1986

Intelligent Systems

MEIS
MICROELECTRONIC &
INFORMATION SCIENCES
CENTER

University of Minnesota
Institute of Technology

Microelectronic and Information Sciences Center

Commitment to Focus

Commitment to Cooperation

INTELLIGENT SYSTEMS REVIEW

November, 1986

MEIS Center
227 Lind Hall
207 Church Street, S.E.
Minneapolis, MN 55455

(612) 624-8005

MEIS Intelligent Systems Group

Overview.

This report summarizes activities of the MEIS Intelligent Systems Group in 1986 and outlines our immediate plans for continuing research.

The Intelligent Systems project has as its principal focus the development of high performance computing systems that incorporate human reasoning techniques. Since many of the scientific, engineering, and management problems for which computer aids are envisioned are currently solved with reasonable competence by human experts, an understanding of the mechanisms of human performance can offer powerful insights into possible computational approaches to these problems. Few existing AI centers, however, have made any serious attempt to combine a formal analysis of possible solutions to problems with an understanding of the functioning of biological systems that have evolved to solve the same problems. Our group at Minnesota has gained prominence in this area by focusing narrowly on the construction of machine representations of highly skilled activities drawn from an analysis of the information processing requirements of a task and biological models of task performance.

The field of Artificial Intelligence is rapidly extending the power and flexibility of modern computers. Working systems exist in areas such as expert problem solving, automated perception, the organization and retrieval of complex data, and natural language understanding. The difficulties and complexities of the problems attacked by researchers in Artificial Intelligence requires expertise in a wide variety of disciplines. Sophisticated software is needed to effectively represent and use complex knowledge bases. Often, specialized computer architectures are called for to efficiently execute AI programs. Psychology provides representations of knowledge and skill based on insight into how people solve similar problems. Electrical and mechanical engineering skills allow the construction of sensors for perception and robots to effect actions. Our interdisciplinary approach to these problems gives us opportunities available in few other AI research centers.

The Intelligent Systems project is a nationally recognized, interdisciplinary Artificial Intelligence research group. One of our members is on the editorial board and another on the advisory board of the IEEE's *Transactions on Pattern Analysis and Machine Intelligence,* one of the most important journals in the field. We are a member of the SUMEX-AIM computer network managed by the Stanford University Heuristic Programming Project. This is a select group of eight Universities funded by NIH to study and develop Artificial Intelligence decision support systems. Group activities are visible through presentations at national and international conferences on Artificial Intelligence and via numerous journal publications.

The Intelligent Systems group is active in two different research areas. Expert System research works to develop models of the processes by which difficult problems can be solved. Such models allow the implementation of computer programs for problem solving in a broad range of application areas. Computer Vision research aims at gaining an understanding of the fundamental principles that underlie vision. With such an understanding, it will be possible to build automatic vision systems for tasks such as manufacturing, vehicle guidance, etc. In addition, the group has been active in the acquisition of essential laboratory facilities. AI research requires specialized, high capacity computing resources. Thus, the development of our Artificial Intelligence Laboratory is crucial to the groups continued success.

Expert Systems -- Strategies for Expert Problem Solving.

In the past few years it has become possible to build computer programs capable of performing tasks previously done only by human experts. These programs, called expert systems, have dramatically changed the way we think about problems ranging from the diagnosis of disease to the design of computer hardware. Although still in its infancy, expert systems technology offers the opportunity for enhancing human productivity through the development of tools that aid in the performance of complex problem solving and decision making tasks.

One aspect of the research of the Intelligent Systems Group is the study of expert reasoning. In this work, we have developed methods for investigating human experts that draw upon the field of cognitive science and lead to computational models for research and development in the fields of artificial intelligence and expert systems. In addition to the development of a framework for conducting research on expert systems, recent research has focused upon two specific areas of interest to both the industrial scientist and the artificial intelligence researcher: (1) the investigating of specific representations employed by human experts in performing ill-structured decision making and problem solving tasks and (2) the development of computer programs that embody these representations as a means of achieving expert performance on a class of tasks.

We have conducted specific studies of expertise in a number of areas including medicine, science, engineering, law and business. In most cases, a better understanding of expert reasoning has lead to the construction of high performance expert systems. This past year, for example, we have collaborated with sponsors from the business and industrial community to investigate expertise in computer hardware fault diagnosis, product development and off-line quality control, chemical reactor design and auditing. These recent research projects provide a model for the kind of collaborative effort we hope to establish as part of the continuing work of the AI Pinnacle research program.

To investigate expertise means to inquire about operative knowledge and how it functions in the performance of a task. Although we assume that this knowledge reflects the constraints of the human information processing system, we do not investigate the role of these constraints in task performance. Rather, we ask what the expert's problem solving method is and how it functions to complete a task. We have divided our investigation of the operative knowledge comprising expertise into three phases.

In the first phase of research our objective is to identify the knowledge necessary for accomplishing a class of tasks. Our methodology here includes structured interviews with experts (and non-experts), the analysis of tasks and records of task performance such as video tapes and audio transcription of thinking aloud comments made by experts while performing tasks, and the systematic observation of the behavior experts in typical work environments. The result of this phase of activity in information processing terms is a statement of what must be computed in order for a task to be done.

The second phase of research is the development of a process theory of the expertise required to perform a task. In its most general sense a process theory is a statement of ingredients or events and how these may be combined to achieve task performance. However, process theory of expertise is different from a process theory of the expert. In the latter case we are testing how symbol structures representing explicit psychological constructs might perform some task of interest, while in the former case we are testing constructs about the domain of expertise. Methodology for the development of process theory of expertise includes the empirical testing of events identified in phase 1 research as well as the explicit simulation of actions to achieve goals, and comparisons between the behavior of computer models and the behavior of human experts (the source of our models of expertise).

In the third phase of research we manipulate the theory developed and tested in phases 1 and 2 to determine its behavior under new, and in some cases extreme conditions. The goal here is to generate predictions that can be tested against the behavior of additional experts and tasks. If we are successful, in this third phase of research we discover conditions under which the structures we have posited may be organized to avoid misconception and error, providing thereby not only useful information to the domain of expertise being represented, but also a stronger test of our theory.

The research we conduct leads to models of expertise in the sense that the knowledge we uncover has been developed by human beings to perform tasks to which they have become highly adapted. The development of models of expertise is important because there are often several kinds of models that will perform a given task. In the case of medical diagnosis, for instance, models based upon Bayes Theorem have performed as well or better than some human diagnosticians. Similarly, in areas such as scheduling or purchasing, mechanical models drawn from the field of operations research may provide perfectly adequate operative knowledge for carrying out a given task. The argument of our research is that when the objective is to build models that will aid human problem solvers in the performance of complex tasks, the development of models of expertise provides an important resource, not only to the field of artificial intelligence, but to the field from which the task is drawn.

