A SHORT INTRODUCTION TO EXPERT SYSTEMS

Jim Clifford

Matthias Jarke

Yannis Vassiliou

September 1983

Center for Research on Information Systems
Computer Applications and Information Systems Area
Graduate School of Business Administration
New York University

Working Paper Series

CRIS #59

GBA #83-97(CR)

Published in Database Engineering, Vol.6, No.4, pp.3-16

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

A SHORT INTRODUCTION TO EXPERT SYSTEMS

Jim Clifford, Matthias Jarke, and Yannis Vassiliou

Computer Applications and Information Systems Area
Graduate School of Business Administration
New York University

It is a generally accepted view among researchers in Artificial Intelligence that the 1980's will witness a tremendous upsurge in the number of successful applications of AI expertise to real-world systems. High on the list of the technologies that are expected to be applied in the marketplace are expert, or knowledge-based, systems. The formation of a number of expert system companies, often in close collaboration with major academic AI research centers, attests to the growing belief in the economic viability of this technology transfer. Although there is yet to be developed a formal theory of what constitutes an expert system, there are some general features that can be identified.

An expert system (ES), by definition, is a computer system which attempts to act like a human expert in some limited application domain. For decades people have certainly been building computer

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

systems that have attempted to be expert in their field of application -- no one has purposefully (unless maliciously) built a system that was intended to bungle its job! There are perhaps two aspects to an expert system that distinguish it from more traditional computer systems: overall architecture, and method of development.

An expert system architecture consists of two interacting components: a "knowledge base" and an "inference engine." The knowledge base contains all of the information that a human expert would normally need to carry out the desired task. This knowledge base itself is usually divided into two sub-components, the first containing specific, or "ground" facts (e.g., "Mary Smith is 35 years old"), and the second containing more general principles, rules, or problem-solving heuristics (e.g., "If a person is single then that person has no spouse"), which come from accumulated empirical observations or technical knowledge of the domain. An important feature of ES's is that both of these knowledge bases are stored declaratively in some assertion language, and not buried somewhere in computer code. This means that the knowledge incorporated into the system is easily accessed by the users, and potentially more easily modified or extended. The second component in an ES is a general purpose inference engine that is capable of making decisions from, answering questions about, and determining the consequences implied by the knowledge that is built into the system.

The other unusual aspect of expert systems is the manner in which they are constructed. The architecture of an ES in a way dictates the often-quoted motto of ES researchers that "in the knowledge lies the

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

power." What this slogan means is that the knowledge base component of an ES contains <u>all</u> of the domain-specific information for the application. In practice, because of the declarative nature of this knowledge base, and the power of the AI languages that have been developed for these systems, this has led to an incremental approach to ES development. Working in small teams of about 3 people, consisting minimally of the domain expert, a programmer, and a knowledge engineer, a small prototype ES is developed, usually in a matter of 2 or 3 months. The system is then successively refined in a process of examining its behavior, comparing it to that of the human expert, and correcting its reasoning processes by modifying its knowledge base. This process continues until the system performs at a level of expertise that approximates that of the human expert. At this point the system is ready for evaluation in the field. However, just as a human expert never stops developing or expanding his/her expertise, the ES is structured to facilitate continued growth and expansion of its capabilities.

In this short paper, some basic aspects of the structure and range of expert system applications are addressed, and directions of current research are indicated. Other comprehensive references on the subject are: [Barr and Feigenbaum 1982, Buchanan 1981, Davis 1982, Duda 1981, Gevarter 1982, Hart 1982, Hayes-Roth 1981, Hayes-Roth et al 1983, Michie 1980, Nau 1983, Stefik et al 1982].

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

1.0 ARCHITECTURE OF EXPERT SYSTEMS

For a long time, artificial intelligence has concentrated on the development of procedural techniques and representations such as heuristic search methods and problem transformation techniques. These have proven too general to solve real world problems in specific domains. Therefore, the focus has shifted to the representation and use of domain knowledge to guide search processes more efficiently.

