3.   AIHANDBOOK.   Documentation of methods, techniques, programs, languages, etc., with tutorial overviews, whose goal is the transfer of knowledge-based systems and other A.I. technology to the broader engineering, scientific, and military community.

Detailed Discussion:

A significant "technology transfer" gap exists in the dissemination of information concerning the concepts, methods, techniques, and programs of artificial intelligence research. The quality of the science is widely recognized (e.g. four of the twelve Turing Award winners of the Association for Computing Machinery are researchers in artificial intelligence). But the use of the science in engineering applications is very small. And the amount of university teaching being done to train new scientists and engineers in the use of the work is also small. Consequently, we were motivated to begin a major work of documentation to fill the gap. The emerging pair of volumes is called the Handbook of Artificial Intelligence. It is being prepared by project scientists and Stanford Computer Science Students. It is a project of low cost and potentially high impact.

The Handbook will consist of approximately two hundred articles on all aspects of artificial intelligence research and development. Each article will be approximately three pages in length; will include brief discussions of the concept, method, technique, program, etc., and a small set of primary source references for further study. In addition, the Handbook will contain numerous "overview" articles summarizing main lines of activity. The entire work will be tied together by an index constructed as a relational network (a "semantic net") so that a user can be led to an article related to his interest even though he did not know that he needed that article.

The outline for the AI Handbook has been put together with the advice and consultation of many AI scientists from laboratories in this country and abroad.

Research Plan:

The Handbook will be produced first in a preliminary edition of two volumes, so that comments, criticism, and correction can be obtained from a wide audience. The preliminary edition of the first volume will be available in Fall, 1977 (so we intend). The preliminary edition of the second volume will be available by Summer, 1978. All the editorial changes necessary, and the entire Handbook preparation for publication will be done during the 1978-79 academic year. More than 150 "first drafts" of articles are already written, and have been made available to ARPA-IPTO for its internal uses.

B.   Production rule representation of knowledge.

1. Particular focus will be placed on the representation of expert problem solving strategies. (i.e., rule-directed behavior at the level of program control.)

Detailed Discussion:

The strategies for problem solving are as important for high performance as the domain specific rules themselves. Some of the strategies that experts use can be expressed in meta-rules, which guide the application and determine the relevance of domain rules. We have already developed the meta-rule formalism for representing knowledge about the reordering and pruning of domain rules. We are proposing to extend that to additional kinds of strategy knowledge. For example, one strategy for dealing with a complex situation is to view it initially as if it were a simple one, find the specific complicating f factors, and deal with those as   special sases. Here is   is "default information" about the uncomplicated situation that allows us to focus on the important special case rules. We propose to use knowledge about commonly expected associations (not expressed in meta-rules but in general "models" of the domain) to guide the application of rules. This work will be carried out in the context of the MYCIN and MOLGEN programs.

Research Plan:

a.   Work with MYCIN's domain experts to formulate at least three models of associated knowledge and expected ("default") values for relevant parameters.

b.   Develop a formalism for encoding this kind of knowledge (6/78).

c.   Isolate the critical times in MYCIN's performance, specifically in the control structure, when strong expectations from models will be most beneficial for guiding the consultation. Modify the control structure to refer to the stored information in models.

d.   Test the cost effectiveness of using models of typical situations in the domain in the context of the three (or more) sample in-depth models (6/79).

2.   Incorporate time-dependent relations into the production rule formalism and control structure, i.e., use time as an integral part of the analysis, where appropriate. Correspondingly, in the dynamic model of the situation (current best hypothesis) built up by a reasoning program, keep a "time line" of relations between events. HASP was the ground-breaking program for this work, and the new work will be oriented toward signal data interpretation and also toward MYCIN.

Detailed Discussion:

One area of improvement for production systems is to uniformly handle rules which interpret a sequence of events over time, as opposed to making conclusions based on a single "snapshot" of the decision making situation. This area of interest will allow the expert to capture the more dynamic situations in his problem domain. The main extension to knowledged-based systems is to allow the expert to specify trends to be searched for in data which has occured up to the present time, to provide a prediction of future trends, and to specify specific milestones which are anticipated. As the processing continues forward in time, the program can compare these anticipated events against the actual data to determine if the previous decision making process was effective. The representation of trends can also be used to "work backwards" to estimate what certain data values might have been at a previous time, but were not then reported.

The major problems include picking both the most effective representation for this type of information and the methods for recognizing and describing trends, and in incorporating the uncertainties in the data or method of data collection into the prediction process. Much of the richness of time oriented analysis is based on mathematical models (e.g. differential equations, simulations, and statistics), which must be integrated with the heuristic information used by the expert to make decisions. The problems of predicting past or future events based on data not known with certainty, or where the current situation must be used to imply how much reliability can be placed on a prediction using the available data must be handled. In a consultation system, it is important to know which type of facts are expected to change, in order to inquire of the user only the minimum facts needed to update the current model of the problem situation.

The major problems include picking both the most effective representation for this type of information and the methods for recognizing and describing trends, and in incorporating the uncertainties in the data or method of data collection into the prediction process. Much of the richness of time oriented analysis is based on mathematical models (e.g., differential equations, simulations, and statistics), which must be integrated with the heuristic information used by the expert to make decisions. The problems of predicting past or future events based on data not known with certainty, or where the current situation must be used to imply how much reliability can be placed on a prediction using the available data must be handled. In a consultation system, it is important to know which type of facts are expected to change, in order to inquire of the user only the minimum facts needed to update the current model of the problem situation.

