Control Issues in Classificatory Diagnosis*

Jon Sticklen, B. Chandrasekaran, and John R. Josephson

Laboratory for Artificial Intelligence Research
Department of Computer and Information Science
The Ohio State University
Columbus, Ohio 43210 USA

Abstract

A good part of medical diagnosis can be modeled as classification problem solver producing a "differential", working in conjunction with an abductive component that performs differential diagnosis by synthesising a "best" composite hypothesis out of the hypotheses in the differential list. Classification problem solving itself can be viewed as having a control component which selects hypotheses to consider, and a decision component associated with each selected hypothesis In this paper we study the family of control regimes that are useful in classificatory problem solving. We start with MDX, a classification system organized as a hierarchical collection of hypothesis specialists, critique its control behavior, and by considering a set of situations involving multiple diseases, show how elements can be added to the control regime in a modular way to handle a large variety of situations.

1. Introduction

As the field of knowledge-based systems is maturing, a clear movement is taking place from relatively flat and uniform representation systems and associated inference regimes, to an appreciation of the multiplicity of knowledge structures and problem solving types that typically play a role in complex real world tasks [4, 6]. There is also an increasing awareness that the problem solving behavior of knowledge-based reasoning systems is best understood at what Marr [13] has called the *information processing /eve/*, or the *knowledge level* as it has recently been called by Newell [15]. E.g., at the implementation language level MYCIN'S diagnostic action can be thought of as backward-chaining, while at the information processing level its activity is best understood as a form of classification. In our group we have persistently emphasized the information - processing level analysis of knowledge-based tasks: in [10], we identified certain aspects of medical diagnosis with classification problem solving, and in [4] we propose a taxonomy of knowledge-based reasoning tasks at the information processing level. Clancey [7| has recently taken a similar perspective and shown that a number of knowledge based systems which appear to be doing different kinds of things at the implementation language level, can be in fact seen to be performing variants of the classificatory task.

Medical diagnostic reasoning is a complex activity, and can be decomposed into a number of different types of problem solving, each with specific kinds of knowledge structures and control regimes. Depending upon the subdomain, the basic task may differ considerably: a knowledge-level statement of the tasks faced by CASNET/Glaucoma [12] or ABEL [16] (viz., put together a "causal" story) is bound to be different from those describing tasks faced by Mycin, Internist or MDX. As a general statement, however, one may say that medical diagnosis is an *abductive* process. Peirce [17], and,following him, Pople [18] have viewed the abductive enterprise as the generation of hypotheses, which, if true, would explain some collection of observed facts.

In medicine the diagnostician is presented with a patient who has a set of signs and symptoms (collection of observed facts), and proceeds to reason toward a set of diseases that can account for the observed signs and symptoms (a set of hypotheses that explain the observed facts)

If a given abductive reasoning task involved the selection of hypotheses from a small set of candidates, then it would be reasonable to directly compare all of the candidates But since the number of potential medical hypotheses is enormous, the medical diagnostic process is observed to consist of two components:

1. the selection of a limited number of candidate diseases, and
2. a decision on which members of the candidate set are required to account for patient signs and symptoms.

The work of Feltovich et al [8] supports such a distinction, and goes on to suggest that much of the difference between novice and expert diagnostic behavior is due to a difference in how the limited set of possibilities is formed (Factor 1 above). In medical terminology, Step 1 is often called *forming the differential* while Step 2 is called *differential diagnosis.*

Classification problem solving is useful in forming the differential. The MDX system [l, 5] can in this sense be thought of as forming a differential, since it produces a list of disease classes into which the signs and symptoms of the case can be classified. NEOMYCIN uses classification explicitly, and a large part of Mycin's work in diagnosis can also be viewed as implicitly classificatory.

The MDX approach, while concentrating on the classificatory part, envisaged a component called the *Overview Critic* for performing Step 2 above. As a result of recent investigation by Josephson, et al. [II] of our group, we now have a first cut theory of how this component should work. The differential diagnosis process is viewed as the process of *assembling a composite hypothesis* from the list of classificatory hypotheses produced by the classificatory component, such that the composite hypothesis "best explains" all the data. This theory is presented in detail in [11].

