# EXPERT SYSTEMS TECHNOLOGY
## A GUIDE

L. JOHNSON
E.T. KERAVNOU

# EXPERT SYSTEMS TECHNOLOGY

## A GUIDE

L. JOHNSON
E.T. KERAVNOU

ABACUS PRESS
Tunbridge Wells ■ London
■ Dover, New Hampshire, USA ■

First Published in 1985 by

ABACUS PRESS

51 Washington Street, Dover, New Hampshire 03820, USA

Abacus House, Speldhurst Road, Tunbridge Wells, Kent TN4 OHU, UK

NOTE: Abacus Press US office now at:
P.B.S, PO BOX 643, Cambridge, MA 02139, USA.



British Library Cataloguing in Publication Data

Johnson, L.
Expert systems technology : a guide.——
(Information technology & systems series)
1. Expert systems (Computer science)
I. Title II. Keravnou, E.T. III. Series
001.64 QA76.9.E96

ISBN 0-85626-446-6

Printed in Great Britain by Nene Litho and bound by Woolnough Bookbinding both of Wellingborough, Northants.

# PREFACE

There is considerable interest in Advanced Information Technology and this book may be a useful guide to all those involved, be they computer scientist, electrical engineer, cognitive scientist or those orientated towards applications. Any of these may feel in need of more deep descriptions of expert systems than are usually available but yet do not need these descriptions entangled with implementation details and the idiosyncratic terminology of the systems designers. These same needs are felt by the neophyte knowledge engineer and hence it is hoped that the committed will find this book a challenging manual. Finally, we have found that our own draft documents have proved useful as a reference work and the book may also serve this purpose.

In this book we describe several expert systems at a conceptual level. We bring out their novel architectural features which gave rise to our selecting them for inclusion in the book. Each chapter has an overview, a description of the system's static structure and a description of its dynamics. We have not used the idiom of the original system constructors but have imposed as much commonality as possible. In this way we hope that the concepts employed, either implicitly or explicitly, by them are more readily discerned and comparisons are easier to make.

We would like to thank the many researchers and others involved in the design of the expert systems, particularly those connected to the systems selected for inclusion in this book, for making the field one of interest and enduring importance. It is worth the effort in getting to know the fruits of their labour. We would like to express our gratitude to the members of the Man-Computer Studies Group (here at Brunel University) who encouraged us to believe that the style and content might be worth wider exposure than our own members.

Our thanks are due to Shirley Hatch who entered the draft copy into the SUPERBRAIN QD under the SPELLBINDER word processing package. We took up the task of editing this raw material into the form seen here. Neither Shirley nor SPELLBINDER are responsible for the errors and omissions that undoubtedly remain.

April 1984
Brunel University, London.

I J & E T K

# CONTENTS

# Chapter 1
# INTRODUCTION

The purpose of this introductory chapter is to orientate the reader and summarize the concepts we use throughout the book. Its purpose is not to introduce these concepts to those innocent of them. The relevant concepts fall under three of the subsections of the introduction. Section 1.1 is a discussion of interactive expert systems. Section 1.2 is a sketch of the main currently available knowledge representation schemes. Section 1.3 is an overview of the inferences employed in the discharge of the tasks typically undertaken by the systems which we have selected for discussion in subsequent chapters. The final section of the introduction previews these systems.

## 1.1 FOCUS OF THE BOOK

A **Knowledge-Based System** is a system which manipulates "knowledge" in order to perform a task or tasks. The knowledge in a knowledge-base, is highly structured symbolic data which represents a model of the relationships between data elements and the uses to be made of them. The performance of a knowledge-based system depends both on the quality of its factual knowledge (structure, completeness, validity, consistency, etc.) and the ways in which this knowledge is applied. (See Johnson and Hartley 1981; Addis and Johnson 1983.)

The field of **Expert Systems** is a subgroup of knowledge-based systems. A prerequisite for applying the expert systems technology to some knowledge domain is the existence of human experts for that domain. The field investigates methods and techniques for constructing Human-Computer Systems encorporating the domain specific knowledge. In a manner of speaking expert systems can thus be said to grasp fundamental domain principles (and weaker general methods), to solve complex problems and to interact intelligibly with the user (see Johnson and Keravnou 1983). And in this way, using the same manner of speaking, an expert system can be said to interpret, diagnose, predict, instruct, monitor, analyse, consult, plan or design.