Research Plan:

Design internal representation (3 months)

Integrate mathematical models (4 months)

Forward and backward prediction under heuristic control (5 months)

Handle uncertainty in the prediction process (6 months)

3. Extend the capabilities of the production rule formalism to include general knowledge of the task domain not easily represented in production rules, such as mathematical relations. Extend the control structure to use this general knowledge in order to direct the reasoning and focus the system's attention on the most relevant subset of rules. Toward similar ends, we will develop further our previous work on the use of meta-rules as guidance for the use of domain rules in a program.

Detailed Discussion:

One very powerful form of reasoning used by a human consultant when he is presented with a new situation is that of classifying and comparing the new situation with "typical" situations he has encountered previously. For example, an auto mechanic who is asked to fix a car which won't start, immediately begins to check for those things which are typically malfunctioning in a "car won't start" situation. That is, he will check the starter, the battery, the supply of gas, etc., and not bother to check the oil level, because he knows that the oil level is not typically at fault in a "car won't start" situation. Sometimes the consultant is unable to classify the new situation as one he has encountered previously, and he must rely upon his basic knowledge of the task domain in order to complete the

consultation. When he is able to classify the new situation as being similar to a "typical" situation, he has, in effect, a model of what to expect in the current consultation. He uses this model of a similar consultation to guide him in performing the current consultation. The model also gives him expectations about what will happen in this new situation because it tells him what typically happens in this situation. And because he can compare his findings with what is "typically" found, he has a way to check the accuracy of the information he obtains in the new situation. We propose to add such a set of models to a rule-based consultation system in order to control the invocation of the rules and to give the system knowledge about what to expect in typical situations. These models will also give us the ability to take appropriate action in response to inconsistent information. That is, given knowledge about what constitutes a "typical" case, atypical findings can be isolated and checked for their validity.

Research Plan:

a. Plan design of models and model-based system (9/77)

b. Add models to data base and implement model-directed consultation system (3/78)

c. Provide consistency checking of information obtained during the consultation

d. Provide explanation and interactive knowledge-acquisition facilities for the model-based system (12/78)

4. Develop a formalism for measuring the power and performance of a program's knowledge base. Create methods for assessing strengths and weaknesses in the knowledge base and for recommending modifications.

Detailed Discussion:

The objective of this research is to develop principles and methods for measuring the performance of certain kinds of AI real world problem solving systems, and to investigate how to use such measurements to improve system performance. This work is motivated by the belief that for an AI system to solve real world problems effectively, it is just as important for it to know how its knowledge is used as it is to know what knowledge it has in the first place. The work is also motivated by the conviction that performance knowledge is essential for conveying a program's scope and limitations to a user, who can then use the program more intelligently. Furthermore, it is hoped that the principles developed here will facilitate the creation of a general laboratory for testing AI methods, whose use will increase our understanding of the field.

This investigation will primarily be concerned with how well the knowledge base of an AI problem solving system is organized and used. The emphasis will be on such factors as the frequency of accessing and grouping information about domain objects; the cost of evaluating and invoking rules characterizing the domain; and the relative merits of strategies for selecting knowledge sources. There are of course other important determinants of system performance, such as the selection of the best internal representation for a data structure, or the selection of the most efficient method of searching a data file. However, the main concern here is to assist expert users in organizing domain knowledge for effective problem solving in conjunction with the AI system.

We will seek to develop interactive methods for a user or designer of a knowledge based system to express system performance criteria; to describe which parts of the system should be monitored, in what detail, and what special events to look for; and to enter heuristics to help the system decide what changes to recommend on the basis of its observations. The system should provide a description of distribution of effort and an estimate of the costs and benefits of further observation of parts of the system.

Research Plan:

a. develop layered information gathering tools — tools which offer a choice of more or less detail for more or less cost — and establish how to use them effectively.

b. establish ways to characterize which parts of a system to observe and what events to detect.

c. develop a means for stating performance criteria (6/78).

d. provide examples of using a description of behavior and a statement of performance criteria to recommend knowledge base changes.

e. investigate ways of estimating the utility of monitoring the system (6/79).

C. Program designs for expert problem solving, with particular focus on planning methods

1. Develop programs and methods for hierarchical planning in task domains in which knowledge is uncertain. Though neither hierarchical planning methods nor methods of inexact inference are entirely new, their confluence is now an important issue.

Detailed Discussion:

31

Significant progress has been made in recent years on methods for both hierarchical planning, producing successively more detailed plans until all factors have been considered, and inexact inference, reasoning with rules whose conclusions can only be "trusted" with a measured degree of certainty. Combining these two powerful problem-solving techniques is now an important research issue for complex, real-world domains. Hierarchical planning in these domains is necessary because of the combinatorial explosion that results when all details are considered at all stages of plan formation. Yet, the rules of these domains make use of uncertain knowledge to make uncertain conclusions, conclusions which may change as further details become available at more refined stages of planning.

The major research question is how to properly modify the basic structure of hierarchical planning to handle the reality of inexact domain rules. Certainly Sacerdoti's idea of postponing decisions involving the ordering of steps in a plan until such a decision is necessary seems a reasonable beginning. The system must be able to handle cases where plans fail at low levels of refinement not only because detailed preconditions of rules cannot be satisfied, but also because the very actions of the rules themselves change when all details are considered. This occurs because the abstractions of rules which are valid at the highest levels of hierarchical planning cannot be expected to provide more than a reasonable guess about the likely action of a rule for which a detailed simulation may be necessary. In simple domains the virtues of hierarchical planning were seen by merely considering successively more rule preconditions at each level of planning, the action of the rule stayed constant. The problem for domains with inexact knowledge is not merely one of considering at each level of planning only those details which are necessary, it is one of considering those details which are reasonable, a task which itself may be rule based.