The observation that human domain experts use domain knowledge as well as meta-knowledge (knowledge about the scope of one's knowledge and knowledge about how to use one's knowledge) efficiently has lead to the idea of extracting knowledge from a human expert into a knowledge base. The knowledge base is therefore at the heart of any expert system. It is a storehouse of knowledge in the form of specific facts and general rules, or in frames of reference that structure the expert's experience and expectations.

To exploit the knowledge, an inference engine is required that relates a problem description to the stored knowledge in order to analyze a certain situation (e.g., in medical diagnosis) or to synthesize a solution for a specific problem (e.g., a computer configuration). Such an inference engine can be a pattern matcher, theorem prover, or network search mechanism customized for one expert system, or it may exist already in the compiler of a corresponding knowledge representation language such as OPS-5 [Forgy 1980], Prolog [Kowalski 1979], or EMYCIN [van Melle 1979]). Even in the latter case, some additional control mechanism may be required to cut down the number of inferences to be made.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

The third major component of an expert system contains a number of user interfaces for various purposes. The two most important seem to be an interface for knowledge acquisition through which the expert or an intermediary can insert, update, and check knowledge in the knowledge base, and an interface through which end-users can get consultation from the expert system. As a windfall profit, the stored expertise can sometimes be made available to train new human experts.

2.0 KNOWLEDGE REPRESENTATION AND INFERENCE PROCEDURES

The knowledge base may require the description of facts about specific objects, relationships, and activities; of classification and generalization hierarchies; of general relationships between object and activity classes; and of meta-knowledge about the scope, importance, precision, and reliability of the stored knowledge. Just as database research has developed multiple representations for specific facts, many techniques exist to represent the more general knowledge required for expert systems.

A "good" knowledge representation should support the tasks of acquiring and retrieving knowledge as well as of reasoning. Factors that have to be taken into account in evaluating knowledge representations for these three tasks include:

1. the naturalness, uniformity, and understandability of the representation;

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

2. the degree to which knowledge is explicit (declarative) or embedded in procedural code;

3. the modularity and flexibility of the knowledge base;

4. the efficiency of knowledge retrieval and the heuristic power of the inference procedure (heuristic power is defined as the reduction of the search space achieved by a mechanism).

Below, four major knowledge representation techniques and their related inference mechanisms will be briefly reviewed. A thorough examination of knowledge representation is given in [Mylopoulos 1980].

2.1 Production Rules

Rules [Davis, Buchanan, and Shortliffe 1977] have been the most popular form of knowledge representation in expert systems. [Chandrasekaran 1983] points out three interpretations of the function of rules in expert systems. First is the interpretation of rules as a programming language. A rule typically has the form

if X then Y.

It can be used in computations in different ways. On one hand, in a data-driven or forward chaining approach, one can try to match a given situation to the condition X in order to infer a possible action Y. On the other hand, one can try to "prove" a hypothesis Y by establishing the preconditions X through further analysis (backward

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

chaining). Combinations of both methods are also sometimes used.

Both approaches require a pattern matching process, perhaps combined with unification (substitute constants or other variables for variables in the pattern to be matched) to identify the applicable rules in a given problem situation. If there is more than one of those, one has to be selected for further processing first. Control structures for rule application can be distinguished by their flexibility of rule choice into irrevocable ("hill-climbing") or tentative, and by the sequence of analysis in depth-first with backtracking or breadth-first with parallel graph search [Nilsson 1980].

Secondly, rules can be used as description tools for problem-solving heuristics, replacing a more formal analysis of the problem. In this sense the rules are thought of as "rules of thumb," incomplete but very useful guides to make decisions that cut down the size of the problem space being explored. These rules are input to an expert system by the human expert, usually iteratively and perhaps by means of an interactive program that guides and prompts the expert to make this task easier, and perhaps does some limited consistency checking.

