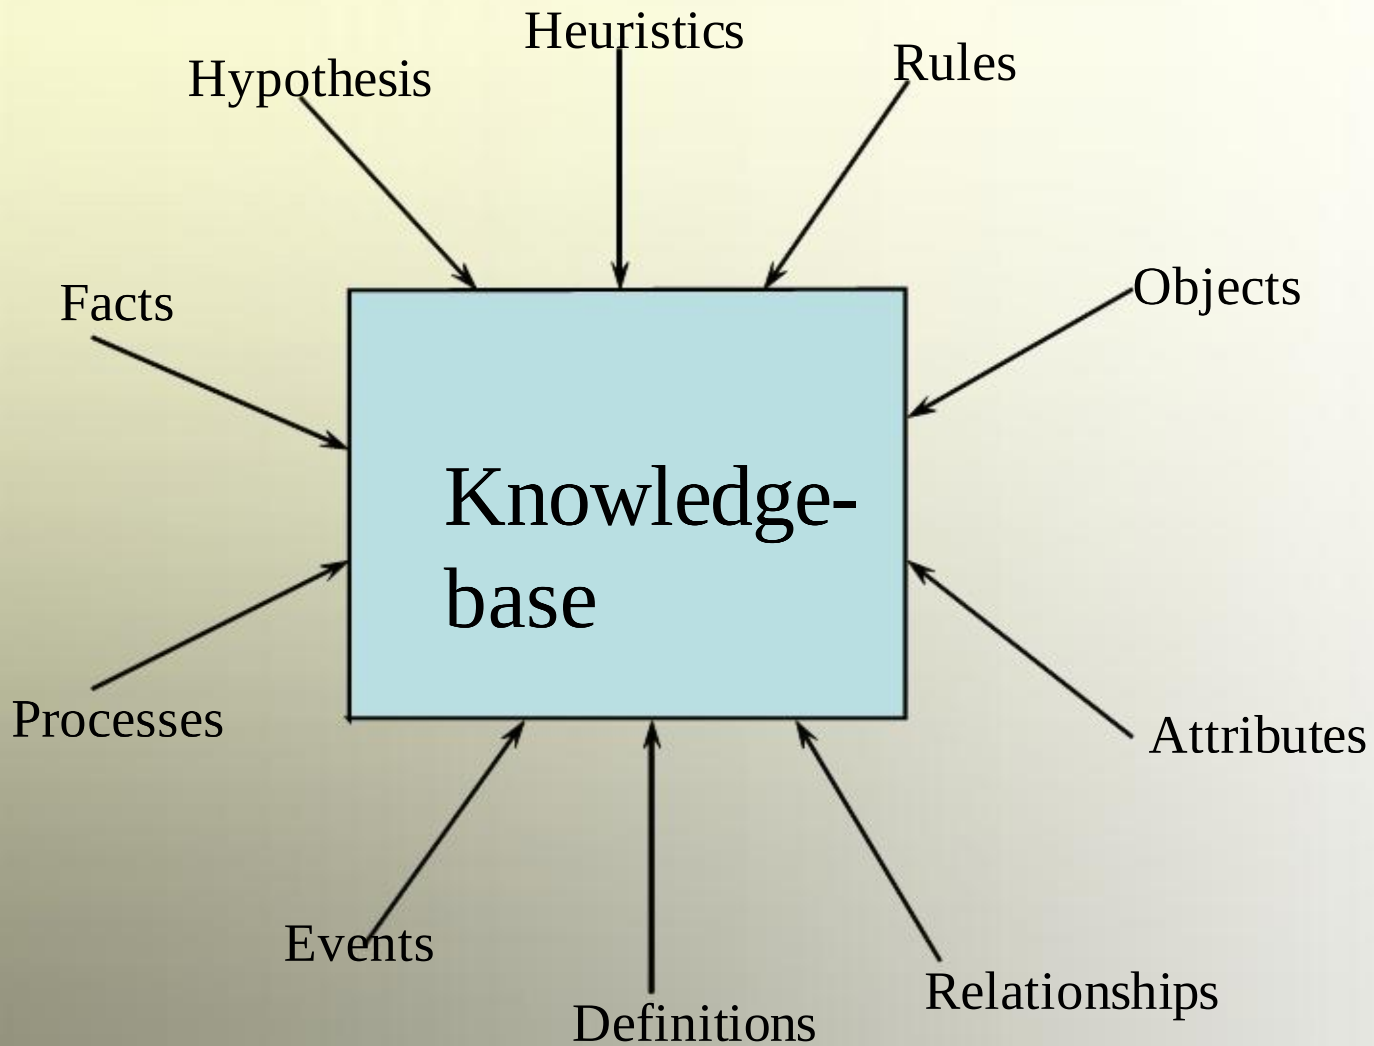

- An expert system is a system that uses human knowledge captured in a computer to solve problems that ordinarily require human expertise.

