SECOND GENERATION KNOWLEDGE ACQUISITION SYSTEMS

Brian R Gaines

Knowledge Science Institute
University of Calgary
Calgary, Alberta, Canada T2N 1N4
gaines@calgary.cdn

ABSTRACT

Knowledge acquisition has developed from manual techniques in the 1970s through a first generation of computer-based tools in the early 1980s to a second generation of integrated knowledge support systems in the late 1980s. This paper reviews the requirements for such second generation systems, places the integration requirements within the framework of support for human knowledge processes, analyses the theoretical foundations of such systems through a model of knowledge validation, and analyses the technical foundations of such system through a model of generations of information technology. It concludes with an overview of the systems architecture of second generation knowledge acquisition systems.

INTRODUCTION

The development of knowledge acquisition systems is facing a critical transition. The zeroth generation of knowledge acquisition techniques in the 1970s were based on human interviewers debriefing experts about their skills. The first generation of computer-based tools in the 1980s were based on the computer-aided application of specific techniques such as construct and conceptual mapping, behavior modeling and text analysis. However, neither manual nor tool-based approaches have yet made a significant impact on the knowledge acquisition bottleneck which continues to impede the development and effective application of knowledge-based systems. Current, second generation research is establishing a new paradigm for knowledge acquisition, one that recognizes the wide diversity of forms and sources of knowledge, its essential dynamism making it subject to contradiction and change, and the constructive nature of knowledge acquisition whereby it is often the knowledge acquisition process itself that creates knowledge.

Second generation knowledge acquisition targeted on applications in the 1990s have to:

- Integrate a wide variety of weak methods encompassing the diversity of forms and sources of knowledge
- Integrate conflicting sources of knowledge to cope with essential contradictions that cannot be replaced by an artificial and misleading 'consensus'
- Integrate knowledge acquisition with system operation to facilitate a heavy activity of continuing system maintenance and upgrade as knowledge changes
- Integrate knowledge representation in the acquisition and performance systems, retaining the flexibility for new representations to be used that more accurately carry the knowledge

Thus, second generation developments have two objectives that may come into conflict (Gaines 1987):

- To integrate closely a variety of techniques and technologies to cope with the diversity of knowledge
- To provide flexibility within the integration framework for new techniques and technologies to be incorporated to cope with the dynamics of knowledge

This paper analyses the integration requirements within the framework of support for human knowledge processes, analyses the theoretical foundations of such systems through a model of

knowledge validation, analyses the technical foundations of such system through a model of generations of information technology, and uses these analyses to develop an architecture for second generation developments.

THE NATURE OF KNOWLEDGE

The diversity of human knowledge processes cannot be over-emphasized—indeed 'anarchy' might be a better term than diversity (Feyerabend 1975). "Knowledge" in the artificial intelligence literature has become associated with formal structures from which inferences can be made. However, much human knowledge is not well-structured for formal representation and computer processing (Gaines 1986). The philosophical dictionary definition of knowledge as 'more than opinion, less than truth' (Runes 1972) encapsulates what can be said. Classical data processing methods take it for granted that the data can be assumed to be 'true' and base their inferences on this. In knowledge-based system we have data that is less than truth, but more than opinion because it is 'expert opinion'. The implications of these notions of untruth (inconsistency, contradiction, change) and expertise (modeling, cognitive authority, reflective equilibrium) are critical to the design of knowledge acquisition systems.

It is interesting to note that in the late 1980s computer-based systems have been developed that are targeted on the representation of very loosely structured knowledge with less emphasis on formal structures from which inferences can be made. *Hypertext* and *hypermedia* systems approach knowledge a different direction to that of expert systems, accepting it as having limited structure and processibility, and adding a limited structure through the links that is intended to aid the user rather than allow computer-based inferencing (Conklin 1986, Gaines & Vickers 1988). They extend knowledge representation to include any media whereby human knowledge can be transferred even if there is no means for analyzing that knowledge or its transfer. For example, a film of a skilled person performing a task which enables another person to perform that task better contains knowledge which can be accessed and transferred through a hypermedia system (Vickers & Kingston 1987).

Hypermedia weaken the formal, computational nature of the current concepts of 'knowledge bases' (Brodie & Mylopoulos 1986) in not requiring that the information system be able to represent and process the knowledge encoded in the media used. However, if the links in hypertext systems are labeled (Trigg 1983) then the possibility of automatic inferencing is enhanced, and if arbitrary processes can be associated with the traversal of links (HyperCard 1987) then the boundary between user-mediated inferencing and computer-mediated inferencing becomes blurred.

The notions of computer-based hypermedia generate the interesting conceptual framework that a knowledge-based system is a medium for carrying human knowledge, and hence its functionality and operation can be related to those of existing media. Figure 1 shows how this framework can be developed to relate human knowledge processes with the relevant techniques used, the pre-computer supporting technology, and the computer-based supporting technology.

The basic level is that of human *experience* as the generative substratum for knowledge processes. In the survival of the species it is important that experience and the lessons learnt from it are shared. Others may be able to learn from experiences more effectively than those who have them. Individuals are finite and expendable but their experience is part of the knowledge economy of the species. Hence the next level is that of *communication* of experience.

However, the process of communication goes beyond sharing of experience because it allows for indirect action upon the world by persuading others to act in one's place, and this needs to be differentiated as a different form of communication, that of *discourse*. Discourse involves the use of communication intentionally to affect the behavior of others and, in particular, to communicate not just experience but *understanding*.

