



DOCUMENT RESUME

ED 242 310

IR 011 066

AUTHOR Clancey, William J.
TITLE The Advantages of Abstract Control Knowledge in Expert System Design. Technical Report #7.
INSTITUTION Stanford Univ., Calif. Dept. of Computer Science.
SPONS AGENCY Office of Naval Research, Arlington, Va. Personnel and Training Research Programs Office.
REPORT NO HPP-83-17; STAN-CS-83-995
PUB DATE Nov 83
CONTRACT N00014-79-0302
NOTE 22p.; For related documents, see IR 011 074 and IR 011 079. Also in the Proceedings of the National Conference on Artificial Intelligence 1983, pages 74-78.
PUB TYPE Reports - Research/Technical (143)
EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Artificial Intelligence; Clinical Diagnosis; Communicable Diseases; *Computer Software; Databases; *Design Requirements; Man Machine Systems; *Problem Solving; Program Effectiveness
IDENTIFIERS Diagnostic Consultation Systems; Domain Specifications; *Intelligent CAI Systems

ABSTRACT

This paper argues that an important design principle for building expert systems is to represent all control knowledge abstractly and separately from the domain knowledge upon which it operates. Abstract control knowledge is defined as the specifications of when and how a program is to carry out its operations, such as pursuing a goal, focusing, acquiring data, and making inferences; domain knowledge is defined as the facts and relations of a knowledge base, such as a knowledge base of medical information. It is noted that a body of abstract control knowledge provides a generic framework for constructing knowledge bases for related problems in other domains, and also provides a useful starting point for studying the nature of strategies. The idea of separating control and domain knowledge is illustrated by discussing knowledge representation on three intelligent computer-aided diagnostic consultation systems, MYCIN, NEOMYCIN, and CENTAUR. The scientific, engineering, and practical benefits of separating control and domain knowledge are outlined and the difficulty of attaining this ideal design is considered. Also provided are a 19-item bibliography and a list of names and addresses of government and private sector research/information centers and personnel concerned with computer-aided instruction. (ESR)

 * Reproductions supplied by EDRS are the best that can be made *
 * from the original document. *

12

ED242310

The Advantages of Abstract Control Knowledge in Expert System Design

by

William J. Clancey

Department of Computer Science

Stanford University
Stanford, CA 94305

U.S. DEPARTMENT OF EDUCATION
NATIONAL INSTITUTE OF EDUCATION
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

This document has been reproduced as received from the person or organization originating it.

Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official NIE position or policy.



R011066

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING	
1. REPORT NUMBER STAN-CS-83-995 Tech. Report #7	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG	
4. TITLE (and Subtitle) The Advantages of Abstract Control Knowledge in Expert System Design		5. TYPE OF REPORT & PERIOD Technical, Nov. 1983	
7. AUTHOR(s) William J. Clancey		6. PERFORMING ORG. REPORT NUMBER STAN-CS-83-995	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of Computer Science Stanford University Stanford, CA 94305		8. CONTRACT OR GRANT NUMBER N00014-79-0302	
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Program Office of Naval Research (Code 458) Arling, VA 22217		10. PROGRAM ELEMENT, PROJECT, AREA & WORK UNIT NUMBERS NR 154-482	
14. MONITORING AGENCY NAME & ADDRESS (if diff. from Controlling Office) ONR Representative - Mr. Robin Simpson Durand Aeronautics Building, Rm. 165 Stanford University, Stanford, CA 94305		12. REPORT DATE November 1983	13. NO. OF PAGES 19
16. DISTRIBUTION STATEMENT (of this report) Approved for public release; distribution unlimited		15. SECURITY CLASS. (of this report) Unclassified	
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from report)		15a. DECLASSIFICATION/OWNGRADING SCHEDULE	
18. SUPPLEMENTARY NOTES Also in the Proceedings of the National Conference on Artificial Intelligence 1983, Pages 74-78. Also in HPP Memo 83-17			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) A poorly designed knowledge base can be as cryptic as an arbitrary program and just as difficult to maintain. Representing control knowledge abstractly, separately from domain facts and relations, makes the design more transparent and explainable. A body of abstract control knowledge provides a generic framework for constructing knowledge bases for related problems in other domains and also provides a useful starting point for studying the nature of strategies.			

DD FORM 1473
1 JAN 73

EDITION OF 1 NOV 65 IS OBSOLETE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

**THE ADVANTAGES OF ABSTRACT CONTROL KNOWLEDGE
IN EXPERT SYSTEM DESIGN**

William J. Clancey

**Department of Computer Science
Stanford University, Stanford CA 94305**

**Contract No. N000C14-79-0302, effective March 15, 1979.
Expiration Date: March 14, 1985
Total Amount of Contract -- \$1,126,897
Principal Investigator, Bruce G. Buchanan (415) 497-0935
Associate Investigator, William J. Clancey (415) 497-1997**

**Sponsored jointly by:
Office of Naval Research and Army Research Institute,
Personnel and Training Research Programs,
Psychological Sciences Division.
Contract Authority No. NR 154-482
Scientific Officers: Dr. Marshall Farr and Dr. Henry Half**

**The views and conclusions contained in this document are those of the authors and should not be
interpret as necessarily representing the official policies, either expressed or implied, of the Office of
Naval Research or the U.S. Government.**

**Approved for public release; distribution unlimited, Reproduction in whole or in part is permitted for
any purpose of the United States Government.**

Abstract

A poorly designed knowledge base can be as cryptic as an arbitrary program and just as difficult to maintain. Representing control knowledge abstractly, separately from domain facts and relations, makes the design more transparent and explainable. A body of abstract control knowledge provides a generic framework for constructing knowledge bases for related problems in other domains and also provides a useful starting point for studying the nature of strategies.