This past year we have conducted research on expertise in four types of problems: model-based reasoning which we are investigating in the domain of computer hardware diagnosis, heuristic design which we are investigating in the area of product development and off-line quality control, qualitative reasoning, which we are investigating in the area of chemical reactor design, and perturbated tasks and Garden Path errors, which we are investigating in the field of auditing. Each of these areas of proposed work is described in turn.

Hardware Diagnosis:

Our original research on models of expertise was carried out in the field of medical diagnosis. In this domain work has been completed in each of the three phases of research described above. The model of expertise upon which the work in medicine is based assumes that a diagnostic problem is solved by recognizing a new case as an instance of a problem that has been solved in the past. This recognition-based approach to diagnosis leads to efficiency in reasoning as well highly successful performance, sometimes surpassing that of the human expert.

More recently, we have applied our methodology from the study of medical diagnosis to the investigation of expertise in computer hardware diagnosis. Based upon work carried out this past year at the phase 1 and 2 levels of research described above, we have discovered that experts in hardware diagnosis seem to reason quite differently from physicians. Rather than relying upon a vocabulary of faults against which to match the data of a current case (as was done by the medical diagnosticians), expert hardware diagnosticians appear to reason from some model of the system being diagnosed. At first we thought this model might be one that represented the physical components of the system. Upon closer study of the expert's performance, however, we discovered that hardware diagnosticians appear to base their problem solving on a representation that we have termed a functional model. Such a model contains knowledge of physical components at a global level (e.g., disk driver, cable, etc.), as well as more detailed levels of representation (e.g., boards and chips). In addition to these physical components, however, the expert's representation uses knowledge of "imaginary " components, such as noises in the environment, which provide useful diagnostic cues. The expert's functional model represents, in fact, a sort of vertical slice through levels of the computer's physical representation. This slice captures just the knowledge that is most diagnostically useful.

Following our discovery of the functional nature of an expert's reasoning in the area of hardware diagnosis, we initiated a joint research project with two local computer companies - IBM and The Control Data Corporation. Each of these companies has supplied problems and experts for our research. We are currently studying the exact composition of expert models for the hardware diagnosis task of field returns on a large mainframe computer. Based upon this work, conducted at the phase 1 level, a computational model VESALIUS has been developed (phase 2).

VESALIUS is a program to diagnose faults in computer hardware. A computer hardware device (like a memory board or CPU board) contains a large number of individual components. VESALIUS must examine the device's symptoms and test results to determine which components have failed. The number of ways each component can fail is small, and we can completely categorize them all. However, since the number of components in the entire device is large, we cannot completely categorize all the ways the device as a whole can fail. VESALIUS therefore makes some assumptions about individual component failures and uses a computational model of the device to see if the consequences of those assumptions are consistent with what has been observed about the device. VESALIUS views diagnosis as a search for a consistent set of these assumptions.

VESALIUS' device models are networks of individual component models. Unlike some related systems, we do not attempt to model the actual operation of the device's components. Rather, we model the expert diagnostician's pragmatic view of how components operate, which generally differs from the way they really operate. This facilitates acquiring knowledge from human experts since our models are described in terms that they understand. More importantly from a computational point of view, we find that the diagnostician's model is often simpler than the real model but still contains enough information to make correct diagnoses.

VESALIUS is implemented using a domain independent rule interpreter, which in turn is implemented in Common Lisp. Individual components and their interconnections are represented by small groups of rules. The design of our rule interpreter borrows most of its ideas from logic programming, but contains some novel features. It is able to introduce new rules by specializing existing rules so that they can be applied more efficiently. It also avoids time consuming backtracking by incorporating an assumption based truth maintenance system based on recent work by Johan de Kleer.

During the coming year we shall investigate the following questions: How does the expert shift attention among competing model based representations (e.g., physical, functional, and fault models)? What kind of explanation for problem solving steps can be provided from the representation employed in carrying out the task? And finally, how does the expert's representation guide the gathering of additional problem solving data.

Product Development and Off-line Quality Control:

This past year we have investigated expertise in product development and what has been recently called off-line quality control. Expertise in both of these areas is based upon the field of industrial statistics and the subfield of experimental design. One of our objectives in this project was to determine whether knowledge acquisition methods developed through extensive study of diagnostic and troubleshooting tasks in medicine and computer engineering could be transferred to a design task. A second objective was to investigate expert representation in a design environment. The research has been carried out in collaboration with statisticians from the University of Minnesota and the 3M Company. Research has been conducted at the phase 1 level and work has begun at phase 2.

In the design task we have investigated, a client is interested in studying the relationships of certain variables in his work environment. A typical problem is to design an empirical study (an experimental design) that optimizes a ratio of cost to information gained. This aspect of design becomes increasingly important when the experiments are expensive to run. This is where expertise of the consultant statistician is applied.

From the clients point of view the problem is to select a set of variables so as to optimize some criterion such as texture of bread (food scientist) or amount of some chemical produced (chemical engineer). When combined with environmental (non-control) variables such as permissible range on temperature setting for baking, for example, experimental design becomes a powerful technique for not only the development of new products in many areas of manufacturing but also in improving a current production process.

From intensive study of thinking aloud protocols collected from experts performing the statistical consultation task, the experts operative knowledge (expertise) has been divided into three parts:

- Knowledge of statistical design categories such as factorial and response surface designs, plus heuristics for selecting from among these alternatives and their variants.

- Knowledge of reasoning tools such as blocking, and heuristics for applying these tools to achieve client goals such as maximum yield at minimum cost.

- Knowledge of statistical concepts such as confoundings, variance and replication and how these must be employed to achieve an answer to the client's problem consistent with the rules of statistical reasoning.

Our investigation of design experts (statisticians) has also revealed a four phase reasoning process that we refer to as design by debugging: gathering data, proposing a design, refining a design, and evaluating the design. Expertise in the process occurs in proposing an initial designs that captures the clients objectives within a given cost-benefit structure. A second source of expertise consists of selecting environmental variables to incorporate into the design such that the resulting design is stable within a range of values of the independent variables. A third, and perhaps most critical form of expertise then consists of modifying the proposed design as the consultant learns more about the exact nature of the client's problem.

An expert system has been developed based upon the above analysis. It accepts input in the form of a set of potential design variables and uses experimental, cost and environmental constraints to recommend a feasible design solution. The system represents various forms of experimental design knowledge in four levels: In the first level, the design level encodes the alternative design possibilities together with the experts' heuristics for considering, rejecting and refining them. The second level termed the Statistical level, is where experts' knowledge of specific statistical concepts such as variability, replication and blocking is represented. The third level corresponds to the domain of experimental design problems. This level incorporates the statistically relevant knowledge about the domain of application. The fourth level corresponds to the user. This level has knowledge about the statistical competency of the user and is used in implementing a user-conscious interface.

Because of similarities between the overall process being employed by our experts in experimental design and the reasoning process employed by design experts in previous work, we have adapted the KEE software tool developed from the MOLGEN project as programming environment. Our current system is based upon a blackboard-like architecture in which constraints are posted and evaluated. An early prototype of the system is currently being evaluated by statisticians from the 3M Company and the University of Minnesota.