Finally, rules have been proposed as in some sense a simulation of the cognitive behavior of human experts. By this claim, rules are not just a neat formalism to represent expert knowledge in a computer but rather a model of actual human behavior.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

A problem with rule-based techniques is the organization of the stored knowledge in a way that permits efficient yet transparent control over the search processes inside the knowledge base. There is currently no satisfactory formal solution to this problem but a number of ad-hoc programming tricks have been developed.

2.2 First-Order Logic

Precise knowledge can be stated as assertions over objects that take the form of first-order predicates with functions and equality [Kowalski 1979]. Logic has the advantage of offering a sound and complete set of inference rules. It is also purely declarative and therefore allows multiple uses of the same piece of knowledge. For inference purposes, predicates are usually transformed in a quantifier-free normal form called clausal form.

As an illustration, Prolog's [McDermott 1980, vanEmde and Kowalski 1976] inference procedure is based on the resolution principle [Robinson 1965]. In order to prove a theorem in clausal form, its negation is added to the set of knowledge clauses or "axioms". If the thus augmented conjunction of clauses can be shown to be contradictory, the theorem has been proved.

A major problem with general first-order logic as a knowledge representation is again the difficulty to express control structures that efficiently guide the use of a large knowledge base. To reduce such problems, practical tools such as the logic programming language Prolog use the subset of definite (Horn) clauses rather than full

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

first-order logic. Furthermore, these clauses are interpreted in a procedural way similar to backward chaining in production rules, leading to a more efficient search process while reducing somewhat the generality of interpretation possible in a nonprocedural interpretation.

2.3 Networks

Semantic networks [Quillian 1968, Brachman 1979, Schubert 1976] seem to be more popular in other AI applications (e.g., natural language processing) than in expert systems. Nevertheless, a number of expert systems rely on network formalisms, among them very large systems such as INTERNIST [Pople 1983], Prospector [Hart et al 1979], and SOPHIE [Brown et al 1981]. Networks are a natural and efficient way to organize knowledge. Nodes describe objects, concepts, or situations whereas arcs define the relevant relationships. Reasoning corresponds to network traversals along the arcs or to pattern matching of problem descriptions and subnets. A large number of exact and heuristic mechanisms exist for these tasks. The disadvantages of this approach stem from the lack of formal semantics making verification of the correctness of reasoning very difficult.

2.4 Frames

Much knowledge is based on experience and expectations adapted from previous situations and general concepts to a specific problem. Frames [Minsky 1977, Schank 1972, 1975, Bobrow 1977] provide a

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

structure to such experiential knowledge by offering so-called slots which can be filled with type descriptions, default values, attached procedures, etc. Frames are a very general and powerful representation form. It may be difficult, however, to specify their meaning precisely as well as to implement them efficiently.

2.5 Multiple Knowledge Representations

It should be clear by now that no one of the knowledge representation methods is ideally suited for all tasks. In very complex systems using many sources of knowledge simultaneously (e.g., speech recognition [Erman et al. 1980]), the goal of uniformity may have to be sacrificed in favor of exploiting the benefits of multiple knowledge representations each tailored to a different subtask. Similar to the interdisciplinary cooperation of several human experts, the necessity of translating among knowledge representations becomes a problem in such cases.

The need for translation also occurs when an expert system is interfaced with other software systems, e.g. database management systems.

3.0 USER INTERFACES

There are at least three distinct modes of interacting with the expert systems that are now being developed: consultation, knowledge acquisition, and training. Of course not every system allows these three types of interaction, nor is this interaction always facilitated

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

by the means of automated tools. Nevertheless the basic expert system architecture that has emerged has shown itself to be capable of at least these modes of interaction. In this section we will give a brief overview of these three interaction types.

3.1 Consultation

The primary mode of interaction is the consultation session, wherein the expert system is used to solve the problem for which it was constructed. There are really two forms that this interaction can take. In the simplest case some member of the user community, not necessarily the expert, presents a problem to the system and requests that the system apply its expertise to generate a solution. Assuming it is capable of understanding the problem statement and then of solving the problem, the system responds to the user with the solution and everyone is happy.

