Towards a generalized model of diagnostic behaviour

E T Keravnou* and L Johnson†

A model of diagnostic behaviour shared by most diagnostic tasks is discussed. The model is, by-and-large, presented at a conceptual (knowledge) level. Its expression at more concrete levels is task dependent. The model's use can be seen as being twofold: first as an aid in the construction of more adequate shells for diagnostic tasks; and second, and more importantly, as the initial (albeit crude) interpretation model held by the knowledge engineer about to embark on the cumbersome task of analysing a particular expert diagnostician's behaviour.

Keywords: diagnostic model, expert system, findings base, hypotheses base

Two approaches to the use of computers in diagnosis can be identified: first, one can use the computer's speed, arithmetic capabilities and storage capacity to solve problems in ways that people cannot; and second, the computer's plasticity can be used to model human problem solving. Combining these approaches offers the prospect of extending the range of diagnostic procedures. The diagnostic model discussed in this paper is intended to be a framework for combining the two fundamental approaches. We believe that only a framework which is a global model of human diagnostic reasoning can form an adequate basis for such a combination. Our global model contains a fundamental division between what Chandrasekaran[1] has called the *generic task*, and the data handling capacities of human beings. In the model, intelligent data handling is performed through a multi-faceted procedure called *Decide-Status:* it is through Decide-Status that access to mathematical models and large amounts of data can be provided in ways that are congruent with the demands of the generic task.

The diagnostic model, however, has a second very important function: it forms part of a well founded and coherent approach to building knowledge-based systems. When embarking upon the analysis of some knowledge-intensive task it is necessary to do so with a general and complete model of the generic task. The model focuses the analysis of the domain more completely, effectively and efficiently than otherwise. If the model is not misdirective at crucial points, feedback from the initial analysis enables the refinement of the model with respect to the specific task, and guides the subsequent stages of analysis (see Johnson and Johnson[2] and Breuker and Wielinga[3]; also, Part One of Reference 4 contains much relevant material). This paper is therefore an indication of a framework in which the major conceptual, technical and methodological issues in the design of knowledge-based systems can be addressed.

*Department of Computer Science, University College London, Gower Street, London WC1E 6BT, UK
†Department of Computer Science, Brunel, The University of West London, Uxbridge, Middlesex UB8 3PH, UK

Paper received 5 July 1988. Accepted 28 March 1989

DIAGNOSTIC CONCEPTS

Diagnosis is probably the field where most of the empirical expert systems work has been undertaken. Medical diagnostic systems encompass a substantial proportion of the pioneering attempts in articulating expertise (e.g. Shortliffe *et al.*[5]; Young[6]). Malfunctioning devices other than the human body have also attracted attention in recent years[7].

In a diagnostic task the important concepts are *findings* and *hypotheses*. Findings (or data) on a problem case is the known information on that case. Findings can refer to the case's history, can express universal truths about any such case instance; or can be statements of malfunctioning. It is the statements of malfunctioning that constitute the direct evidence, while other findings constitute contextual or circumstantial evidence. Each finding is associated with a temporal aspect, e.g. *past, recently, currently*. Hypotheses are statements of possible explanations of the malfunctioning: an established hypothesis becomes a finding. It may be that the only way to establish a particular hypothesis is by establishing the effectiveness of an associated rectification action (treatment). A treatment is effective only if it turns the system's behaviour back to normal (assuming no side-effects from the treatment). Thus, treatment is another very important concept[8]. This concept, however, is not of central concern here.

0950–7051/89/030165–13 $03.00

A general model of diagnostic behaviour can only be adequate if it is a model at a sufficiently high level of abstraction, induced from the analysis of a sufficient sample of the relevant task population. We have studied the results of researchers analysing diagnostic behaviour[9–14], made an analysis of diagnostic systems[15], and had experience in building an actual diagnostic system[16]. Our abstraction and integration involved the highlighting of divergences between the practices adopted by diagnosticians and the dynamics of existing diagnostic systems.

REPRESENTING REASONING KNOWLEDGE

The maintainability, explanatory power, flexibility and adaptability of an expert system all depend on the quality of the knowledge representation structures for the incorporated expertise. The completeness of the system's knowledge depends on the expressiveness of the chosen knowledge structures. Domain knowledge is of two types: factual knowledge (i.e. knowledge about domain entities and their interrelationships), and reasoning knowledge (i.e. knowledge of how to use the factual knowledge to generate knowledgeable behaviour). Generally, factual knowledge is domain-specific, while reasoning knowledge can apply to more than one domain, e.g. the same diagnostic principle applies to more than one diagnostic domain. A knowledge structure for the reasoning knowledge is discussed here, which we view as a generic structure.

Task analysis

A generic task (see Chandrasekaran[1]) we decompose into a set of simpler possible tasks. These in turn are decomposed until directly executable tasks are reached. Non-decomposable tasks are those that manipulate factual knowledge and case-specific information. In a diagnostic domain it is these primitive tasks that progress the construction of the diagnostic picture (see below). The tree representing the decomposition of a task into subtasks is called the *task analysis tree*. A task analysis tree makes the functionality of the overall task more explicit. Our analysis of knowledge-intensive diagnostic tasks resulted in a greater articulation of this notion of task analysis. This is the framework in which we represent reasoning knowledge.

Analysing knowledge-intensive tasks

A knowledge-intensive task is decomposed into a set of subtasks that collectively constitute the means available to complete the task (e.g. the Decide-Status task mentioned in the section below on reasoning within a findings base is decomposed into the four cooperating subtasks, *viz* matcher, inferencer, generalizer-restrictor and default-reasoner). To achieve a given global task, instantiations of the specific subtasks are repeatedly selected and executed until a termination condition for the task is satisfied. The selection of a subtask instantiation is geared by case-specific information on which the selection conditions for the subtasks are applied. The selection conditions for a subtask form its *logical basis* in the context of the particular task. Thus, a strategic principle is encompassed in the link:

$$\text{task} \xrightarrow{\text{logical basis}} \text{subtask}$$

It might be concluded from this that the rule scheme is the most appropriate for representing strategic knowledge. However, we would resist this conclusion. The essence of a strategic principle is the underlying logical basis. Unless the logical bases are, therefore, adequately explicated (and subsequently reasoned with) in the chosen representational structure, this structure will not be suitable. We found that rules are not suitable for representing strategic knowledge: rules are an inflexible representation structure, in the sense that they leave many assumptions governing their application as implicit[17], thus preventing their violation when these assumptions cease to hold. A logical basis for a subtask (always in the context of some task) consists of the conditions that must be true for the subtask to be selected (*enabling condition*) assuming that another set of conditions are not true (*disabling condition*). In addition, disabling assumptions can be violated through a set of *relaxation conditions*. Making the components of a logical basis (enabling, disabling and relaxation conditions) explicit is of paramount importance in knowledge-intensive processing, where case-specific information is incomplete, imprecise, and continually changing. A condition that is considered to be valid