The purpose of this paper is to discuss the control issues in the classificatory component of diagnostic reasoning. We will briefly review the MDX approach, discuss the essential aspects of its current control strategy, some of the difficulties faced by it, and discuss extensions to the MDX viewpoint by systematically

*This work has been supported by NSF grant MCS-8103480, and Biomedical Computing and Information Processing Training Grant NLM 5T15LM07023-05.

considering examples of increasing difficulty, especially examples involving multiple diseases.

2. The MDX Approach

The hallmarks of the MDX approach to date have been

- a cognitively plausible, hierarchical decomposition of diagnostic domain knowledge into a number of cooperating, classificatory specialists, each representing a diagnostic hypothesis, [2|
- a reliance on the *individual* specialist's ability to determine *on a local basis* the applicability of the diagnostic hypothesis (represented by the specialist) to the current case [3, 19], and
- an overall control regime call*establish/refine* in which classification is begun at the top of the classification hierarchy, and at each step *establishes* a specialist before going on to *refine* the specialist by examining less general subordinate specialists [10].

Figure 1 shows a fragmentary high level level decomposition of medical knowledge.

```
                Internist
                    |
        --------------------
        |           |      ...
      Heart       Liver
        |           |
      ------      ------------------
      | ...       ...  |          |
ValvularDisease      Cirrhosis  Cholestasis
```

Figure 1: Sample MDX Specialist Hierarchy

Suppose that a patient problem consisting of cholestatic liver disease is presented to the MDX-style system in Figure 1. The control of the system will proceed in a top down fashion: Internist will first *establish,* then *refine* itself. The establish step consists of Internist utilizing locally available domain knowledge to determine if the patient is sick at all. The refine step consists in Internist sending to its subordinate specialists (Heart, Liver, ...) the composite message "establish/refine". For the hypothesised case, only Liver will be established. Hence all subordinates to Heart, ... will not be considered. The Liver specialist, in contrast, will establish and subsequently refine itself. Finally, Cholestasis will receive an establish/refine message and establish itself. Note that when a hypothesis is rejected, all its successors are also rejected: this is the pruning power of hierarchical classification. Also, typically, many hypotheses may be "suspended"; i.e., may not have enough data to positively establish them or rule them out. It will be in general combinatorially prohibitive to explore all of them further.

2.1. Two Sets of Control Issues in Classification Systems

In classification problem solving in the style of MDX, we can see *two distinct kinds of control activity:*

1. invocation of different classificatory hypotheses for evaluation. (Control issue: What are the potential paths of invocation, and what are the potential tasks that may be requested?) We shall refer to this type of control activity as *classificatory control.*
2. once a hypothesis is invoked, making decisions about the relevance or applicability of the hypothesis to the case at hand, e.g., reject, establish, decide what data it can account for, etc. (Control issue: What knowledge sources can help in this, and how and when to invoke them?)

For 2 above, the current version of MDX uses a fairly straightforward pattern-matching mechanism using "compiled" knowledge to map from a subset of patient data to qualitative decisions about the hypothesis. But there is no reason why this task cannot use other kinds of knowledge and problem solving such as causal reasoning, or table look up, or resolution theorem proving for that matter. (In fact, the *DART* system of Genesereth et al. [9]does precisely that: its higher level architecture is classificatory, closely following the design component/subcomponent hierarchy, while the presence of a component fault is determined by a theorem prover working with the test values and the axioms that define the component.) Thus, *how* a decision is made within a hypothesis, so to speak, is irrelevant from the point of view of *classificatory control,* as long as the information that is needed is available after an invoked hypothesis has completed its action.

Thus the control issues in 1 above are issues *intrinsic* to classification problem solving, while control issues in decisions within each hypothesis are issues for a different type of problem solving. Our concern in this paper is with the former. We call this set of issues *classificatory control issues.*

2.2. Classificatory Control in MDX

It is convenient to study control in distributed systems such as MDX, which are implemented as a community of specialists which coordinate their work by exchanging messages, by examining the constraints on message channels and the contents of the message types. There is no implication that classificatory problem solving need necessarily be implemented in this way, or that our discussion of control is relevant only for classificatory problem solvers implemented in this way. The message language is merely a convenient device to talk about which hypothesis can invoke which others, requesting or providing which information.

Classificatory control in MDX was provided by the following elements:

1. The hypothesis hierarchy sets invocation paths, or message channels, as between parents and children. A hypothesis may invoke its children in order that the corresponding hypotheses may be evaluated for their presence or absence

   (message types: (establish) and (refine))

   or it may invoke its parent reporting on the results of its activity

   (message type: (establishedAt *<confidence level>))* .