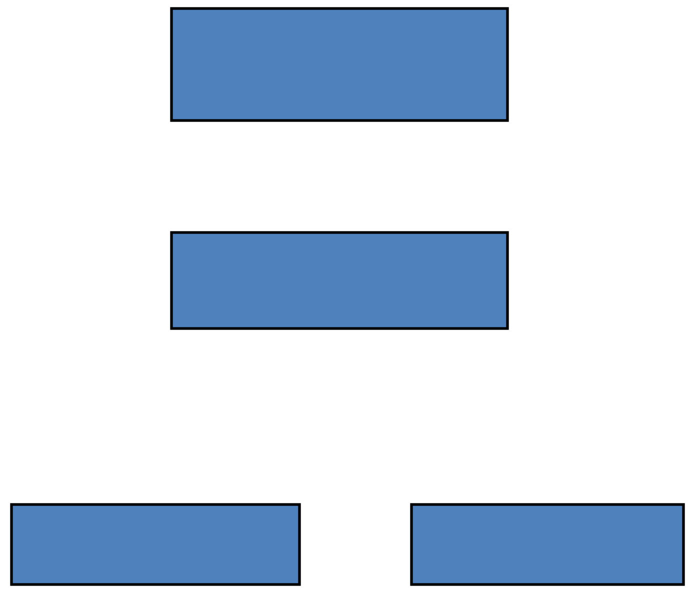

Expert System Structure

In theory we could take the knowledge base & use different inference mechanisms, or take the inference mechanism & use different KBs.