Research Plan:

a. A goal of the first six months should be to categorize the variety of problems which do result when planning in complex domains where knowledge is tentative and uncertain.

b. Develop rule-based systems which provide heuristic control at the problem points in hierarchical planning--expand the notion of domain dependent critics as introduced by Sacerdoti. This process should take between 1 and 1 1/2 years.

c. Test these systems in a variety of problem domains. This will involve methods for expert entry and explanation of the rules described in (b). The process of testing and expanding will start after 1 year and continue through the second year (6/79).

D. Heuristic Rule Acquisition, by automatic rule formation methods and semi-automatic model-directed rule extraction

1.   Extend the interactive production rule acquisition system to handle acquisition of other types of information about the task domain, i.e., domain-specific models other than rule-based models. Use the domain- specific models to check for inconsistencies in new rules, i.e., exploit the inevitable redundancy to enforce logical consistency among rules.

Detailed Discussion:

The information about "typical" situations, discussed in section B above, will be very helpful fo reasoning and problem solving. It is necessary to be able to acquire this information from experts first, however, before it can be used. This will require giving the production rule system the same acquisition and explanation capabilities for these "models" as it has for iis rules.

The current knowledge-acquisition facility of the consultation system will be extended so that these models can be acquired through interaction with the user. The explanation capabilities of the system will also be extended and improved so that "higher-level" explanations of system performance will be provided. For example, we could explain which hypothesis, i.e., which model, is being tried during a given consultation, and we could say why it is being tried.

Research Plan:

a. Define a vocabulary to be used in the acquisition of models.

b. Implement the basic Knowledge-Acquisition program based on this vocabulary.

c. Provide good human engineering features for Knowledge-Acquisition. For example, simple spelling errors should be tolerated and careful prompts should be used when communicating with the user.

d. Extend the explanation facility so that "higher-level", contextual explanations can be handled.

2. Develop programs for acquiring one particular type of non-rule-based auxiliary model of a domain: an expert's taxonomy of the objects in his domain. This taxonomy is a necessary precondition for applying rules that mention the objects.

Detailed Discussion:

In many domains a strict hierarchical relationship exists among an important subset of the objects in that domain. Experts in these domains often form that subset of objects into a taxonomy with important consequences for rules affecting or using those objects. The major idea is that objects can be understood and accessed through this hierarchical relationship; for example, if no rules can be found which specifically mention object A, then perhaps relevant rules can be found which do mention its parent in the taxonomy, object B. Also, a taxonomical classification makes acquisition and explanation affecting objects simpler and more efficient; misunderstood objects can be clarified by showing their relationship to other, better understood ones.

An automated system for extracting from an expert his taxonomy of a complex domain should provide interactive facilities to allow the expert to make completely clear the relationships between objects in the taxonomy. It is important for a rule-based system using the taxonomy to understand if these relationships are based on some underlying theory of the domain, or if they simply reflect a particular expert's view of the pathways through which objects should be accessed. The system should allow diverse types of - relationships within the taxonomy and should make expansion and modification a relatively simple process for even a large taxonomy.

Research Plan:

a.   Initially develop an interactive framework for describing and modifying an expert's taxonomy of objects in his domain. This should be accomplished within six months.

b.   Combine this taxonomy acquisition system into a variety of knowledge-based systems to test its flexibility and utility. This will be an ongoing process over a several year period.

c.   Develop methods for describing the domain theory that underlies the taxonomic classification along with means for testing the supplied taxonomy to check if its organization really does represent that theory. This should be accomplished in some generality within two years.

3. Develop a general representation for checking the form of new knowledge entered about objects and relations in the domain. These so-called "schema" can also guide the acquisition of knowledge.

Detailed Discussion:

One of the most successful systems to address the knowledge acquisition issue is the MYCIN/TEIRESIAS system developed at Stanford. MYCIN/TEIRESIAS developed techniques for the

acquisition of domain knowledge in the form of rules and objects. MYCIN/TEIRESIAS has approximately four hundred production ("If ... then ...") rules in its knowledge base. Problem solving context and expectation models are used to guide the process of acquiring these rules from a user. MYCIN/TEIRESIAS also has several hundred objects in its vocabulary (eg. gramstain, nutrient, and culture). To guide the acquisition of these objects which have varied structure, a schema, that is a description of structure, is kept for each type of object. When a new instance of an object is acquired, MYCIN/TEIRESIAS decomposes the acquisition process into a number of small steps corresponding to the components of the object. A schema for schemata (the "Schema-schema") gives a description of schemata in the knowledge base. This makes it possible to acquire a new kind of object by first acquiring a new schema for it using the schema-schema, and then acquiring an instance using the new schema.