2. Each hypothesis, after its establishment, may use any knowledge it has to <u>order</u> its successors for invocation. This ordering may be based on likelihood, risk, etc.
3. A hypothesis may also have knowledge corresponding to which hypothesis, one or more levels below its immediate successor, may be considered at that time in order to

Although the control issues of classification may be viewed independently of the message passing paradigm, at the level of cognitive organisation it is very useful to understand a complex cognitive agent in terms of the interactions of many simpler agents via well constrained message types and communication channels; but this is another issue.

increase problem solving efficiency, i.e., it may use context-dependent *suggestion rules.* These provide economy of effort by suggesting hypotheses which are highly plausible, but which would not be immediately considered within the purely hierarchical regime.

(message type:
  (refine (suggest *<subordinate specialis>))).*

The introduction of suggestion rules can potentially cause difficulties in control if improperly used. In the hierarchical control regime, when a hypothesis is invoked for the purpose of establishing it, it can be assumed that its parent has been established. This provides a context for its establishment activity. But the suggestion rule mechanism will typically result in the invocation of hypotheses whose parents have not been established, and it is possible that without this context, the establishment activity of the invoked hypothesis will be erroneous. Thus the use of the suggestion rule mechanism will need to be restricted to hypotheses that are able to arrive at decisions without significant dependence on the context provided by the parent.

2.3. Limitations of the MDX Control Regime

Three kinds of limitations are worth mentioning at this point.

1 As a classification problem solver, MDX outputs a simple list of disease hypotheses which are considered to be likely As mentioned earlier, the MDX approach is currently being extended to provide a differential diagnosis component, which assembles the best composite hypothesis out of the output list from the classifier using information about what those hypotheses can explain The availability of this abduction component can contribute to better classification problem solving by providing the following control possibilities:

   As mentioned in the section on how MDX works, it is often the case that many hypotheses are suspended; i.e., there is not enough data to establish or reject them. Typically it will be quite prohibitive to "expand" all of the suspended hypotheses by considering their successors in the hopes that they may be able to find data to reject or establish themselves. However, after the assembler has put together a composite hypothesis based on the current output list from the classifier, it can produce a list of manifestations that still remain to be explained. (For how this is done, see [11].) If this list is empty, the problem is "done". If not, a selection can be made of the suspended hypotheses, on the basis of which of them can potentially explain the remaining unexplained findings, and only this subset need now be further explored. Thus the abduction assembly machine and the classifier can work in close collaboration, and the former can provide focus to the classification problem solver.

2. Lack of Establish Knowledge for Some Hypotheses: In some domains, the only way to establish or reject a classificatory hypothesis, call it H, is by invoking its successors to see if any of them can be established, or whether all of them can be rejected. This might be either because there simply is no domain knowledge available about H, or because data at the level of H are not available in the case at hand. In the former case, this fact can be explicitly made available in H, which can then be made to invoke its successors (ordering them if information for that is available). In the case where data are not available for the case at hand, the hypothesis in question would have been suspended, and the approach in 1 above could be useful.

3. Limitations of hierarchical invocation for multiple diseases: Selective communication/invocation outside the hierarchical channels is often needed in the case of multiple diseases, where the decision status of one classificatory hypothesis may be needed for making decisions about another hypothesis. Problems of this sort will be discussed in the next sections. They provide some of the most interesting challenges in the control of classification problem solving

3. Cases Of Multiple Patient Diseases

For purposes of illustration, in this section we will present examples of medical situations in which multiple disease hypotheses are required to fully account for patient data. Successful solutions for cases involving multiple disease hypotheses have been difficult to achieve in most medical AI diagnostic systems.

3.1. Multiple Independent Diseases

Referring to Figure 1, suppose that a patient has both valvular disease and cholestasis. The output list from MDX of possible patient diseases would contain both valvularDisease and cholestasis, each typically accounting for different findings. Thus the control component of MDX responsible for differential formation is adequate for this task.

3.2. When One Hypothesis Needs The Establishing Status Of Another - Example: "Secondary To" Diseases

Consider two classificatory specialists, Spl and Sp2. Many times in order for Spl to establish itself, it is useful for the establishing status of Sp2 to be available. Note that given this situation, it is not necessary for Spl to know the detailed problem solving that Sp2 used to determine its establishing status; i.e., how Sp2 determined its establishing status.