If the user is unhappy with the solution, uncertain as to its validity, or desirous of an explanation of "why" or "how" the system has reached its conclusion, the user can typically enter into a second form of the consultation mode of use and request an explanation of the steps that the system has followed to achieve the generated result. In most cases this explanation takes the form of a formatted presentation of the chain of rules that were activated by the inference engine in reaching the solution. This explanatory capability is a major advantage over more conventional systems, and is facilitated by the architectural feature of a clear separation between the knowledge base and the inference mechanism.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

3.2 Knowledge Acquisition

A second form of interaction with the expert system is the knowledge acquisition process, wherein the knowledge and heuristics used by the human expert in the problem-solving task are transferred into the knowledge base of the expert system. This dialogue is the least understood process in the expert system paradigm. In most systems this interaction is not automated, but rather is mediated by a "knowledge engineer" [Feigenbaum 1980] whose job it is to (a) pick the brains of the human expert for the knowledge, principles, and heuristics used to solve the problem at hand, and (b) translate this communicated information into the form(s) required by the representation language(s) within which the expert system is being implemented.

There are very few guidelines available for how to facilitate this process. It is generally recognized that this is a long and tedious process, requiring good conceptual and communication skills, considerable patience, and experience. Moreover, it is this knowledge acquisition process which is iterative, continuing throughout not only the development of the system but during all of its useful life. On the other hand, this is another touted advantage of expert systems over conventionally engineered systems -- the ability to grow and learn, thereby providing the opportunity to continually improve performance. While this expendability is certainly enhanced by the isolation of the knowledge base, it is clearly not always a simple task to expand the limits of a system's expertise.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

Much research is currently being devoted to techniques for at least partially automating the knowledge-acquisition process. Such systems as AGE [Nii and Aiello 1979], KAS [Duda et al 1979], TEIRESIAS [Davis and Lenat 1982], EXPERT [Weiss and Kulikowski 1979], HEARSAY-III [Erman et al 1980], etc. have all attempted to provide a framework within which the system can guide the expert in communicating his/her expertise to the system. Much work remains to be done in this area, both in the development of automated tools for the existing paradigms of problem solving, and in the more basic research into the understanding of the very nature of human problem-solving strategies and abilities.

3.3 Training

A final form of interaction with the expert system occurs when the system is used as a training tool to teach new human experts the problem-solving skills embodied in its knowledge base. Relatively few systems have been used in this mode. However, such systems as SOPHIE [Brown et al 1981] have demonstrated that the existence of a clearly formulated, central repository of expertise provides a solid foundation for the development of such computer-based teachers as a fortuitous side-effect.

4.0 CURRENT STATUS - A POTPOURI OF EXPERT SYSTEMS

There is as yet no well-developed theory of problem-solving techniques, no theory of problem space complexity comparable, say, to

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

a theory of database query complexity. Moreover, or perhaps partially in consequence, the development of expert systems is still more of an art than a science. It is therefore difficult to find a concrete opinion held about these systems by a reputable researcher in the field whose opposite is not held by another researcher equally as reputable. Nevertheless, a consensus is beginning to emerge as to the characteristics of problem domains appropriate for the technologies that exist today. Recent surveys of expert systems [Davis 1982, Gevarter 1982, Nau 1983] have emphasized a number of characteristics to look for in a problem domain before considering it as a candidate for current expert system technology, and have identified a number of considerations involved in the development of such a system.

Foremost among these characteristics are the selection of an appropriate domain, and the availability of a human expert. The most successful domains seem to be those wherein the expertise is based on experience of associations, rather than causal links or use of structural information. Equally important is a close collaboration and active participation of the human expert throughout the entire system development process. Other considerations frequently mentioned are: the necessity for an experienced "knowledge engineer," the efficacy of a quick (3 months?) development of a first system prototype to test the feasibility of the initial problem-structuring ideas, and an average development time of 5 years, regardless of the number of people on the project.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

In Table 1, examples of expert systems are presented. As can be seen, expert systems have been built for several domains, which include Medicine, Geology, Chemistry and Physics, Mathematics, and Computers (both software and hardware). Among these, R1, Macsyma, and the Dipmeter Advisor are widely used in commercial environments.