### Expert Consultant Systems

Currently most expert systems engage in a dialogue with the user; the computer acting as a "consultant". The computer system suggests options on the basis of its knowledge and the symbolic data supplied by the user. A dialogue is usually terminated when a decision or a recommendation is reached. In this book we consider only **Expert Consultant Systems** which designate this large subgroup of expert systems. The term, "consultant", makes explicit the most important and distinguishing feature of these systems; i.e. the fact that they are **interactive** computer systems. Figure 1.1 gives the taxonomy of Computer Systems from the perspective of the book.

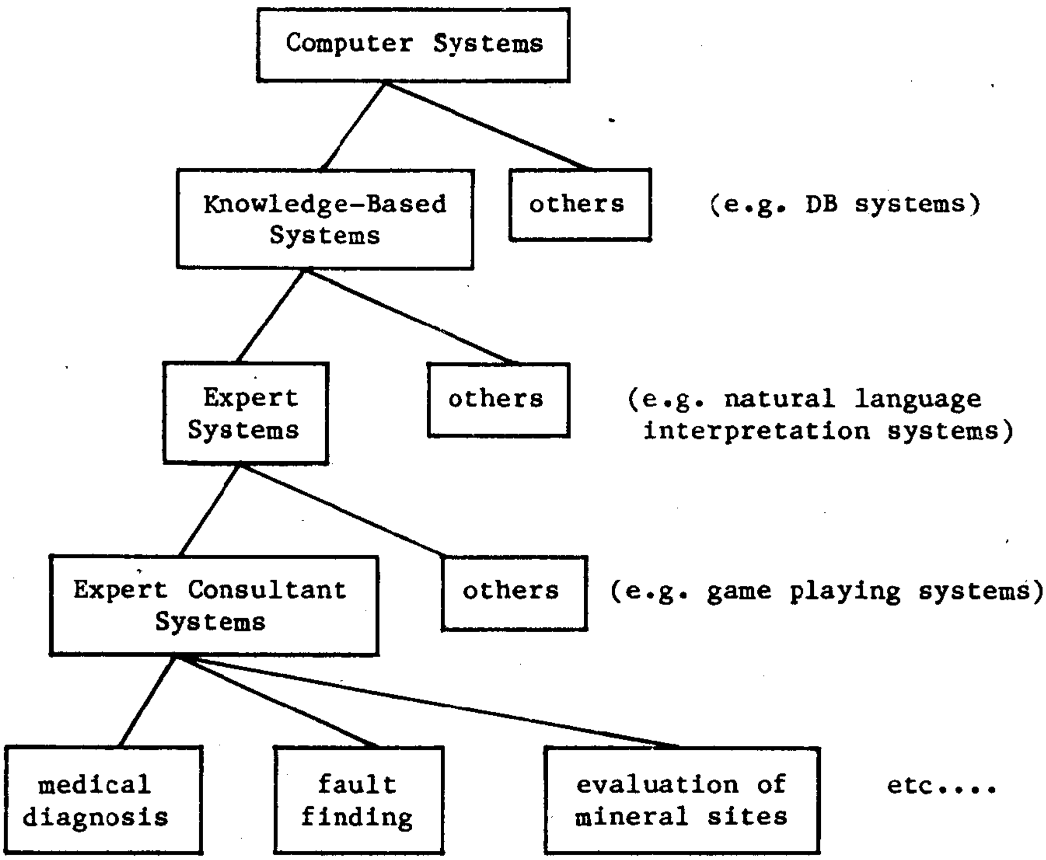


**Fig. 1.1**: Taxonomy of Computer Systems from the perspective of Knowledge-Based Systems.

The vital process of a consultant system is the process of moving from known items of information (or "seen" concepts) to unknown information (or "unseen" concepts). The user of an expert consultant system has "observed" some particular state of affairs, within the domain of the system's expertise and submits these observations to the system. Examples of such states of affairs are: a sick person, a faulty machine, an earth region (for evaluation of mineral sites), a malfunctioning business environment. Based on the observations the system makes inferences and suggests new routes of investigation which will yield high grade information. An interaction continues until the system finds the most likely explanation of the observations. Once a likely explanation is reached the system may go on to compile recommendations.

**Modes of Interaction**

There are two basic forms of interaction: 1) user initiated 2) computer initiated.

In the **user initiated** mode of interaction the system is restricted to respond to user requests only. The accuracy of any conclusion or recommendation reached by the system is constrained by the amount of input provided by the user. The decision to reach more accurate conclusions or recommendations lies entirely with the user. If the user is not satisfied with the current system output then he/she can formulate his/her next request based on a careful examination of the output.

In the **computer initiated** mode of interaction the user is restricted to respond to system requests only. The system starts by a set task (usually this will be a very general task like "Compile the best therapy regime for this patient") and in the light of this, the system requests input that will enable it to accomplish the task. The process is not determinate for the system makes inferences on input and requests further input in the light of these inferences.