	Relevant Techniques	*Supporting Technology*	*Computing Technology*
Experience	Perception, Action	Instruments, Tools	Games, Simulation
Communication	Mimicry, Representation, Language	Painting, Writing, Phonograph, Television	Word Processing, Graphics, Hypermedia
Discourse	Rhetoric, Communicative Action	Books, Newspapers, Journals, Telephone	Electronic Mail, Structured Hypermedia
Practical Reasoning	Logic, Induction, Analogy	Jurisprudence Systems, Uncodified 'Commonsense'	Semantic Nets, Conceptual Graphs
Formal Reasoning	Formal Logic	Predicate Calculus, Sequent Calculi	Relational Databases, Expert Systems
Mathematical Reasoning	Set Theory, Category Theory	Calculators, Metatheorems, Proof Techniques	Number Crunchers, Theorem Provers

Fig.1 Knowledge processes, techniques and support

The next level of *practical reasoning* is concerned with the creation of knowledge from experience, and its application to the generation of experience—that is, determination of the meaning, significance, consequences, actions and principles based on experience.

The next two levels, of *formal* and *mathematical* reasoning, are separated from practical reasoning because of their fundamental role in the development of modern science and technology, and their significance in computing. Formal reasoning provides foundations for precise deductive processes and mathematical reasoning applies this to the structures underlying arithmetic and algebraic systems.

The techniques and support systems underlying these knowledge processes are shown in three columns in Figure 1:

• Experience
—Relevant Techniques: human interaction with worlds through perception and action
—Supporting Technology: instruments for enhancing perception and tools for enhancing action
—Computing Technology: the creation of artificial experience through games and simulation
• Communication
—Relevant Techniques: demonstration and mimicry of behavior, representation of experience through media, description of experience through language
—Supporting Technology: painting, writing, phonograph and television as media for capturing experience
—Computing Technology: word processing, graphics and hypermedia as media for representing and presenting experience
• Discourse
—Relevant Techniques: rhetorical techniques for persuasion and communicative action to secure agreement over an understanding
—Supporting Technology: rhetorical and interactive media such as books, newspapers, journals and the telephone
—Computing Technology: electronic mail supporting conversation and structured hypermedia supporting the presentation of arguments
• Practical Reasoning
—Relevant Techniques: informal deductive logic, inductive and abductive inference, and analogical pattern recognition
—Supporting Technology: surprisingly little—jurisprudence through the legal system to support justice in socially important practical reasoning, but primarily left to uncodified 'commonsense'
—Computing Technology: knowledge representation and inference through semantic nets and conceptual graphs
• Formal Reasoning
—Relevant Techniques: formal logic based on axiom systems
—Supporting Technology: various calculi for formal logics such as predicate and sequent calculi
—Computing Technology: relational databases and expert systems
• Mathematical Reasoning
—Relevant Techniques: set theory and category theory
—Supporting Technology: calculators for arithmetic, metatheorems and proof techniques for algebraic systems
—Computing Technology: number crunchers and spreadsheets for arithmetic, theorem provers for algebraic systems

Looking down the right hand column of Figure 1 we can see the full spectrum of computer-based techniques supporting human knowledge processes, and the need for integration becomes very apparent. A knowledge support system must be capable of providing any combination of graphics, text processing, hypermedia, number crunching, databases, communications, theorem proving, databases, and semantic nets and expert system technologies. These are not highly separated functions but rather multiple aspects of human knowledge processing.

It is particularly interesting to note the weaknesses in the support of discourse and practical reasoning processes. These are the levels at which the colloquial concept of 'expertise' is highly applicable. However, expert systems technology has developed primarily at the level of formal reasoning which is a very narrow subdomain of human reasoning processes.

THE VALIDATION OF KNOWLEDGE

The analysis of the previous section may be taken a stage further by considering the processes by which knowledge can be validated, that is rationally developed from the realm of opinion towards that of truth. Typical criteria used for distinguishing knowledge from opinions are: observation; reliable evidence; rigorous application of method; the use of a precise language; testing by experiment; power of anticipation; formalization; ability to proceed to applications; usefulness of these applications; accountability to certain principles or to acquired knowledge; impartiality; accepted views of specialists (Franck 1979).

These criteria can be portrayed systematically using Popper's (1968) distinctions between three 'worlds': "If we call the world of things—of physical objects—the first world and the world of subjective experience the second world we may call the world of statements in themselves the third world (world 3)." This definition can be expanded in terms of the distinctions involved:

- World 1 arises from the assertion that there exist *physically necessary* distinctions

—this defines a world of *physical systems*

—a key concept is that of *actuality*, that the distinctions made are actual

—the conceptual framework involved is that of *epistemology*, that the distinctions have to be drawn from the actual world

—truth in this world is by *correspondence*, that the distinctions made correspond to the properties of physical systems

—inference is *causal*, based on chains of physical cause and effect

- World 2 arises from the assertion that there exist *humanly chosen* distinctions

—this defines a world of *mental systems*

—a key concept is that of *agency*, that distinctions have a maker

—the conceptual frameworks involved are those of *psychology and sociology*, that distinctions exist in the mental worlds of individuals and groups

—truth in this world is *performative*, that the distinctions chosen will be continue to be chosen