Prominent researchers in the area (e.g. [Davis 1982]) see future expert systems departing from simple rules and uniform knowledge representations, to causal models employing multiple representations that concentrate on the understanding and description of "structure" and "function".

4.1 Expert Systems And Database Management

There have been several research efforts to combine expert system technology with that of database management systems. Historically, knowledge-based techniques were first applied at the query language level (e.g., natural language). Systems like RENDEZVOUS, LADDER and KLAUS [Haas and Hendrix 1980] have successfully employed knowledge-bases to disambiguate and process English queries to and about databases. In addition, formal specification languages like TAXIS [Mylopoulos et al 1980] have been proposed for the design of databases and, more generally, information systems. Knowledge-based technology may also be used in such database topics as, query optimization [King 1981], transaction management (e.g., constraint maintenance), and data representation [Jarke and Vassiliou 1983, Vassiliou et al 1983].

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

SYSTEM	FUNCTION	DOMAIN	REF	SOME UNDERLYING PRINCIPLES
Casnet	Consulting	Medicine	66	Production rules. Causality. Semantic net
Internist	Consulting	Medicine	50	Forward/backward chaining. Frames.
KMS	Consulting	Medicine	52	Conditional probabilities.
MDX	Consulting	Medicine	8	Hierarchical, subproblem formation.
Mycin	Consulting	Medicine	59	Backward chaining. Exhaustive search.
Puff	Consulting	Medicine	33,49	Backward chaining. Exhaustive search.
AQ11	Diagnosis	Plant Diseases	9	Multiple-valued logic.
Dipmeter Advisor	Exploration	Geology	14	Causality.
Prospector	Exploration	Mineral	16	Backward chaining. Semantic net.
R1	Configuration	Computer	31	Forward chaining. No backtracking. Subproblem formation. Pattern match.
EL	Analysis	Circuits	60	Forward chaining. Backtracking. Constraint propogation.
SOPHIE	Troubleshoot	Electronics	5	Multi-knowledge representation.
Molgen	Planning	DNA Exper.	36	Forward/backward chaining. Hierarchical, subproblem formation.
Macsyma	Manipulation	Math	43	Pattern match.
AM	Formation	Math	34	Forward Chaining. Generate, test.
Dendral	Generation of hypotheses	Chemistry	21,35	Forward Chaining. Generate, test.
SYNCHEM2	Organic Synth.	Chemistry	23	Multi-representation. Subproblem formation
Hearsay	Interpretation	Speech Recognition	1	Forward/backward chaining. Multi-representation.
Harpy	Interpretation	Speech Recognition	37	Forward chaining
Crysalis	Interpretation	Crystallo-graphy	17	Event Driven. Generate, test.
Noah	Planning	Robotics	54	Backward chaining. Subproblem formation.
Abstrips	Planning	Robotics	55	Back-chaining. Hier. sub-problem formation
VM	Monitoring	Medicine	19	Event Driven. Exhaustive Search.
Guidon	CAI	Medicine	10	Event Driven.

TABLE 1: Example Expert Systems

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

A more recent research topic is that of coupling ESs with DBMSs. To date the applications that have been chosen for expert systems have had the property that their knowledge base of rules has been relatively small (around 1000 rules is common) and their base of specific facts has been considerably smaller, usually data pertaining to a single problem case and obtained interactively during system execution. In almost all cases, then, these knowledge bases have been implemented directly in main memory. The work of Kunifuji and Yokota for the Fifth Generation Computer project, and that of [Vassiliou,, Clifford, and Jarke 1983] attempt to apply the ES paradigm to a problem characterized by the existence of a large database of specific facts which the expert must access in order to perform successfully.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

References

1. Balzer, R., Erman, L.D., London, P., and Williams, C., "HEARSAY-III: A Domain-Independent Framework for Expert Systems," Proc. 1st Natl. Conf. of the Amer. Assoc. for Artificial Intelligence, Palo Alto, 1980, pp. 108-110.