1. Introduction

The quality of a knowledge base depends not only on how well it solves problems, but also how easily its design allows it to be maintained. Easy maintenance--the capability to reliably modify a knowledge base without extensive reprogramming--is important for several reasons:

- Knowledge-based programs are built incrementally, based on many trials, so modification is continually required, including updates based on improved expertise;
- A knowledge base is a repository that other researchers and users may wish to build upon years later;
- A client receiving a knowledge base constructed for him may wish to correct and extend it without the assistance of the original designers.

A knowledge base is like a traditional program in that maintaining it requires having a good understanding of the underlying design. That is, you need to know how the parts of the knowledge base are expected to interact in problem solving. Depending on the representation, this includes knowing how default and judgmental knowledge interact, whether rule clauses can be reordered, when attached procedures are applied, how constraints are inherited and ordered, etc. One way to provide this understanding is to have the program explain its reasoning, using an internal description of its own design (Davis, 1976), (Swartout, 1977). However, problems encountered in understanding traditional programs--poorly structured code, implicit side-effects, and inadequate documentation--carry over to knowledge-based programming and naturally limit the capabilities of explanation programs. For example, a knowledge base might arbitrarily combine reasoning strategies with facts about the domain. Implicit, procedurally-embedded knowledge cannot be articulated by an explanation system (Swartout, 1981), (Clancey, 1983) and is not visible to guide the program maintainer (see (Ennis, 1982) for an entertaining study of this problem).

This paper argues that *an important design principle for building expert systems is to represent all control knowledge abstractly, separate from the domain knowledge it operates upon.* This idea is illustrated with examples from the NEOMYCIN system (Clancey, 1981). There are many scientific,

engineering, and practical benefits. The difficulty of attaining this ideal design is also considered.

2. What is Abstract Control Knowledge?

"Control knowledge" specifies when and how a program is to carry out its operations, such as pursuing a goal, focusing, acquiring data, and making inferences. A basic distinction can be made between the facts and relations of a knowledge base and the program operations that act upon it. For example, facts and relations in a medical knowledge base might include (expressed in a predicate calculus formulation):

(SUBTYPE INFECTION MENINGITIS)
-- "meningitis is a kind of infection"

(CAUSES INFECTION FEVER)
-- "infection causes fever"

(CAUSES INFECTION SHAKING-CHILLS)
-- "infection causes shaking chills"

(DISORDER MENINGITIS)
-- "meningitis is a disorder"

(FINDING FEVER)
-- "fever is a finding"

Such a knowledge base might be used to provide consultative advice to a user, in a way typical of expert systems (Duda and Shortliffe, 1983). Consider, for example, a consultation system for diagnosing some faulty device. One typical program operation is to select a finding that causes a disorder and ask the user to indicate whether the device being diagnosed exhibits that symptom. Specifically, a medical diagnostic system might ask the user whether the patient is suffering from shaking chills, in order to determine whether he has an infection. The first description of the program's operation is *abstract*, referring only to domain-independent relations like "finding" and "causes"; the second description is *concrete*, referring to domain-dependent terms like "shaking-chills" and "infection". ("Domain-independent" doesn't mean that it applies to every domain, just that the term is not specific to any one domain.)

The operation described here can be characterized abstractly as "attempting to confirm a diagnostic hypothesis" or concretely as "attempting to determine whether the patient has an infection." Either description indicates the *strategy* that motivates the question the program is asking of the user. So in this example we see how a strategy, or control knowledge, can be stated either abstractly or concretely. The following two examples illustrate how both forms of control knowledge

might be represented in a knowledge base.

2.1 An Implicit Refinement Strategy

In MYCIN (Shortliffe, 1976), most knowledge is represented as domain-specific rules. For example, the rule "If the patient has an infection and his CSF cell count is less than 10, then it is unlikely that he has meningitis," might be represented as:

PREMISE:

```
($AND (SAME CNTXT INFECTION)
      (ILESSP (VAL1 CNTXT CSFCELLCOUNT) 10))
```

ACTION:

```
(CONCLUDE CNTXT INFECTION TYPE MENINGITIS TALLY -700)
```

The order of clauses is important here, for the program should not consider the "CSF cell count" if the patient does not have an infection. Such clause ordering in all rules ensures that the program proceeds by top-down refinement from infection to meningitis to subtypes of meningitis. The disease hierarchy cannot be stated explicitly in the MYCIN rule language; it is implicit in the design of the rules. (See (Clancey, 1983) for further analysis of the limitations of MYCIN's representation.)

CENTAUR (Aikins, 1980) is a system in which disease hierarchies are explicit. In its representation language, MYCIN's meningitis knowledge might be encoded as follows (using a LISP property list notation):

```
INFECTION
  MORE-SPECIFIC ((disease MENINGITIS)
                (disease BACTEREMIA)...)
  IF-CONFIRMED (DETERMINE disease of INFECTION)

MENINGITIS
  MORE-SPECIFIC ((subtype BACTERIAL)
                (subtype VIRAL)...)
  IF-CONFIRMED (DETERMINE subtype of MENINGITIS)
```

In CENTAUR, hierarchical relations among disorders are explicit (meningitis is a specific kind of infection), and the strategies for using the knowledge are domain-specific (after confirming that the patient has an infection, determine what more specific disease he has). This design enables CENTAUR to articulate its operations better than MYCIN, whose hierarchical relations and strategy are procedurally embedded in rules.

However, observe that each node of CENTAUR's hierarchy essentially repeats a single strategy--try

to confirm the presence of a child disorder--and the overall strategy of top-down refinement is not explicit. Aikins has *labeled* CENTAUR's strategies, but has not stated them abstractly. By representing strategies abstractly, it is possible to have a more explicit and non-redundant design. This is what is done in NEOMYCIN.

In NEOMYCIN domain relations and strategy are represented *separately* and strategy is represented abstractly. A typical rule that accomplishes, in part, the abstract task of attempting to confirm a diagnostic hypothesis and its subtypes is shown below.