In the coming year we will conduct research in three areas: First we shall investigate expertise in optimal design, although a number of designs would satisfy a given clients' needs, only one will be optimal. Expertise in optimal experimental design is relatively recent and in some cases is still in the research phase within the field of applied statistics. We shall work with experts at the University of Minnesota to identify knowledge and heuristics of optimal design. The second area of investigation concerns the development of more explicit methods of knowledge acquisition in the design area. Thus far we have relied extensively upon protocol studies of experts solving simulated client problems. In the coming year we shall employ what we have termed elsewhere perturbated tasks (see discussion below) in order to elicit deeper levels of expert reasoning. In addition, we shall study consultants interacting with actual clients in order to understand the role of the consultation process. Finally, we shall investigate the kind of explanations and justifications offered by the expert both to other experts and to the design client. Our objective here is to discover the kind of representation that supports expert communication regarding the nature of problem solving activity as well as problem solving itself.

Chemical Reactor Design and Auditing.

The next two areas of research, chemical reactor design and auditing are both still at the phase 1 level of research. Therefore, we will describe only the basic motivation for the research, and our plans for the comming year.

Our reason for investigating expertise in the field of chemical engineering is that a preliminary analysis of experts performing a reactor design task convinced us that problems in this area are solved by a process of classification rather than design. This observation prompted us to investigate similarities and differences between expertise in classification as we have observed it in diagnostic tasks (medicine, hardware diagnosis) and expertise in classification in design.

In collaboration with two faculty members in the Chemical Engineering department at the University of Minnesota, as well as experts in the 3M Company, we have continued a pilot study of expertise in chemical reactor design. Preliminary results from protocol studies of individuals solving simulated design tasks indicate that experts appear to estimate aspects of the design problem graphically rather than working out the exact form of solution by means of an appropriate set of differential equations. This characteristic of estimating rather than calculating exact results has led to a consideration of models of qualitative reasoning as a means of describing the expert's knowledge.

The major question of interest in the work in chemical engineering is how experts employ qualitative reasoning techniques so as to simplify the complexity of a design problem and at the same time maintain a representation that is powerful enough to result in a successful solution. In the coming year, we will investigate expertise in reactor design, especially in the setting where chemical engineers seek the advice of industrial statisticians. This will allow us to develop a chemical engineering knowledge base for the off-line control system project described above. Our goal is to describe in detail the nature of the consultation process that occurs between the expert statistician and a client designer.

Finally, we are investigating expertise in the field of auditing in the task of second partner review. In this work, we are applying the methodology developed during the design and testing of our earlier Galen diagnostic system. In this methodology we employ knowledge acquisition tasks in which the knowledge of a given set of experts "almost applies." The rationale for the use of such tasks is that when we give experts tasks to which their knowledge is well adapted we learn more about the task and its demands than we do about the expertise that accomplishes it.

The basis of our methodology is the development of tasks which contain unusual or anomalous combinations of cues, but which are nevertheless real examples of problems encountered by experts in the domain under investigation. In medicine the tasks with unusual combinations of cues (to which we have given the name Garden Path) tend to occur naturally, though infrequently. In auditing such tasks are often designed purposely to mislead the expert. They are called fraud, and form the basic tasks investigated in this research.

In the coming year, we will investigate how expert auditors detect fraud and error in an auditing case when they are playing the role of a second partner reviewing the work of another in-charge partner. During this past summer we constructed three experimental tasks based upon actual audit cases. One of the major accounting firms in the Twin Cities area has agreed to provide experts for the study. Our major objective will be to develop a model of expertise for the audit review task. In developing this model we will draw upon work previously performed by Rayman Meservy (1985) and reported in last years AI Pinnacle review. In the next phase of research (phase 2), an expert system will be developed and tested.

Knowledge-Based Decision Support for Professionals.

Last year we reported on the design of the program BRAMBLE. In BRAMBLE we implemented a computational model of processes of reasoning by analogy. The program was different from many typical knowledge-based systems in that it included explicit models of contextual knowledge and the role of contextual knowledge in analogical reasoning.

Work on the BRAMBLE project has continued during the past year. Parts of the computer program have been redesigned or refined. The principle effort on the project during the year has been the development of a plan to implement a larger knowledge base for BRAMBLE and to design experiments to test its effectiveness.

The underlying premise in this project is that a decision support system will be a more effective tool if it "thinks like the user." A continuing line of research by members of the AI Pinnacle demonstrates that expert reasoning processes in a task domain tend to follow certain *lines of reasoning*. We believe that some aspects of these lines of reasoning serve as a conventional form of communication between professionals in a field and as a medium for problem solution. Thus an expert system or decision support tool that uses these same lines of reasoning will seem like a natural tool to the professional user.

An example of this can be seen in legal problem solving. Attorneys rarely think of a legal problem in terms of fixed legal rules which can be applied to a factual situation. Rather they look to analogical cases to derive rules that fit the particular context of a problem. Thus a rule-based expert system which attempts to model the law as a set of fixed rules is simply not "thinking like a lawyer." On the other hand, knowledge-based decision support tools that aid the attorney in identifying analogous cases and in mapping the consequences of an analogy fit the actual problem solving processes that the attorney is trained to use.

Having developed a computational model that supports representation of contextual knowledge and the use of that knowledge, we now have a tool that can be used to model problem solving processes that seem to place special emphasis on the use of context. In order to test this tool we need to implement a knowledge base so that the program becomes an expert system that can perform problem solving tasks. Furthermore, if the knowledge for the program is modeled from the actual knowledge used by human experts, then the program becomes a simulation of the expert problem solving process. This is important not because we are trying to create a psychological model of the user: rather we are concerned with designing a program which displays lines of reasoning which are familiar to and understandable by users trained in the problem domain.

Thus we are faced with a problem of knowledge acquisition: until we have a testable knowledge base the BRAMBLE system remains only theory. Sources of knowledge that will be investigated in the coming year include publications in a domain, consulting experts, and observation of experts performing problem solving tasks. Since BRAMBLE is particularly concerned with the role of the contextual knowledge that a problem solver takes into account, our knowledge acquisition will concentrate upon data collected from experts solving actual client problems.

In 1985 Bonnie Pechtel completed a study of legal reasoning. In her study, Pechtel used verbal protocols collected from two dozen attorneys engaged in an experimental problem solving task which she had designed. The task was the analysis of proposed corporate acquisition with a view to advising a client concerning legal issues raised by the acquisition. Pechtel analyzed her data in order to identify lines of reasoning and problem solving processes. Her subjects included both expert and novice attorneys, and the subjects were selected from different areas of legal specialization.

In the coming year, we will reanalyze the Pechtel data in order to empirically identify knowledge structures used by the subjects in her study. Traces of subjects' reasoning in this study show the importance of context setting activities and the way in which the context focuses and directs reasoning processes. From this data we will infer the structure of the underlying contextual knowledge structures for incorporation of this knowledge into BRAMBLE.