2. Barr, A., Feigenbaum, E.A., The Handbook of Artificial Intelligence, Volume 2, William Kaufmann, Inc., Los Altos, CA, 1982.

3. Bobrow, D.G., and Winograd, T. "An Overview of KRL, a Knowledge Representation Language," Cognitive Science, Vol. 1, No. 1, 1977, pp. 84-123.

4. Brachman, R. "On the Epistemological Status of Semantic Networks," Associative Networks: Representation and Use of Knowledge by Computer, N.V. Findler, ed., Academic Press, 1979, pp. 3-50.

5. Brown, J., Burton, R, deKleer, J., "Pedagogical Natural Language and Knowledge Engineering Techniques in SOPHIE I, II, and III", Intelligent Tutoring Systems, Sleeman et al (eds), Academic Press, 1981.

6. Buchanan, B.G., "Research on Expert Systems," Stanford University Computer Science Department, Report No. STAN-CS-81-837, 1981.

7. Chandrasekaran, "Expert Systems: Matching Techniques to Tools", Artificial Intelligence Applications for Business (W.Reitman, ed.), Ablex, to appear 1983.

8. Chandrasekaran et al., "An Approach to Medical Diagnosis Based on Conceptual Structures," Proc. Sixth Int'l Joint Conf. Artificial Intelligence, 1979, pp. 134-142.

9. Chilausky, R., Jacobson, B., and Michalski, R.S. "An Application of Variable-Valued Logic to Inductive Learning of Plant Disease Diagnostic Rules," Proc. Sixth Annual Int'l Symp. Multiple-Valued Logic, 1976. Tokyo, Aug., 20-23, 1979, pp. 645-655.

10. Clancey, W.J., "Dialogue Management for Online Tutorials", IJCAI 6, 1979, pp.155-161.

11. Davis, R., "Expert Systems: Where are we? and Where do we Go from Here?", Massachusetts Institute of Technology, AI MEMO No. 665., June, 1982.

12. Davis, R., and Lenat, D.B. Knowledge-Based Systems in Artificial Intelligence, New York: McGraw-Hill, 1982.

13. Davis, R., Buchanan, B., and Shortliffe, E. "Production Rules as a Representation for a Knowledge-Based Consultation Program," Artificial Intelligence, Vol. 8, No. 1, 1977, pp. 15-45.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

14. Davis, R., et al., "The Dipmeter Advisor: Interpretation of Geological Signals," Proc. Seventh Int'l Joint Conf. Artificial Intelligence, 1977, pp. 1030-1037.

15. Duda, R.O., Gaschnig, J.G., "Knowledge-Based Expert Systems Come of Age," Byte, Vol. 6, No. 9, Sept. 81, pp. 238-281.

16. Duda, R.O., Gaschnig, J.G., and Hart, P.E., "Model Design in the PROSPECTOR Consultant System for Mineral Exploration," in Expert Systems in the Micro-Electronic Age, D. Michie ed., Edinburgh Univ. Press, Edinburgh, 1979, pp. 153-167.

17. Engelmen, R, Terry, A., "Structure and Function of the CRYSALIS System", IJCAI 6, 1979, pp.250-256.

18. Erman, L.D., et al., "The Hearsay-II Speech-Understanding System: Integrating Knowledge to Resolve Uncertainty," Computing Surveys, Vol. 12, No. 2, June 1980, pp. 213-253..

19. Fagan, L.M. "VM: Representing Time-Dependent Relations in a Medical Setting," Ph.D. dissertation, Comp. Sci. Dept., Stanford Univ., Stanford, Ca., 1980.

20. Fahlman, S.E. NETL: A System for Representing and Using Real-World Knowledge, MIT Press, Cambridge, Mass., 1979.

21. Feigenbaum, E., Buchanan, G., and Lederberg, J. "Generality and Problem Solving: A Case Study Using the DENDRAL Program," Machine Intelligence 6, D. Meltzer and D. Michie, eds., Edinburgh University Press, 1971, pp. 165-190.