<Domain Knowledge>

INFECTION
CAUSAL-SUBTYPES (MENINGITIS BACTEREMIA ...)

MENINGITIS
CAUSAL-SUBTYPES (BACTERIAL VIRAL ...)

<Abstract Control Knowledge>

TASK: EXPLORE-AND-REFINE
ARGUMENT: CURRENT-HYPOTHESIS

METARULE001

IF the hypothesis being focused upon
has a child
that has not been pursued,
THEN pursue that child.

(IF (AND (CURRENT-ARGUMENT \$CURFOCUS)
(CHILDOF \$CURFOCUS \$CHILD)
(THNOT (PURSUED \$CHILD)))
(NEXTACTION (PURSUE-HYPOTHESIS \$CHILD)))

NEOMYCIN uses a deliberation/action loop, for deducing what it should do next. *Metarules*, like the one shown above, recommend what task should be done next, what domain rule applied, or what domain finding requested from the user (details are given in (Clancey, 1981) and (Clancey and Bock, 1982) and are not important here). The important thing to notice is that this metarule will be applied for refining any disorder, obviating the need to "compile" redundantly into the domain hierarchy of disorders how it should be searched. When a new domain relation is declared (e.g., a new kind of infection is added to the hierarchy) the abstract control knowledge will use it appropriately. That is, we separate out what the domain knowledge is from how it should be used.

Metarules were first introduced for use in expert systems by Davis (Davis, 1976), but he conceived

of them as being domain-specific. In that form, principles are encoded redundantly, just like CENTAUR's control knowledge. For example, the principle of pursuing common causes before unusual causes appears as specific metarules for ordering the domain rules of each disorder.

The benefits of stating metarules abstractly are illustrated further by a second example.

2.2. An Implicit Question-Asking Strategy

Another reason for ordering clauses in a system like MYCIN is to prevent unnecessary requests for data. A finding might be deduced or ruled out from other facts available to the program. For example, the rule "If the patient has undergone surgery and neurosurgery, then consider diplococcus as a cause of the meningitis" might be represented as follows.

```
PREMISE: ($AND (SAME CNTXT SURGERY)
              (SAME CNTXT NEUROSURGERY))
ACTION: (CONCLUDE CNTXT COVERFOR DIPLOCOCCUS TALLY 400)
```

We say that the surgery clause "screens" for the relevance of asking about neurosurgery. Observe that neither the relation between these two findings (that neurosurgery is a *type of* surgery) nor the strategy of considering a general finding in order to rule out one of its subtypes is explicit. An alternative way used in MYCIN for encoding this knowledge is to have a separate "screening" rule that at least makes clear that these two findings are related: "If the patient has not undergone surgery, then he has not undergone neurosurgery."

```
PREMISE: ($AND (NOTSAME CNTXT SURGERY))
ACTION: (CONCLUDE CNTXT NEUROSURGERY YES TALLY -1000)
```

Such a rule obviates the need for a "surgery" clause in every rule that mentions neurosurgery, so this design is more elegant and less prone to error. However, the question-ordering strategy and the abstract relation between the findings are still not explicit. Consequently, the program's explanation system cannot help a system maintainer understand the underlying design.

In NEOMYCIN, the above rule is represented abstractly by a metarule for the task of finding out new data.

<Domain Knowledge>

(SUBSUMES SURGERY NEUROSURGERY)
(SUBSUMES SURGERY CARDIACSURGERY)

<Abstract Control Knowledge>

TASK: FINDOUT
ARGUMENT: DESIRED-FINDING

METARULE002

IF the desired finding
is a subtype of a class of findings and
the class of findings is not present in this case,
THEN conclude that the desired finding is not present.

```
(IF (AND (CURRENT-ARGUMENT $SUBTYPE)  
(SUBSUMES $CLASS $SUBTYPE)  
(THNOT (SAMEP CNTXT $CLASS)))  
(NEXTACTION  
(CONCLUDE CNTXT $SUBTYPE 'YES TALLY -1000)))
```

This metarule is really an *abstract generalization* of all screening rules. Factoring out the statement of relations among findings from how those relations are to be used produces an elegant and economical representation. Besides enabling more-detailed explanation, such a design makes the system easier to construct and more robust.

Consider the multiple ways in which a single relation between findings can be used. If we are told that the patient has neurosurgery, we can use the subsumption link (or its inverse) to conclude that the patient has undergone surgery. Or if we know that the patient has not undergone any kind of surgery we know about, we can use the "closed world assumption" and conclude that the patient has not undergone surgery. These inferences are controlled by abstract metarules in NEOMYCIN.

The knowledge base is easier to construct because the expert needn't specify every situation in which a given fact or relation should be used. New facts and relations can be added in a simple way; the abstract metarules explicitly state how the relations will be used. The same generality makes the knowledge base more robust. The system is capable of making use of facts and relations for different purposes, perhaps in combinations that would be difficult to anticipate or enumerate.

3. Studying Abstract Strategies and Structural Relations

In NEOMYCIN, domain findings and disorders are related in the way shown above, and there are approximately 75 metarules that constitute a procedure for doing diagnosis. Besides abstract domain relations, such as SUBSUMES, NEOMYCIN's metarules reference:

- Knowledge about metarules and tasks: (static) the argument of a task, whether metarules are to be applied iteratively, the condition under which a task should be aborted, (dynamic) whether a task completed successfully, whether a metarule succeeded or failed, etc.
- Domain problem-solving history: the active hypotheses, whether a hypothesis was pursued, cumulative belief for a hypothesis, rules using a finding that are "in focus", a strong competitor to a given hypothesis, etc.

These concepts form the vocabulary for a model of diagnosis, the terms in which expert behavior is interpreted and strategies are expressed.