—inference is *conventionalist*, based on accepted chains of reasoning that are chosen within a culture

- World 3 arises from the assertion that there exist *formal systems* of distinctions

—this defines a world of *coherent systems*

—a key concept is that of *abstraction*, that systems of distinctions have their own existence

—the conceptual framework involved is that of *ontology*, that some distinctions are made independent of any criteria other than intrinsic structure

—truth in this world is by *coherence*, that the systems of distinctions have an acceptable intrinsic structure

—inference is *structuralist*, based on chains of reasoning that conform to the structure

Figure 2 shows the interactions between a knowledge-based systems and these three worlds in ways that may be used to validate it:

- *Prediction* is the capability to use knowledge to model the world and predict its behavior
- *Action* is the capability to use knowledge to effectively control the world
- *Explanation* is the capability to use knowledge to explain the basis of predictions and actions
- *Understanding* is the capability to accept new knowledge through discourse
- *Derivation* is the capability to deduce knowledge from coherent knowledge systems
- *Creation* is the capability to add knowledge to coherent knowledge systems

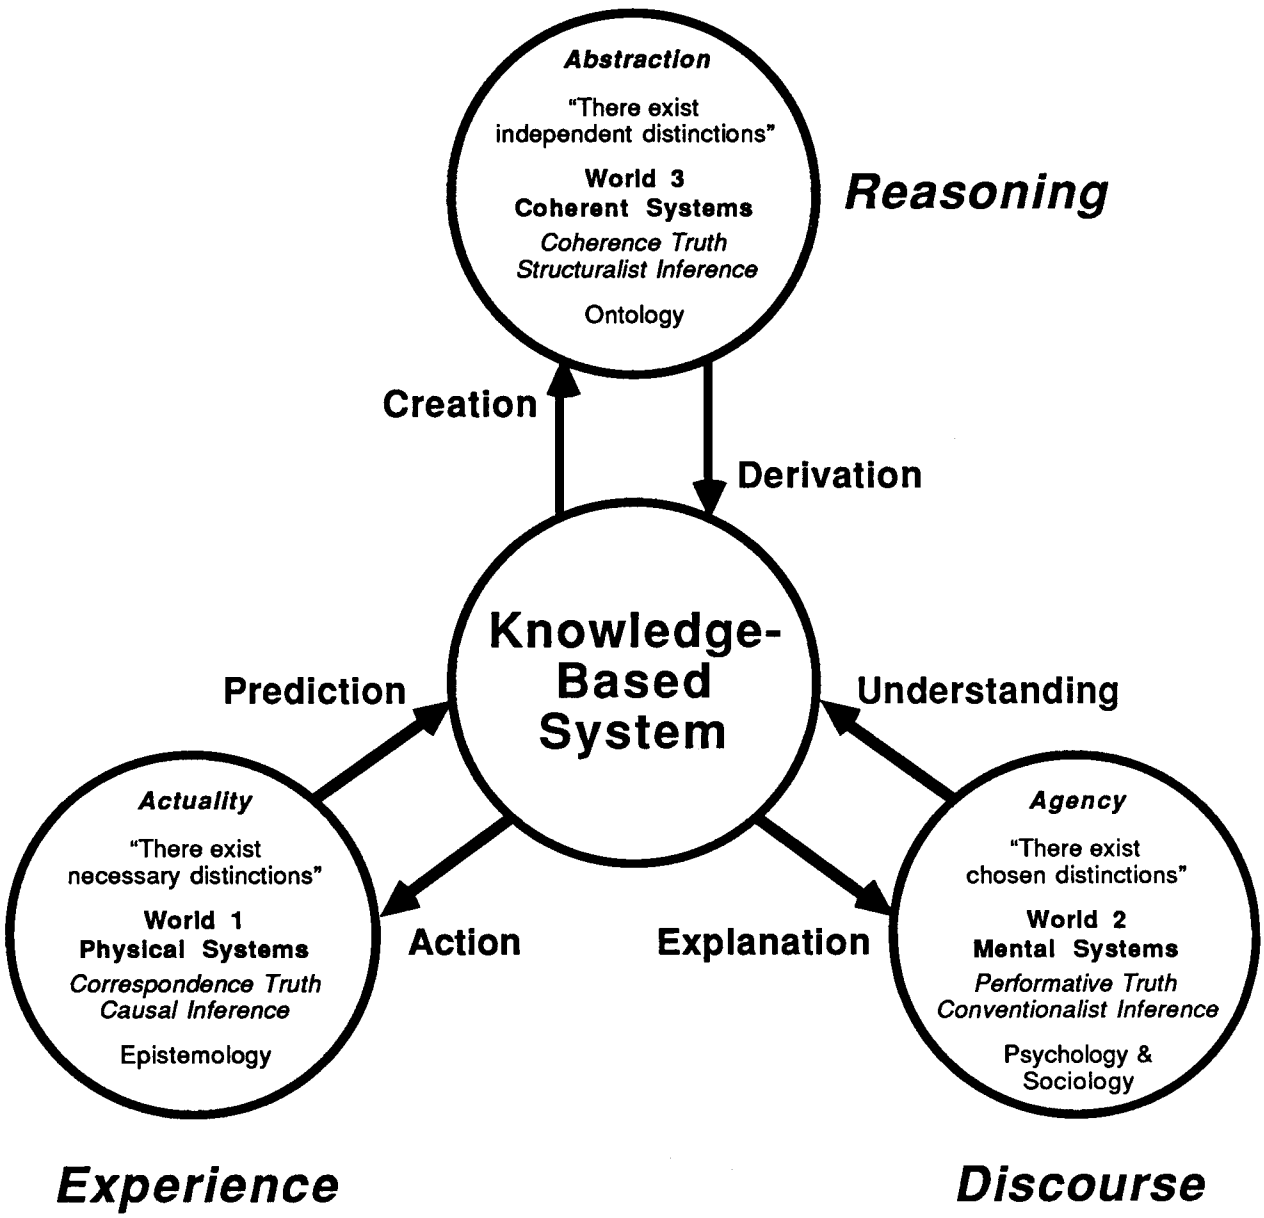

Fig.2 Interaction of a knowledge-based system with 3 worlds

For example, the MYCIN/TEIRESIAS system (Shortliffe 1976, Davis & Lenat 1982) exhibits all these activities to some limited extent. It predicts microbial infections from clinical data and recommends appropriate action. It gives explanations of these predictions and is able to understand new knowledge given in natural language. It derives its knowledge from a coherent set of rules and is able to create new rules from linguistic input. Analysis of the strengths and weaknesses of this and related systems in these various areas has been the major source of progress in the development of knowledge-based systems (Hasling, Clancey & Rennels 1984, Chandrasekaran & Mittal 1984).

Figure 2 also illustrates how the 3 worlds may be used to give foundations to the distinctions between levels in Figure 1, with experience in world 1, communication and discourse in world 2, and reasoning in world 3. The 3 worlds model provides a theoretical framework for the development of second generation integrated knowledge acquisition systems.

It is apparent in Figure 2 that the links from the knowledge-based system to world 1 have been studied in depth in the methodologies and philosophies of science. The links to world 3 have been studied in depth in the methodologies and philosophies of deductive and inductive inference. Both these domains of study have been used as major input to the development of computer-based systems. However, the links to world 2 which are studied in depth in the psychology, sociology and philosophy of discourse are the weakest in current knowledge-based systems. Explanation and understanding will be crucial areas of development in second generation knowledge acquisition systems.

The sociological basis of explanation and understanding has been developed in depth by Habermas (1984) through his concept of *communicative action.* Habermas distinguishes four primary components of communication acts: *strategic* in which one person treats another as an entity whose behavior is manipulable through language; *normative* in which the communication conforms to shared norms and values; *dramaturgical* in which one person presents their self to others; and *communicative* in which people attempt to reach voluntary agreement on their situation for the sake of cooperating. Habermas argues for communicative action as the foundation of a healthy society.