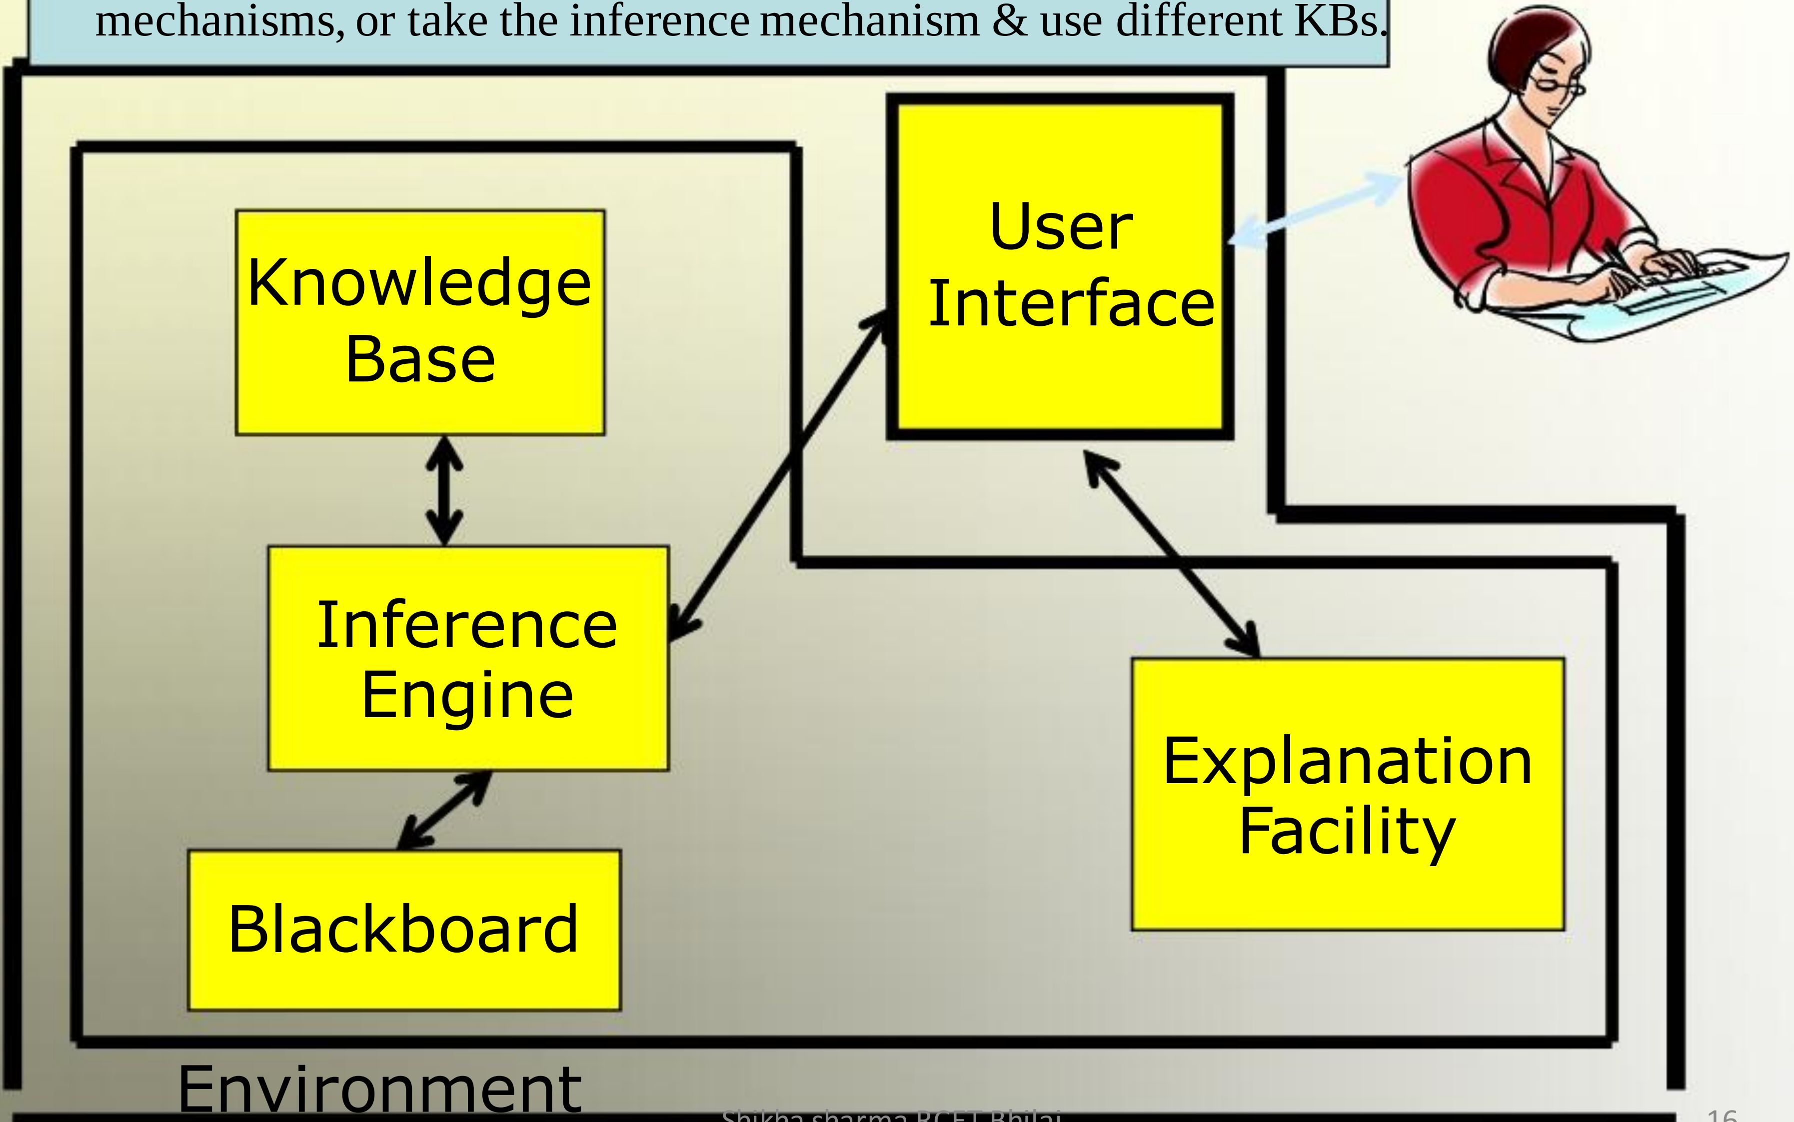

KBS architecture and components

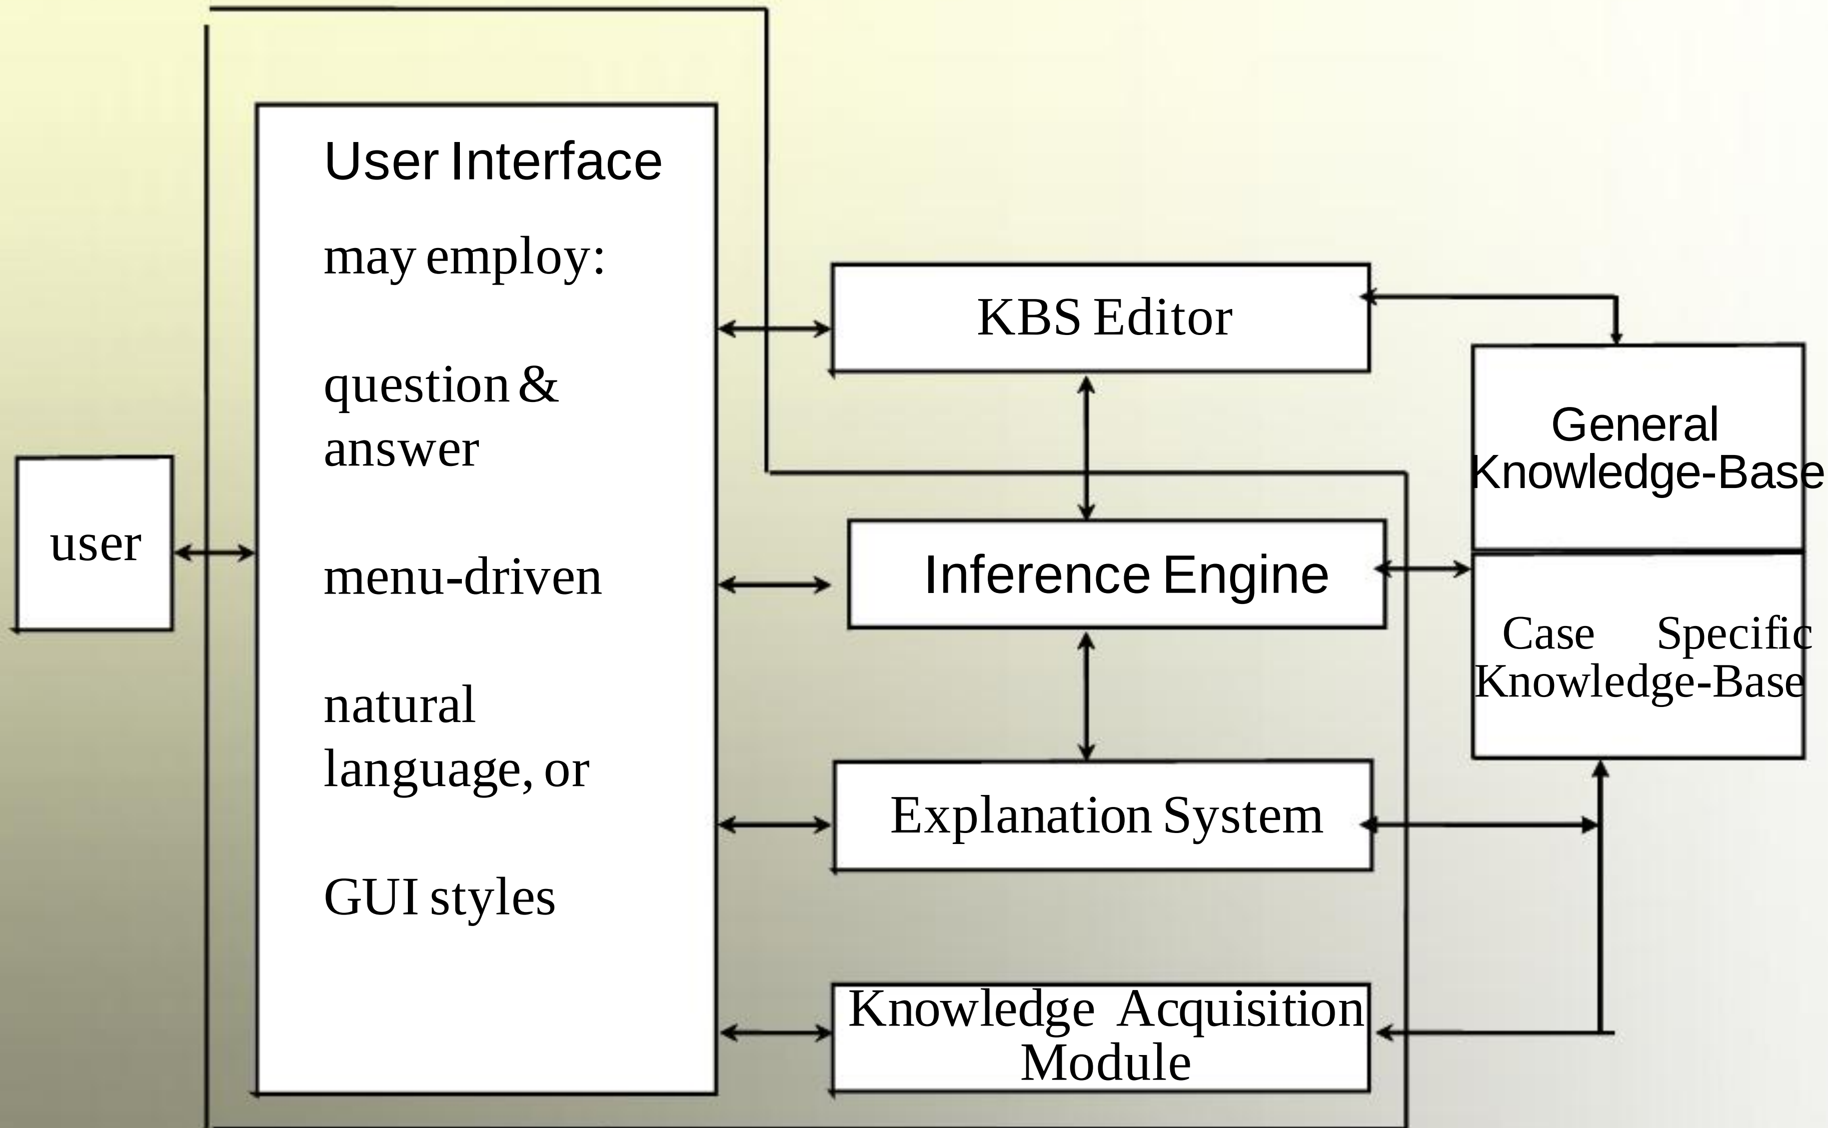

Knowledge representation formalisms & Inference

KR	Inference
* Logic	Resolution principle
* Production rules	backward (top-down, goal directed)
	forward (bottom-up, data-driven)
* Semantic nets & Frames	Inheritance & advanced reasoning

Expert System Shells

- Separate the mechanisms for making inference from the rule base
- Facilitate the entry of rules by non-programmers