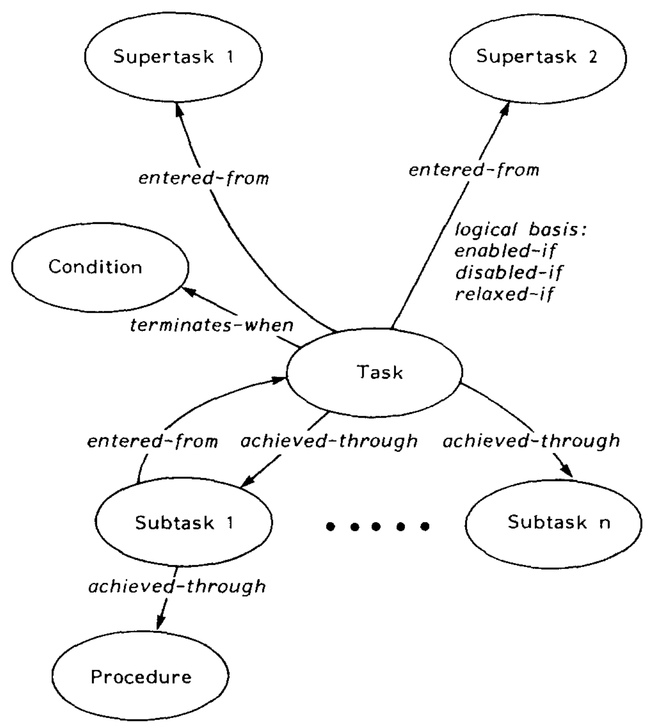

Figure 1. Analysing reasoning knowledge. A task is linked to its subtasks via achieved-through *links. A primitive task is achieved-through the procedure implementing it. Non-primitive tasks are linked to their* termination conditions. *A task is linked to its invocation contexts (other tasks) via* entered-from *links which are qualified with the respective logical bases. A task can be invoked from numerous points. Some tasks are entry (global) tasks. If there is only one entry task then we have a tangled hierarchy. Entry tasks form the roots of strategic explanation trees (see discussion below)*

may cease to be valid, or an unknown disabling condition may become true, etc. In such a dynamic environment new strategic choices are frequently made. Figure 1 gives a network representation, in associative form, of this analysis.

Representation structure

Tasks are represented declaratively in terms of frames (the format of which is given in Reference 16). As the reasoning knowledge is represented declaratively, an interpreter embodying the above semantics is required. The *task interpreter* selects, instantiates and executes tasks until the global task is achieved. In addition, the task instantiations are linked into a strategic explanation tree that constitutes the means for the generation of strategic explanations.

The justification of diagnostic hypotheses through the tracing of their status transitions is discussed below. Such transitions are directly attributed to primitive tasks that form leaf nodes in a strategic tree. By traversing the tree upwards from these nodes, the entire rationale behind a status transition can be obtained. (For further details on the analysis, representation and interpretation of strategic knowledge, see Reference 16.)

ABSTRACT ARCHITECTURE

The relationships within the domain of findings and the relationships within the domain of hypotheses are far more numerous and complex than their interrelationships. However, during the initial stages of diagnostic activity, the shifts of reasoning from the findings base to the diagnostic (hypotheses) base are crucial, as they establish the initial context for the problem-solving activity. Clancey[18] refers to such initial reasoning jumps as 'heuristic inferences' — a non-hierarchical and non-definitional connection between concepts of distinct classes. Although such inferences are very nearly categorical, as their name implies they are not infallible. Examples are the *constrictor* associations in CADUCEUS[19] and the *trigger* associations in PIP[20], PUFF[21] and NEOMYCIN[22].

The model of diagnostic behaviour proposed here has an abstract architecture that separates the two bodies of factual knowledge along the lines of conceptual cleavage in the diagnostic concepts (see Figure 2). At this level of analysis the model resembles that proposed by Chandrasekaran and Mittal[23]. However, at more refined levels of analysis and on more concrete issues, the two models differ significantly.

The *findings reasoner* operates on the conceptual organization of the general findings knowledge to make intelligent inferences on the available case-specific information. Such inferencing could be of a commonsense nature, deducing that a child is a non-smoker, for example, or it could be based on specialist knowledge, e.g. deriving suitable qualitative abstractions from the quantitative results of laboratory tests entails knowledge of the units of measurement used. The *diagnostician* generates and refines hypotheses by operating on the conceptual organization of the general hypotheses knowledge. A critical function of the diagnostician is in assessing what additional information would be needed for the diagnostic activity to progress. The *diagnostic*

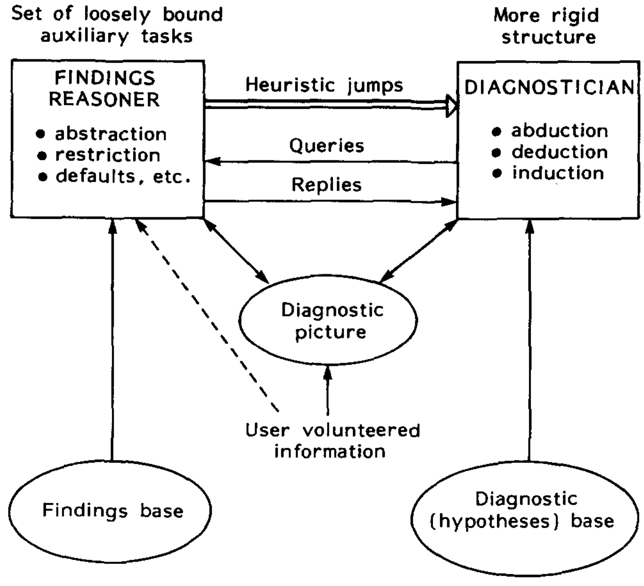

Figure 2. Abstract architecture overview, representing the separation of the findings knowledge and the diagnostic knowledge. The reasoning processes within each of these components and the communication protocols between them are not represented

picture is the data structure that holds the user-supplied case-specific information and the results of the operations of the findings reasoner and diagnostician. The diagnostic picture for a particular domain could be of considerable structural complexity. As this data structure essentially holds the instantiations of the general knowledge that apply in the particular case, its complexity would depend on the complexity of the organization of the knowledge bases. The communication between the findings reasoner and the diagnostician could be restricted to message passing through the diagnostic picture, but this is an implementation rather than a conceptual issue.

The separation of the data handling function from the main diagnostic function is seen as the first step towards the explication of the total functionality of a diagnostic task. An important feature of the proposed architecture is the distribution of the overall reasoning among a number of knowledge-intensive tasks (see above). The top level distribution is between the diagnostic and data handling tasks. The findings reasoner comprises a set of loosely bound auxiliary tasks (see below), while the tasks comprising the diagnostician are rigidly bound together into a cooperating whole. The distribution of reasoning as such enhances the modularity of the overall system, and by implication, its extensibility. Furthermore, and importantly, this allows the construction of a knowledge engineering environment with acceptable skeleton system tools.

REASONING WITHIN A HYPOTHESIS BASE

A diagnostic task can be characterized by its breadth (very narrow, e.g. dealing with one disease; or very broad, e.g. dealing with internal medicine) and its depth (very shallow, e.g. reasons from the clinical knowledge; or very deep, e.g. reasons from the pathophysiological knowledge). Although at a high level of abstraction every diagnostic process can be captured within a uniform framework (see below and Reference 24), at less abstract levels individual processes can diverge significantly in the way these high level forms of inference are carried out. This is reflected in the wide spectrum of representation structures that are encountered in the knowledge bases of diagnostic systems.

There are three forms of diagnostic inference[25]: *abduction*, the process of generating hypotheses; *deduction*, the process of testing hypotheses; and *induction*, the process of evaluating hypotheses. Figure 3 qualifies diagnostic steps from this perspective. A hypothesis could be a simple assertion, a collection of property value associations, possibly with certainty measures, an associative network of simpler hypotheses, or a combination of these. The same hypothesis could be described at different levels of abstraction and within the same system; hypotheses could be of different types, e.g. in a medical system some hypotheses are primary etiologies, others are syndromes or intermediate states, etc. By definition, hypotheses can not be directly established through observation, and thus they need to be inferred or gradually pieced together from observations and inferences.

Diagnostic inquiries

Abductive diagnostic steps

Abductive diagnostic steps generate hypotheses of two types: non-contextual and contextual:

non-contextual: ⟨findings⟩ –⟩⟨hypothesis⟩

contextual: ⟨hypothesis⟩ * ⟨findings⟩ –⟩⟨hypothesis⟩

contextual: ⟨hypothesis⟩ –⟩⟨hypothesis⟩

Non-contextual abductive steps are Clancey's (simplest) heuristic jumps. They function to set up the initial context for the problem-solving activity. In the proposed model, initial context formation is a collaborative activity between the findings reasoner and the diagnostician (see below). The hypotheses thus generated are usually few and general. Non-contextual abductive steps, therefore, function to significantly constrain the range of possible explanations of the malfunctioning behaviour.