Finally, a verification study (phase 2) will be conducted. In this study, the verification output of BRAMBLE will be tested against the reasoning of expert attorneys to see the extent to which it has accomplished our goals of incorporating expert knowledge and expert lines of reasoning.

Expert Systems - Expert System Shells.

Expert system technology aims to represent and apply knowledge obtained from a specialist in a problem domain. Since most of the operational code can be separated from the domain specific knowledge, one program can be written to handle rule bases from several domains. Thus a system can be developed for a new domain simply by changing the rules that the operational system, now called an expert system shell, handles. We have developed an expert system shell called AGNESS (A Generalized Network-based Expert System Shell) as a successor to our previous shell, GENIE. AGNESS uses a computation network rather than a production rule base and supports values of any well-defined data type, the Merit questioning scheme, an explanation facility, and expert-defined inference methods. AGNESS was designed to provide greater representational power, simplified knowledge specification, and smaller storage requirements. Moreover, representational structures are shared whenever possible, and some redundancies of the previous system are reduced to increase storage efficiency. We are also in the process of developing a logic-based language, called PEIRCE, that can be used for describing and implementing expert classification systems.

In AGNESS, each general proposition is represented by a node in the network. A value is associated with each node in a given *context*. For example, given a node AGE1 representing the general proposition of a person's age and the context (Steve), the two together have an associated value, say 24, which is stored in a database separate from the computation network. Since domain specific knowledge is relatively fixed, it is represented directly in the computation network. User-supplied and problem-specific knowledge, on the other hand, is more volatile, and so is stored in a separate database.

The triple made up of a node, context, and value is called a *datum*. Edges in the network represent possible dependencies of one datum's value on that of another; each edge connects an antecedent with a consequent. A node may be linked to arbitrarily many nodes and arbitrarily many nodes may be linked to it. Associated with each node is a function -- the *inference method -- that takes the values of the antecedent data and generates the values of the consequent data. Nodes with no consequents are top nodes; those with no antecedents are bottom nodes. The ultimate consequent datum in the problem at hand is the topic. The topic may be associated with an interior just as well as a top node since a particular problem instance may be concerned with only part of the computation network.*

AGNESS provides a forward chaining mechanism -- the *propagation process* -- and a backward chaining mechanism -- the *questioning process.* Propagation is invoked each time a new value is added to the database. It updates the database so that the values of the consequents of the modified datum, the values of their consequents, and so on, up to the values of the affected top data, are consistent with the new data.

Questions directed to the user -- the questioning process -- present the major bottleneck for most expert systems. Older consultation systems generally follow an exhaustive depth-first network traversal to direct the consultation process. AGNESS, however, like its predecessors, assigns every datum whose node is marked *askable* a weight, called its merit value, to represent the ratio of the expected change in the value of the topic to the cost of supplying a value for the datum. Thus merit is a measure of the utility of requesting information. (Questions are never asked for unaskable nodes.) AGNESS is able to direct its consultation sessions more efficiently, using a best-first strategy, because the user is asked a question with highest merit value at each stage in the consultation. This not only substantially reduces interaction time by avoiding unnecessary nodes, but also assures the user that the most critical questions are asked first in a time-critical task.

Two factors contribute to a datum's merit: its self-merit and edge-merits along the path from the topic. The self-merit of a datum is an attribute of the corresponding node, so is unaffected by change of context. It is an expert-defined approximation of the ratio of the expected change in the value of a datum to the expected cost of determining the value. An edge-merit value for an antecedent-consequent pair is found by evaluating the partial derivative of the consequent's assignment function with respect to the value of the antecedent in question; the assignment function of a node is a procedure for deriving the value of a consequent datum from the values of its antecedent data. The product of the datum's self-merit and the various edge-merits along the path from the topic to the node gives the merit value of that node.

Progress.

When the computation network has the form of a directed acyclic graph in general, rather than simply the form of a tree, merit calculation is complicated because a node may have more than one consequent. For then there may not be a unique path from the topic to a node, and we defined merit in terms of the edge-merits along THE path from the topic. This problem has been solved, however, and the solution incorporated into AGNESS. The solution involves computing the merit of a node on more than one path in terms of its self-merit, the merits of its consequents, and the edge-merits of the edges to its consequents, and then comparing the merits of nodes at various levels.

In most realistic applications, the values supplied to an expert system are not known with certainty. This uncertainty is propagated through the system to all values depending, however remotely, on the supplied values. Furthermore, various inference methods themselves introduce uncertainty. The method we have considered to model uncertainty is classical probability theory; in particular, we have considered propagating probability distributions in a computation network. At any step, distributions with known parameter values are associated with the antecedents, and one must infer distributions with explicit parameter values for the consequents; in general, the consequents' distributions need not be of the same family as the antecedents'. Bayesian probabilities have been used in the propagation.

The theory has been worked out to extend AGNESS to vector-valued computation networks. The difficulty due to child-node vectors of different dimensions is overcome by multiplying each by an appropriate transition matrix to furnish vectors of a common dimension. Assignment functions are defined to take vector arguments and compute vector-valued results. Edge-merit may then be defined in terms of the Jacobian matrix (the matrix analogue of the derivative) of a differentiable assignment function. This scheme has been extended to allow probabilistic vector values, assignment functions, and edge-merits; the extension allows for either probability distributions, described by family and parameter values, or discrete probabilities.

AGNESS may be used as the first phase in a general resource allocation program. The general problem is to assign limited resource items to a set of tasks. In the first phase, AGNESS determines the effectiveness of each individual resource as applied to each prospective task. The second phase then evaluates composite allocations, that is, possible plans for the use of all resources in performing all tasks. Previously, this second phase was accomplished by traversing and pruning an allocation tree. Recently, however, we have shown that simpler cases of this phase may be done significantly more efficiently by use of the Hungarian Assignment Algorithm from operations research. This algorithm has been implemented in Common Lisp.

Explanation increases confidence in the system's conclusions and is invaluable in building and debugging the rule base. Early expert systems included two basic explanation queries: why a question is being asked and how a value was computed. AGNESS includes an augmented set of explanation queries designed to increase information availability while remaining sufficiently simple to be practical in most of today's expert system shells. These queries are available during the questioning process and whenever the user may view a database entry, as when the system presents the conclusions. As part of the explanation facility, a graphics interface was built to allow the user to view the computation network as a collection of nodes and edges and to ask for information about any node.