In the context of expert consultant systems the desirable mode of interaction is a **mixed-initiative** one, whereby the initiative switches from the one basic mode to the other (both the user and the computer prompt each other).

## 1.2 KNOWLEDGE REPRESENTATION SCHEMES

Knowledge can be represented in schemes that lend themselves to implementation on a computer. The predominant schemes are: Predicate Calculus, Associative Networks, Frames and Rules. This section gives a very brief description of these schemes.

On the whole schemes are Janus faced; the one face looking towards human understanding the other looking towards computer implementation. Hence we will distinguish between Knowledge

Representation Schemes and Knowledge Representation Languages -- **schemes** look towards human understanding and the methodologies for knowledge elicitation etc., **languages** are implementations of schemes and hence questions of control and efficiencey arise. One finds that certain languages have been directly influenced by a scheme. A point that we wish to make is that Predicate Calculus can be thought of as both a scheme and as the basis for a language. When conceived as a language Predicate Calculus may be used, as indeed any language can, to implement any of the schemes. The argument for treating Logic Programming as a basis for work in expert systems, is one to do with having a high level language with a clear semantics. One can accept, or reject, these arguments independently of accepting, or rejecting, that Predicate Calculus is the only scheme in which to consider knowledge representation.

To aid us to make our description of the schemes we use a very trivial body of factual knowledge which we represent in them.

> "A person suffering from a disease belonging to the category of diseases, A, exhibits symptoms X and Y. B, C, and D are such diseases. A person suffering from B in addition to symptoms X and Y, exhibits symptom Z as well".

The conceptual structure of the above knowledge is very simple: we have a taxonomy of diseases and each node of this taxonomy is empirically associated with some symptoms; a category of diseases is associated with those symptoms shared by all its specialisations (in other words a symptom associated with some element of the taxonomy is inherited by all its descendant elements (property inheritance)).

### 1.2.1 Predicate Calculus

The (first order) Predicate Calculus can be used as a basis of a knowledge representation scheme (see Nilsson (1980), chps 4 and 5). Below we express our body of knowledge in this idiom.

```
1.   For all diseases x and y,
          if x is the category of y, then
               for all symptoms s,
                    if s is associated with x, then
                         s is associated with y.
```

becomes

$$1.\quad \forall x\ \forall y((\mathrm{DISEASE}(x) \wedge \mathrm{DISEASE}(y) \wedge \mathrm{CATEGORY}(x,y)) \Rightarrow (\forall s((\mathrm{SYMPTOM}(s) \wedge \mathrm{ASSOCIATED\text{-}WITH}(s,x)) \Rightarrow \mathrm{ASSOCIATED\text{-}WITH}(s,y))))$$

```
2.   For all persons p and diseases x,
          if p suffers from x, then
               for all symptoms s,
                    if s is associated with x, then
                         p exhibits symptom s.
```

becomes

$$2.\quad \forall p \ \forall x((PERSON(p) \land DISEASE(x) \land SUFFERS\text{-}FROM(p,x)) \Rightarrow ( \forall s((SYMPTOM(s) \land ASSOCIATED\text{-}WITH(s,x)) \Rightarrow EXHIBITS(p,s))))$$

```
3., 4., 5., 6.    --- is a disease;     7,8,9   --- is a symptom;
10., 11., 12.     --- is the immediate category of --- ;
13., 14., 15.     --- is associated with --- ;
```

become

```
 3.  DISEASE(A)        4.  DISEASE(B)         5.  DISEASE(C)
 6.  DISEASE(D)        7.  SYMPTOM(X)         8.  SYMPTOM(Y)
 9.  SYMPTOM(Z)       10.  CATEGORY(A,B)     11.  CATEGORY(A,C)
12.  CATEGORY(A,D)              13.  ASSOCIATED-WITH(X,A)
14.  ASSOCIATED-WITH(Y,A)       15.  ASSOCIATED-WITH(Z,B)
```

These sentences, in the idiom, can be said to represent the knowledge. The specific taxonomy of diseases and the knowledge concerning the specific symptoms associated with these diseases are represented in sentences 3-15. In sentence 1, there is the explicit expression of "property inheritance" and this facilitates the representation of empirical associations between symptoms and diseases. We only need to explicitly express the associations that are particular to each disease — the rest can be deduced.