A contextual abductive step occurs when a hypothesis is generated in the context defined by another hypothesis (the two hypotheses could be of different types, thus having another case of heuristic jumps). Such steps are qualified either as traversal of a taxonomy or as lateral shifts. In the former case, the generated hypothesis is either a refinement, or a generalization, to the initial given hypothesis, and in the latter case, an opponent, or a complement, to this hypothesis. Generalization shifts broaden the range of currently entertained explanations, and refinement shifts constrain this range even further. Thus, taxonomic shifts of the former type are more likely to occur at the initial stages of the process, and of the latter type at subsequent stages. Opposing hypotheses share expectations, giving rise to situations where the actual presence of the one may be confused for the presence of the

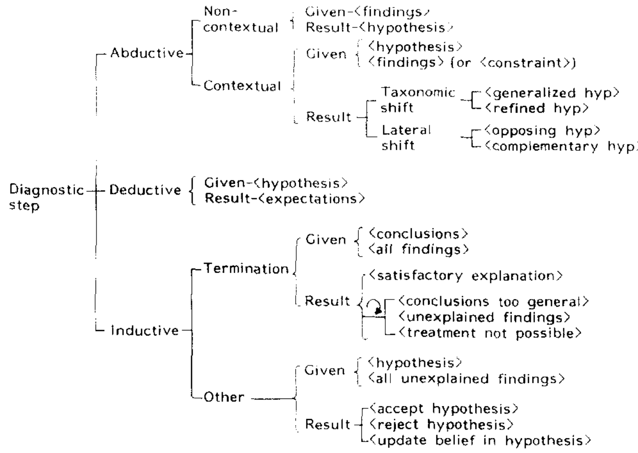

Figure 3. Qualifying diagnostic steps. { : joint selection; [: exclusive selection; ⟨ : selection may be repeated

other. The given findings in such abductive steps are not expected on the currently entertained hypothesis, but on an opponent of it. Complementary hypotheses are related via a causality, or a complication-of, or some other sort of an association relationship. Thus, the given findings point to the direction of the particular complementary hypothesis, the objective being that the two hypotheses together will give a more accurate and complete picture of what is causing the malfunctioning than either one would do on its own. Groups of complementary hypotheses constitute complex or composite hypotheses. Contextual abductive steps function to piece together the components of more complex and more diagnostically complete (global) hypotheses. Contextual abductive steps do not necessarily involve findings: a hypothesis can directly point to a complementary hypothesis.

Contextual abductive diagnostic steps can take place within the scope of inductive diagnostic steps (see below). For example, consider an inductive step concerned with matching a hypothesis' expectations against the reported findings to decide whether the hypothesis can be concluded. If an expectation violation is thus detected, which points in the direction of an opposing hypothesis, a contextual abductive step has occurred.

Deductive diagnostic steps

Given a hypothesis, deductive diagnostic steps decide which findings follow by necessity from the hypothesis. Thus, actions for testing the unobserved expectations, on the hypothesis, can be requested. In the proposed model the acquisition of new information is again a collaborative activity between the diagnostician and the findings reasoner (see below).

Inductive diagnostic steps

An inductive inference is from hypotheses to overall best explanation, or termination of the diagnostic process. Usually a satisfactory explanation is when the critical abnormal findings are accounted for by the concluded hypotheses and, more importantly, when the explanation permits the undertaking of 'treatment' procedures. It is only after treatment procedures have been undertaken that we can really test the correctness of the induced explanation of the problem. For example, in the case of hardware devices, a satisfactory explanation of a problem would indicate a set of replaceable/repairable units as being faulty. If after replacing/repairing the particular units the problem still exists, then the given explanation was incorrect. Of course, if each contending hypothesis gives rise to the same treatment there is no justification in trying to resolve these hypotheses any further.

All three forms of inference should be present in a diagnostic system. We found that evaluating existing systems from this abstract but complete model of diagnostic inference gives rise to many useful insights. It seems that of the three forms of inference, deduction is the one understood best, and induction is the one understood least. Although diagnostic systems have been built that only employ deduction (e.g. MYCIN[26]), neither abduction nor induction can suffice on their own. Knowledge engineers agree that inductive processes are the ones most difficult to analyse. How expert diagnosticians evaluate hypotheses, and how they decide when to stop are still largely unanswered questions. The PIP and INTERNIST-I[27] systems have very *ad hoc* termination criteria, and the designers of ABEL[28] accept this for their system.

Focusing and information acquisition

The focusing heuristics are inherently inductive, and the information acquisition heuristics are inherently deductive in nature. However, focusing and information acquisition are intimately related (and critical) aspects of a diagnostic process. The former decide on the part of the hypothesis space on which to focus next, and the latter decide on the information which is required next. Figure 4 gives a complete abstraction of a diagnostic process from this perspective. To be complete, the analysis of a particular diagnostic task must therefore reveal the heuristics for the initial context formation, the focusing heuristics and the information acquisition heuristics. MYCIN has a simple parameter-value language. The system reasons in a backward-chaining fashion, as such goals are always more general than the premises. The whole notion of hypotheses is alien to MYCIN (see Clancey[29]).

At the heart of a diagnostic process' focusing aspect lie the means for evaluating the promise of hypotheses. As mentioned above, eliciting these means for a particular domain could prove to be a very challenging undertaking. Existing diagnostic systems have been criticized for erratic changes in focus which are not akin to the corresponding human reasoning. As focusing and information acquisition are so intimately related, problems with the one aspect in a particular system may in fact be related to errors in the other aspect, i.e. focusing errors can propagate themselves as information acquisition problems, and *vice versa*. This fact should be borne in mind during the evaluation stages of a diagnostic system. For example, MYCIN's information acquisition problems are due to the lack of focusing, and PIP's focusing problems are to a large extent dependent on the system's shortcomings in its information acquisition.

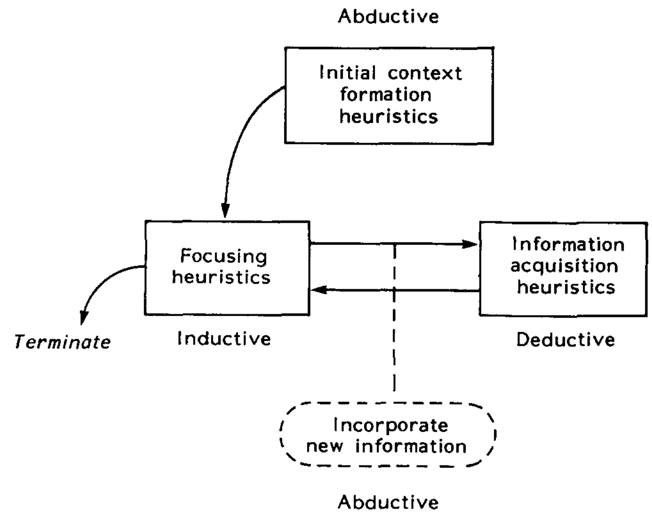

Figure 4. High level abstraction of diagnostic process, showing the intimate link between focusing and information acquisition

REASONING WITHIN A FINDINGS BASE

The reasoning within a findings base has a central process concerned with deciding whether some item of information is true, false or unknown. This reasoning is involved in revealing conflicts among case-specific findings, deciding whether an abnormal finding can be trivially explained and hence ignored, whether it is sensible, in the light of the things known, to ask a user a question, or to request an action for eliciting additional information. We call the central process the Decide-Status process[30]. This process uses a number of structures defined over the findings. Findings are grouped into hierarchies which are associated and labelled in various ways. The basic organization of findings is shared by most knowledge domains: it is usual to classify (finding) concepts; to associate attributes with findings; and to identify property dependencies among findings. The same findings reasoner can be shared by a number of diagnosticians. The procedures operating on the findings base are not, in fact, restricted to the findings in a diagnostic domain.