22. Forgy, C.L. The OPS5 User's Manual, Tech. Report Carnegie-Mellon University, 1980.

23. Gelernter, H.L., Sanders, A.F., Larsen, D.L., Agarwal, K.K., Boivie, R.H., Spritzer, G.A., and Searleman, J.E., "Empirical Explorations of SYNCHEM," Science v.197, 1977, pp.1041-1049.

24. Gevarter, William, B., "An Overview of Expert Systems", National Bureau of Standards Report, NBSIR 82-2505, 1982.

25. Haas, N., and Hendrix, G.G., "An Approach to Acquiring and Applying Knowledge", Proceedings of the First National Conference on AI, Stanford University, 1980, pp.235-239.

26. Hart, P.E., "Directions for AI in the Eighties," Sigart Newsletter, No. 79, Jan. 1982, pp. 11-16.

27. Hart, P.E., Duda, R.O., and Einaudi, M.T. A Computer-Based Consultation System for Mineral Exploration, Tech. Report, SRI International, Menlo Park, Calif., 1978.

28. Hayes-Roth, F., "AI The New Wave - A Technical Tutorial for R&D Management," (AIAA-81-0827), Santa Monica, CA: Rand Corp., 1981.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

29. Hayes-Roth, F., Waterman, D.A., and Lenat, D.B. Building Expert Systems,Addison-Wesley, Reading, MA, 1983.

30. Jarke, M., Vassiliou, Y., "Coupling Expert Systems with Database Management Systems", Artificial Intelligence Applications for Business (W.Reitman, ed.), Ablex, to appear 1983.

31. King, J.J., "Query Optimization by Semantic Reasoning", Rep.no. CS-81-857, Computer Science Dept., Stanford University, 1981.

32. Kowalski, R.A., Logic for Problem Solving. North-Holland, New York, 1979.

33. Kunz, J.C., et al, "A Physiological Rule-based System for Interpreting Pulmonary Function Test Rules", Stanford Memo HPP-78-19, Stanford University, Computer Science Department, November, 1982.

34. Lenat, D.B. "AM: An Artificial Intelligence Approach to Discovery in Mathematics as Heuristic Search," Ph.D. dissertation Memo AIM-286, AI Laboratory, Stanford Univ., Stanford, Ca., 1976.

35. Lindsay, R.K., et al., Applications of Artificial Intelligence for Organic Chemistry: The DENDRAL Project, New York: McGraw-Hill, 1980.

36. Martin, N., et al., "Knowledge-Base Management for Experiment Planning in Molecular Genetics," Proc. Fifth Int'l Joint Conf. Artificial Intelligence, 1977, pp. 882-887.

37. Lowerre, B.T., "The HARPY Speech Recognition System," Ph.D. dissertation, Comp. Sci. Dept., Carnegie-Mellon Univ., Pittsburgh, Pa., 1976.

38. McDermott, D. "The PROLOG Phenomenon," Sigart Newsletter, No. 72, July 1980, pp. 16-20.

39. McDermott, D., "R1: A Rule Based Configurer of Computer Systems", Artificial Intelligence, 19(1), September, 1982.

40. McDermott, J., and Steele, B, "Extending a Knowledge-Based System to Deal with Ad Hoc Constraints," Proc. Seventh Int'l Joint Conf. Artificial Intelligence, 1981, pp. 824-828.

41. Michie, Donald, "Knowledge-based Systems," University of IL at Urbana-Champaign, Report 30-1601, June 1980.

42. Minsky, M. "A Framework for Representing Knowledge," The Psychology of Computer Vision, P.H. Winston, ed., McGraw-Hill, New York, 1975, pp. 211-277.

43. Moses, J. "Symbolic Integration: The Stormy Decade," Comm. ACM, Vol. 14, No. 8, 1971, pp. 548-560.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

44. Mylopoulos, J. "An Overview of Knowledge Representation," Proc. Workshop Data Abstraction, Databases, and Conceptual Modeling, June 1980, pp. 5-12.