Within this scheme, knowledge is only represented as an unstructured sequence of "independent" sentences. For example, there is no specification that sentences 3-6 and 10-12 together represent "a taxonomy of diseases". However, this grouping reflects the application of the sentences in a problem solving context and there is a need to group them explicitly to form higher level knowledge units (see Brachman et al 1983).

## 1.2.2 Associative Networks

An **associative network** is a collection of nodes and arcs (see Findler 1979). Nodes represent terms (names of physical entities, situations, places, processes, events, n-ary relationships (n >= 2), etc.) and arcs represent binary relationships or arguments of n-ary relationships (n >= 2). Figure 1.2 gives the associative network representation of our body of knowledge.

KEY:

s: subset-of; ds: disjoint-subset-of;
de: distinct-element-of; aw: associated-with.

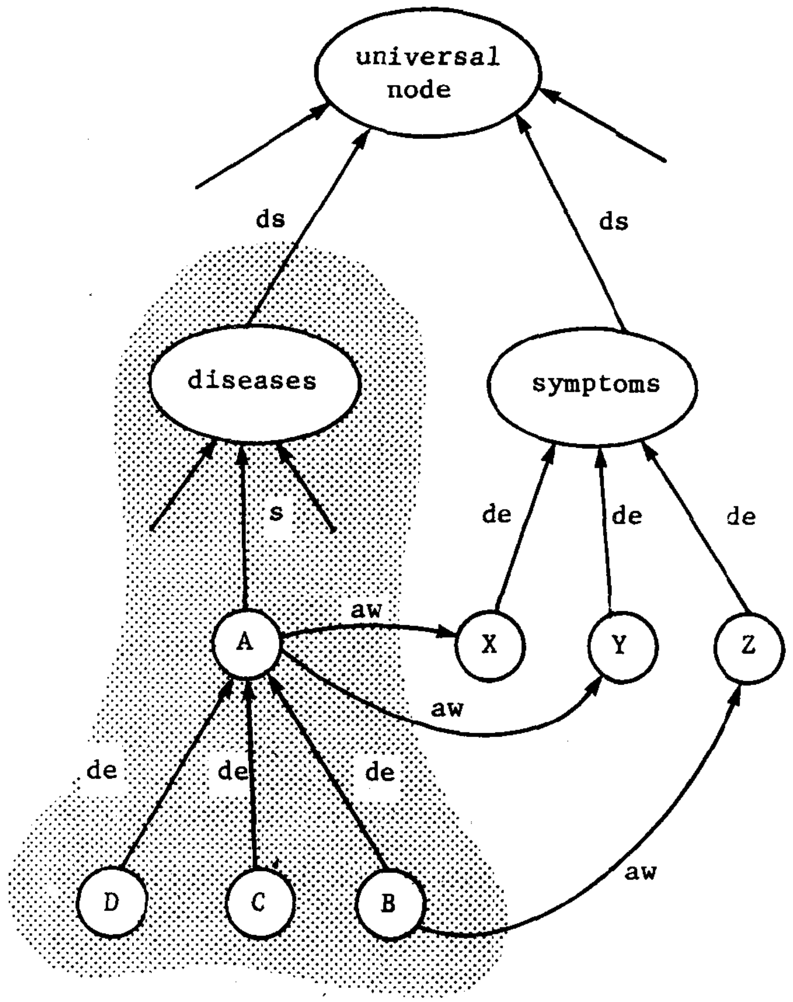


The associative networks scheme captures the given conceptual knowledge structure adequately; the taxonomy of diseases is represented through the "s", "ds" and "de" arcs and the empirical associations between symptoms and diseases are represented through the "aw" arcs.