One can characterize the move towards knowledge-based systems in computing as being one of developing the possibilities for communicative action by basing the systems on overt, high-level data structures. 'Overt' means that the structures are separate from, and not embedded in the technical code of the computer. 'High-level' means that they are understandable by non-computer specialists having a knowledge of the relevant knowledge domain. Communicative action is the capability to provide understandable explanations of data-processing using these overt, high-level structures, and to understand specifications of data-processing in them.

Popper's concept of world 3 is particularly interesting from the viewpoint of explanation and understanding since he sees world 3 as one of *objective knowledge* where the primary criterion is one of *fallibilism*—there must be tests that might show the knowledge to be false. However, this criterion which establishes the possibility of a critical community, also establishes the need for knowledge to be *understandable* so that it can be tested by the community (Notturno 1985). This is consistent with Habermas' analysis of communicative understanding provided the distinction between truth and falsity is taken to include validation in worlds 2 and 3 as well as world 1.

Pask (1980) has developed a similar framework to Habermas as a foundation for computer-based learning. His *conversation theory* singles out as the critical event in discourse the occasions on which two participants agree that they have effectively exchanged a concept and validated its mutual similarity. Such *agreement over an understanding* punctuates the discourse into *strict conversations* whose structure and processes can be analyzed in depth as coherent units. Pask recognizes the diversity of validity tests that can be applied to understanding such as the ability to: *produce* knowledge transferred; to *reproduce* knowledge transferred from an alternative derivation; to *apply* knowledge transferred to problem solving; and, recursively, to *teach back* knowledge transferred.

The significant common theme in Habermas', Popper's and Pask's analyses is that the discourse should provide for as much exploration of a knowledge structure and its validity as the parties to it require in order to establish agreement or disagreement. For a knowledge-based system this implies the capability to expose links over as wide a breadth and in as much depth as the user requires. For knowledge acquisition systems this implies that a minimal translation of expert skills into computational form is inadequate. The acquisition phase has to be able to explore the full richness of the knowledge structure encompassing linkages which, although not computationally useful, may be vital to explanation and understanding.

GENERATIONS OF INFORMATION TECHNOLOGY

To examine the technological support for second generation integrated knowledge acquisition systems it is necessary to have some model of the wide range of technologies underlying computing. The computer itself is only one component in an elaborate and highly differentiated infrastructure which has grown through a succession of generations of computers, each of which represents a major change in computing technology (Gaines 1984). During the eight year span of each computing generation, revolutionary changes have taken place that correspond in magnitude to those taking place over some hundred years or more in other industries. The first generation commenced in 1948 and we are now in the sixth generation spanning 1988 through 1995.