User-volunteered information is processed by the findings reasoner to:

- locate and resolve conflicts;
- remove red herrings;
- and to locate potential triggering information.

Human competent diagnosticians are very quick to spot conflicts in the user-supplied evidence. Exhaustively locating conflicts among a set of findings could be a computationally costly operation. A procedure for this would be to try to decide the status of each finding from a consistent subset of the findings, initialized to the empty set. If the finding is true (in which case it is not a very informative finding) or unknown, then it is consistent with the 'axiom' set, and thus becomes a member of it. If the finding is false then we have an inconsistent set, and the user needs to be queried to resolve the conflict. The above procedure, although it does not display the parallelism inherent in the corresponding human reasoning process, would locate data inconsistencies, if they exist.

A potential red herring is a finding that can both be either trivially explained or attributed to some serious malfunctioning. The findings base must indicate such findings and associate them with their trivial explanations, themselves findings. To locate potential red herrings from among the user-supplied information, the findings reasoner may have to perform suitable abstractions or refinements on the data. For example, one finding could be 'a blood pressure of 15 was read'. The findings reasoner has to translate that to 'a high blood pressure was observed', a potential red herring since it may have the trivial explanation of 'the person concerned was under significant stress'. Once a red herring is located, the findings reasoner has to ensure that none of its trivial explanations holds for the particular case.

Locating potential trigger findings is a crucial process in a diagnostic task, as such information tends to significantly constrain the diagnostician's focus of attention. Trigger links associate groups of findings with, usually, classes of diagnostic hypotheses which they either implicate or rule out. It is these trigger links that enable 'heuristic jumps'. The findings reasoner does not perform these jumps: what it does do, though, is to locate from among the user information which triggers could be instantiated. Here again it has to perform suitable abstractions over the data. The competence with which this and the previous function are performed depends on how well the findings knowledge is 'practically indexed' (see Clancey[18]) to enable the efficient and effective derivation of such abstractions.

The results of the operations of the findings reasoner are entered into a global database or blackboard which is referred to here as the case-specific diagnostic picture. It is up to the diagnostician to decide on the basis of its knowledge and the diagnostic picture which potential triggering information is worth pursuing further, or whether a 'red herring' is truly a red herring.

Deciding the truth status of findings

The central Decide-Status process is carried out by four cooperating tasks: a *matcher* that reasons from attribute types and other constraints (where necessary employing computational procedures) to translate sentences to contextually equivalent expressions; an *inferencer* that reasons from the finding dependencies; a *generalizer/restrictor* that reasons from the hierarchies of finding concepts and employs the principle of non-monotonicity; and a *default-reasoner* that uses assumptions to reach useful conclusions when every other route has failed. Instances of these tasks implicate new findings that entail the queried finding, and which subsequently themselves become queried. This dynamically generated network of findings is referred to as the *implications network*. When the truth status of a finding in the implications network is decided, this information is propagated along the network. The implications network is a component structure of the diagnostic picture. It forms the blackboard for the Decide-Status auxiliary task, while the diagnostic picture is the blackboard for the overall diagnostic system. Nesting of blackboards is common in complex, non-deterministic knowledge-intensive processing. (For a fuller description of these procedures see Reference 31.)

The case-specific findings are partitioned according to their temporal aspects. The Decide-Status process must be able to reason temporally. For example, suppose that the case-specific information indicates that the patient underwent heart surgery three months ago, and that he developed respiratory problems a week after surgery. If the queried finding is whether the patient had respiratory problems in the near past, Decide-Status should be able to say 'yes'; to reach this decision it must 'understand' that three months ago is near past, and that near past plus a week is still near past. Processes capable of temporal reasoning are said to be *deep* processes, as opposed to shallow processes that have no such capabilities. Thus, to be effective Decide-Status needs to incorporate deep reasoning. In contrast, a particular diagnostician's reasoning can be entirely shallow and still be effective.

Constructing a plan of information acquisition actions

The acquisition of new information is a collaborative activity between the diagnostician and the findings reasoner. The diagnostician performs the high level information acquisition tasks, i.e. decides which additional information is required. This information could be expressed as a list of findings, or it could be structured as a knowledge intensive decision tree (c.f. ABEL[28]). The findings reasoner performs the low level information acquisition tasks, i.e. decides which are the appropriate actions, for yielding the required information, in the particular case, and the order of their execution. In addition, the findings reasoner must be able, if necessary, to provide instructions to the user regarding the execution of actions. The findings reasoner would then inform the diagnostician of the truth status of the requested findings, or of the reasons they were unobtainable. (The information acquisition function in the ARBY diagnostic model[32] is distributed between a hypothesis generator and an interaction frames manager.)

Managing case-specific information

Often a consultation consists of a sequence of temporally distinct sessions with the diagnostician. For example, the patient is prescribed some treatment and asked to report in a fortnight's time, or the behaviour of a hardware device suffering from transient faults needs to be observed over a period of time. To be able to continue a diagnostic process at a later time, the contents of the diagnostic picture need to be saved. This information will be used to instantiate the diagnostic picture prior to the next session. (For the sake of simplicity the database that holds the snapshots of the diagnostic picture for the problem case is not shown in Figure 2.) We see the managing of such information as a function of the findings reasoner. In particular, the findings reasoner must ensure that the temporal aspects of findings are maintained. Further, as at the next session only the most critical information need be brought back to the diagnostician's attention, the findings reasoner must be able to select such findings and to provide summaries of groups of other findings; information left behind could be retrieved from secondary memory if and when it is required. Information audit is a very complex task indeed, with an associated database of significant organizational complexity. Although here this task is being categorized under the findings reasoner, its complexity for a particular diagnostic process could justify its incorporation as the third major component of the overall system.

Queries and replies

Once the initially volunteered information is processed by the findings reasoner, control is passed to the diagnostician. Subsequently, volunteered information will temporarily revert control back to the findings reasoner. Also, the diagnostician needs to be able to directly request the services of the findings reasoner. The most common query is 'What is the status of this finding with respect to the case?' Such a query has 'true', 'false' and 'unknown' as replies. Other queries are 'Is this constraint satisfied by that finding instance?', with 'yes', 'no', 'can be satisfied' and 'unknown' as replies; 'Is this finding a potential red herring?', with 'yes' and 'no' as replies; and 'What is the value of something?', with a value and 'unknown' as replies.

Findings base

The conceptual structure of a findings base can be accurately captured in terms of frames for the finding subjects[33]. Procedural attachments and/or declarative rules can encode the various finding (subject) interdependencies. The overall function of the findings reasoner is open-ended, and this is reflected in the more general nature of the content of the findings base. The findings base can include general commonsense knowledge, specialist domain knowledge, or even deep knowledge on domain concepts. Deep knowledge allows the prediction of new information from known information or the explication of unstated constraints on case-specific findings. Such diverse sources of knowledge would be needed by the central function of the findings reasoner, Decide-Status (see above). Thus, Decide-Status is an open-ended operation, and can always be extended with the addition of new structures defined over the findings knowledge.

Knowledge duplications

Diagnostic subtasks and auxiliary tasks can share knowledge. Tasks that share knowledge are strongly coupled by virtue of their close collaboration in achieving some function. Knowledge duplications, if present, should be made explicit to ensure that knowledge updates are properly propagated across the boundaries of the knowledge bases.