**Fig. 1.2:** Illustrating the scheme of associative networks.

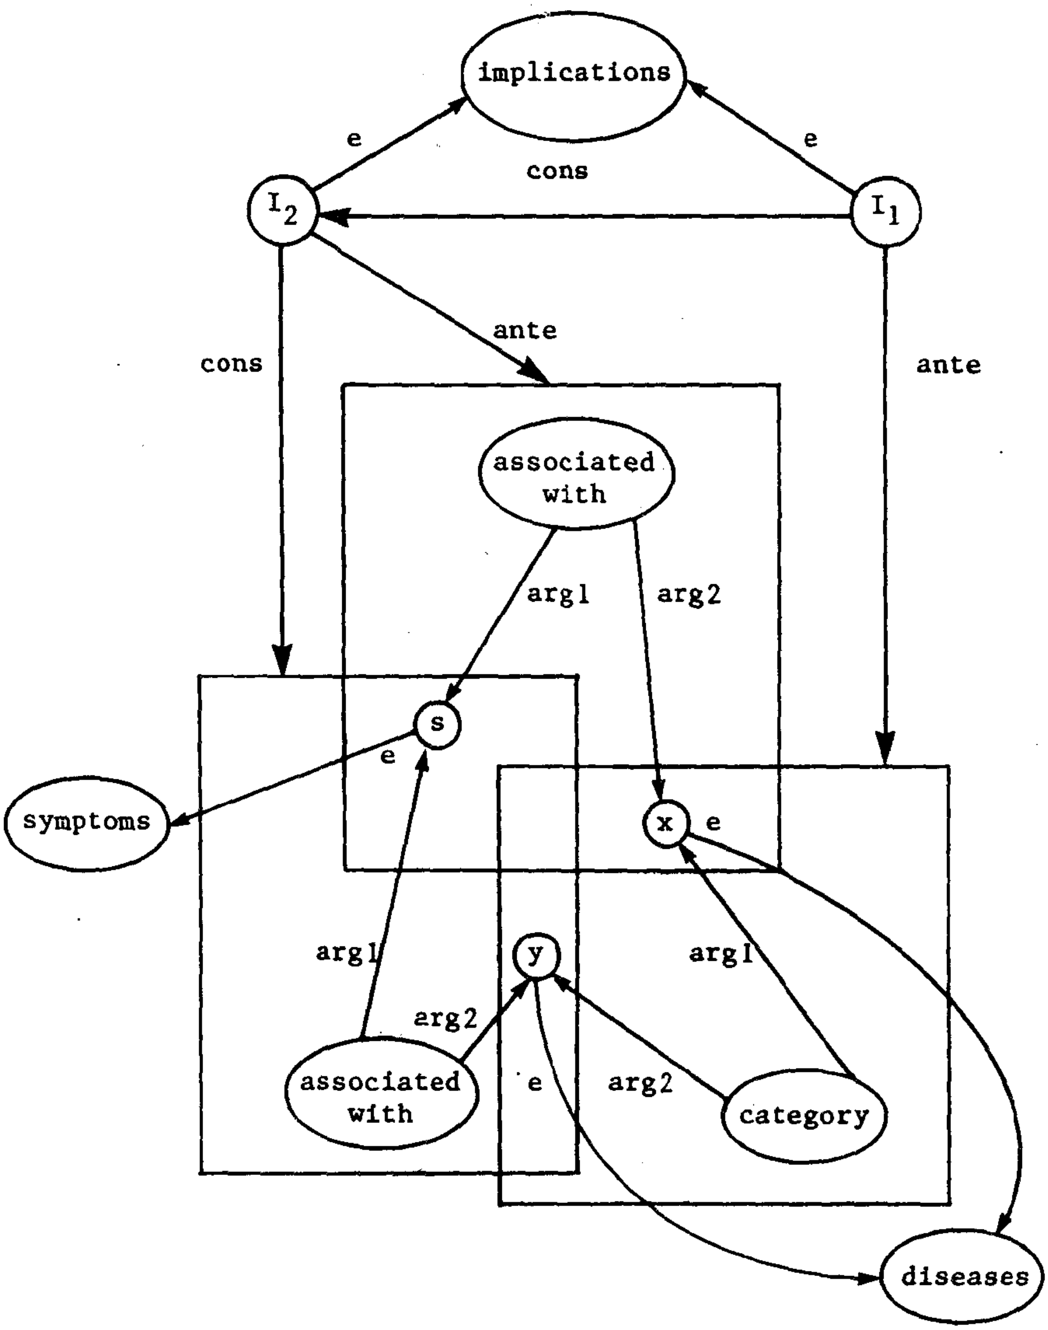


Partitioned associative network representation of the sentence "Any disease inherits the symptoms associated with its category". $I_2$ is an implication node with arcs connecting it to antecedent and consequent. Similarly $I_1$ is another implication node whose consequent is the implication represented by $I_2$.

Fig. 1.3

The network concept is extended by **partitioning** groups of nodes and arcs into a kind of super node (Hendrix, 1979). Figure 1.3 gives a partitioned associative network. This illustrates the point that the boxed areas represent units which are the supernodes of a super binary predicate.

### 1.2.3 Frames

A **frame** (Minsky 1975) is a structure consisting of a network of nodes and relations, used for representing a situation or topic stereotype. Attached to the frame is information about **how to use** the frame, what to expect to happen, and what other frames it might be appropriate to move to in certain circumstances. In short the scheme permits the co-existence of the factual knowledge and the reasoning knowledge that manipulates it. Some aspects of the frame are fixed; these are **slots** that are initially filled with "**default**" assignments containing information which holds unless new information displaces them. (See figure 1.4.)