The generation structure of information technology can itself be analyzed in terms of the knowledge flows involved. Like other basic technologies (Marchetti 1981), the growth of computing has an infrastructure comprised of the envelope of learning curves in successive underlying technologies (Ayres 1968). Figure 3 shows this structure through the zeroth through fifth generations projected to the sixth and seventh. Each technology has a learning curve in which a *breakthrough* leads to a phase of: *replication* in which the results are duplicated; *empiricism* in which rules for design are derived from experience; *theory* in which generative principles are derived; *automation* in which the theory is made operational; and *maturity* when the knowledge becomes proven, widely available and widely used (Gaines & Shaw 1986).

At the base is the learning curve for different electronic device technologies: 1940 zeroth generation using relays; 1948 first using vacuum tubes; 1956 second using transistors; 1964 third generation using using integrated circuits; 1972 fourth using large-scale integration; 1980 fifth using VLSI; 1988 sixth projected to use ultra large-scale integration with some 10 million transistors on a chip; 1996 seventh projected to use grand-scale integration with 1,000 million on a chip. The definition of generations in terms of electronic device technology captures some important aspects of computing such as cost decreases, size decreases, power increases, and so on. However, it fails to account for the qualitative changes that have given computing its distinct character in each generation. These appear through the tiered succession of learning curves of higher level technologies based on the lower level developments.

Now the generations may be defined in terms of major qualitative changes:

- the initial breakthrough in *electronic device technology* leading to the zeroth generation is placed at 1940 about the time of the Atanasoff and Berry experiments with tube-based digital calculations
- the first generation breakthrough was the introduction by Mauchly and Von Neumann of stored program and subroutine concepts around 1947 which led to the transition from the ENIAC to the EDVAC designs—this detached computing as a separate discipline from electronics by substituting software for hardware in a *virtual machine architecture*
- the second generation breakthrough was to bridge the gap between machine and task through the development of *problem-oriented languages* such as FORTRAN around 1956
- the third generation breakthrough was to bridge the gap between the computer and the person with the development of interactive time-shared computers in 1964 allowing close *human-computer interaction*
- the fourth generation breakthrough was in the early 1970s with developments in expert systems based on knowledge-processing—such *knowledge-based systems* have the capability to store information through its inter-relations and make inferences about its consequence
- the fifth generation breakthrough was in 1980 with developments in machine learning and *inductive inference systems*

- the growth of robotics will provide the next breakthroughs in which goal-directed, mobile computational systems will act autonomously to achieve their objectives—the breakthrough into the sixth generation era commencing in 1988 will be one of *autonomous activity systems*
- the interaction between these systems will become increasingly important in enabling them to cooperate to achieve goals—the seventh generation era commencing in 1996 will be one of *socially organized systems*