In summary, the functions of the findings reasoner are: to decide the truth status of a queried finding from the known findings (this is the central operation, and requires the ability to reason temporally); to locate and resolve conflicts in user supplied information; to locate potential triggers and red herrings; to devise a global plan for the acquisition of a set of information requirements; and to manage case-specific information so that a consultation can progress over a number of distinct sessions. The last function may need to be extended to a fully fledged information audit subsystem, capable of managing all historical information on a case, and not just the information relevant to a single consultation.

EVALUATING HYPOTHESES: QUANTITATIVE OR QUALITATIVE BELIEFS?

Here we show how the diagnostic model allows us to address an important issue in the design of expert systems. It is argued that the evaluation of hypotheses should be primarily qualitative. Quantitative methods, where used, should be secondary from an architectural point of view. Thus, our framework is a suitable basis for integrating qualitative reasoning and quantitative methods.

People do not use anything closely resembling mathematical methods in making diagnoses[34,35]. Our postulate is that hypotheses are ranked from a single direct (global) perspective, and a number (possibly zero) of contextual (local) perspectives, especially when one is dealing with a hypothesis space partitioned into interrelated subspaces. The global belief in a hypothesis indicates the strength of the belief in the hypothesis from the available direct (hard) evidence. A contextual belief in a hypothesis indicates the strength of the belief in the hypothesis in relation to other hypotheses (of the same or different types). The split in direct and contextual belief measures is exhibited by the CASNET system[36] and advocated by the designers of the ABEL[28] system.

Rule scheme lends itself to quantitative methods

It seems that quantitative methods have evolved from diagnostic systems that are primarily rule-based. The concept of a production rule is meant to capture any conditional association in a uniform way. The rule scheme, therefore, lends itself naturally to uniform numeric manipulations. This is not so for semantically richer knowledge structures; qualitative methods would seem more appropriate for these.

Those advocating quantitative methods for evaluating hypotheses do not claim that these procedures are the result of an analysis of the corresponding reasoning processes of the experts. A notable exception is the claim that the underlying assumption of MYCIN's method, *viz* that a piece of evidence cannot both favour and disfavour a hypothesis, is a well known paradox in physicians' reasoning[37]. Of course, this paradox could equally well constitute the basis for a qualitative method. What is perhaps interesting to note is the fact that a lot of emphasis is placed on how to train the experts to 'think' in terms of the particular quantitative methods, and thus to provide the necessary numeric values. Surely, if such methods were to capture the reasoning of the experts, such training would be unnecessary? The thrust of our argument is not which methods (quantitative or qualitative) yield better performance, but which methods are nearer to the experts' thinking. Our philosophy is that expert systems should model human competence, and should not aim to outperform human expertise[38].

Most quantitative methods used in diagnostic expert systems are rooted in probability theory. Their underlying assumptions (e.g. hypotheses are independent) are pragmatically imposed. Assuming a production rule implementation, such a method can be summarized as: the associations captured in the rules are assigned numeric values, usually from the interval $[-1, 1]$ (these values are predetermined and do not change during a consultation). Hypotheses are dynamically assigned similar values in the context of a problem. The basic procedures for this, as exemplified by systems like MYCIN[37] and PROSPECTOR[39] are:

- How are antecedent concept (evidence) uncertainties propagated over rule (associational) uncertainties to yield belief measures for the consequent concepts (hypotheses) from the perspective of the particular evidence?
- How are these individual conditional beliefs in some hypothesis combined to yield an overall belief measure for the hypothesis?

The hypotheses are, therefore, ranked on a finite scale ranging from some value (usually -1) denoting 'rejection' to some other value (usually 1) denoting 'confirmation'. Thus, quantitative methods reduce a hypothesis space to a linear configuration, where hypotheses types and their interdependencies are hidden (see also Cohen and Grinberg[40] for a similar conclusion). This could not be acceptable for semantically rich hypotheses spaces; a single number cannot capture a complex, interplanar situation.

PIP's qualitative method

Humans use a finite and rather small set of belief values for qualifying hypotheses. Consider the following low level reasoning constructs which were observed being used by an expert:

1 If this were the case then I would *eliminate* that possibility.
2 If this were the case then you must *definitely* go for that.
3 If this were the case then possibility X is *more likely* than possibilities Y or Z.
4 If this were the case then any acceptable solution *must satisfy* that.
5 If this were the case then a possibility that *satisfies* that would be *more likely*.
6 If you were to *accept* X then Y is *more likely* than Z.
7 If you were to *accept* this then any solutions for that *must satisfy* this condition.

Reasoning constructs 1 to 5 and 6 and 7 respectively relate findings to hypotheses and hypotheses to hypotheses (possibly in another space). The expert was expressing these relations through qualitative terms like *eliminate, definitely, more likely, must satisfy*, etc., instead of numbers. The PIP, NEOCRIB[31] and CASNET systems[36] use methods that are more or less qualitative in nature. PIP's method is outlined below as an illustrative example of these.

In PIP, those hypotheses that are strongly believed are the *active* ones. Hypotheses that are not yet strongly believed, but which have somehow been pointed at (usually through an active hypothesis), are *semi-active* ones. This qualification signifies the fact that although they are not in the immediate focus space, they are in the surrounding vicinity and ready to enter it. The other belief statuses are *concluded*, for those hypotheses that themselves become findings (facts), and *inactive*, for the general hypothesis pool. A *rejected* belief status is not used. PIP's method is qualitative because the transitions between the belief statuses are related to domain heuristics, and by implication to domain structures, not because it employs four rather than an infinite number of belief values. These heuristics provide the basis for justifying the transitions. PIP also employs a secondary quantitative method: hypotheses can be assigned matching scores. If a hypothesis' matching score is above a given threshold, the hypothesis is concluded, and if below another given threshold, the hypothesis is removed from further consideration. The cut-off values are used to convert a quantitative scale to a qualitative scale. The choice of these values is quite significant—in the absence of the qualitative criteria this choice could even be critical.

Skeleton for a qualitative method

What is being offered here is not a fully fledged qualitative method for evaluating hypotheses, but a skeleton for supporting the acquisition of the relevant method for a specific task. (We believe that what is suggested here bears application outside the diagnostic field.) The point is that such a method would depend on a complete understanding of the particular diagnostic reasoning processes, and thus domain structure. In contrast, quantitative methods treat a hypothesis space syntactically.

We propose the discrete belief scale, *concluded, active, semi-active, suspended, rejected* and *inactive* (for the general hypothesis pool). At any instant a hypothesis will have a single global (direct) belief status, and possibly a number of contextual belief statuses. The need for a concluded status and a rejected status is intuitively obvious. Active hypotheses should be strongly suggested, while semi-active hypotheses should simply be suggested. Hypotheses are activated through non-contextual and/or contextual abductive steps (see above). Hypotheses are semi-activated through contextual abductive steps. The suspended belief status is a kind of special status. A hypothesis moved to this status still retains its prior status (active or semi-active). Such transitions are the results of deductive diagnostic steps that were not successful in acquiring the required information. This information could refer to direct evidence (in favour or against) the hypothesis; it could refer to direct evidence (for or against) a hypothesis that is contextually related to the suspended hypothesis; or it could directly refer to the hypothesis, but its acquisition is being imposed by a contextually related hypothesis (see the discussion on constraints above). In the former case the hypothesis is directly suspended, and in the latter cases it is contextually suspended. A suspended status for a hypothesis, therefore, signifies that the information required to cause a transition from the hypothesis' current belief status is unobtainable. The suspended status simply implements a mechanism for temporarily removing hypotheses from the immediate focus. The unknown required information should still remain in the diagnostician's immediate focus space. These items of information can be viewed as signals on which the suspended hypotheses are waiting. As incoming user information is processed by the findings reasoner, the matching of signals against new information should be a function of the findings reasoner.