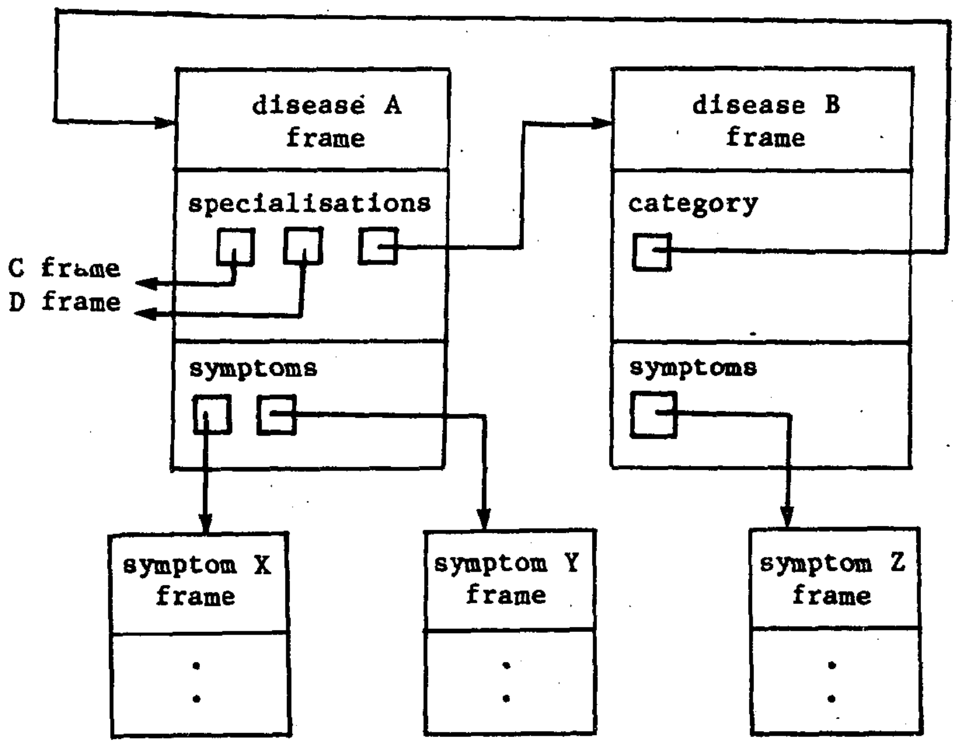


**Frames also adequately capture the given structure; the taxonomy of diseases is represented through the "specialisation" and "category" slots and the empirical associations between symptoms and diseases are represented through the "symptom" slots.**

**Fig. 1.4**

## 1.2.4 Production Rules

Consider the following relationships in the given body of knowledge:

patient suffers from disease A ⇛
patient exhibits symptoms X and Y

patient suffers from disease B ⇛
patient exhibits symptoms X, Y and Z

The knowledge is represented in rule form as follows:

$R_1$: **If** the patient exhibits symptom X **and**
the patient exhibits symptom Y
**then** the patient is likely to be suffering from disease A

$R_2$: **If** the patient exhibits symptom X **and**
the patient exhibits symptom Y **and**
the patient exhibits symptom Z
**then** the patient is likely to be suffering from disease B

Since the first two clauses of the antecedent of rule $R_2$ represent the consequent of rule $R_1$, $R_2$ could be reexpressed as:

$R_2^*$: **If** the patient suffers from A **and**
the patient exhibits symptom Z
**then** the patient is likely to be suffering from disease B

Other rules which do not constitute reversions of deductive relationships but could also be included in the knowledge-base are:

$R_3$: **If** the patient suffers from A **and**
the patient does not exhibit symptom Z
**then** the patient is not likely to be suffering from disease B

$R_4$: **If** the patient suffers from A **and**
the patient does not suffer from B
**then** the patient could be suffering from disease C

$R_5$: **If** the patient suffers from A **and**
the patient does not suffer from B
**then** the patient could be suffering from disease D

$R_6$: **If** the patient suffers from A **and**
the patient does not suffer from C **and**
the patient does not suffer from D
**then** the patient could be suffering from disease B

The scheme of rules, as used above, does not adequately capture the given conceptual knowledge structure. The taxonomy of diseases is not explicitly represented through the rules; it is implicitly represented by repeating the same conditions in the antecedents of rules (e.g. the condition "the patient suffers from A" is present in the antecedent of every rule) and by including clauses that restrict competing hypotheses (e.g. "B is present", "C is present", and "D is present", i.e. the specialisations of A, represent a set of mutually exclusive hypotheses).