An unexpected effect is that there is no more backward chaining at the domain level. That is, the only reason MYCIN does backward chaining during its diagnostic (history and physical) phase is to accomplish top-down refinement and to apply screening rules. This is an important result. By studying the hundreds of rules in the MYCIN system, factoring out domain relations from control knowledge, we have greatly deepened our understanding of the knowledge encoded in the rules. There are two specific products: a *body of abstract control knowledge* that can itself be studied, as well as applied in other problem domains, and a *language for representing knowledge about disorders* (in terms of causality, subtype, etc.). We call these abstract relations *structural relations*.

Structural relations are a means for indexing domain-specific knowledge: They *select* hypotheses to focus upon, findings to request, and domain inferences that might be made. As such, structural relations constitute the organization, the *access paths*, by which strategies bring domain-specific knowledge into play. For example, the metarules given above mention the CHILDOF and SUBSUMES relations. METARULE001 looks for *the children of* the current hypothesis in order to pursue them; METARULE002 looks for *a more general finding* in order to ask for it first.

These relations constitute the language by which the primitive domain concepts (particular findings and disorder hypotheses) are related in a network. *Adding a new strategy often requires adding a new kind of structural relation to the network.* For example, suppose we desire to pursue common causes of a disorder before serious, but unusual causes. We must partition the causes of any disorder according to this distinction, adding new relations to our language--COMMON-CAUSES and

SERIOUS CAUSES.

Similarly, the applicability of a strategy depends on the domain. For example, a strategy might give a problem domain all findings might be equal with how to search a deep hierarchy of deep or shallow, making the strategies inapplicable to explicate structural relations. On this basis, a set of strategies, referring to structural properties

Lenat has found a similar relationship between his program for discovering new heuristic heuristics in EURISKO depends on breaking slots that the program can inspect and modify concepts whose representation is refined (FACTORS slot to every number). The program that collects relations among entries of an

4. Given the Benefits, Can it

An initial reaction might be that for some abstract strategies--all facts and relations and the procedure for confirming any given discovery (some device) might be completely situation-specific. It would appear to be an unusual kind of domain where principles can be applied over and over again.

Teaching and learning are made incredibly difficult from problem to another. Domains with a strong structure might be like this. Perceptual skills rely on a large amount of data; they are not good candidates for

We also know that in many domains, for solving routine problems. These procedures are important for organization management, equipment control, etc. It is important to recognize that these procedures

causal, temporal, and spacial interactions, problem-solving goals, abstract principles of design, diagnosis, etc. Except where a procedure is arbitrary, there must be some underlying rationale for the selection and ordering of its steps. Knowing this rationale is certainly important for reliably modifying the procedure; such procedures are often just prepared plans that an expert (or a user following a program's advice) may need to adapt to unusual circumstances. At one level, the rationale can be made explicit in terms of an abstract plan with its attendant domain structural relations; a redundant, compiled form can be used for efficient routine problem solving.

In theory, if the rationale for a procedure or prepared plan can be made explicit, a program can reconstruct the procedure from first principles. This approach has two basic difficulties. First, the procedure might have been learned incrementally from case experience. It simply handles problems well; there is no compiled-out theory that can be articulated. This problem arises particularly for skills in which behavior has been shaped over time, or for any problem in which the trace of "lessons" has been poorly recorded. The second difficulty is that constructing a procedure from first principles can involve a great deal of search. Stefik's (Stefik, 1980) multi-leveled planning regime for constructing MOLGEN experiments testifies to the complexity of the task and the limited capabilities of current programs. In contrast, Friedland's (Friedland, 1979) approach of constructing experiment plans from skeletal, abstract plans trades flexibility for efficiency and resemblance to human solutions. While skeletal plans may sometimes use domain-specific terms, as precompiled abstract procedures they are analogous to NEOMYCIN's tasks.

Importantly, the *rationale for the abstract plan* itself is not explicit in any of these programs. For example, NEOMYCIN's metarules for a given task might be ordered by preference (alternative methods to accomplish the same operation) or as steps in a procedure. Since the constraints that suggest the given ordering are not explicit, part of the design of the program is still not explicit. For example, the abstract steps of top-down refinement are now stated, but the sense in which they constitute this procedure is not represented. (Why should pursuing siblings of a hypothesis be done before pursuing children?) As another example, the task of "establishing the hypothesis space" by expanding the set of possibilities beyond common, expected causes and then narrowing down in a refinement phase has mathematical, set-theoretic underpinnings that are not explicit in the program. Similarly, Stefik's abstract planning procedure of "least-commitment" is implicit in numeric priorities assigned to plan design operators (Clancey, 1983). Automatically constructing procedures at this high level of abstraction, as opposed to implicitly building them into a program, has been explored very little.

Even within the practical bounds of what we make explicit, it might be argued that representing

procedures abstractly is much more difficult than stating individual situation-specific rules. This might differ from person to person; certainly in medicine some physicians are better than others at stating how they reason abstractly. A good heuristic might be to work with good teachers, for they are most likely to have extracted the principles so they can be taught to students.

There is certainly an initial cost whose benefit is unlikely to be realized if no explanation facility is desired, only the original designers maintain or modify the knowledge base, or there is no desire to build a generic system. But even this argument is dubitable: a knowledge base with embedded strategies can appear cryptic to even the original designers after it has been left aside for a few months. Also, anyone intending to build more than one system will benefit from expressing knowledge as generally as possible so that lessons about structure and strategy can speed up the building of new systems.

The cost aside, it appears that there is no way to get strategic explanations without making domain relations explicit and stating strategies separately. This was the conclusion of Swartout, who was led to conclude that an automatic programming approach, as difficult as it first seemed, was a natural, direct way to ensure that the program had knowledge of its own design (Swartout, 1981). That is, providing complete explanations means understanding the design well enough to derive the procedures yourself.