Figure 5 gives the possible (direct or contextual) status transitions for a hypothesis. Each transition is attributed to some reasoning chain that resulted in performing the particular manipulations on the domain knowledge and case information. In summary: abductive diagnostic steps are responsible for the transitions from the inactive to the active and semi-active statuses; inductive diagnostic steps are responsible for the transitions from the active and semi-active statuses to the rejected and concluded statuses; and deductive diagnostic steps are responsible for the transitions to the suspended status. Thus, understanding the reasoning behind all possible status transitions implies an understanding of the entire diagnostic process.

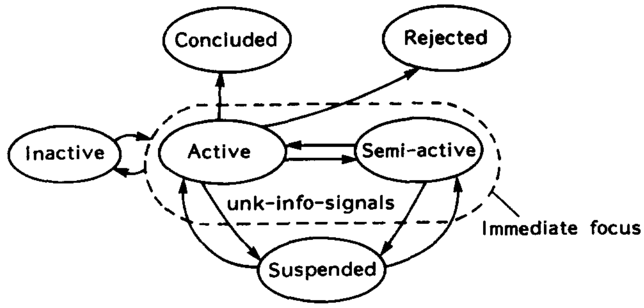

Figure 5. Hypotheses status transitions

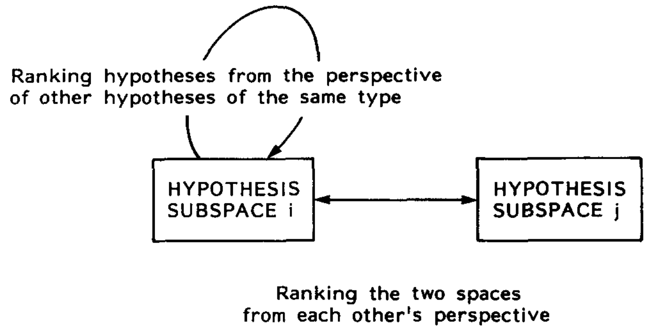

Figure 6. Contextual beliefs function to interrelate hypotheses spaces. Hypotheses suggest other hypotheses or impose constraints to be satisfied by other hypotheses

Assigning global beliefs

The global belief in some hypothesis must reflect the belief in the hypothesis, essentially due to direct factual information, i.e. case-specific findings. Findings either point to specific hypotheses or define constraints that must be satisfied by a hypothesis as a condition for its generation. Hypotheses, globally concluded, become case-specific findings.

Assigning contextual beliefs

Suppose we have a causally related hypothesis space (c.f. CASNET[36]). Hypotheses (generated or still inactive) can be evaluated from the perspective of the generated ones: how promising is this hypothesis given the hypotheses already generated? Thus, hypotheses of a given space can be contextually evaluated within their own space. Alternatively, some hypothesis subspace can be evaluated from the perspective of some other hypothesis subspace. For example, a hypothesis in one subspace could define a constraint to be satisfied by the hypotheses of the other subspace (such a situation arises in 'synthesis' problem-solving, where one part of the solution defines constraints for another part). If the former hypothesis is finally accepted, then a hypothesis from the other subspace must satisfy this constraint to be acceptable. The hypotheses thus generated are evaluated according to how well they satisfy the posted constraints. If a number of constraints (from one or more sources) are being posted, the evaluation may need to take into consideration the importance of the sources of the various constraints.

Contextual beliefs function to interrelate hypotheses spaces (see Figure 6). Contextual rankings are ignored if the ranking hypothesis is globally rejected. However, if the hypothesis is globally concluded the relevant contextual beliefs of the associated hypotheses should be raised to the global level, as the former hypothesis is now a finding. Contextual beliefs can guide the acquisition of additional information. For example, suppose that a hypothesis is globally suspended (at a semi-active status), but is concluded from the perspective of another hypothesis which is globally semi-active. Pursuing the latter hypothesis may therefore lead to the conclusion of both hypotheses.

Combining procedures

The term *combining procedure* is used to refer to the function of combining together all the evidence relevant to a hypothesis, to give an overall ranking for the hypothesis. The distributed nature of the proposed qualitative method precludes the use of a uniform combining procedure, as in quantitative methods. A hypothesis has a global belief, determined by the available direct evidence, and a number of contextual beliefs determined by its associations to other generated hypotheses. Thus, there is an immediate differentiation between the direct and the contextual evidence. A hypothesis that is only suggested by other hypotheses has an uninstantiated global status. An established (concluded) hypothesis becomes a finding, and thus direct evidence for its related hypotheses. The procedure that combines all the direct evidence relevant to a hypothesis (to yield the hypothesis' overall global belief) computes the match between the expectations on the hypothesis and the observations.

The separation between the direct and contextual evidence aids the focusing and information acquisition. A hypothesis may not look very promising from the perspective of direct evidence, but very likely from the perspective of contextual evidence. Direct and contextual evidence are semantically different, and should be treated as such. Pursuing only hypotheses with direct evidence (even strong direct evidence) generates a behaviour that lacks the parallelism inherent in multi-fault cases. Hypotheses would be pursued sequentially, hoping that the multi-fault picture will gradually be pieced together by promoting the consideration of hypotheses related to previously concluded hypotheses (which now constitute direct evidence for them). This is an important criticism of the INTERNIST-I system[27], where contextual evidence is essentially ignored. A hypothesis with strong contextual evidence, i.e. one that is strongly interconnected with other hypotheses, would be a more promising focus point from a hypothesis with some direct evidence only. Further, the separation between direct and contextual evidence eases the processing of finding revocations. As finding information is not propagated beyond their immediate associated hypotheses, only the global beliefs of the associated hypotheses for the revoked finding need to be modified. A procedure for combining direct and contextual evidence, if required, is envisioned to embody knowledge-intensive, domain-specific methods.

The discussion above applies both for simplex and compound hypotheses. The situation for simplex hypotheses is illustrated in Figure 7a, and the corresponding situation for compound hypotheses in Figure 7b. Referring to Figure 7a, the links from the direct evidence to the hypothesis, H, are of the same associational nature. However, the links from the contextual evidence to the hypothesis could represent different relationships. Contextual evidence could be of varying types.

A compound hypothesis has components that are hypotheses at a lower level of abstraction. Usually such components are mutually interrelated in some sense. However, the most general fault model is where there are multiple dependent and/or independent faults, thus allowing for compound hypotheses to consist of independent islands of interrelated components (see Figure 7b). Usually compound hypotheses evolve as the diagnostic process progresses. At the end of a consultation one of these hypotheses must completely explain the diagnostic problem. The direct evidence relevant to a compound hypothesis is the combination of the direct evidence for its component hypotheses. The contextual evidence for a compound hypothesis is split into intra- and inter-contextual evidence. Either of these may (partly) depend on the contextual evidence for the component hypotheses. The intra-contextual evidence must reflect the quality (plausibility) of the hypothesis based on criteria such as simplicity, completeness, coherence, etc. The inter-contextual evidence reflects how the hypothesis interrelates with other hypotheses of the same or different type (same definition as for simplex hypotheses).