The goal of the PEIRCE project is to design a logic-based language that can be used for describing and implementing expert classification systems. The first implementation of this language augmented traditional first order logic with a modal operator for belief. Modal operators indicate attitudes toward a proposition other than the normal default of asserting truth. Thus if we just write **p**, we assert that proposition p is true, whereas if we write **Bp**, the modal operator **B** indicates that proposition p is believed but not necessarily true. Observations input to the system were represented as non-modal propositions. Hypotheses, and inferences from them, were represented by belief propositions. Since hypotheses are often incorrect, it was necessary to assure that the logic could detect and manage inconsistencies without allowing the derivation of any arbitrary proposition (as would be the case with standard first order logics). To this end, many of the ideas of *relevance* logic were incorporated into the system. Relevance logic disallows the *fallacies of relevance* that first order logics allow, among which is the inference from *p* and *not p* to any arbitrary proposition *q*.

While testing the first implementation, two problems became apparent. First, the language still did not contain enough different modal concepts to represent theories of classification adequately. Second, additional control structures were required to prevent computation of useless facts. Correcting these deficiencies required reformulation of the logical theory underlying the system, and this has been the major effort of the past year. The new theory uses multiple modal operators to represent different kinds of assumptions that can be made, including assumption by hypothesis, by observation, by inference, by logical definition, by terminological definition, by historical association, and by causal association. Better control is achieved by isolating incompatible alternative hypotheses so that they are no longer used together to make inferences about "impossible" events.

Plans.

Thus far, the major application of AGNESS has been to implement the clinical expert system ETA (Exercise Test Analyzer). The cases studied come from the Program on the Surgical Control of Hyperlipidemas (POSCH), a study of the effect of reduced cholesterol in heart attack victims. ETA outperformed both the average POSCH cardiologist and a multiple regression model used to analyze the data. Use of an expert system has the advantage that the system can use rules approximating the reasoning process followed by a clinician and can explain the process in a useful way. As a benefit of using AGNESS in particular, both numeric and nonnumeric values, along with expert-defined inference methods, are used; this makes ETA extremely natural and easy to follow. Work has been resumed on ETA, and will continue in the coming year.

We also have begun a knowledge engineering and expert system project, with emphasis on knowledge acquisition, in the domain of cystic fibrosis. Dr. Warwick of the University of Minnesota Hospitals will be the domain expert. Currently the expert looks at home reports and previous medical records to determine whether a patient should visit the clinic immediately or wait for his regular quarterly visit.

We shall incorporate into AGNESS the theoretical work done for handling uncertainty, namely, the adaptation of classical probability theory to the sort of computation network found in AGNESS. This method will also be extended to cover merit questioning. Also to be incorporated is the theory developed for vector-valued computation networks. We shall as well combine our implementation of the Hungarian Assignment Algorithm with AGNESS to form a resource allocation shell. In the process of implementing these various features, we shall hone the notions involved, experiment with variations, and possibly extend the theoretical work.

The Merit scheme should be made practical for a large computation network, say of 7,000 or more nodes. This will involve circumscribing the part of the network considered in both the propagation and questioning processes. The Merit scheme should also be made practical for a large -- 50,000 or more items -- database. We have already addressed these areas, and will continue our investigations.

Future work in explanation will concentrate on two areas. First, we are investigating more extensive ways of using the Merit scheme to improve explanation. For example, merit based explanations should give insight into the relative importance of data to different hypotheses. Second, we are working on explanation systems that are more advanced than the queries referred to above. These explanations will be based more on causal relationships than on the rules themselves and will justify the system's conclusions with domain principles rather than merely describe the reasoning processes leading to these conclusions.

We plan to reimplement PEIRCE according to the new theory mentioned above. The result will be a new form of *truth maintenance* system that will support multiple modalities and isolate incompatible but competing *possible worlds*. A possible world is simply an alternative internally consistent model of the subject or event being classified. This system will then be used as a basis for implementing an expert classification system.

Computer Vision.

Automated visual processing is needed for tasks ranging from manufacturing to autonomous vehicle guidance. Sensor-based robots use vision and other techniques to allow closed-loop control of manipulators. Automated inspection techniques utilize image processing procedures to provide efficient and cost effective flaw detection. Surveillance systems must processes large quantities of image data, looking for particular features. Autonomous navigation requires vision systems capable of recognizing obstacles, comparing the visual environment to prestored maps, and locating objects relevant to the vehicles assigned goals.

Most of the computer vision systems currently in commercial use have been built for special purpose applications. These systems are engineered for individual tasks. Often, cost and processing speed requirements limit the technology that is utilized. Binary imagery and silhouette-based processing are still common, restricting visual analysis to two-dimensional object boundaries. Our long-range goal is the development of general-purpose vision systems. Such systems must embody techniques for interpreting visual data applicable to a broad class of situations. While improvements in processing power and algorithm development will help, in large part this generality must come from a better understanding of the basic principals of vision.

Most available computer vision methods apply to single, static images. More recently, attention has been directed at the analysis of time varying imagery. Much of the information required in computer vision applications is intrinsically dynamic. Moreover, even static information such as the location of obstacles or the shape of objects can be easier to determine if either the camera or object is moving. As a result, our work currently concentrates on techniques based on visual motion, though the problem of integrating information from other visual and non-visual sources is obviously a crucial factor in understanding perceptual information. Motion-based computer vision is much less subject to problems with ambiguity than methods that deal with one image frame at a time. Image dynamics obviously provide information about scene dynamics such as sensor trajectory. Visual motion is also, however, a rich source of information about scene geometry. Methods already exist for passive ranging based on motion. In theory, motion can also be used to determine such properties as surface orientation, though significant practical limitations exist. Current efforts involve a variety of motion-based techniques for finding object boundaries, determining spatial relationships, passive ranging, and the recognition of certain types of objects and events.

Progress.

Work is continuing on problems associated with *motion-based segmentation.* Visual motion can be used both to find and to interpret object boundaries. Discontinuities in optical flow are necessarily due to surface boundaries or discontinuities in depth in the scene. (Optical flow is the image plane velocity of the projection of points on visible surfaces.) Thus, detected edges in flow necessarily correspond to important properties of scene geometry, where as edges in properties such as luminance can be due to a wide variety of scene properties. Because our approach is based on understanding the three-dimensional scene structure leading to an edge in the image, it allows the determination of important three-dimensional properties of the associated scene surfaces. Visual motion can be used to distinguish between occlud*ing* and occlud*ed* surfaces at an edge. The method is equally applicable to optical flow edges or those due to other image properties such as discontinuities in brightness. Occlusion boundaries arise due to geometric properties of the occluding surface, not the occluded surface. Thus, while the shape of the edge provides significant information on the structure of the occluding surface, it says little or nothing about the structure of the surface being occluded. This technique may make it possible to link image regions corresponding to a partially occluded object and to produce descriptions of object boundaries that are less affected by occlusion.

Figure 1 shows a simple example. Each line in the figure corresponds to an occlusion boundary. Two regions are shown, corresponding to two object surfaces. The object to the right is in front of and partially occluding the object to the left. The object to the left is in turn occluding a portion of the background. A traditional segmentation techniques applied to the portion of the image corresponding to the left object would produce an image region as in figure 1b. Such a region would be difficult to recognize as nearly half of the region boundary does not actually correspond to the object. In fact, the only portions of the contour of the left region that provide information about the silhouette of the object are those which correspond to the surface occluding some other surface. Motion-based segmentation allows the recognition of such boundary segments. In figure 1c, the occluding boundaries of the left object are shown as dark lines, while other boundaries, recognized by the occlusion analysis, are shown as dashed lines. A segmentation such as shown in figure 1c is likely to be quite a bit easier to recognize than that shown in figure 1b.