NEOMYCIN's factoring of knowledge into domain and strategic knowledge bases is comparable to the input requirements of Swartout's automatic programming system. However, NEOMYCIN interprets its domain knowledge, rather than instantiating its abstract strategies in a compiled program. (Maintaining the separation is important so the metarules can be used in student modeling (London and Clancey, 1982).) Moreover, NEOMYCIN's strategies are abstract, unlike the domain-specific "principles" used in Swartout's program. This design decision was originally motivated by our desire to replicate the kind of explanations given by teachers (Hasling, 1983). However, we now realize that representing control knowledge abstractly has engineering and scientific benefits as well.

5. Advantages of the Approach

The advantages of representing control knowledge abstractly can be summarized according to engineering, scientific, and practical benefits:

• Engineering.

- The explicit design is easier to debug and modify. Hierarchical relations among findings and hypotheses and search strategies are no longer procedurally

embedded in rules.

- Knowledge is represented more generally, so we get more performance from less system-building effort. We don't need to specify every situation in which a given fact should be used.
- The body of abstract control knowledge can be applied to other problems, constituting the basis of a generic system, for example, a tool for building consultation programs that do diagnosis.
- **Science.** Factoring out control knowledge from domain knowledge provides a basis for studying the nature of strategies. Patterns become clear, revealing, for example, the underlying structural bases for backward chaining. Comparisons between domains can be made according to whether a given relation exists or a strategy can be applied.
- **Practice.**
 - A considerable savings in storage is achieved if abstract strategies are available for solving problems. Domain-specific procedures for dealing with all possible situations need to be compiled in advance.
 - Explanations can be more detailed, down to the level of abstract relations and strategies, so the program can be evaluated more thoroughly and used more responsibly.
 - Because strategies are stated abstractly, the program can recognize the application of a particular strategy in different situations. This provides a basis for explanation by analogy, as well as recognizing plans during knowledge acquisition or student modelling.

Representing control knowledge abstractly moves us closer to our ideal of specifying to a program WHAT problem to solve versus HOW to solve the problem (Feigenbaum, 1977). Constructing a knowledge base becomes a matter of declaring knowledge relations. HOW the knowledge will be used needn't be simultaneously and redundantly specified.

An analogy can be made with GUIDON (Clancey, 1979) (Clancey, 1982), whose body of abstract teaching rules make the program usable with multiple domains. Traditional CAI programs are specific to particular problems (not just problem domains) and have both subject matter expertise and teaching strategies embedded within them. The separation of these in GUIDON, and now the abstract representation of strategies in NEOMYCIN, is part of the logical progression of expert systems research that began with separation of the interpreter from the knowledge base in MYCIN. The trend throughout has been to state domain-specific knowledge more declaratively and to generalize the procedures that control its application.

Another analogy can be made with database systems that combine relational networks with logic programming (e.g., see (Nicolas, 1977)). To conserve space, it is not practical to explicitly store every relation among entities in a database. For example, a database about a population of a country might record just the parents of each person (e.g., (MOTHEROF \$CHILD \$MOTHER) and (FATHEROF \$CHILD \$FATHER)). A separate body of *general derivation axioms* is used to retrieve other relations (the *intensional database*). For example, siblings can be computed by the rule:

```
(IF (AND (PERSON $PERSON)
         (MOTHEROF $PERSON $MOTHER)
         (PERSON $PERSON2)
         (MOTHEROF $PERSON2 $MOTHER))
    (SIBLING $PERSON $PERSON2))
```

Such a rule is quite similar to the abstract metarules that NEOMYCIN uses for deducing the presence or absence of findings. NEOMYCIN differs from database systems in that its rules are grouped and controlled to accomplish abstract tasks. Only a few of NEOMYCIN's metarules make inferences about database relations; most invoke other tasks, such as "ask a general question" and "group and differentiate hypotheses." Moreover, the knowledge base contains judgmental rules of evidence for the disorder hypotheses. These differences aside, the analogy is stimulating. It suggests that treating a knowledge base as an object to be inspected, reasoned about, and manipulated by *abstract procedures*--as a database is checked for integrity, queried, and extended by general axioms--is a powerful design principle for building expert systems.

Acknowledgments

This research has been supported in part by ONR and ARI Contract N00014-79C-0302. Computational resources have been provided by the SUMEX-AIM facility (NIH grant RR00785).

References

- Aikins J. S. *Representation of control knowledge in expert systems*, in *Proceedings of the First AAAI*, pages 121-123, 1980.
- Clancey, W. J. Tutoring rules for guiding a case method dialogue. *The International Journal of Man-Machine Studies*, 1979, 11, 25-49. (Also in Sleeman and Brown (editors), *Intelligent Tutoring Systems*, Academic Press, 1982).
- Clancey, W. J. and Letsinger, R. *NEOMYCIN: Reconfiguring a rule-based expert system for application to teaching*, in *Proceedings of the Seventh IJCAI*, pages 829-836, 1981. (Revised version to appear in Clancey and Shortliffe (editors), *Readings in medical artificial intelligence: The first decade*, Addison-Wesley, 1983).
- Clancey, W. J. GUIDON. In Barr and Feigenbaum (editors), *The Handbook of Artificial Intelligence*, chapter Applications-oriented AI research: Education. William Kaufmann, Inc., Los Altos, 1982. (Revised version to appear in the *Journal of Computer-Based Instruction*, 1983).
- Clancey, W. J. The epistemology of a rule-based expert system: A framework for explanation. *Artificial Intelligence*, 1983, 20(3), 215-251. (Also to appear in Buchanan and Shortliffe (editors), *Rule-based expert systems: The MYCIN experiments of the Stanford Heuristic Programming Project*, Addison-Wesley, 1983).
- Clancey, W. J. and Bock, C. *MRS/NEOMYCIN: Representing metacontrol in predicate calculus*. HPP Memo 82-31, Stanford University, November 1982.
- Davis R. *Applications of meta-level knowledge to the construction, maintenance, and use of large knowledge bases*. HPP Memo 76-7 and AI Memo 283, Stanford University, July 1976.
- Duda, R. O. and Shortliffe, E. H. Expert systems research. *Science*, 1983, 220, 261-268.
- Ennis, S. P. *Expert systems: A user's perspective of some current tools*, in *Proceedings of the Second AAAI*, pages 319-321, August, 1982.
- Feigenbaum, E. A. *The art of artificial intelligence: I. Themes and case studies of knowledge engineering*; in *Proceedings of the Fifth IJCAI*, pages 1014-1029, August, 1977.
- Friedland, P. *Knowledge-based experiment design in molecular genetics*, in *Proceedings of the Sixth IJCAI*, pages 285-287, 1979.
- Hasling, D. W. *Abstract explanations of strategy in a diagnostic consultation system*. (To appear in the *Proceedings of AAAI-83*).
- Lenat, D. B. The nature of heuristics. *Artificial Intelligence*, 1982, 19(2), 189-249.
- London, B. and Clancey, W. J. *Plan recognition strategies in student modeling: prediction and*