45. Mylopoulos, J., Bernstein, P., and Wong, H., "A Language Facility for Designing Database Intensive Applications", ACM, 5, 1980, pp.185-207.

46. Nau, D.S., "Expert Computer Systems", Computer, February, 1983, pp.63-85.

47. Nii, H.P., and Aiello, N., "AGE (Attempt to Generalize): A Knowledge-BasedProgram for Building Knowledge-Based Programs," Proceedings of the Sixth International Conf. on Artificial Intelligence, (IJCAI-79), Tokyo, Aug., 20-23, 1979, pp. 645-655.

48. Nilsson, N.J. Principles of Artificial Intelligencej, Tioga, Palo Alto, Calif., 1980.

49. Osborn, J., et al., "Managing the Data from Respiratory Measurements," Medical Instrumentation, Vol. 13, No. 6, Nov. 1979.

50. Pople, H., "Knowledge Based Expert Systems: The Buy or Build Decision", Artificial Intelligence Applications for Business (W.Reitman, ed.), Ablex, to appear 1983.

51. Quillian, M.R., "Semantic Memory", in Semantic Information Processing, M. Minsky (ed), 1968, pp.227-270.

52. Reggia, J., et al., "Towards an Intelligent Textbook of Neurology," Proc. Fourth Annual Symp. Computer Applications in Medical Care, 1980, pp. 190-199.

53. Robinson, J.A., A Machine Oriented Logic Based on the Resolution Principle, JACM, 1965, Vol.1, No.4, pp.23-41.

54. Sacerdoti, E.D., "Planning in a Hierarchy of Abstraction Spaces," Artificial Intelligence Vol. 5 No. 2, 1974, pp. 115-135.

55. Sacerdoti, E.D., A Structure for Plans and Behavior, American Elsevier, New York, 1977.

56. Schank, R.C. "Conceptual Dependency: A Theory of Natural Language Understanding." Cognitive Psychology, Vol. 3, No. 4, 1972.

57. Schank, R.C. Conceptual Information Processing, North-Holland, New York, 1975.

58. Schubert, L.K. "Extending the Expressive Power of Semantic Nets," Artificial Intelligence, Vol. 7, No. 2, 1976, pp. 163-198.

59. Shortliffe, Computer Based Medical Consultation: MYCIN, New York, American Elsevier, 1976.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97

60. Stallman, R.M., and Sussman, G.J. "Forward Reasoning and Dependency-Directed Backtracking in a System for Computer-Aided Circuit Analysis," Vol. 9, Artificial Intelligence, 1977, pp. 135-196.

61. Stefik, M., et al., "The Organization of Expert Systems: A Prescriptive Tutorial", XEROX, Palo Alto Research Centers, VLSIi-82-1, January 1982.

62. van Emden, M.H., and Kowalski, R.A., "The Semantics of Predicate Logic as a Programming Language," J. ACM, Vol. 23, No. 4, 1976.

63. van Melle, W. "A Domain-Independent Production Rule System for Consultation Programs," Proc. Sixth Int'l Joint Conf. Artificial Intelligence, 1979.

64. Vassiliou, Y., Clifford, J., Jarke, M., "How does an Expert System Get Its Data?", NYU Working Paper CRIS#50, GBA 82-26 (CR), extended abstract in Proc. 9th VLDB Conf., Florence, October 1983.

65. Weiss, S.M., and Kulikowski, C.A. "EXPERT: A System for Developing Consultation Models," Proc. Sixth Int'l Joint Conf. Artificial Intelligence, 1979, pp. 942-947.

66. Weiss, S.M., et al., "A Model-Based Method for Computer-Aided Medical Decision-Making," Artificial Intelligence, Vol. 11, No. 2, 1978, pp. 145-172.

Center for Digital Economy Research
Stern School of Business
Working Paper IS-83-97