Motion-based segmentation uses visual motion to determine geometric properties of the scene. We are also investigating methods for determining temporal scene properties. One important function of a vision system is to recognize the presence of moving objects in a scene. If the camera is stationary and illumination constant, this can be done by simple techniques which compare successive image frames, looking for significant differences. If the camera is moving, the problem is considerably more complex. We have developed a number of techniques for detecting and locating moving objects based on imagery from a (possibly) moving camera.

If the parameters of camera motion are known (e.g. from inertial guidance systems or other sensors), a computationally simple and reliable technique can be used to find moving objects. If the exact camera motion is not known, the problem is significantly harder. Often, however, other simplifying assumptions can be exploited. One such technique we have developed applies to situations in which the camera is actively tracking an object of interest. Though the tracked object will remain stationary in the image, motion of the object with respect to

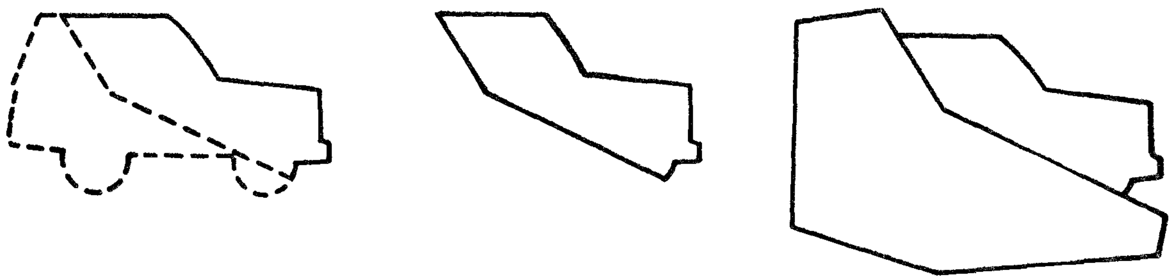

(a) Segments (b) Left region (c) Silhouette

Figure 1

the environment can be detected based on the visual motion of the area surrounding the object in the image.

Techniques similar to those used for motion-based segmentation are also relevant for determining other aspects of spatial organization. Identification of occluding and occluded surfaces relates directly to the relative depth of surfaces. For example, it may be straightforward to find that surface **A** is in front of surface **B**, even though the actual positions of **A** and **B** are difficult or impossible to determine. Qualitative descriptions of geometry most often are phrased in terms of relationships between geometric entities (e.g. the *in front* relationship above). We call such descriptions specifications of *topological spatial organization* and distinguish them from more precise specifications of *topographical spatial organization.*

Much of the existing work in estimating scene geometry uses cues such as stereo, shading, texture, and motion to determine precise positional properties -- principally 3-D location and/or surface orientation. These methods are seldom able to deal with the variability, complexity, and noise encountered in real imagery. Robustness can be substantially improved if the analysis need only provide a qualitative description of a scene. While a qualitative analysis may at first seem much less desirable than a quantitative analysis, this is not necessarily the case. Substantially greater reliability is possible if only qualitative descriptions are required. In addition, in unfamiliar environments, goals and maps are likely to be stated topologically, rather than topographically. (Eg. "turn left at the second junction.")

These ideas have led to the development of a preliminary model of a process we refer to as *inexact vision.* The basic argument in support of the inexact vision approach is that vision systems cannot determine scene geometry precisely under any but the most trivial of circumstances. Furthermore, efforts at determining precise geometry lead to complex, numerically unstable algorithms. Increased robustness is possible if the requirements for precision are reduced by designing vision algorithms which only determine a partial specification of scene properties.

The majority of the work in computer vision over that last ten years has been plagued by two problems: 1) unrealistic constraints are required in order to invert the image formation process, and 2) even when these assumptions are made, the problems remain numerically ill-conditioned and highly prone to error. Both problems are associated with the ambiguity inherent in basing interpretations on local image features over surfaces. The inexact vision approach emphasizes the computation of less specific scene properties that can be determined reliably, even when presented with noisy input data. One key to this approach is to develop a representation methodology for describing imprecise geometric information. The use of a range of measurement scales seems a good first step towards dealing with this issue. In addition, investigations of alternate measurement scales provides an increased understanding of the problems which arise when integrating information derived from alternate visual cues, as these alternate sources of information are often most naturally represented in different measurement scales.

Planned Work.

The development of the inexact vision methodology is a difficult task and will continue for some time. The immediate questions to be dealt with are: 1) the applicability of our ideas to non-motion vision, 2) the characterization of measurement scales for multi-dimensional data, and 3) general principles for combining visual cues represented in different measurement scales. In the longer term, both a *competence-based* and a *task-based* analysis of visual systems will be required. Competence-based analysis involves the determination of what is possible to compute and, equally important, the accuracy of such computations. Task-based analysis focuses on what information actually needs to be computed at each stage of processing. For example, to determine the potential three-dimensional information that can be recovered from image motion, we must understand the characteristics and limitations of the motion estimation techniques. The characteristics of the motion estimates determine in large part the success of later stages of processing. Methods of motion analysis must not rely of estimates that are impossible to obtain.

To provide a well defined task within which to investigate automated vision, we are initiating a major new effort to study *perception for navigation.* Intelligent machines must be able to maneuver effectively within their environment. In all but the most simple of circumstances, this maneuvering requires sophisticated navigational skills in order to determine an appropriate path for motion and to make sure that the intended motion is actually carried out. If perfect information about the world were available, then navigation would consist of a fairly straightforward problem in geometric reasoning. In practise, complete information about the structure of the environment is seldom available. Furthermore, mechanical imprecision makes it impossible to precisely predict the effect of a particular action. As a result, navigation requires closed loop control in which perceptual systems are used to gather information about the world and the effectiveness of plans for moving through that world.

The first step in the navigation project focus on *map orientation* -- the establishment of a correspondence between a map and visual observations of the terrain described by the map.

The map orientation task starts with one or more images taken from a particular vantage point and a topographic map including most or all of the terrain viewed in the images. The problem is to find: 1) key topographic features on the map and the corresponding visual features in the imagery, 2) key image features and the corresponding topographic features on the map, and 3) the location on the map of the observation point and the direction of view for each image.

Map orientation is a good problem within which to investigate the geometric representations needed by a sophisticated perceptual system.. A map represents precise geometric properties, while image cues are more relevant to relationships (e.g. in-front-of) than to precise position. In addition, map orientation requires significant change of perspective from a down-looking map to horizontal-looking imagery.