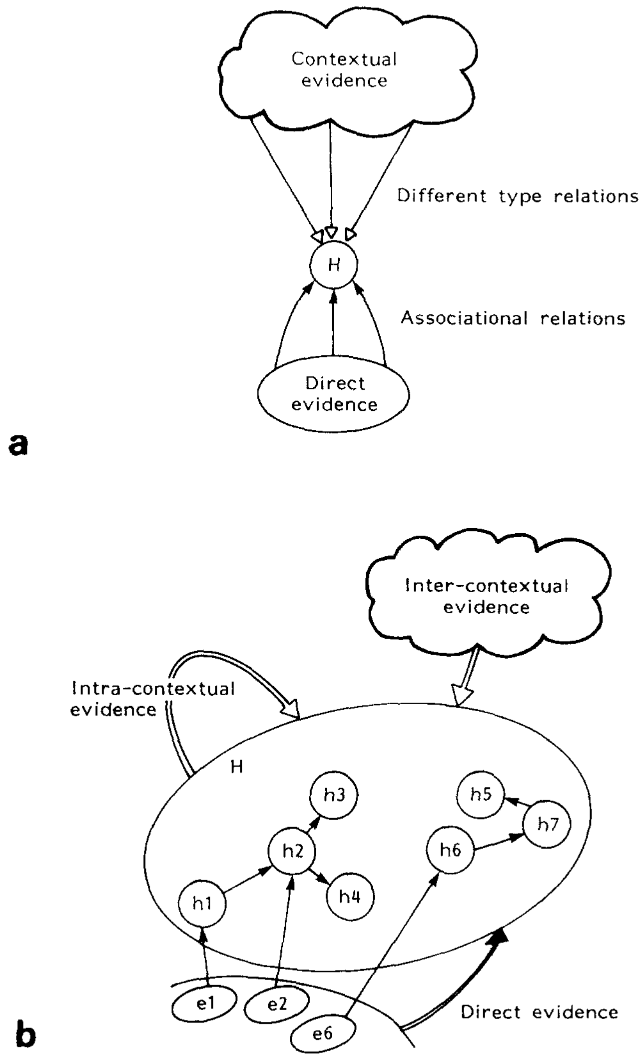

Figure 7. Direct and contextual evidence for hypotheses. (a) Simplex hypotheses, (b) Compound hypotheses

Constraint satisfaction

Hypotheses can be defined implicitly through constraints. A constraint is taken to be a qualified assertion, e.g. '*must-not* cause teeth discolouration' could be a constraint for a treatment hypothesis. Qualifiers could include *must-not*, *rather-not*, *rather-yes* and *must*. The 'match' between a hypothesis and an assertion could be given by the three-valued scale, *no*, *possible* and *yes*, where the middle value denotes that the assertion can be satisfied by the hypothesis, possibly at an extra cost. The belief status for a hypothesis is determined by the particular match given the specified qualifier (see Table 1). If the constraint corresponds to direct evidence, then the belief status is at the global level. If, on the other hand, the constraint is defined by another hypothesis, then the belief status is a contextual one.

Provided that a previously generated hypothesis already satisfies the constraint to a specified required level (i.e. activation or semi-activation levels), no new hypotheses need be generated. Otherwise, the most general hypothesis that satisfies the constraint to the required level must be generated. The 'matching' function of the constraint satisfaction procedure could, by and large, be carried out by the findings reasoner.

Justifying hypotheses

It is essential for a diagnostic system to be able to adequately justify to the user its belief in some hypothesis. Within the proposed qualitative framework, rich justifications can be generated by tracing a hypothesis' status transitions from the perspectives of direct and contextual evidence. Recall that a status transition is justified in terms of the reasoning task that gave rise to it. Thus, the quality of explanations would depend on the representational adequacy of the reasoning tasks (see above).

Table 1. Constraint satisfaction

Belief-status		*Qualifier* Must-not	Rather-not	Rather-yes	A-must
	No	Active	Semi-active	Semi-active	Reject
Assertion	Possible	Active	Active	Semi-active	Semi-active
type	Yes	Reject	Semi-active	Semi-active	Active

CONCLUSION

A model of diagnostic behaviour has been presented at a sufficiently high level of abstraction to cover most diagnostic tasks. The model separates the data handling function from the main diagnostic function.

The data (findings) reasoner deals with conflicts in user evidence, the identification of potential triggers and red herrings, instructing the user in carrying out information acquisition actions and the management of case information (historic or current). In this respect, deciding the truth status of a finding from the known findings is a central operation. To be fully adequate, the findings reasoner must be capable of deep temporal reasoning.

The diagnostic function deals with the generation and evaluation of hypotheses. A diagnostic step is abductive, deductive or inductive in nature. Focusing and information acquisition are critical aspects of a diagnostic process that entail the evaluation of current hypotheses. A qualitative method for evaluating hypotheses that separates the direct and contextual evidence relevant to a hypothesis is being proposed. Direct evidence is derived from case-specific information, while contextual evidence is derived from the hypothesis' interrelationships to other hypotheses. A hypothesis' (direct or contextual) beliefs are from the qualitative range—inactive, semi-active, active, suspended, concluded and rejected. A hypothesis can be justified by tracing the transitions in its belief.

In a domain, a knowledge of entities and their interrelationships is factual knowledge, and knowledge of how to use the factual knowledge to generate knowledgeable behaviour is reasoning knowledge. Factual knowledge can be represented in schemes that combine rules and frames. An analysis of reasoning knowledge that can be captured in generic frame structure has been presented here. This structure accurately captures the semantics of the reasoning knowledge as transpired from our analysis. According to the proposed structure, reasoning knowledge is statically represented as a tangled task hierarchy. Dynamically, tasks are instantiated and executed. All the task instantiations required to achieve a global task are related into a tree, referred to as the strategic explanations tree. Therefore, in the proposed diagnostic model the reasoning is distributed among a number of knowledge-intensive tasks. This brings out the total functionalilty of the system, including the bases for passing control from one task to another (control knowledge).

Figure 8 overviews the components of a diagnostic picture. The diagnostic picture is the placeholder for the case-specific information and the results of the operations of the findings reasoner and the diagnostician. The strategic trees have been included as components of the diagnostic picture in Figure 8. Note that although there is only one tree for the diagnostician, there are numerous trees for the findings reasoner. Findings reasoner trees correspond to nodes in the diagnostician tree.

REFERENCES

1 **Chandrasekaran, B** 'Generic tasks in knowledge-based reasoning: high level building blocks for expert system design' *IEEE Expert* Vol 1 No 3 (Fall 1986) pp 23–30

2 **Johnson, L and Johnson, N E** 'Knowledge elicitation involving teachback interviewing' in **Kidd, A L (ed)** *Knowledge Acquisition for Expert Systems* Plenum Press, UK (1987)

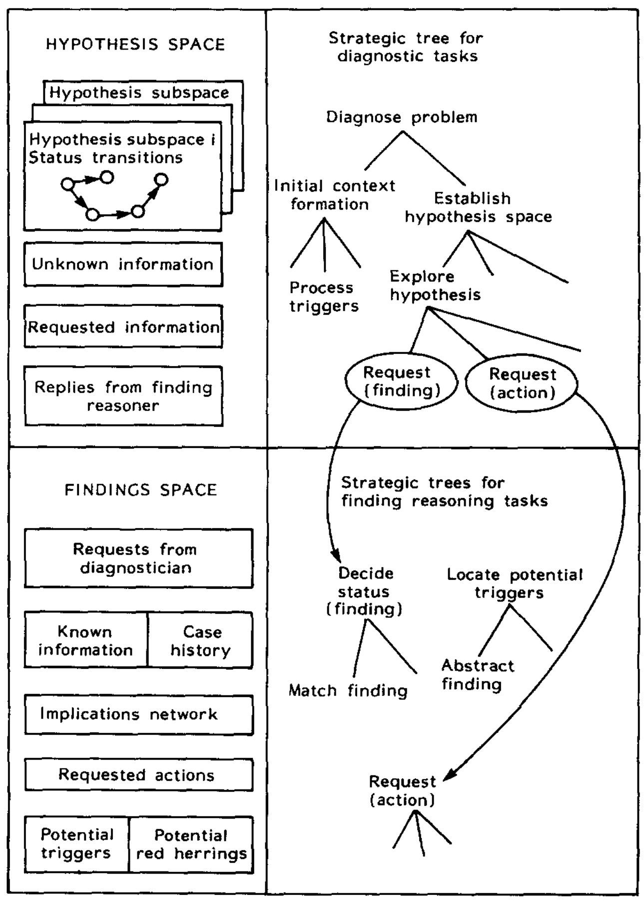

Figure 8. Diagnostic picture overview

3 **Breuker, J A and Wielinga, B** 'Use of models in the interpretation of verbal data' in **Kidd, A L (ed)** *Knowledge Acquisition for Expert Systems* Plenum Press, UK (1987)