Several extensions to the ideas in MYCIN/TEIRESIAS are needed to build a powerful problem solving system. The part of the medical diagnostic task represented in this system is described in terms of rules of a very limited format. Some common kinds of operations and strategies are awkward to express within this constrained format. For example, iteration can be expressed only by introducing some extra dummy variables. Following the analogy of using schemata to describe different types of objects, it is proposed that rule schemata be developed to describe different types of rules. (MYCIN/TEIRESIAS may be viewed as having only one type of rule which has an "if" component and a "then" component with associated experts for guidance in filling the two components. This "schema" for rules in MYCIN/TEIRESIAS is built into the system.) We propose to develop a variety of schemata to cover a broad range of domain operations and strategies. Just as an object can be decomposed into it component objects, a rule could be decomposed into its component subactions. In the proposed system, a user could select from a menu of standard artificial intelligence strategies and the system would guide the acquisition of the strategy rule using appropriate schemata from the knowledge base.

In summary, the acquisition process for a strategy rule is broken down into a number of small and manageable steps. The schema used to guide the acquisition process inherits many of its specifications for creating the rule from its ancestors in the schemata hierarchy. It is suggested that this process can be used to help prevent required entries from being forgotten when a new rule is acquired. Much of the structure of a rule can be filled in automatically - for example, the iterative loop in the Means-ends analysis example. Tests on the sets of acceptable values for the components of instances can be built into the schemata as a further check on the correctness of what a user enters. The goal of this process of assisted acquisition is to

make the acquisition of domain specific strategy rules as painless and bug-free as possible.

Research Plan:

a. Develop the vocabulary of schemas that extends the MYCIN/TEIRESIAS capabilities to iterative statements.

b. Design a hierarchy of schemas and programs for "inheriting" or transfering knowledge from one level of the hierarchy to others.

c. Design the "mean" of AI strategies and instantiate all schemas below at least one of them.

4. Additional experiments will test the relative power of rule formation under different learning conditions.

Detailed Discussion:

Given a body of data from which rules are to be formed, together with a basic approach to rule induction, there remains a range of ways in which the data may be utilized, which differ in the degree of parallelism involved in the examination of instances. At one extreme are methods in which rules are formed and refined in a sequence of steps, each step involving the examination of one new instance. At the other extreme are methods which involve a single-pass rule formation process, using all available data. There are, of course, many intermediate possibilities. We propose to investigate, within the Meta-DENDRAL framework, whether some of these methods are optimal in the sense of yielding rules of comparatively high quality with the expenditure of comparatively little computing effort. It is hoped that the investigation will lead us to some general insights concerning the optimal utilization of data in automatic rule formation.

Research Plan:

a. Develop and implement one or more procedures for updating an evolving set of rules on the basis of newly examined data. These procedures will make use of existing capabilities of the RULEGEN and RULEMOD programs, and will make possible the implementation of a variety of schemes for data utilization, as described above.

b. Select and implement a representative subset of the class of data utilization schemes indicated above, and test their performance in the application area of mass spectrometry.

c. Describe in a technical report these experiments, their results, and the lessons learned.

E.   Explanation by program of its line-of-reasoning

Current capabilities are exemplified by the interactive explanation capability (in English) of the MYCIN system, and non-interactive capability in HASP.

1.   Continue program development to extend explanation capabilities to complex procedures in addition to production rules. Also extend facility to allow explanation programs to be provided with additional knowledge about the purposes of requests,thereby allowing them to produce the "most relevant" explanations of lines of reasoning.

Detailed Discussion:

Our previous work demonstrated the explanation capability of rule- based expert systems [Comp. Biomed. Res. ppr, and AJCL microfiche]. We propose to extend the advantages of an explanation system to programs whose knowledge is procedurally embedded. The key concepts we will explore are: 1) an "event history" which records decisions made by the program in the form of traces, 2) an "event structure" which abstractly represents the steps and strategies of a procedure, and 3) an interactive question-answerer which can selectively retrieve traces from the event history, guided by the organization of the event structure. For the sake of illustration we will use an example from antibiotic therapy selection in the MYCIN program.   One relevant question at the end of a consultation with MYCIN is "Why did you recommend streptomycin for the streptococcus infection?". In order to answer this question the program needs to keep a record of the major conceptual reasoning steps taken by the various procedures in the therapy selection algorithm at the end of a MYCIN consultation. It must also be able to structure these "events" so that only the relevant steps will be retrieved in answer to the user's question.

Research Plan:

a. Develop a model for types of problem-solving methods which might be explained by this scheme. Is there a broad class of programs to which it might be applicable? (8/77)

b. Develop a general schema for the event structure which can be used to abstractly represent the class of programs determined in 1. (10/77)

2. Extend the explanation capabilities of production systems to allow the explanation of strategy decisions embodied in meta-rules.

Detailed Discussion:

Strategy decisions may sometimes be more important than decisions about more specific steps in the reasoning. We propose to extend the capabilities of the MYCIN explanation system to phrase answers to questions in terms of the general strategies as well as in terms of individual rules.

For example, the answer to the question "Why did you ask me about X?" may sometimes refer to the strategy decision in which X was relevant, e.g., "Because one of the strategies for finding out about Q is only relevant when you know X." That is, the answer may be given in terms of the strategy rules as well as the individual domain rules. Again, the question "Why didn't you use rule R?" may refer to the general strategy under which the program was operating at the time -- "Because my strategy rule S said that R would not be useful in this context." Finally, general questions about the domain, and not about the specific consultation, may refer to strategies. E.g., "In general how do you conclude the identity of infecting organisms?" This may be answered by mentioning the individual domain rules, as MYCIN does now, or by mentioning some of the general strategies for reasoning about this parameter.