As an example, consider this image:

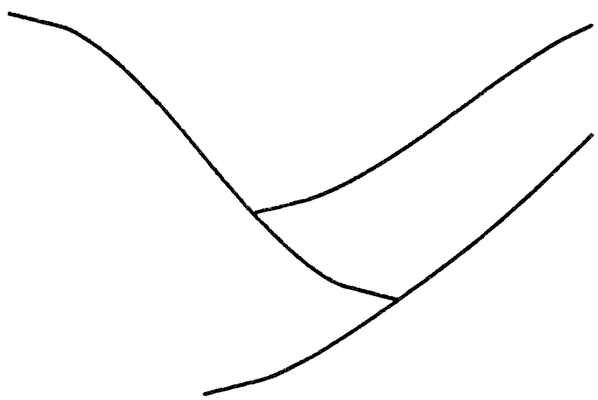

the corresponding topographical map:

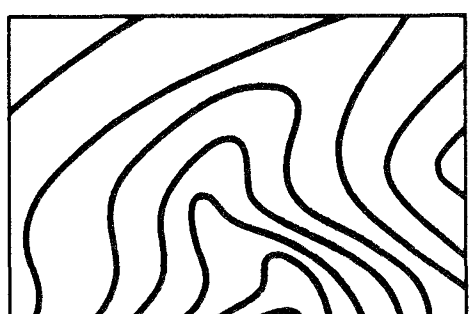

and the intermediate representations likely to be necessary to establish correspondence:

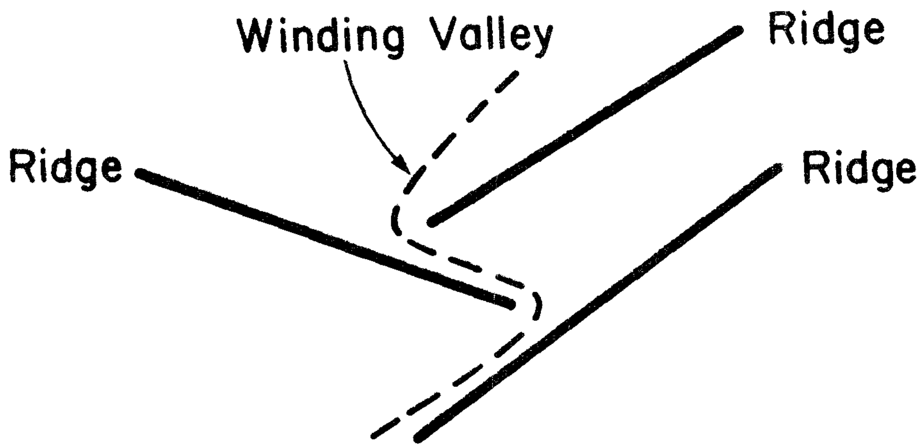

Map orientation requires effective integration of low-level vision, models of spatial organization, and problem solving. Considerably more problem solving competence is needed than for visual tasks so far investigated. Multiple matching directions must be supported. Image features can be searched for in a map. Map features can be searched for in an image. Both map and image features can be searched for in some intermediate representation. Much of this search will require efficient hypothesize-and-test matching strategies.

Finally, we are continuing our study of the fundamental properties of visual motion that make it possible to determine the three dimensional structure of a scene under view. This effort is aimed at the discovery of both new methods for using motion information, and ways of improving the reliability and effectiveness of existing techniques.

Artificial Intelligence Laboratory.

The Artificial Intelligence Laboratory provides computational facilities for research in artificial intelligence. Laboratory equipment has been purchased with major funding provided by MEIS, NSF, and AFOSR, along with an equipment donation from SUN Microsystems. Space is provided by the Computer Science Department. Operational support is assisted by the systems staff of the Computer Science Department.

Artificial intelligence research requires high performance computers capable of running specialized software. Highly interactive systems are needed to develop complex AI systems. Specialized peripherals are also often required. The laboratory supports research in both computer vision and expert systems. Vision research requires cameras and graphics displays closely coupled to high performance computers. Work in expert systems requires Lisp machines and associated, specialized software tools. The Artificial Intelligence Laboratory provides these important resources.

Current facilities include 8 Sun-3 workstations, 3 Sun-2 workstations, a Sun-3 file server, 1,280 Mb of disk capacity, 4 graphics displays including one with a frame grabbing capability, 2 Dage-MTI high quality video cameras and a writable optical video disk to assist in the digitization and replay of time-varying imagery. Common Lisp is supported, along with a variety of our own and other vendor's software tools.

Curriculum Development.

The Intelligent Systems group has recently completed a major revision of courses in artificial intelligence taught by the Computer Science Department. These revisions significantly expand the courses offered, providing an integrated curriculum for students concentrating in AI. The core of the new AI curriculum is a three quarter sequence intended for advanced undergraduates and graduate students. The first two quarters give an overview of all important aspects of the field. The third quarter is a laboratory course on AI programming techniques. These courses are complemented by three regularly scheduled graduate level courses and several specialty courses offered on an ad hoc basis. Courses on computer vision, expert systems, and AI techniques for robotics are offered each year. In addition, research seminars and special topics courses are frequently given. The AI curriculum is supported by courses in pattern recognition, cognitive psychology, perception, and linguistics taught by a variety of departments.

Other Activities.

Intelligent Systems Group is involved in a wide variety of technology transfer activities aimed at facilitating communications between the group and others interested in artificial intelligence. One of the most effective of these activities has been a *visiting scientist* program in which industrial researchers have worked directly with the Intelligent Systems Group on problems of common interest. In the last year, personnel from 3M, CDC, and Honeywell have participated. In May, the group sponsored a day-long symposium providing an overview of much of the important AI work currently underway in the Twin Cities area. Professor Lotfi Zadeh of the University of California-Berkeley delivered the keynote address, "Formalization of Common Sense Reasoning and Knowledge Representation". This was followed by presentations by ten industrial scientists from Control Data Corp., Honeywell, Inc., 3M Company, and Sperry. It proved to be a day of fruitful exchanges and its success has led to plans for a similar symposium next year. Finally, two colloquia series have been held this year. In the spring, group members made presentations on recent research accomplishments. In the fall, four nationally prominent researcher were brought to the University for presentations and informal interactions with group members, students, and industrial scientists.

Investigators.

Three principal investigators direct the research activities of the Intelligent Systems Group:

William Thompson, Computer Science Department.

Professor Thompson is actively investigating problems in computer vision and in expert systems. His work in vision is focused on discovering the processes involved in perceiving the spatial relationships between objects and surfaces. Significant results have been achieved in describing how motion provides information about spatial organization. For the last six years, Thompson has also been active in the study of expert systems. Together with Paul Johnson, he has developed a formal model of diagnostic reasoning capable of expert level performance.

Paul Johnson, Department of Management Sciences.