- Provide reuse for what would otherwise be redundant code across expert systems

Expert System Components

- Inference Engine
  - Forward or Backward-chaining
  - Conflict resolution algorithms
- Rule-base
  - IF-THEN rules
- Database
  - Current state on which IF-THEN rules are applied.
- Explanation Facilities
  - An important advantage rule-based expert systems hold over other types of AI

Inference Engines

- Forward-chaining
  - Submit current data to all rules
  - Rules make conclusions, which in turn, generate new data
  - “Inference Chains” result from initial data and the data generated in conclusions.
- Backward-chaining
  - Try to prove a conclusion by working backwards from ways to prove it.

Forward-chaining Example (A,B,E, and D are given)

- If Y and D then Z
- If X and B and E then Y
- If A then X
- If C then L
- If L and M then N

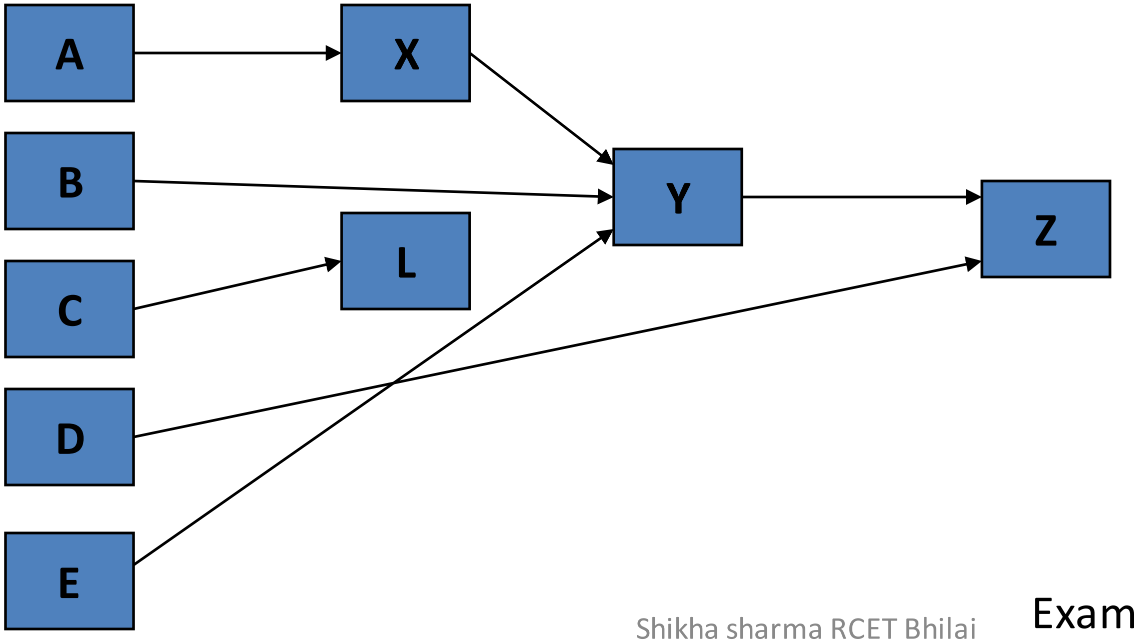

Example from Negnevitsky

Backward Chaining

- Prolog uses backward chaining
  - Work backward from the goal.
  - Check rules that can provided the desired goal.