In medicine an example of the control problem posed in the above paragraph takes place when two disease hypotheses stand in a secondary to relation. At a pathophysiological level, there is often an expression of medical knowledge that one disease can either

1. be caused by another via some disease process, or
2. is temporally preceded by another statistically.

In compiled diagnostic knowledge, such a relation between two diseases is called the *secondary to* relation. Although at the compiled level, there will not be any knowledge about *how* the one disease causes the other, or *why* there is a statistically significant temporal ordering, still the knowledge that the one is *secondary to* the other can be useful in diagnosis. In most instances, if A is *secondary to* B, then knowing that A is established or rejected can directly contribute to a determination of A.

In Figure 2 we have expanded the diagnostic hierarchy of Figure 1 to include one disease hypothesis (CardiacClrrhosis)

It may be possible to to tag that part of the establishing knowledge that is needed for context setting of underlying specialists and test it before considering the suggested specialist. In the original MDX implementation [14], all the specialists in the path towards the suggested specialist were required to *establish.*

that is *secondary to* another (TricuspidInsufficiency).* Suppose that the original exploration path in a given case was as follows.

```
Internist -> Liver
          -> Cirrhosis
          -> CardiacCirrhosis
```

The hypothesis CardiacCirrhosis cannot establish unless the hypothesis it is secondary to, TricuspidInsufficiency, can first establish.

```
                  Internist
                      |
            ----------------
            |           |   ...
          Heart       Liver
            |           |
       ---------    ------------------
       |    ...    ...  |          |
  ValvularDisease    Cirrhosis  Cholestasis
     |                  |          ...
   ----------       ---------
   |   ...          |   ...
TricuspidInsufficiency  CardiacCirrhosis
```

Figure 2: Example Of "Secondary To" Hypothesis

The inverse relation to *secondary to* is *can have complication.* For example,

TricuspidInsufficiency *canHave Complication* CardiacCirrhosis.

Suppose the exploration path of the diagnostic system of Figure 2 was

```
Internist -> Heart
          -> ValvularDisease
          -> TricuspidInsufficiency.
```

In this example, the establishing values of either tricuspid insufficiency or of cardiac cirrhosis can have direct bearing on the establishment of the other.

In [10], the problem of a given classificatory specialist needing to know the establishing value of a second specialist is solved by positing a *blackboard* in which the decision status of hypotheses are written as soon as they are available. The solution also involved *parallel invocation* of successors of a hypothesis that has been established. Thus both of the exploration paths above, viz., from Internist to CardiacCirrhosis and from Internist to Tricuspid Insufficiency would be performed in parallel if the patient had both the diseases. But since TricuspidInsufTiciency can be established independent of CardiacCirrhosis, it will not need to refer to the status of the latter. However the latter hypothesis can obtain from the blackboard the information regarding the status of TricuspidInsufflciency. If, say, in a particular case heart is ruled out, the rejection of its grandchild, TricuspidInsufTiciency, can also be inferred on the blackboard.

In a serial implementation, there will be some problems that would need special attention. If the Internist to CardiacCirrhosis path had been preferentially pursued, either because of a suggestion rule or because Liver had a higher ordering than heart as more likely, there would be a need for the other line of exploration to be activated. There are some straightforward ways this can be done, and we forego further discussion of it here.

*Medically, cardiac cirrhosis can be *secondary to* a number of diagnostic alternatives under the heart category. For simplicity here, we will discuss only tricuspid insufficiency.

It is also possible to solve this version of the multiple disease situation by explicitly providing an invocation path from CardiacCirrhosis to TricuspidInsufTiciency; i.e., when the former hypothesis is invoked it checks the status of the latter by directly invoking it, with the message *Establish.* However, this solution requires additional work, since, as in our earlier discussion on suggestion rules, TricuspidInsufTiciency cannot in general use its decision knowledge without knowing that its parent has been established. Eventually at some point the Heart specialist's establishment status will need to be known. It is fairly straightforward to come up with additional message types to invoke the appropriate chain.

To reiterate, although we have motivated this section by explicating the control capabilities necessary to handle the *secondary to* example, these same elements of control will be needed to cope with <u>any</u> situation in which one diagnostic specialist needs knowledge of the establishing values of other specialists in order to undertake its own establishment.