description, in *Proceedings of the Second AAAI*, pages 335-338, 1982.

Nicolas, J. M. and Gallaire, H. Data base: Theory vs. interpretation. In H. Gallaire and J. Minker (editors), *Logic and data bases*, pages 33-54. Plenum Press, New York, 1977.

Shortliffe, E. H. *Computer-based medical consultations: MYCIN*. New York: Elsevier 1976.

Stefik, M. J. *Planning with constraints*. PhD thesis, Computer Science Department, Stanford University, 1980.

Swartout, W. R. *A digitalis therapy advisor with explanations*. Technical report 176, MIT Laboratory for Computer Science, February 1977.

Swartout W. R. *Explaining and justifying in expert consulting programs*, in *Proceedings of the Seventh IJCAI*, August, 1981. (Also to appear in Ciancay and Shortliffe (editors), *Readings in medical artificial intelligence: The first decade*; Addison-Wesley, 1983).

Navy

- 1 Robert Ahlers
Code M711
Human Factors Laboratory
NAVTMAEQUIPCEN
Orlando, FL 32813
- 1 Dr. EJ Alken
Navy Personnel RSD Center
San Diego, CA 92152
- 1 Dr. Meryl S. Baker
Navy Personnel RSD Center
San Diego, CA 92152
- 1 Dr. Robert Blanchard
Navy Personnel RSD Center
San Diego, CA 92152
- 1 Liaison Scientist
Office of Naval Research
Branch Office, London
Box 39
FPO New York, NY 09510
- 1 Dr. Richard Cantone
Navy Research Laboratory
Code 7510
Washington, DC 20375
- 1 Dr. Stanley Collyer
Office of Naval Technology
800 N. Quincy Street
Arlington, VA 22217
- 1 CDR Mike Curran
Office of Naval Research
800 N. Quincy St.
Code 270
Arlington, VA 22217
- 1 Dr. Jude Franklin
Code 7510
Navy Research Laboratory
Washington, DC 20375
- 1 Dr. Mike Gaynor
Navy Research Laboratory
Code 7510
Washington, DC 20375
- 1 LT Steven D. Harrie, MSC, USN
RFD 1, Box 243
Blair, VA 24149
- 1 Dr. Jim Mollan
Code 14
Navy Personnel R & D Center
San Diego, CA 92152
- 1 Dr. Ed Hutchins
Navy Personnel RSD Center
San Diego, CA 92152
- 1 Dr. Norman J. Kerr
Chief of Naval Technical Training
Naval Air Station Memphis (75)
Millington, TN 38054
- 1 Dr. Peter Kincaid
Training Analysis & Evaluation Group
Dept. of the Navy
Orlando, FL 32813
- 1 Dr. William L. Maloy (02)
Chief of Naval Education and Training
Naval Air Station
Pensacola, FL 32508
- 1 CAPT Richard L. Martin, USN
Commanding Officer
USS Carl Vinson (CVN-70)
FPO New York, NY 09558

- 1 Dr. Joe McLachlan
Navy Personnel RSD Center
San Diego, CA 92152
- 1 Dr. George Moeller
Director, Behavioral Sciences Dept.
Naval Submarine Medical Research Lab
Naval Submarine Base
Groton, CT 06349
- 1 Dr. William Montague
NPRDC, Code 13
San Diego, CA 92152
- 1 Library, Code P201L
Navy Personnel RSD Center
San Diego, CA 92152
- 1 Technical Director
Navy Personnel RSD Center
San Diego, CA 92152
- 6 Commanding Officer
Naval Research Laboratory
Code 2627
Washington, DC 20390
- 1 Office of Naval Research
Code 433
800 N. Quincy Street
Arlington, VA 22217
- 6 Personnel & Training Research Group
Code 442PT
Office of Naval Research
Arlington, VA 22217
- 1 Office of the Chief of Naval Operations
Research Development & Studies Branch
OP 115
Washington, DC 20350
- 1 LT Frank C. Petho, MSC, USN (Ph.D.)
CNET (N-432)
NAS
Pensacola, FL 32505
- 1 Gary Poock
Operations Research Department
Code 5597
Naval Postgraduate School
Monterey, CA 93940
- 1 Dr. Robert G. Smith
Office of Chief of Naval Operations
OP-987H
Washington, DC 20350
- 1 Dr. Alfred P. Smode, Director
Training Analysis & Evaluation Group
Dept. of the Navy
Orlando, FL 32813
- 1 Dr. Richard Sorensen
Navy Personnel RSD Center
San Diego, CA 92152
- 1 Dr. Frederick Steinhafer
CNO - OP115
Navy Annex
Arlington, VA 20370
- 1 Roger Weissinger-Baylon
Department of Administrative Sciences
Naval Postgraduate School
Monterey, CA 93940
- 1 Mr John E. Wolfe
Navy Personnel RSD Center
San Diego, CA 92152
- 1 Dr. Wallace Wulfeck, III
Navy Personnel RSD Center
San Diego, CA 92152