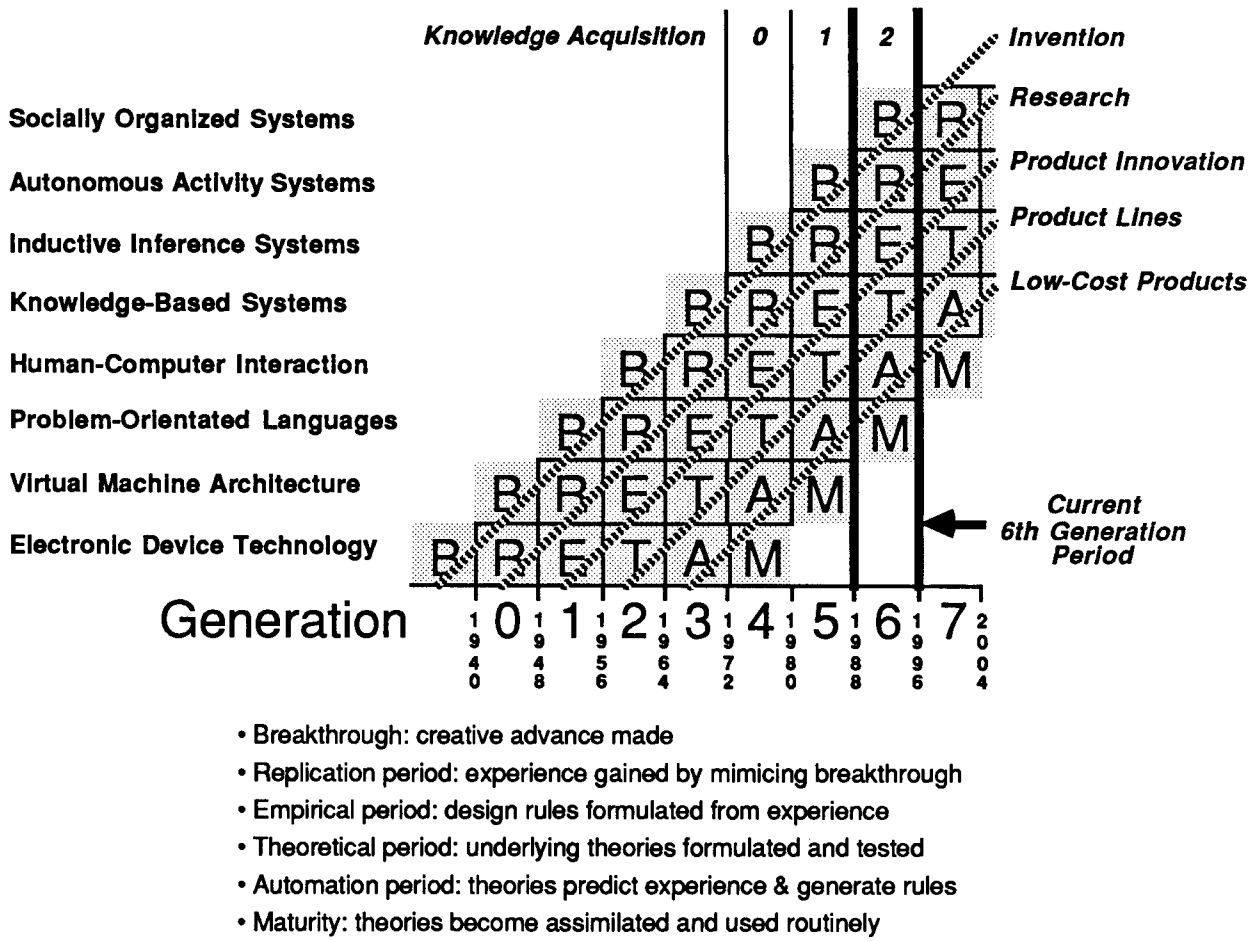

Fig.3 Knowledge acquisition in the Infrastructure of Information Technology

The model of the development of computing described above is important in any evaluation of the impact of computing. Each generation represents a revolution in technology with a qualitatively different impact. Each generation subsumes the capabilities of that preceding it, providing very much better facilities at very much lower cost, and adding new capabilities not possessed by the previous generations. The availability of these new capabilities in a generation is linked to the stage of the learning curve for each technology in that generation as shown in Figure 3:

- *Invention* is focused at the **BR** interface where new breakthrough attempts are being made based on experience with the replicated breakthroughs of the technology below;
- *Research* is focused at the **RE** interface where new recognized breakthroughs are being investigated using the empirical design rules of the technologies below;
- *Product Innovation* is focused at the **ET** interface where new products are being developed based on the empirical design rules of one technology and the theoretical foundations of the technologies below;

- *Product Lines* are focused at the **TA** interface where product lines can rest on the solid theoretical foundations of one technology and the automation of the technologies below;
- *Low-cost Products* are focused at the **AM** interface where cost reduction can be based on the the automated mass production of one technology and the mature technologies below.

For example, in the fourth generation (1972-79):
- **BR**: recognition of the knowledge acquisition possibilities of knowledge-based systems led to the breakthrough to inductive-inference systems;
- **RE**: research focused on the natural representation of knowledge through the development of human-computer interaction, e.g. the Xerox Star direct manipulation of objects;
- **ET**: experience with the human-computer interaction using the problem-oriented language BASIC led to the innovative product of the Apple II personal computer;
- **TA**: the simplicity of the problem-oriented language RPG II led to the design of the IBM System/3 product line of small business computers;
- **AM**: the design of special-purpose chips allowed the mass-production of low-cost, high-quality calculators.

In the fifth generation (1980-87):
- **BR**: recognition of the goal-seeking possibilities of inductive inference systems is leading to the breakthrough to autonomous-activity systems in robotics;
- **RE**: research is focused on learning in knowledge-based systems;
- **ET**: the advantages of the non-procedural representation of knowledge for human-computer interaction led to the innovative designs of the Visicalc spread-sheet business product and the Lisp-machine scientific product;
- **TA**: the ease of human-computer interaction through a direct manipulation problem-oriented language led to the Apple Lisa/Macintosh product line of personal computers;
- **AM**: the design of highly-integrated language systems has allowed the mass-production of low-cost, high-quality software such as Turbo Pascal.

The generations of knowledge acquisition systems are shown at the top of Figure 3, and the current sixth generation of information technology is also marked. It coincides with the second generation of knowledge acquisition systems and this four generation gap indicates an analogy between the state of knowledge-based systems development now and that of data processing 32 years ago. We do indeed have a similar situation to the 1950s in computing when software development was new with problems of shortage of skills, unrealistic schedules and budget over-runs.

The detailed infrastructure for the development of the second generation of knowledge acquisition systems may be seen in Figure 3. At the three lowest levels, the hardware machines architectures and languages are all mature technologies posing no problems. Human-computer interaction has entered its phase of automation with the increasing availability of machine-independent interactive systems such as X-windows. Knowledge-based systems have entered the theoretical development phase with increasing depth and structure to knowledge representation and inferencing schema. Inductive inference systems have entered the empirical phase where there are beginning to be available a range of usable tools. Autonomous activity systems are still in the research phase and the lack of experience and foundations will restrict the development of knowledge acquisition systems for goal-directed, planning activities during this generation.

Thus overall, we should expect the emphasis in second generation knowledge acquisition systems to be on the use of high-speed graphic workstations, programmed in high-level languages, using high-quality user interfaces to provide effective knowledge representation, some inductive inferences and little or no capability to deal with planning.

A SYSTEMS ARCHITECTURE FOR THE SECOND GENERATION

A general framework for a systems architecture for the second generation of knowledge acquisition systems may be developed based on the analysis of the previous three sections.

1 The distinction between the acquisition phase and the performance phase will begin to break down so that one is talking of a generation of *knowledge support systems* allowing users to switch modes between being knowledge sources and knowledge appliers.