Backward Chaining Example

- If Y and D then Z
- If X and B and E then Y
- If A then X
- If C then L
- If L and M then N

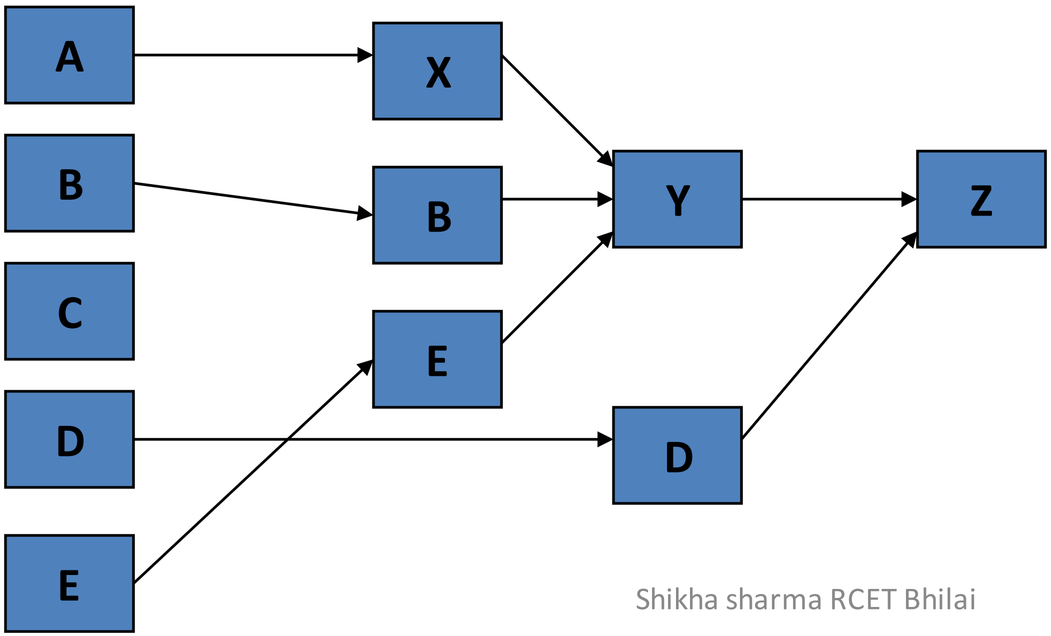

Forward or Backward Chaining?

- What do experts use?
- Are we trying to prove a particular hypothesis?
  - Backward chaining
- Are we trying to find all possible conclusions?
  - Forward chaining
- What does the rule set look like?
  - Could be either one or a combination of both.

Conflict Resolution

- What happens when two rules provide conflicting conclusions?
  - If it has feathers then it is a bird
  - If it can't fly then it is not a bird
  - What if has feathers, but can't fly?

Conflict Resolution Methods

- Use rule-order as an implied priority
  - The first rule to provide an answer is used.
- Assign a priority to each rule, the rule with the higher priority is sustained.
- Longest Matching Strategy uses the rule with the most specific information.
  - If it cannot fly and has feathers then it is a bird.
- Certainty-based conflict resolution
  - Measures of certainty are provided for data and rules. Most certain rule is sustained.

Life Cycle for Developing Expert Systems

- Problem Definition
- Knowledge Acquisition
- Knowledge Representation
- Prototype system
- Operational system
- Knowledge base maintenance

Problem Definition

- The essential problem is selecting an appropriate domain:

  - the problem must require some type of specialized knowledge, if there are human "experts" this criteria is probably satisfied

  - must not be overly large: define the problem fairly narrowly.

  - in business organizations, it should a problem that is handled often enough that an investment is expected to have some payoff: the once every 5 years sort of problem going to payoff.

Knowledge Acquisition

- " the transfer and transformation of potential problem-solving expertise from some knowledge source to a program.”

  \- Buchanan 1983.

- machine learning - building capabilities into the system that allow it to learn from what it is doing.
  - the problem of induction - how many instances must be observed before it can be added to the knowledge base as "true"

Knowledge Acquisition (cont.)

- knowledge elicitation - extract the knowledge from the human expert, through some means
  - direct - interaction with the human expert interviews, protocol analysis, direct observation, etc.
  - indirect - utilize statistical techniques to analyze of data and draw conclusions about the structure of the data.

Knowledge Representation

- A method to represent the knowledge you are eliciting and/or learning.
- Several major methods -rules, bayes nets, frames
- Strengths and weaknesses for each.
- None is completely dominant.
- Trent is to build heterogeneous systems, that 's what experts are.