Marine Corps

- 1 H. Willies Greenup
Education Advisor (2031)
Education Center, MCDEC
Quantico, VA 22134
- 1 Special Assistant for Marine
Corps Matters
Code 100M
Office of Naval Research
800 N. Quincy St.
Arlington, VA 22217
- 1 DR. A.L. SLAFKOSKY
SCIENTIFIC ADVISOR (CODE RD-1)
HQ, U.S. MARINE CORPS
WASHINGTON, DC 20380

Army

- 1 Technical Director
U. S. Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22339
- 1 Dr. Beatrice J. Parr
U. S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333
- 1 Dr. Harold P. O'Neill, Jr.
Director, Training Research Lab
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333
- 1 Commander, U.S. Army Research Institute
for the Behavioral & Social Sciences
ATTN: PERI-32 (Dr. Judith Orsatti)
5001 Eisenhower Avenue
Alexandria, VA 20333
- 1 Joseph Peotks, Ph.D.
ATTN: PERI-1C
Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333
- 1 Dr. Robert Sasmor
U. S. Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333
- 1 Dr. Robert Wisler
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Air Force

- 1 U.S. Air Force Office of Scientific
Research
Life Sciences Directorate, V
Bolling Air Force Base
Washington, DC 20332
- 1 Mr. Raymond E. Christel
AFRL/PHO
Brooks AFB, TX 78235
- 1 Bryan Dallas
AFRL/LRT
Lowry AFB, CO 80230

BEST COPY AVAILABLE

1 Dr. Genevieve Haddad
Program Manager
Life Sciences Directorate
AFOSR
Bolling AFB, DC 20332

1 Dr. John Teaguey
AFOSR/ML
Bolling AFB, DC 20332

1 Dr. Joseph Yasutake
AFHRL/LRT
Lowry AFB, CO 80230

Department of Defense

12 Defense Technical Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC

1 Military Assistant for Training and
Personnel Technology,
Office of the Under Secretary of Defense
for Research & Engineering
Room 3D129, The Pentagon
Washington, DC 20301

1 Major Jack Thorpe
DARPA
1400 Wilson Blvd.
Arlington, VA 22209

Civilian Agencies

1 Dr. Patricia A. Butler
NIE-38N Bldg, Stop # 7
1200 19th St., NW
Washington, DC 20208

1 Dr. Susan Chipman
Learning and Development
National Institute of Education
1200 19th Street NW
Washington, DC 20208

1 Edward Esty
Department of Education, OERI
MS 40
1200 19th St., NW
Washington, DC 20208

1 Dr. Arthur Melmed
724 Brown
U. S. Dept. of Education
Washington, DC 20208

1 Dr. Andrew R. Molner
Office of Scientific and Engineering
Personnel and Education
National Science Foundation
Washington, DC 20550

1 Dr. Everett Palmer
Mail Stop 239-3
NASA-Ames Research Center
Moffett Field, CA 94035

1 Dr. Mary Stodderd
C 12, Mail Stop B296
Los Alamos National Laboratories
Los Alamos, NM 87545

1 Chief, Psychological Research Branch
U. S. Coast Guard (G-P-1/2/TP42)
Washington, DC 20593

1 Dr. Edward C. Weiss
National Science Foundation
1800 G Street, NW
Washington, DC 20550

1 Dr. Frank Withrow
U. S. Office of Education
400 Maryland Ave. SW
Washington, DC 20202

1 Dr. Joseph L. Young, Director
Memory & Cognitive Processes
National Science Foundation
Washington, DC 20550

Private Sector

1 Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

1 Dr. Patricia Baggett
Department of Psychology
University of Colorado
Boulder, CO 80309

1 Mr. Avron Barr
Department of Computer Science
Stanford University
Stanford, CA 94305

1 Dr. Manucha Birenbaum
School of Education
Tel Aviv University
Tel Aviv, Ramat Aviv 69978
Israel

1 Dr. John Black
Yale University
Box 11A, Yale Station
New Haven, CT 06520

1 Dr. John S. Brown
XEROX Palo Alto Research Center
3133 Coyote Road
Palo Alto, CA 94304

1 Dr. Glenn Bryan
6208 Poe Road
Bethesda, MD 20817

1 Dr. Jaime Carbonell
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

1 Dr. Pat Carpenter
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

1 Dr. William Chase
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

1 Dr. Micheline Chi
Learning E & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 Dr. Michael Cole
University of California
at San Diego
Laboratory of Comparative
Human Cognition - 0003A
La Jolla, CA 92093

1 Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

1 Dr. Lynn A. Cooper
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 ERIC Facility-Acquisitions
4833 Rugby Avenue
Bethesda, MD 20814

1 Professor Ruvan Feuerstein
EWERI Ramot Karmon 6
Bet Hakarem
Jerusalem
Israel

1 Mr. Wallace Faurzeig
Department of Educational Technology
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02238

1 Dr. Dexter Fletcher
WICAT Research Institute
1875 S. State St.
Orem, UT 22333

1 Dr. John E. Fredariken
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

1 Dr. Don Gentner
Center for Human Information Processing
University of California, San Diego
La Jolla, CA 92093

1 Dr. Dedra Gentner
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02138

1 Dr. Robert Glaser
Learning Research & Development Center
University of Pittsburgh
3939 O'Hara Street
PITTSBURGH, PA 15260