This lack of clear role definition in second generation systems has other interesting consequences since it merges the concepts of knowledge acquisition and computer-assisted learning, distinguished only by the intended direction of flow of knowledge. Figure 4 shows this in terms of the knowledge flows in a knowledge support system. Clearly, each flow involves "agreement over an understanding" and hence two-way interaction. Clearly, also, the three roles shown can be within one individual acting sometimes to add knowledge to the system, sometimes to take its advice and sometimes to learn from it.

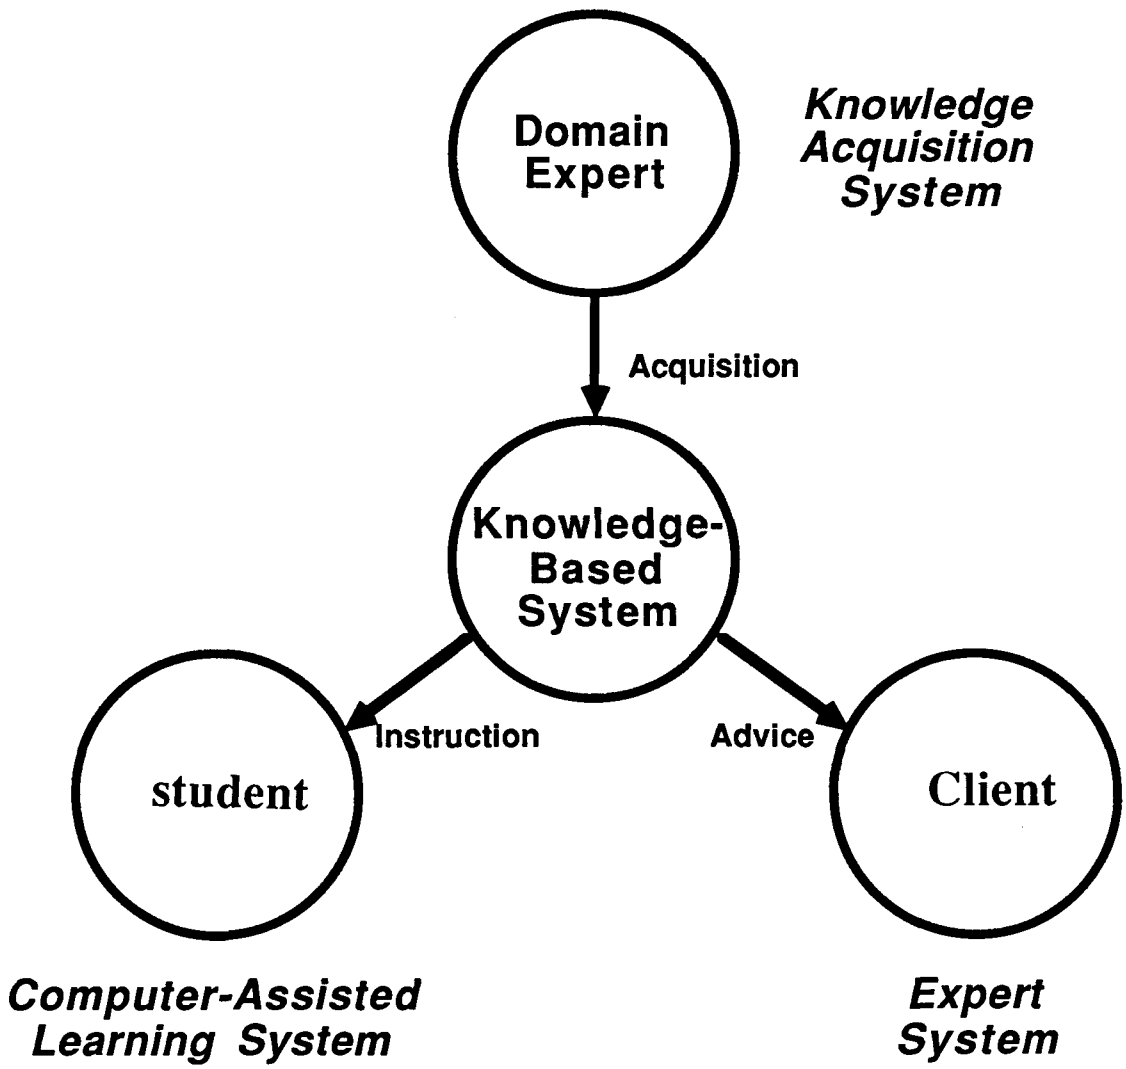

Fig.4 Knowledge flows in a knowledge support system

2 Acquisition will not be dominated by the requirement for computer-based inferencing but will encompass knowledge that can be interpreted only by the user as well as that interpretable by the computer.

In particular, acquisition will be as much for explanation and understanding as it is for inferencing.

3 A basic data item in knowledge support systems will be an uninterpreted knowledge representation objects encompassing text, graphics, sound, videodisc sequences, and so on.

The capability for objects to establish their identity, meaning, significance, denotation, and so on through direct interaction with the user is very important. The computer-based system does not have to understand the processes involved, only that they may have occurred.

4 Hypermedia-like links will be added to these data objects, some of which will be only user interpretable, but others of which will be computer-interpretable for inferencing.

This structure of objects, attributes and relations encompasses hypertext, semantic networks, conceptual graphs, and repertory grids, among other possibilities.