Research Plan:

a. Extend the history list mechanism of the MYCIN program to accept information about the strategies that MYCIN is carrying out during a consultation.

b. Upgrade the question-answering capabilities of the MYCIN program to answer general questions about the program's reasoning in terms of strategy rules. An important problem to be solved is recognizing the need to frame an answer in terms of strategies rather than individual rules.

II. Research in Support of Other IPTO R&D Programs

A. Rule-based methods in the interpretation of image data

Many of the rule-based methods previously developed by our project, and to be studied and developed in this period, can be applied to assist the interpretation of image data, in at least the following ways:

...   modularity/flexibility. Qualitative rules of interpretation can be changed easily when existing rules are deemed inadequate by the image analysis expert.

... acquisition. New rules can be inserted into the program under guidance of a rule acquisition system (similar to MYCIN's) to extend performance to new or changed circumstances.

... explanation. The line-of-reasoning being used by the program to interpret the image can be exhibited for the human analyst to understand and evaluate.

This important application of our work is not now being made, and we propose to orient a portion of our effort in this direction. A.I. researchers in IPTO's Image Understanding program are doing significant work primarily in "front-end" image processing and procedure-based interpretation; but not in rule-based methods such as exist in MYCIN, DENDRAL, and HASP.

Because the start-up costs in money and scientific energy to understand the real problems in image analysis and build base-level software are so great, we propose to develop our prototype system in a domain in which we have considerable expertise, excellent data, extensive software, and (last but not least) some other financial support from NSF. This domain is protein crystallography, involving the interpretation of three-dimensional images of the electron density of a protein molecule in a crystal.

The rule-based system to be developed in this prototype will be engineered to be as portable as possible to facilitate its application to the photo-interpretation prototypes being built by the Image Understanding projects, e.g., those at the Stanford AI Laboratory, SRI, and Carnegie Mellon University.

This proposed activity has additional plausibility because of the following technical circumstances:

... the HASP rule-based program structure developed by some of our staff is applicable to this problem.

... as part of our current activity, to test the generality of that HASP program structure, we have already mapped the basic ideas in the HASP program onto the crystallographic data interpretation problem. Hence we have a running start on the problem.

Detailed Discussion:

For the reasons discussed above, we propose to develop rule-based methods for interpreting 3-D images, within the context of protein crystallographic data. Our ultimate goal is to transfer the rule-based methodology to other IPTO R&D programs in image interpretation. In this respect, preliminary discussions have been held with Prof. Raj Reddy of CMU. Prof. Reddy is interested in our work and eager to collaborate. The technology

transfer is expected to be bi-directional in this effort: the Image Understanding Projects stand to benefit by employing a rule-based structure in their programs, and we expect to benefit from their expertise in front-end processing of the primary image data.

The interpretation of an electron density image, derived from the reduction of X-ray crystallographic data, is a necessary step in understanding the 3-D structure of proteins and other macromolecules. When crystallographers use the term "electron density map" they usually have in mind some pictorial representation of the electron density defined over a certain region of 3-space (usually some fraction of the unit cell of the crystal). By carefully studying the map the experienced interpreter can find features which allow him to infer important features of the image: approximate atomic locations, molecular boundaries, groups of atoms, the backbone of the polymer, etc. An "interpretation" is a model of the molecular structure which conforms to the electron density map and is also consistent with his knowledge of protein chemistry, stereochemical constraints and other available chemical and physical data (e.g., the amino acid sequence).

Our goal is to build an intelligent assistant to the image analyst, i.e., a computational system that can generate and verify its own structural hypotheses, and explain the steps taken in the interpretive process. This capability implies (1) a representation of the 3-D image data suitable to machine interpretation, (2) a substantial knowledge base appropriate to the task domain, (3) a wide assortment of model building algorithms and heuristics, in order to achieve acceptable performance, and (4) a flexible, rule-based system for incorporating both the knowledge sources and the control strategy. An elaboration of these points has been recently documented (Engelmore and Nii, HPP-77-2). The fourth point, concerning knowledge representation and utilization, is of primary concern here, because of its applicability in other image understanding contexts.

The expert analysts who interpret these images move continually across a large field of basic facts, special features of the data and implications of the partial model already built, looking for any and all opportunities to add another piece to their structure. There are several desiderata to working in this "opportunistic" mode of hypothesis formation: (1) the inference making rules and the strategies for their deployment should be separated from one another, (2) the rules should be separated from the mechanics of the program in which they are embedded, and (3) the representation of the hypothesis space should be compatible with the various kinds of hypothesis generating rules available. (The hypothesis structure represents an a priori established plan for problem solving.) The modularity of such a

system allows users to add or change rules for manipulating the data base, as well as to investigate different solution strategies, without having to make major modifications to the system. These issues, which have general applicability beyond the specific task domain of protein crystallography, are discussed further in Appendix A.

The formal and informal procedures which comprise our knowledge sources are expressed as rules, as in the DENDRAL, HASP and MYCIN systems. These rules are collected into sets of rules, each set being appropriate to use on a particular class of events. The events generally reflect the level on which the inference is being made, which in turn reflects the level of the detail of the model. The correspondence between event classes and rule sets is established by another set of rules, the event rules. The event rules thus form a second layer of rules which direct the system's choice of knowledge sources for a given event, reflecting the system's knowledge of what it knows. (A similar set of rules, the job rules, perform the same role when the system operates in goal-driven mode.) Maintaining the rule-based structure affords a flexibility in choosing different combinations of knowledge sources to work together, without having to make any changes in the knowledge sources themselves. Thus, yet a higher level knowledge source, the strategy rules, can manipulate the events in order to choose the appropriate combination of KSs suited to a particular stage or state in the solution hypothesis. Examples of strategy rules are the merging of two or more events to produce a new event that can lead to further progress, or shifting from event-driven to goal-driven mode.