Professor Johnson's research focuses on the investigation and formalization of expertise in professional and technical fields including science and engineering, medicine, management, and law. He heads the GALEN project that is developing computational models for fault diagnosis and that has built a high performance diagnostic expert system at Minnesota. Johnson's current work addresses two major barriers to the development of more effective expert systems: 1) the need for tools and techniques for identifying and transferring expertise from human to computer environments, and 2) the lack of an understanding of efficiency mechanisms employed by human experts in order to avoid exhaustive search of large data bases in solving relatively unstructured problems.

James Slagle, Computer Science Department.

Professor Slagle is interested in all aspects of artificial intelligence, including automated deduction, heuristic search, clustering pattern recognition, learning, robotics, multipurpose problem solving, symbolic mathematics, and especially expert consultant systems. Slagle's earlier work in heuristic search led the way to modern symbolic mathematics systems. He is currently particularly interested in control strategies and explanation strategies in expert systems.

Other University faculty are affiliated with the Intelligent Systems Group through their participation in joint research activities and/or their use of AI Lab facilities. Current affiliated faculty include John Carlis, Maria Gini, and Ting-Chuen Pong in Computer Science, Mostafa Kaveh and Harry Wechsler in Electrical Engineering, Albert Yonas in Child Development, Christopher Nachtsheim in Management Science, and Glen Berryman in Accounting.

Graduate Students.

Martha Arterberry, Charles Butler, Keith Bellairs, Alice Chan, Bharat Charan, Tass Dimitropoulos, John Doyle, Albert Esterline, Claude Fennema, Gregory Frascadore, James Held, Ian Horswill, Erach Irani, Karim Jamal, Chung Lee, Paul Krueger, Ruth Meyer, James Moen, Madhu Nigam, Bruce Petrick, David Pogoff, Marius Poliac, Nancy Reed, Steven Savitt, Kent Spackman, Elizabeth Stuck, Nikolas Vasillas, Rand Whillock, Michael Wick, Lee Zimmerman, Imran Zulkernan.

Research Support.

External research support is provided by NSF, AFOSR, NICHD, NIH, CDC, IBM, and Sun Microsystems.

Recent Publications.

K.A. Smith, B. Farm, and P.E. Johnson, "Surface: A Prototype Expert System for Selecting Surface Analysis Techniques," *Proceedings of the IEEE Phoneix conference on Computers and Communications,* March 1985.

J. Slagle and H. Hamburger, "An Expert System for a Resource Allocation Problem," *Comm. of the ACM,* Vol.28, No.9, Sept.(1985), pp.994-1004.

J.A. Larsen, W.F. Kraemmerer, K.L. Ryan, J. Slagle and W.T. Wood, "ATOZ -- A Prototype Intelligent Interface to Multiple Information Systems," *Proceedings of the Conference on the Future of Command Languages: Foundations for Human-Computer Communication,* Rome, Italy, Sept. 1985.

R.L. Madarasz and W.B. Thompson, "Recognition of Moving Objects Using Feature Signatures," *IEEE Transactions on Pattern Analysis and Machine Intelligence* November 1985.

A.D. Bailey, Jr., G.L. Duke, P.E. Johnson, R.D. Meservy, and W.B. Thompson, "Auditing, Artificial Intelligence, and Expert Systems," *Decision Support Systems: Theory and Application,* A.B. Whinston, ed., North Holland Publications, 1985.

P.E. Johnson, "Cognitive Models of Expertise," *Proceedings of the USC Symposium on Expert Systems for Auditing,* February 1986.

W.B. Thompson, "Comments on 'Expert' Vision Systems: Some Issues," *Computer Vision, Graphics, and Image Processing,* April 1986.

W.B. Thompson and J.K. Kearney, "Inexact Vision," invited paper, *Proceedings Workshop on Motion: Representation and Analysis,* May 1986.

J.M. Long, J. Slagle, M.R. Wick, J.P. Matts, A.S. Leon, and the POSCH ECG Panel, "Using an Expert System to Automate the Element of Clinical Judgment in Data Analyses," *Proc. of the Seventh Annual Meeting of the Society for Clinical Trials,* May 1986, Montreal.

J. Slagle, M.R. Wick, and M.O. Poliac, "AGNESS: A Generalized Network-Based Expert System Shell," *Proc. of the National Conference on Artificial Intelligence,* August 1986.

J. Slagle, J.M. Long, M.R. Wick, J.P. Matts, and A.S. Leon, "The ETA Project: A Case Study of Expert Systems for Analysis of Serial Clinical Trial Data," *Proc. of Medinfo 86: 5th World Conference on Medical Informatics,* October 1986.

J.K. Kearney, R. Sedlmeyer, W.B. Thompson, M. Adler, and M. Gray, "An Analysis of Software Complexity Measures," *Communications of the ACM* November 1986.

J.K. Kearney and W.B. Thompson, "Bounding Constraint Propagation for Optical Flow Estimation," in *Motion Understanding: Robot and Human Vision,* W.N. Martin and J.K. Aggarwal, eds., Kluwer Press Academic Publishers, 1986.

J. Slagle, "Alpha Beta Pruning," *Encyclopedia of Artificial Intelligence,* S. Shapiro (ed.), John Wiley and Sons, Inc., New York, 1986.

J. Slagle and M. Gini, "Problem Reduction," *Encyclopedia of Artificial Intelligence,* S. Shapiro (ed.), John Wiley and Sons, Inc., New York, 1986.

J. Slagle and M. Gini, "Pattern Matching," *Encyclopedia of Artificial Intelligence,* S. Shapiro (ed.), John Wiley and Sons, Inc., New York, 1986.

J. Slagle and H. Hamburger, "Heterogeneous Discrete Expenditure for Diminishing Returns," *Naval Research Logistics Quarterly,* John Wiley and Sons, Inc., New York, 1986.

J. Slagle, J. Long, M. Wick, Matts, J., and A. Leon, "An Expert System Treadmill Exercise ECG Test Analysis," *Proc. of the 1986 ACM Annual Computer Science Conference,* 1986, p. 421.

J. Slagle, J. Long, M. Wick, Matts, J., and A. Leon, "Expert Systems in Medical Studies -- A New Twist," *Proc. of SPIE -- The International Society for Optical Engineering,* Vol. 635, *Applications of Artificial Intelligence III,* 1986, p.25.

J. Slagle, M.W. Gaynor, and H. Hamburger, "A Decision Support System for Fire Support Command and Control," *High Technology Initiatives in C3I,* Stephen J. Andriole (ed.), Washington D.C., AFCEA International Press, 1986.

P.E. Johnson, J.B. Moen, and W.B. Thompson, "Garden Path Errors in Medical Diagnosis," in *Computer Expert Systems,* L. Bolc and M.J Coombs, eds., Springer-Verlag (in press).

R. Meservy, A. Bailey, and P.E. Johnson, "Internal Control Evaluation: A Computational Model of the Review Process," *Auditing: A Journal of Practice and Theory* (in press).

P.E. Johnson, "Problem Solving," in *Encyclopedia of Science and Technology,* McGraw-Hill (in press).