3.3. When Reconsideration Is Called For - Example: Multiple Diseases With Additive Symptoms

In this section, we will consider patient situations that can be successfully explained only by hypothesizing two diseases, each associated with the same patient sign being higher than the normal range or lower than the normal range. For example, again referring to Figure 2, suppose that a patient has both cardiac cirrhosis and cholestasis. Also assume that cardiac cirrhosis can typically account for an elevation of up to 3 mg/dl of bilirubin in the blood, and that the presence of cholestasis can account for bilirubin in excess of 5 mg/dl. Finally, suppose that a blood test of the patient shows 9 mg/dl of bilirubin.

A physician's understanding of this case might be captured in outline by:

1. Relying on relevant patient signs and symptoms, it appears lik*ely* that the patient has cholestasis.
2. It also appears *likely* that the patient has cardiac cirrhosis, although only part of the bilirubin is explainable by the hypothesis of cardiac cirrhosis alone. The problem with not accounting for the total observed bilirubin reduces the overall confidence in this hypothesis somewhat.
3. But a coherent understanding of the current patient state is that she has *both cardiac cirrhosis and cholestasis.* By deciding that the patient state is best characterized by these multiple diseases:
    a. the total bilirubin elevation is explained as the sum of contributions associated with cardiac cirrhosis (3 mg/dl) and cholestasis (6 mg/dl), and
    b. the confidence that the patient has cardiac cirrhosis is restored.

Bilirubin is a blood chemical that is cleared by the liver. Elevated bilirubin is associated with many liver diseases. The medical situation for the example shown is complicated here by the need to do a base line study to determine how much bilirubin can be accounted for by cardiac cirrhosis. For simplicity, we assume that the base line study shows that for this case of cardiac cirrhosis, 3 mg/dl is a reasonable elevation.

Reasoning Steps 1 and 2 can be construed as normal establishing behavior for disease hypotheses. In Step 1, cholestasis is established at +2 (likely), because the signs and symptoms associated with cholestasis are matched relatively well in the patient data. In Step2, cardiac Cirrhosis is established at +1 (possible). Signs and symptoms associated with cardiac cirrhosis match the patient data well too, except that the bilirubin is too high.

Reasoning Step 3 requires the addition of new control elements. Step 3a requires that a *higher level* specialist integrate results from the subordinates cholestasis and cardiac Cirrhosis. In particular, those two subordinates must return information to liver to enable an integration.

If cholestasis returns to liver the result

```
[(establishedAt +2)
 (canAccountFor
    (bilirubin 9))]
```

and cardiacCirrhosis returns to cirrhosis and subsequently to liver the result

```
[(establishedAt +1)
 (canAccountFor
    (bilirubin 3))]
```

then liver will have the necessary information for understanding that the elevated bilirubin can be explained by hypothesising multiple diseases, assuming that liver knows that bilirubin levels combine additively.

Step 3b is the most interesting step above. The local establishing knowledge of cardlacClrrhosls yielded an establishing value of +1 because even though most knowledge "fit" the patient data, the value of bilirubin was higher that would be expected in a case of cardiac cirrhosis only. But after reasoning Step 3a, we see that the high reading for bilirubin can be understood as coming from two sources. Hence, cardiacCirrhosis need account for only part of the elevated bilirubin.

Step 3b is then a re-establishment of cardiacCirrhosis with the assumption that only part of bilirubin value need be accounted for. This step requires that a new message type:

```
[re-establish
 (assume ...)].
```

In the case above, we want cardiacCirrhosis and cholestasis to receive the messages

```
to cardiacCirrhosis.
    [re-establish (asanas (bilirubin 3))]. and


to cholestasis:
    [re-establish (assuae (bilirubin 6))].
```

On receiving this message, cardiacCirrhosis will establish in the usual method, *except* the value of bilirubin will be assumed 3. Sending the above message to cholestasis was crucial. We at least suspected that by forcing cardiacCirrhosis to assume a bilirubin level of 3, that its establishing value would be raised. But it also had to be demonstrated that by lowering the bilirubin value that cholestasis sees to 6, that the multiple disease hypothesis would remain viable.