We thus have a completely rule-based control structure, employing three distinct levels of rules (or knowledge): the specialist, commonly called the knowledge sources, the event processing rules (or job processing rules), representing knowledge about the capabilities of the specialist, and the strategy rules which know when to use all available knowledge to solve the problem. Although this pyramidal structure of rules and meta-rules could continue indefinitely, the flexibility of knowledge deployment offered by our three-tiered system would appear to be sufficient for this problem solving system. Similar ideas in a simpler context have been explored by Davis (HPP-76-7) for the MYCIN system.

Research Plan:

We propose the following tasks and dates of completion:

1. Continue the implementation of the rule-based image interpretation system in the crystallographic context. Sub-tasks include:

a. Implementation of additional knowledge sources for matching templates of predicted sub-structures with primary and/or processed image data. (3/78)

b. Partitioning of the KSs into functionally related sets of rules. (3/78)

c. Development and implementation of new strategy rules for directing the utilization of the KSs. (6/78)

2.Implement an explanation subsystem, using the history list feature   presently   incorporated,   which   interfaces conveniently with the user. (6/78)

3. Develop contacts with at least one IPTO Image Understanding group, in order to understand their particular problems and associated technical isssues. (9/77)

4. Select one of the Image Understanding groups for collaboration, and begin working with it, in order to orient our prototype development towards its particular needs. (6/78)

5. Begin the transfer of ideas and software from the prototype system to the collaborating Image Understanding project. (6/79)

B. Distributed Sensor/Computer Networks

We propose a two-year effort of relatively small size to study issues in signal interpretation, artificial intelligence methods, and computer organization raised by the IPTO initiative in Distributeed Sensor/Computer Networks.

The general problem to be studied is an inference problem of interpreting and integrating the sensory information being received by an array of sensors/computers geographically distributed, for which there is no designated central processing agent.   Programs resident at the local sensor/ computer would compute a "local" understanding, communicating it to neighbors as appropriate and calling upon neighboring computational assistance as needed.   The local sensor/computer stations may be integrating more than one type of instrumental data.

Our experience with this type of problem arises partly from our work on the geographically distributed, multi-sensor (but central processing) HASP task; and partly from our planning for multi-instrument data interpretation in analytic chemistry (arising from the work on DENDRAL,and currently under funding consideration by NIH).

The problem described above is inextricably tied up with two other problems, only one of which is being studied elsewhere

in the program. The first is one of conceiving of how the methods
of artificial intelligence research that are currently well
understood can be mapped into a highly- parallel asynchronous
computing environment. Attempts will be made to study this
problem for (at least) heuristic search methods, rule-based
inference methods and graph matching. The second is one of
conceiving of a flexible enough architecture of small computers
and communications to support the requirements of the AI methods,
especially in their application to the interpretation of
geographically distributed signals. We have begun initial and
relatively abstract studies of both of these problems.

We define "distributed processing" to mean processing that
involves one or both of the following attributes: physical
decomposition of the processor into relatively independent
processor nodes, and functional decomposition of the top level
task into relatively independent subtasks. There are several
applications, particularly distributed sensor networks and
distributed command and control systems, in which the confluence
of distributed processing and AI techniques appears attractive.
We propose to study and develop techniques for symbolic inference
in a distributed processing environment; and to discover the
constraints (if any) placed on a distributed processor
architecture by these applications. These techniques are based
on an emerging formalism we have been developing, which we call
the CONTRACT NET. In a contract net, individual tasks are dealt
with as contracts that exist between pairs of processor nodes
that take on the roles of contract manager and contractor
respectively. In any such system with distributed executive
control, there must be a method for choosing the small number of
tasks that can be processed at a given time instant from the
large number generally available. To this end, we wish to explore
the concept of a "priority description" which is intended to
offer a more detailed specification of task characteristics and
importance than the usual integer ratings.

Our approach to the architectural problem is toward a much
more loose coupling between the computers of the net (i.e., much
more local processing,much less communication) than is implicit
in the design of the c.m star machine architecture being studied
at Carnegie-Mellon. A possible and important by-product of this
study could be a computational route to extreme speed-ups in
symbolic computation arising from the collective processing
activity of large numbers of inexpensive microcomputers. The slow
speed of symbolic computation, in contrast to numeric
calculation, has been a major hurdle in the development of AI
science and AI applications.

Joining us in these studies have been a member of the staff
of the Canadian Defense Research Establishment Atlantic (a
Stanford student) and staff member of the Norwegian Defense
Research Establishment (a Stanford post-doctoral Fellow). We

have had some indications of possible industry support for the research on distributed architectures. We propose three tasks:

1. Interact with the other investigators in the IPTO Distributed Sensor Network project to understand and help to further refine the scientific and practical problems therein; and produce a report (or series of reports) on the methods available from artificial intelligence research that might be applicable to the development of a total system. Our methodology will be computer simulation (not hardware building). The sense data interpretation problem that will give specificity to our study will be an image data interpretation problem, probably the same image data interpretation problem discussed in (a.) above (the crystallographic image data interpretation problem) for reasons of economy of effort and funds.