Knowledge Representation

- A method to represent the knowledge about the domain
- Three major symbolic methods:
  - rules
  - semantic objects
  - logic
- Although a shell contains a way to represent knowledge, shell selection should be influenced by the matching the representation to the knowledge in the domain.
- Knowledge must be coordinated, so that the knowledge base is consistent.

Prototype system

- Typically use an "incremental" development approach to an expert system.

  - Build an initial prototype and adjust and expand - Allow the expert to interact with the prototype to get feedback

- Reevaluate if the project should be continued, if major redesign (knowledge representation) is necessary, or to go ahead.

Build Operational System & Knowledge base maintenance

- Once The actual system is built

  - New rules can be continually added and old ones refined/ removed.

- This is a tricky process, but there does not seem to be much literature on it.
- One characteristic of an Expert system should be maintainability, so the ability to add/change/delete rules is essential.

A Representation: First-Order Logic

- Constants: Mr_Smith, Dr._Jones, anemia
- Variables: X, Y
- Functions: Address(X), Age(Y)
- Predicates: Diagnosis(X, anemia); Male(Y); Patient(Z)
- Negation: ¬Male(X); ¬Name(X, Smith)
- Connectors:

- Conjunction (AND): Patient(X) ∍ Male(X) - Disjunction (OR): Doctor(X) (Nurse(X) - Logical implication: Female(X) ® ¬Male(X) • Quantifiers:

- Universal quantifier: X (Patient(X) (Doctor(X)) - Existential quantifier: □Y (Patient(Y) ∍ Name(Y, Jones))

From Yuval Shahar, “Frame-Based Representations and Description Logics” Temporal Reasoning and Planning in Medicine

Alternatives Ways of Modeling

- X has Diabetes:
  - Diabetes (x)
  - Has_Diagnosis (x, "Diabetes") - Has (x, "Diagnosis", "Diabetes")
- Trade off between efficiency and expressiveness
  - Has (x, y, "Diabetes")

Relationship Of this K to a DB

- Representing patient X has Diabetes in a table:
  - Diabetes (x)
    - A table called Diabetes with column (s) identifying patient x and a column of the value of Diabetes (x) - Has_Diagnosis (x, "Diabetes")
      - A table called Diagnosis with column (s) identifying patient x, and diagnosis y and a column of the value of Has_Diagnosis (x, y)
  - Has (x, "Diagnosis", "Diabetes")
      - A table called observation with column (s) identifying patient x, observation type y and observation value z and a column of the value of Has (x, y, z)

"Production" Rule sets

- Experts typically form sets of rules to apply to a given problem
- Set of rules reflects the skill of the expert on a topic; use different rule sets to reflect problem-solving competence of expert
- Need a strategy to know when to apply them ie use meta rules
- Rule sets often represented in a tree-like structure with most general, strategic rules at the top of the tree; most specific rules at leaf nodes
- Adopts a top-down approach to problem-solving, where rule sets only used when appropriate;
  - reflects human approach divide and conquer
  - eases modular development
  - each module may use different representation and reasoning techniques (say for body system)

Rules & Decision Tree Example

Q1: Test is

Depth First

not included

included

Q2: the Panel cost is

Q3: the panel cost is

2K 3K 4K

>=6K <5K

c3: 10% chance Rule 7

c2: 70% chance Rule 8

Q3: the cost is

Q3: the Cost is

Q3: the cost is

< >=x <y

c1: 30% chance Rule 1

c2: 70% chance Rule 2

c1: 30% chance Rule3

c2: 70% chance Rule 4

c1: 30% chance Rule5

c2: 70% chance Rule6

Easy to find a recipe to turn this into a rule representation.

Knowledge acqusition

Medical Knowledge (Adjusting to Situations)

Biochemical lab rules
Go from simple, modular to confusing complications

From "Toward Situated Knowledge Acquisition"
Tim Menzies,
Int. J of Human Computer Studies, 1998

A. Originally

```
RULE(22310.01)
IF  (bhthy or
     utsh_bhft4 or
     vhthy) and not on_t4
         and not surgery
         and (antithyroid or
              hyperthyroid)
THEN DIAGNOSIS("...thyrotoxicosis")
```