4 **Keravnou, E T and Johnson, L** *Competent Expert Systems: A Case Study in Fault Diagnosis* Kogan-Page, UK (1986)

5 **Shortliffe, E H, Buchanan, B G and Feigenbaum, E A** 'Knowledge engineering for medical decision making: a review of computer-based clinical decision aids' *Proc. IEEE* Vol 67 No 9 (1979) pp 1207–1224

6 **Young, D W** 'A survey of decision aids for clinicians' *Br. Med. J.* Vol 285 (November 1982) pp 1332–1336

7 **Keravnou, E T** 'Automated reasoning in fault diagnosis and verification' *Int. J. Syst. Res. & Info. Sci.* Vol 1 No 4 (1986) pp 269–285

8 **Whitbeck, C** 'What is diagnosis? a preface to the investigation of clinical reasoning' *Metamedicine* Vol 2 (1981) pp 319–329

9 **Elstein, L D, Shulman, L A and Sprafka, S A** *Medical Problem Solving: An Analysis of Clinical Reasoning* Harvard University Press, USA (1978)

10 **Feltovich, P J, Johnson, P E, Moller, J H and Swanson, D B** 'LCS: the role and development of medical knowledge in diagnostic expertise' in **Clancey, W J and Shortliffe, E H (eds)** *Readings in Medical AI: The First Decade* Addison-Wesley, UK (1984)

11 **Johnson, P E, Duran, A S, Hassebrock, F, Moller, J, Prietula, M, Feltovich, P J and Swanson, D B** 'Expertise and error in diagnostic reasoning' *Cognitive Sci.* Vol 5 (1981) pp 235–283

12 **Kassirer, J P, Kuipers, B J and Gorry, G A** 'Towards a theory of clinical expertise' *Am. J. Med.* Vol 73 (1982) pp 251–259

13 **Pauker, S G and Szolovits, P** 'Analyzing and simulating taking the history of the present illness: context formation' in **Schneider, W and Sagvall Hein, A-L (eds)** *Computational Linguistics in Medicine* North Holland, Netherlands (1977) pp 109–118

14 **Swanson, D B, Feltovich, P J and Johnson, P E** 'Psychological analysis of physician expertise: implications for design of decision support systems' MEDINFO77 — *Proc. 2nd World Conf. on Medical Informatics* Toronto, Canada (1977) pp 161–164

15 **Johnson, L and Keravnou, E T** *Expert Systems Technology: A Guide* Abacus Press, UK (1985)

16 **Johnson, L and Keravnou, E T** 'Analysing, representing and interpreting expert strategic knowledge' in *Cybernetics & Systems* Reidl Publishing Co., USA (1986) pp 743–750

17 **Dhar, V and Pople, H E** 'Rule-based versus structure-based models for explaining and generating expert behaviour' *Commun. ACM* Vol 30 No 6 (1987) pp 542–555

18 **Clancey, W J** 'Heuristic classification' *Artif. Intell.* Vol 27 (December 1985) pp 289–350

19 **Pople, H E** 'Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnostics' in **Szolovits, P (ed)** *Artificial Intelligence in Medicine* Westview Press, USA (1982) pp 119–185

20 **Pauker, S G, Gorry, G A, Kassirer, J P and Schwartz, W B** 'Toward the simulation of clinical cognition: taking a present illness by computer' *Am. J. Med.* Vol 60 (1976) pp 981–995

21 **Aikins, J S, Kunz, J C, Shortliffe, E H and Fallat, R J** *PUFF: An Expert System for Interpretation of Pulmonary Function Data* Stanford report STAN-CS-82-931, Stanford University, USA (1982)

22 **Clancey, W J and Letsinger, R** 'NEOMYCIN: reconfiguring a rule-based expert system for application to teaching' *Proc. Int. Joint Conf. on Artif. Intell.* Vancouver, Canada (1981) pp 829–835

23 **Chandrasekaran, B and Mittal, S** 'Conceptual representation of medical knowledge for diagnosis by computer: MDX and related systems' *Adv. Comput.* Vol 22 (1983) pp 217–293

24 **Johnson, L** 'Analysis of the diagnostic task' *Int. J. Syst. Res & Info. Sci.* Vol 1 No 4 (1986) pp 225–236

25 **Peirce, C S** 'Illustrations of the logic of science, sixth paper—deduction, induction, hypothesis' *Pop. Sci. Month.* Vol 1 (1878) pp 470–482

26 **Shortliffe, E H** *Computer-based Medical Consultations: MYCIN* Elsevier, USA (1976)

27 **Miller, R A, Pople, H E and Myers, J D** 'INTERNIST-I: an experimental computer-based diagnostic consultant for general internal medicine' *New England J. Med.* Vol 307 No 8 (1982) pp 468–476

28 **Patil, R S** *Causal Representation of Patient Illness for Electrolyte and Acid-base Diagnosis* MIT/LCS/TR-267 Technical Report, Massachusett Institute of Technology, USA (1981)

29 **Clancey, W J** 'From Guidon to Neomycin and Heracles in twenty short lessons: ORN final report 1979–1985' *AI Magazine* Vol 7 No 3 (Fall 1986) pp 40–60

30 **Keravnou, E T and Johnson, L** 'Intelligent handling of data by integration of commonsense reasoning' *Knowl.-Based Syst.* Vol 1 No 1 (December 1987) pp 32–42

31 **Keravnou, E T and Johnson, L** 'NEOCRIB: an expert fault finding system that articulates the competence of field engineers' *Proc. 3rd Int. Conf. on Fault Tolerant Comput. Syst.* Brehenhaven, FRG (1987)

32 **McDermott, D and Brooks, R** 'ARBY: diagnosis with shallow causal models' *Proc. AAAI-82* Pittsburg, PN, USA (1982) pp 370–372

33 **Johnson, L** 'Organizing a findings (data) base for use in a competent automated diagnostician' *Int. J. Syst. Res. & Info. Sci.* Vol 1 No 2 (1985) pp 91–112

34 **Clancey, W J** 'Acquiring, representing and evaluating a competence model of diagnostic strategy' in **Chi, M, Glaser, R and Farr, M (eds)** *Contributions to the Nature of Expertise* Stanford University, USA (1985)

35 **Kidd, A L and Cooper, M B** 'Man-machine interface issues in the construction and use of an expert system' *Int. J. Man-Mach. Stud.* Vol 22 No 2 (1985) pp 91–102

36 **Weiss, S M** *A System for Model-based Computer-aided Diagnosis and Therapy* PhD thesis, Department of Computer Science, Rutgers University, USA (1974)

37 **Shortliffe, E H and Buchanan, B G** 'A model of inexact reasoning in medicine' *Math. Biosci.* Vol 23

(1975) pp 351–379
38 **Johnson, L** 'The need for competence models in the design of expert consultant systems' *Int. J. Syst. Res. & Info. Sci.* Vol 1 No 1 (1985) pp 23–26
39 **Duda, R, Hart, P and Nillson, N** 'Subjective Bayesian methods for rule-based inference systems' *Proc. Nat. Comput. Conf. (Vol 45)* (1976) pp 1075–1082
40 **Cohen, P R and Grinberg, M R** 'A framework for heuristic reasoning about uncertainty' *Proc. Int. Joint Conf. on Artif. Intell.* Karlsruhe, FRG (1983) pp 355–357