1 Dr. Marvin D. Glock
217 Stone Hall
Cornell University
Ithaca, NY 14853

1 Dr. Joseph Goguen
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

1 Dr. Daniel Gopher
Faculty of Industrial Engineering
& Management
TECHION
Haifa 32000
ISRAEL

1 DR. JAMES G. GREENO
LRDC
UNIVERSITY OF PITTSBURGH
3939 O'HARA STREET
PITTSBURGH, PA 15213

1 Dr. Barbara Hayes-Roth
Department of Computer Science
Stanford University
Stanford, CA 94305

1 Dr. Earl Hunt
Dept. of Psychology
University of Washington
Seattle, WA 98105

1 Dr. Marcel Just
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

- 1 Dr. David Kieras
Department of Psychology
University of Arizona
Tucson, AZ 85721
- 1 Dr. Walter Rentsch
Department of Psychology
University of Colorado
Boulder, CO 80302
- 1 Dr. Stephen Kosslyn
1236 William James Hall
33 Kirkland St.
Cambridge, MA 02138
- 1 Dr. Pat Langley
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
- 1 Dr. Jill Larkio
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213
- 1 Dr. Alan Lesgold
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260
- 1 Dr. Jim Levin
University of California
at San Diego
Laboratory of Comparative
Human Cognition - 0903A
La Jolla, CA 92093
- 1 Dr. Michael Levine
Department of Educational Psychology
210 Education Bldg.
University of Illinois
Champaign, IL 61801
- 1 Dr. Marcie C. Linn
Lawrence Hall of Science
University of California
Berkeley, CA 94720
- 1 Dr. Don Lyon
AFRL/OT (UDRI)
Williams AFB, AZ 85225
- 1 Dr. Jay McClelland
Department of Psychology
MIT
Cambridge, MA 02139
- 1 Dr. James E. Miller
Computer Thought Corporation
1721 West Plano Highway
Plano, TX 75075
- 1 Dr. Mark Miller
Computer Thought Corporation
1721 West Plano Highway
Plano, TX 75075
- 1 Dr. Tom Moran
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304
- 1 Dr. Allen More
Behavioral Technology Laboratories
1845 Elena Ave., Fourth Floor
Redondo Beach, CA 90277
- 1 Dr. Donald A. Norman
Cognitive Science, C-015
Univ. of California, San Diego
La Jolla, CA 92093
- 1 Dr. Jesse Orlesky
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA 22311
- 1 Prof. Seymour Papert
20C-139
Massachusetts Institute of Technology
Cambridge, MA 02139
- 1 Dr. Nancy Pennington
University of Chicago
Graduate School of Business
1101 E. 58th St.
Chicago, IL 60637
- 1 DR. PETER POLSON
DEPT. OF PSYCHOLOGY
UNIVERSITY OF COLORADO
BOULDER, CO 80309
- 1 Dr. Fred Rief
Physics Department
University of California
Berkeley, CA 94720
- 1 Dr. Lauren Resnick
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15211
- 1 Dr. Jeff Richardson
Denver Research Institute
University of Denver
Denver, CO 80208
- 1 Mary S. Riley
Program in Cognitive Science
Center for Human Information Processing
University of California, San Diego
La Jolla, CA 92093
- 1 Dr. Andrew M. Rose
American Institutes for Research
1055 Thomas Jefferson St. NW
Washington, DC 20007
- 1 Dr. Ernst Z. Rothkopf
Bell Laboratories
Murray Hill, NJ 07974
- 1 Dr. William B. Rouse
Georgia Institute of Technology
School of Industrial & Systems
Engineering
Atlanta, GA 30332
- 1 Dr. David Rueschert
Center for Human Information Processing
Univ. of California, San Diego
La Jolla, CA 92093
- 1 Dr. Michael J. Samet
Perceptronics, Inc
6271 Verbal Avenue
Woodland Hills, CA 91364
- 1 Dr. Roger Schank
Yale University
Department of Computer Science
P.O. Box 2158
New Haven, CT 06520
- 1 Dr. Walter Schneider
Psychology Department
603 E. Daniel
Champaign, IL 61820
- 1 Dr. Alan Schoenfeld
Mathematics and Education
The University of Rochester
Rochester, NY 14627
- 1 Mr. Colin Sheppard
Applied Psychology Unit
Admiralty Marine Technology Est.
Teddington, Middlesex
United Kingdom
- 1 Dr. E. Wallace Sineko
Program Director
Manpower Research and Advisory Services
Smithsonian Institution
801 North Pitt Street
Alexandria, VA 22314
- 1 Dr. Edward L. Smith
Bolt Beranek & Newman, Inc.
30 Moulton Street
Cambridge, MA 02138
- 1 Dr. Richard Snow
School of Education
Stanford University
Stanford, CA 94305
- 1 Dr. Elliott Soloway
Yale University
Department of Computer Science
P.O. Box 2158
New Haven, CT 06520
- 1 Dr. Kathryn T. Spoehr
Psychology Department
Brown University
Providence, RI 02912
- 1 Dr. Robert Sternberg
Dept. of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520
- 1 Dr. Albert Stevens
Bolt Beranek & Newman, Inc.
10 Moulton St.
Cambridge, MA 02138
- 1 David E. Stone, Ph.D.
Hazeltine Corporation
7680 Old Springhouse Road
McLean, VA 22102
- 1 Dr. Kikumi Tatsuoka
Computer Based Education Research Lab
252 Engineering Research Laboratory
Urbana, IL 61801
- 1 Dr. Perry W. Thorndyke
Perceptronics, Inc.
545 Middlefield Road, Suite 140
Menlo Park, CA 94025
- 1 Dr. Douglas Touss
Univ. of So. California
Behavioral Technology Labs
1845 S. Elena Ave.
Redondo Beach, CA 90277
- 1 Dr. Kurt Van Lehn
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304
- 1 Dr. Keith T. Wascourt
Perceptronics, Inc.
545 Middlefield Road, Suite 140
Menlo Park, CA 94025
- 1 William B. Whitten
Bell Laboratories
2D-610
Holmdel, NJ 07733

BEST COPY AVAILABLE

I Dr. Thomas Wickens
Department of Psychology
Frans Hall
University of California
405 Hilgarde Avenue
Los Angeles, CA 90024

I Dr. Mike Williams
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 04304

I Dr. Joseph Wohl
Alphatech, Inc.
2 Burlington Executive Center
111 Middlesex Turnpike
Burlington, MA 01803