A partial message trace for the problem discussed in this section is shown in Figure 3.

```
from liver
 to cardiacCirrhosis:      (establish/refine)

from cardiacCirrhosis
 to liver:                 [{(cardiacCirrhosis
                               established-at +I)
                             (can-account-for
                               (bilirubin 3))}|

from liver
 to cholestasis:           (establish/refine)



from cholestasis
 to liver:                 [{(cholestasis
                               established-at +2)
                             (can-account-for
                               (bilirubin 9))}|

from liver
 to cardiacCirrhosis:      (re-establish
                             (assume bilirubin 3))

from cardiacCirrhosis
 to liver:                 [{(cardiacCirrhosis
                               established-at +2)
                             (can-account-for
                               (bilirubin 3))}|

from liver
 to cholestasis:           (re-establish
                             (assume bilirubin 6))



from cholestasis
 to liver:                 [{(cholestasis
                               established-at +2)
                             (can-account-for
                               (bilirubin 6))}|
```

Figure S: Message Trace For Multiple Disease With Additive Symptoms

Thus we have expanded the control aspects to include:

1. message channels: a specialist may send a re-establish message to *any* subordinate specialist.
2. message types:
   a. can-account-for: used to indicate what the specialist once established can account for of the observed patient signs,
   b. re-establish-assume: used to force a specialist to assume a certain patient sign and symptom.

3.4. When Reconsideration Is Called For - Example: Multiple Diseases With Canceling Symptoms

In this section, we consider patient states which can be explained by multiple disease hypotheses in which the presence of both diseases cancel some particular patient sign or symptom.

Consider for example a patient who has both osteo myelitis secondary to a staph aureus infection, and a generalized gram negative sepsis.

The osteo myelitis is typically associated with an increased blood platelet count, while the gram negative sepsis is associated with a decreased platelet count Hence the patient suffering from both diseases will typically have a *normal platelet count.*

A critical point in understanding this situation is to realize that the establishment of any specialist will in general depend on *many* patient signs and symptoms. Thus both the hypotheses osteo myelitis and gram negative sepsis will be established, although both should "point out" that expectations for the patient platelet count are not realized.

The situation here is conceptually very similar to the one discussed in the preceding section on multiple diseases with additive symptoms. There, the specialist reported how much of an observed anomaly could be accounted for; here the specialist must report that some expected sign or symptom is *not* observed. To capture this idea, we introduce the message type

```
(expected-but-not -found
   ...).
```

For example, the specialist representing the osteo myelitis hypothesis should return

```
[(established-at +1)
 (expected-but-not-found
   (plateletCount high))],
```

and the gram negative sepsis specialist should return

```
[(established-at +1)
 (expected-but-not-found
   (plateletCount low))].
```

As in the case with multiple diseases with additive symptoms, once a possible multiple disease hypothesis has been composed by a specialist *above* osteoMyelitis and gramNegSepsis, then both will be forced to re-establish assuming appropriate platelet count. For example, in subsequent processing, osteomyelitis and gramNegSepeis will receive the messages

```
to osteoMyelitis: [re-establish
                      (assume (plateletCount high))],
```

and

```
to gramNegSepsis: [re-establish
                      (assume (plateletCount low))]
```

In this section, to handle cases of multiple diseases with canceling symptoms, we have added to the set of message types the possibility

```
(expected-but-not-found ...)
```

4. Discussion

This paper is based on identifying classification as a generic problem solving activity with a set of control issues that are particular to it. By considering a number of examples, we have shown how we can add control elements to the control structure of MDX, a hierarchical classification-based diagnostic system, to deal with situations where multiple, dependent classificatory hypotheses are appropriate; e.g., in medical diagnosis when a patient has multiple diseases. These extensions have been described in terms of two sets:

- the set of possible paths by which control can go from consideration of one classificatory hypothesis to another,
- the kinds of operations that can be performed on hypotheses and their effects on the control behavior of the problem solver.

The existence of other generic tasks in our framework suggests that a similar analysis of the relevant control issues for them would also be fruitful.

There is another aspect of the thrust of our work that is worth remarking on. The control problem *intrinsic to classification* was defined as that which delineates the possible invocation paths that will be available, and the requests that can be made along those paths. Thus the language of *hypotheses as agents* seemed a natural one since it provided the right level of analysis to capture the essence of classification problem solving. This approach may be more than a metaphor, however. It may provide a way to discuss the *knowledge level architecture* of classification, and by extension, other kinds of generic problem solving as well. If that is true, the analysis method we have utilized would be an alternative to Newell's proposals in [15], in which he is concerned precisely with how an agent can be described *independent of implementation.*

5. Acknowledgments

We wish to thank Drs. Jack Smith and John Svirbely, Pathology Department, Ohio State University, who, by giving us many examples of reasoning in the medical domain, provided the raw data that any theory of classificatory medical diagnosis must account for. Our thanks in addition to Tom Bylander for critical readings of earlier drafts. Finally, we wish to express our thanks to the reviewers for their comments.