2. Prepare a report on the realization of at least two of the well-understood AI methods in a distributed computing environment. Again , the methodology will be computer simulation but the tasks chosen for study will range more broadly than just signal data interpretation.

3. Prepare report(s) on studies of various highly parallel asynchronous architectures suited to symbolic computation. The methodology will involve both abstract analytic studies and computer simulation. We do not at this time intend to construct any network until we understand the issues from a theoretical point of view.

Research Plan:

1. Complete design of CONTRACT NET formalism for expressing control of problem solving in a distributed network, and incorporate the completed design into a simulation program.

2. Report on the variations in control that can be achieved through the use of priority descriptions.

3. Expand the simulator to allow exploration of the suitability of different processor node interconnection alternatives for a distributed processor (6/78).

4. Test the CONTRACT NET formalism on a variety of simple, yet representative problems, using the simulator, and report results.

5. Report on the problems raised by the requirement for global information in a distributed network in which there is limited memory at each processor node (6/79).

Appendix A.   Extended Discussion of Proposed Rule-based System for Image Understanding

In this appendix we amplify some issues on knowledge representation and utilization in the context of the proposed 3-D image understanding research.

Representation of Knowledge in the Image Interpretation System

The problem of representing a large and diverse body of knowledge, in a form which will allow it to be used cooperatively and efficiently in the search for plausible hypotheses, is of central concern to this research.   The system currently under development draws upon many concepts which have emerged in the design of other large knowledge-based systems, e.g., the use of production rules (as in DENDRAL, HASP, etc.) and a multi-level space of hypotheses (as in Hearsay-II).

Knowledge consists of facts, algorithms and heuristics (rules of good guessing). Facts required for protein structure inference are general physical, chemical, stereochemical and crystallographic constraints. Typical factual knowledge stored in the system includes molecular structure and chemical properties of the twenty amino acid building blocks.   These facts are encoded as tables or in property lists attached to specific structural entities.

Algorithms and heuristics comprise the formal and informal knowledge which generate and/or verify hypothesis elements. We have been guided by two general principles in the representation of the knowledge sources:

1) decompose identifiable areas of knowledge into elementary units, each of which increments the hypothesis when specified preconditions are met.

2) represent the elementary units as situation-action rules.

To illustrate, consider the relatively simple example of heavy atom location.   This subproblem is decomposed into two independent parts: 1) inferring the presence of heavy atoms and 2) determining their spatial locations. These two independent parts are represented as two separate KSs, invoked under different conditions. A rule contained in the first KS is:

IF the composition list contains a cofactor of type heme,
THEN:
    1) create a superatom node of type heme in the model plane,
    2) create an atom node of type iron in the model plane,
    3) create membership links between the iron and the heme,
    4) put "cofactor_posited" on the event-list,
    5) put "heavyatom_posited" on the event-list.

45

Note that several actions may be performed for a given situation.
An action may itself be a situation-action rule, and may be
iterative. Also note that at least one of the actions of each
rule is to place a token on an event-list. The event-list is used
by the interpreter, discussed in the next section, to determine
what to do next, i.e., which set of knowledge sources will be
invoked after the current event has been processed.

Control Structure for the Image Interpretation System

There are several choices of control structure faced by the
designer of a knowledge-based system. Basically the choices are
among points on a spectrum, at the extremes of which are goal-
driven and event-driven systems. In a goal-driven system (of
which MYCIN is a well-known example (Shortliffe, HPP-74-2)) the
rule interpreter selects a rule which concludes with the goal
being sought. In our system, we might imagine having such a goal
rule as follows:

IF
1) the topological description is complete, and
2) the coordinates of all atoms in the structure are assigned,
   and
3) the structure satisfies stereochemical constraints, and
4) the structure is consistent with the electron density map,
   and
5) the structure is consistent with auxiliary chemical data,
THEN:
signify that a model has been completed.

The interpreter would then attempt to verify each of the premises
in the goal rule. To do that, other rules would be selected
whose conclusions (the right-hand sides) verified the premises
under consideration and the interpreter would attempt to verify
the premises of these rules, and so on, working through the list
of rules in this recursive fashion. The program's focus of
attention is determined by the current rule whose premises are
being evaluated. Many levels of recursion may occur before a
rule is reached which is relevant to the current state of the
system. A goal-driven monitor is attractive, in that it pursues
a logical chain of reasoning, in which the purpose of each move
is clearly revealed by the tree of subgoals.

A vital requirement for the success of a goal-driven
control structure is the relative independence of the subgoals.
Although some aspects of the protein image interpretation problem
can be solved in this fashion, the independence requirement is
not generally satisfied. For example, the location of each atom
in the molecule must obey stereochemical constraints with respect
to neighboring atoms, so that one cannot simply determine each
atom's position independently.

An alternate way to focus attention is to employ an event-driven control structure. In this scheme the current state of the hypothesis space determines what to do next. The monitor continually refers to a list of current events - the event-list mentioned in the rules discussed above - which is used to trigger those knowledge sources most likely to make further headway. As a knowledge source makes a change in the current hypothesis, it also places a symbol on the event-list to signify the type of change made. Thus as events are drawn from the event-list for processing, new events are added, so that under normal conditions the monitor always has a means for choosing its next move. An explanation capability is provided by keeping a list of events processed, their predecessor and successor events, and the KSs invoked.