B. Same rule, 9 years later

```
RULE(22310.01)
IF  ((((T3 is missing)
       or (T3 is low and
           T3_BRD is low))
       and TSH is missing
       and vhthy
       and not (query_t4 or on_t4 or
                or surgery or tumour
                or antithyroid
                or hypothyroid
                or hyperthyroid))
     or ((((utsh_bhft4 or
            (Hythe and T4 is missing
             and TSH is missing))
           and (antithyroid or
                hyperthyroid))
          or  utsh_bhft4
          or  ((Hythe or borthy)
               and T3 is missing
               and (TSH is undetect
                    or TSH is low)))
         and not on_t4 and not
             (tumour or surgery)))
   and (TT4 isnt low or T4U isnt low)
THEN DIAGNOSIS("...thyrotoxicosis")
```

Disadvantages of Production Knowledge

Difficult to maintain for Very Large-KB

- One reason is addition of new, contradictory knowledge. Consider

Rule 1. IF it is raining

THEN not (weather is sunny)

Rule 2. IF location is Florida

THEN not (weather is cloudy)

Rule 3. IF it is late afternoon

THEN weather is sunny or weather is cloudy

FACTS: it is late afternoon   location is Florida

Conclude?????

Maintenance is to ADD   Rule 4. IF it is late afternoon AND location is Florida

THEN it is raining

Some observe that RB development never ends….KE is a continuous process…..

KBS as real-world problem solvers

- Problem-solving power does not lie with smart reasoning techniques nor clever search algorithms but

domain dependent real-world knowledge - Real-world problems do not have a well-defined solutions in literature

- Expertise not laid down in algorithms but are domain dependent rules-of-thumb or heuristics (cause-and-effect)
- KBS allow this knowledge to be represented in computer & solution explained
  - These are not “logical”

MYCIN

MYCIN - begun in 1972

- Consultation system assist internists in diagnosis and treatment of infectious diseases: meningitis & bacterial septicemia
- When patient shows signs of infectious disease, culture of blood and urine set to lab (>24hrs) to determine bacterial species • Classified as a "production- rule" system, depth-first, backward chaining.

- Given patient data (incomplete & inaccurate) MYCIN gives interim indication of organisms that are most likely cause of infection & drugs to control disease
  - Uses certainty factors to handle incomplete and uncertain information, included the "how" and "why" capabilities that are now considered essential, defining characteristics of Expert Systems.
- Drug interactions & already prescribed drugs taken into account
- Able to provide explanation of diagnosis (limited)
  - Thoroughly documented in Buchanan and Shortliffe Rule Based Expert Systems, Addison- Wesley, Reading, Mass., 1984.

Top-level goal rule

IF there is an organism which requires therapy, and

consideration has been given to the possibility of additional organisms requiring therapy THEN compile a list of possible therapies, and determine the best therapy in this list.

THERAPY rule

IF the identity of the organism is Pseudomonas

THEN I recommend therapy from among the following drugs: 1 - COLISTIN (.98)

2 - POLYMYXIN (.96) 3 - GENTAMICIN(.96) 4 - CARBENICILLIN (.65) 5 - SULFISOXAZOLE (.64)

THERAPY rule

- The number with each drug is the akin to the probability that a Pseudomonas will be sensitive to the named drug.
- To select the actual therapy, the drugs on the list are screened for contra-indications and to minimize the number of drugs administered, while maximizing sensitivity.

Typical RB Exercise: Write Rules by Diagnosis

- Write rules for patients with the following diagnoses (one at a time):

  - diabetes mellitus

  - heart failure

  - myocardial infarction - benign prostatic hyperplasia

K Engineer compares notes and leads discussion on integration.

Evaluation of MYCIN

- In 1974, an initial study of MYCIN was conducted where five experts approved 72% of MYCIN's recommendations on 15 actual cases.
- The system was improved and in 1979 MYCIN was again compared to experts.

MYCIN's Performance Compared to Human Experts

MYCIN	52	Actual Therapy	46
Faculty-1	50	Faculty-4	44
Faculty-2	48	Resident	36
Inf. Dis fellow	48	Faculty-5	34
Faculty-3	46	Student	24

Ratings by 8 experts on 10 cases

Perfect score = 80

MYCIN is not currently in use:

- Knowledge base is incomplete, does not cover a full spectrum of infectious diseases.
- computing power was not available in most hospital wards.
- MYCIN's development lead to the development of "EMYCIN" - for "Empty MYCIN".
  - To demonstrate this capability, they developed "EMYCIN", the first shell.
  - The developers of MYCIN believed that the programming approaches they used in MYCIN could be applied to other domains.

What makes an ES feasiblefeasible ?

1 The need justifies cost.

2 The (human) expertise* is not available in all situations where it is needed.

3 The problem may be solved using symbolic reasoning techniques.

4 The domain is well structured and does not require common sense reasoning.

5 The problem may not be (better) solved using other (traditional) computing methods.

6 Cooperative and articulate experts exist. 7 The problem is of proper size and scope. This is relative to resources and evolving technology.