## 1.3 INFERENCE

Here we report on the forms of inference employed in the reasoning components of the majority of the publicized expert systems. The actual mechanization for implementing these two forms of inference depends on the scheme/s used for representing the knowledge captured within the systems. These schemes we divide into two categories. One category is the rule-based schemes and the other is a catch-all category (associative nets and frames).

### 1.3.1 Three Stages of Inquiry

In this section we introduce Peirce's (1931) terminology and his notions of the changing form of inference process of inquiry. The inference forms are characterized as having certain requirements as a precondition for the validity of the inferences.

(1) We observe some puzzling phenomena and by abduction arrive at a certain hypothesis H.

(2) We deduce experimental consequences of H; experimental consequences are propositions of the form "If a procedure of a certain kind is carried out, a result of a certain kind will be observed".

(3) We carry out experiments from ($E_1$ .... $E_n$) (finite). There are two cases:

- (i) Suppose we find that, say, $E_3$ is false. Then we infer that H as it stands is false, though we may be able to give a modified version $H^*$ from which $E_3$ does not follow.
- (ii) Suppose ($E_1$ .... $E_n$) are all true. Then we conclude by **induction** that either H, or some modified version of H, is the true explanation of the phenomenon.

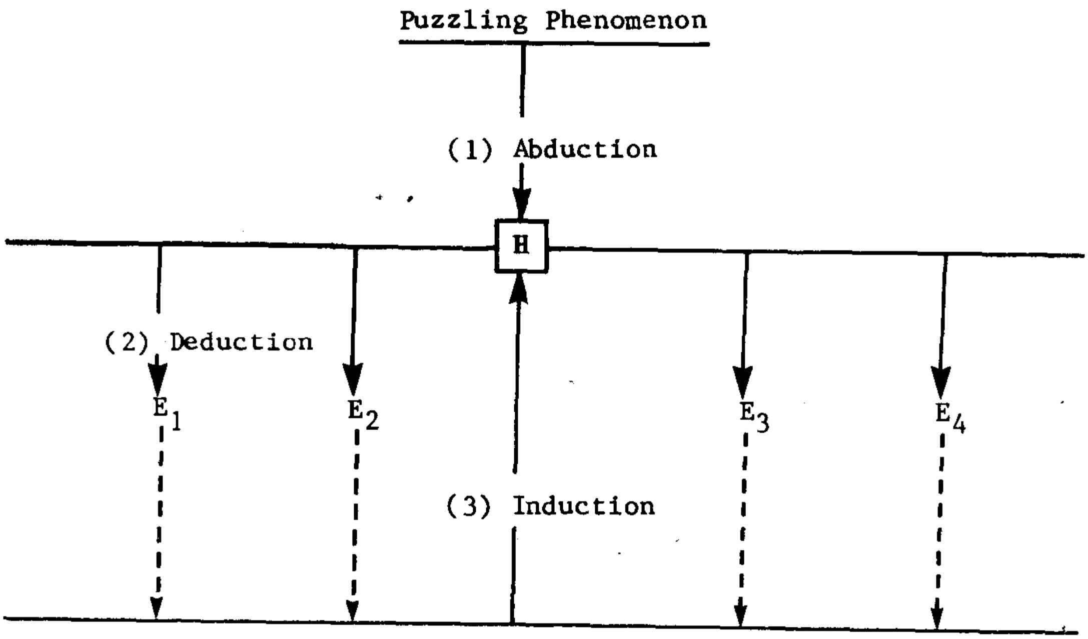


**Requirements for abduction**

(i) The hypothesis must be such that some experimental consequences can be deduced from it ("pragmatic requirement").

(ii) The hypothesis must explain the puzzling phenomenon, hence it must be deducible from the hypothesis that such a phenomenon would occur.

(iii) A hypothesis which, if false, could be easily falsified is to be preferred.

(iv) An initially plausible hypothesis is to be preferred.

**Requirements for deduction**

(i) ($E_1$ .... $E_n$) must follow by necessity from H.

**Requirements for induction**

(i) "Fair sampling" requirements: these relate to the choice from all the possible experiments of those to be actually carried out.

(ii) "Predesignation": we must decide what hypothesis we are testing **before** making our observations.