The system we are currently developing operates in both goal-driven and event-driven modes (with an emphasis on the latter), and is closely related to the HASP system design. The normal iterative cycle of problem solving uses the event-list to trigger knowledge sources, which create or change hypothesis elements and place new events on the event-list. Thus the system's behavior is "opportunistic" in that it is guided primarily by what was most recently discovered, rather than by a necessity to satisfy subgoals. The choice of an event-driven control structure as the primary mode of operation is based partly on subgoal interdependence, partly on efficiency in selecting appropriate knowledge sources and partly on conformity with the structure modeling process normally employed by protein crystallographers.

BIBLIOGRAPHY

Heuristic Programming Project Memos Since 1964

HPP-64-1
J. Lederberg,
"DENDRAL-64-A System for Computer Construction,
Enumeration and Notation of Organic Molecules
as Three Structures and Cyclic Graphs",
(technical reports to NASA, also available from
the author and summarized in (12)).
    (1a) Part I. Notational algorithm for tree
        structures (1964) CR.57029
    (1b) Part II. Topology of cyclic graphs (1965) CR.68898
    (1c) Part III. Complete chemical graphs; embedding
        rings in trees (1969)

HPP-64-2
J. Lederberg,
"Computation of Molecular Formulas for Mass
Spectrometry", Holden-Day, Inc. (1964).

HPP-65-1
J. Lederberg,
"Topological Mapping of Organic Molecules",
Proc. Nat. Acad. Sci., 53:1, January 1965,
pp. 134-139.

HPP-65-2
J. Lederberg,
"Systematics of Organic Molecules, Graph
Topology and Hamilton Circuits. A General outline of
the DENDRAL system." NASA CR-48899 (1965)

HPP-65-3
AIM-30, AD785056
Edward A. Feigenbaum, Richard W. Watson,
"An Initial Problem Statement for a Machine
Induction Research Project", (April 1965).

HPP-66-1
AIM-38, AD785066
Donald A. Waterman,
"A Filter for a Machine Induction System",
(January 1966).

HPP-66-2
AIM-46, CS-50, PB176761
Staffan Persson,
"Some Sequence Extrapolating Programs: a Study
of Representation and Modeling in Inquiring
Systems",
Ph.D. Thesis in Computer Science, (U.C. Berkely September 1966).

HPP-66-3
AIM-47
Bruce Buchanan,
"Logics of Scientific Discovery",
Ph.D. Thesis in Philosophy
(Michigan State University December 1966).

HPP-67-1
AIM-49
Georgia Sutherland,
"DENDRAL -- a Computer Program for Generating and Filtering Chemical
Structures", (February 1967).

HPP-67-2
J. Lederberg,
"Hamilton Circuits of Convex Trivalent Polyhedra (up to 18 vertices),
Am. Math. Monthly, (May 1967).

HPP-67-3
AIM-54
Joshua Lederberg, Edward A. Feigenbaum,
"Mechanization of Inductive Inference in Organic Chemistry",
(August 1967). Also in Formal Representations For Human Judgment,
B. Klainmuntz, Editor, Wiley, 1968.

HPP-68-1
J. Lederberg,
"Online Computation of Molecular Formulas from Mass Number".
NASA CR-94977 (1968)

HPP-68-2
AIM-62
Bruce G. Buchanan, Georgia Sutherland, E. A. Feigenbaum
"Heuristic DENDRAL: a Program for Generating Explanatory
Hypotheses in Organic Chemistry", (July 1968).
Also in Machine Intelligence-4, Edinburgh Univ., Press, 209-254, 1969.

HPP-68-3
AIM-67 AD680487
Edward A. Feigenbaum,
"Artificial Intelligence: Themes in the Second Decade",
Proceedings of IFIP International Congress, Edinburgh, (August 1968).

HPP-68-4
AIM-74, CS-118, AD681027
Donald Waterman,
"Machine Learning of Heuristics",
Ph.D. Thesis in Computer Science, (Stanford University December 1968).

HPP-68-5
E. A. Feigenbaum and B. G. Buchanan, "Heuristic DENDRAL: A
Program for Generating Explanatory Hypotheses in Organic Chemistry,"
in Proceedings, Hawaii International Conference on System Sciences,
B. K. Kinariwala and F. F. Kuo (eds.), University of Hawaii Press, 1968.

HPP-69-1
AIM-80, AD685612
Georgia Sutherland,
"Heuristic DENDRAL: a Family of LISP Programs", (March 1969).

HPP-69-2
AIM-99
Bruce G. Buchanan, G.L. Sutherland, E.A. Feigenbaum,
"Toward an Understanding of Information Processes of
Scientific Inference in the Context of Organic Chemistry",
(September 1969). Also in Machine Intelligence-5, Edinburgh Univ., Press,
1970.

HPP-69-3
AIM-102
Donald A. Waterman,
"Generalization Learning for Automating the Learning
of Heuristics", J. Artificial Intelligence Vol. L, No. 1/2
(July 1969).

HPP-69-4
AIM-104
Joshua Lederberg, Georgia Sutherland, Bruce G. Buchanan,
Edward A. Feigenbaum
"A Heuristic Program for Solving A Scientific Inference Problem:
Summary of Motivation and Implementation", (November 1969).
Proceedings of IV Systems Symposium, Case Western Reserve Univ.,
Springer-Verlas, 1970.

HPP-69-5
J. Lederberg,
"Topology of Molecules", in the Mathematical Sciences -
A Collection of Essays, (ed.) Committee on Support of
Research in the Mathematical Sciences (COSRIMS), National
Academy of Sciences - National Research Council, M.I.T.
Press, (1969).