References

1. B. Chandrasekaran, F. Gomes, S. Mittal, and J. W. Smith. An Approach to Medical Diagnosis Based on Conceptual Structures. Proceedings of IJCAI 6, 1979.

2. B. Chandrasekaran. Decomposition of Domain Knowledge into Knowledge Sources: The MDX Approach. Proc. 4th Nat. Conf. Canadian Society for Computational Studies of Intelligence, Saskatchewan, May, 1982.

8. B. Chandrasekaran, S. Mittal, and J. W. Smith. Reasoning with Uncertain Knowledge: The MDX Approach. Proc. 1st Ann. Joint Conf. of the American Medical Informatics Association, May, 1982.

4. B. Chandrasekaran. "Towards a Taxonomy of Problem-Solving Types". *AI Magazine 4,* 1 (Winter/Spring 1983), 9-17.

5. B. Chandrasekaran and S. Mittal. Conceptual Representation of Medical Knowledge for Diagnosis by Computer: MDX and Related Systems. In *Advances in Computers,* M. Yovits, Ed., Academic Press, 1983, pp. 217-293.

6. William J. Clancey. NEOMYCIN: Reconfiguring a Rule-Based Expert System for Application To Teaching. Proceedings of IJCAI 7, 1981.

7. William J. Clancey. Classification Problem Solving. Proceedings of the AAAI, 1984, pp. 49-55.

8. Paul J. Feltovich, Paul E. Johnson, James H. Moller, and David B. Swanson. LCS: The Role And Development of Medical Knowledge in Diagnostic Expertise. In William J. Clancey and Edward H. Shortliffe, Ed., *Readings In Medical Artificial Intelligence,* Addison-Wesley, 1984, Chap. 12, pp. 275-319.

0. Michael R. Genesereth. Diagnosis UsingHierarchical Design Models. Proceedings of the AAAI, 1982, pp. 278-283.

10. F. Gomes and B. Chandrasekaran. "Knowledge Organisation and Distribution for Medical Diagnosis". *IEEE Transactions on Systems, Man, and Cybernetics SMC-11,* 1 (January 1981), 34-42.

11. John R. Josephson, B. Chandrasekaran, and J.W. Smith. Assembling the Best Explanation. Proceedings of the IEEE Workshop on Principles of Knowledge-Based Systems, IEEE Computer Society, Dec, 1984.

12. Casimir A. Kulikowski and Sholom M. Weiss. *AAAS Selected Symposia.* Volume 1: Representation of Expert Knowledge for Consultation: The CASNET and EXPERT Projects. In *Artificial Intelligence In Medicine,* Peter Ssolovitis, Ed., Westview Press, 1982, ch. 2, pp. 21-54.

13. D. Marr. "Early Processing of Visual Information". *Phil. Trans. Roy. Soc. 197* (, 1976).

14. S. Mittal. *Design Of A Distributed Medical Diagnosis And Data Base System.* Ph.D. Th., Computer and Information Science Department, The Ohio State University, 1980.

15. A. Newell. "The Knowledge Level". *AI Magazine Summer* (1982), 1-19. Presidential Address.

16. Ramesh S. Patil. *Causal Representation of Patient Illness For Electrolyte And Acid-Base Diagnosis.* Ph.D. Th., M.I.T., October 1981.

17. C.S. Peirce. Abduction and Induction. In Buchler, Ed., *Philosophical Writings of Peirce,* Dover, 1955, Chap. 11, pp. 150ff..

18. H. Pople. On The Mechanisation Of Abductive Logic. Proceedings of IJCAI 8, 1983, pp. 147-152.

10. Jon Sticklen, B. Chandrasekaran, J. W. Smith, and John Svirbely. A Comparison Of The Diagnostic Subsystems of MDX and MYCIN. Workshop On Priciples Of Knowledge Based Systems, IEEE, December, 1984, pp. 205-212.