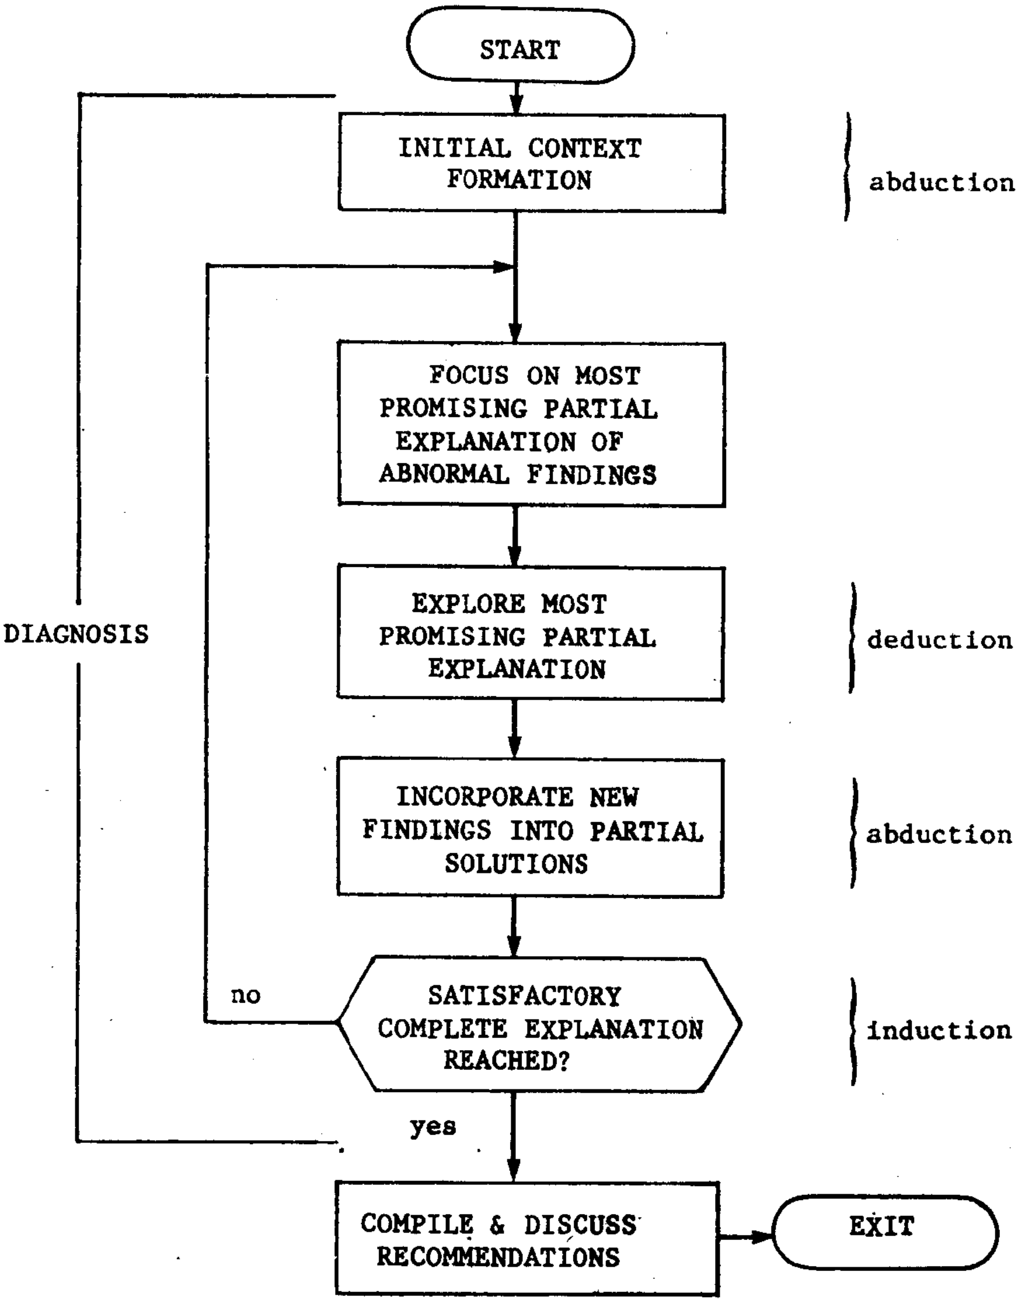
START
INITIAL CONTEXT FORMATION
abduction
FOCUS ON MOST PROMISING PARTIAL EXPLANATION OF ABNORMAL FINDINGS
EXPLORE MOST PROMISING PARTIAL EXPLANATION
deduction
INCORPORATE NEW FINDINGS INTO PARTIAL SOLUTIONS
abduction
SATISFACTORY COMPLETE EXPLANATION REACHED?
induction
no
yes
DIAGNOSIS
COMPILE & DISCUSS RECOMMENDATIONS
EXIT

Fig. 1.5