5 Some links will have associated validation procedures (methods) that test the validity of a link in terms of associated structures and hence may also be used to derive such structures from the link.

This association of validity tests with links allows the introduction of abstraction hierarchies and inference rules within the object-link framework.

The notion of a validity test is recursive since it introduces further knowledge within the system which itself will be presented by objects and links and subject to validity tests. In practical systems the infinite regression possible will be truncated at whatever level is acceptable to users. Further knowledge acquisition will be performed as acceptance criteria change.

6 A major part of the system will be devoted to link integrity maintenance which will be the prime dynamic to the generation and control of knowledge within the system.

In particular, structures when edited will reflect changes into associated structures. This is a weakness in most current knowledge acquisition schema in that they are unidirectional in deriving computational knowledge from expert data and do not reflect edits in the knowledge base back to the knowledge sources.

7 The hardware architecture will be a network of high speed graphic workstations connected through communications networks and utilizing database and computational servers.

In most cases the hardware will be fairly heterogeneous.

8 The software architecture will be based on an object-oriented language, with objects dynamic at run-time, providing the overall system control.

This will be able to invoke processes in other languages, however, so that the overall software architecture will also be heterogeneous.

None of these requirements is technically infeasible currently. Second generation knowledge acquisition systems are possible with existing hardware, software and knowledge technologies.

ACKNOWLEDGEMENTS

I am grateful to many colleagues for discussions over the years that have influenced this paper. In particular I would like to thank John Boose, Jeff Bradshaw, Bruce Conrad, Debbi Leishman, Gordon Pask, Mildred Shaw and Joan Vickers, for access to their own research and many stimulating discussions.

REFERENCES

Ayres, R.U. (1968). Envelope Curve Forecasting. **Technological Forecasting for Industry and Government.** pp. 77-94. New Jersey: Prentice Hall.

Brodie, M.L. & Mylopoulos, J. (Eds) (1986). **On Knowledge Base Management Systems.** New York: Springer Verlag.

Chandrasekaran, B. & Mittal, S. (1984). Deep versus compiled knowledge approaches to diagnostic problem solving. Coombs, M.J. (Ed). **Developments in Expert Systems.** pp.23-34. London: Academic Press.

Conklin, J. (1986). A survey of hypertext. MCC Technical Report STP-356-86. Austin:Microelectronics & Computer Technology Corporation (October).

Davis, R. & Lenat, D.B. (1982). **Knowledge-Based Systems in Artificial Intelligence.** New York: McGraw-Hill.

Feyerabend, P. (1975). **Against Method**. London: NLB.

Franck, R. (1979). Knowledge and opinions. Nowotny, H. & Rose, H. (Eds). **Counter-movements in the sciences**. pp.-. Dordrecht, Holland: Reidel.

Gaines, B.R. (1984). Perspectives on fifth generation computing. **Oxford Surveys in Information Technology, 1**, 1-53.

Gaines, B.R. (1986). Foundations of knowledge engineering. Bramer, M.A. (Ed). **Research and Development in Expert Systems III.** pp.13-24. Cambridge: University Press.

Gaines, B.R. (1987) Advanced expert system support environments. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Second AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.8-0-8-14. Banff (October).

Gaines, B.R. & Shaw, M.L.G. (1986). A learning model for forecasting the future of information technology. **Future Computing Systems, 1**(1), 31-69.

Gaines, B.R. & Vickers, J.N. (1988). Design considerations for hypermedia systems. **Microcomputers for Information Management**, 5(1), 1-27 (March).

Habermas, J. (1981). **The Theory of Communicative Action: Reason and the Rationalization of Society**. Boston: Beacon Press.

Hasling, D.W., Clancey, W.J. & Rennels, G. (1984). Strategic explanations for a diagnostic consulting system. Coombs, M.J. (Ed). **Developments in Expert Systems**. pp.117-133. London: Academic Press.

HyperCard (1987) **HyperCard User's Guide.** Cupertino, California: Apple Computer.

Marchetti, C. (1981). Society as a learning system: discovery, invention and innovation cycles revisited. Report RR-81-29 (November). Laxenburg, Austria: International Institute for Applied Systems Analysis.

Notturno, M.A. (1985). **Objectivity, Rationality and the Third Realm: Justification and the Grounds of Psychologism: A Study of Frege and Popper**. Dordrecht, Holland: Martinus Nijhoff.

Pask, G. (1980). Developments in conversation theory—Part I. **International Journal of Man-Machine Studies, 13**(4), 357-411 (November).

Popper, K.R. (1968). Epistemology without a knowing subject. Van Rootselaar, B., Ed. **Logic, Methodology and Philosophy of Science III**. pp. 333-373. Amsterdam, Holland: North-Holland Publishing Co.

Runes, D.D. (1972). **Dictionary of Philosophy**. London: Peter Owen.

Shortliffe, E.H. (1976). **Computer-Based Medical Consultations: MYCIN**. New York: Elsevier.

Trigg, R.H. (1983). **A Network-Based Approach to Text Handling for the Online Scientific Community**. PhD Thesis, University of Maryland (University Microfilms 8429934).

Vickers, J.N. & Kingston, G.E. (1987). Modelling the master coach: building an expert system for coaching. Norrie, D., Ed. **Learning in Future Education: Proceedings of International Conference on Computer Assisted Learning in Post-Secondary Education**. pp.207-212. University of Calgary.