Gaines, B. R. (1993). Representation, discourse, logic and truth: situating knowledge technology. Conceptual Graphs for Knowledge Representation 36-63.

Representation, Discourse, Logic and Truth: Situating Knowledge Technology

Brian R Gaines

Knowledge Science Institute, University of Calgary
Calgary, Alberta, Canada T2N 1N4
gaines@cpsc.ucalgary.ca

Abstract. The cumulative impact of developments in information technology during the past fifty years is resulting in major qualitative changes in applications of the technology, from information processing to knowledge processing. These changes not only involve new technologies, such as those for knowledge representation and acquisition, but also new social impacts such as the establishment of major scientific communities whose primary mode of discourse is through the internet. The conceptual graphs cg list server is a major example of the effective use of the internet to coordinate the intellectual activities of a widely dispersed research community. The PEIRCE project is an outstanding example of the effective use of the internet to initiate, design and coordinate development of a major software system. This paper examines these activities from a number of different perspectives. It places the cg list server and PEIRCE project in the context of the development of the infrastructure of information technology. It compares the project with other developments in knowledge representation, and the cg list server with related activities. Finally, it looks beyond PEIRCE at what will have been achieved, and what issues will remain to be addressed, once this project, and related ones, have been completed.

1 Introduction

The Knowledge Science Institute (KSI) was established in 1985 to study phenomena of the knowledge economy, and to contribute to the development of knowledge-based systems. It was predicated on the emergence of new forms of scholarly cooperation mediated through electronic networks that would tend to bypass existing institutional structures and forms of communication. It projected that knowledge representation schema and knowledge acquisition tools developed originally in research on artificial intelligence would become significant in the operationalization of scholarship and the communication of research. The KSI has developed models of the infrastructure of information technology [61, 66], of the structure and function of knowledge processes in society [60, 63], of the social and technical factors underlying scholarly knowledge dissemination through books, journals and internet [69, 70], of the quantitative relations between data and knowledge [59, 67], of the architectures of knowledge acquisition systems [75, 76], and of the relations between various forms of knowledge-based system [62]. It tracks a number of activities on the internet relating to new modes of scholarly communication, including news groups, list servers, electronic journals, collaborative projects, and archives of reports and papers.

The PEIRCE project initiated through the cg list server has been of particular interest since it involves so many of the central issues in the development of knowledge science:

1 The activity on the cg list server is an important instance of the use of the internet to coordinate the intellectual activities of a widely dispersed research community.

2 The PEIRCE project is an outstanding example of the effective use of the internet to initiate, design and coordinate development of a major software system.

3 The PEIRCE software is a major example of a new generation of high performance knowledge representation systems now becoming available.

4 More generally also, conceptual graphs are significant as the basis for an operational knowledge science and technology encompassing natural language, formal language, visual language and the wide range of reasoning processes that may be based on them [140].

5 More generally still, the conceptual graphs research community draws heavily on the seminal ideas of C S Peirce whose philosophical work in general is foundational for modern knowledge science.

It is rare for any theory to be able to make legitimate claim to the breadth encompassed by conceptual graphs, and unique for that theory to be available in operational form. Part of the reason for this is the lack of any unified theory of human knowledge processes. At one extreme we may see them as the anarchical pragmatic survival mechanisms of a biological entity comprised of an autopoietic neural structure that operates as a distributed network of physically discrete, coordinated but autonomous agents. At the other extreme we may see them as a progressive approximation to a rational process for modeling, predicting, controlling, and creating the universe. The first perspective emphasizes the chaotic, serendipitous, situated nature of knowledge processes, and sees language and overt knowledge as by-products of the coordination process. The second perspective emphasizes the deterministic, logical, universal nature of knowledge processes, and sees mathematical structures as the ultimate expression of knowledge with natural language as a first approximation. These perspectives are not mutually exclusive, and both are encompassed within the rich framework of Peirce's pragmatism that takes as its domain all the phenomena of the universe and human existence.

This paper addresses the foundations, operationalization, applications and other aspects of conceptual graphs as a knowledge science and technology situated in issues of representation, discourse, logic and truth. It examines these issues from a number of different perspectives, placing the cg list server and PEIRCE project in the context of the development of the infrastructure of information technology, comparing the PEIRCE project with other developments in knowledge representation, comparing the cg list server with related activities, and looking beyond PEIRCE at what will have been achieved, and what issues will remained to be addressed, once this project, and related ones, have been completed.

2 The Infrastructure of Information Technology

It is not coincidence that major changes should be occurring in both the dissemination and the representation of knowledge in the 1990s. Information technology shows strong patterns of growth linked to the medium-term business cycle, and has a complex infrastructure in which the availability of one information technology enables growth in others [73]. The 1980s have been an era in which low cost personal computers and

workstations have become ubiquitous, as has their communication linkage through local and wide area networks. The 1990s is an era of transition from information processing to knowledge processing. All the phenomena outlined in the previous section have depended on the universal availability of workstations and networks at the end of the 1980s to support the development of knowledge processing in the 1990s.

2.1 The Learning Curves of Information Technology

A foundational perspective on trends in information technology may be gained by examining the *learning curves* that characterize innovation and diffusion in all technologies and determine the dynamics of technological substitution [11, 99]. Logistic curves have been found to be a useful model of the introduction of new knowledge, technology or product in which growth takes off slowly, begins to climb rapidly and then slows down as whatever was introduced has been assimilated. Such curves arise in many different disciplines such as education, ecology, economics, marketing and technological forecasting [50, 142]. From the learning curves for information technology it is possible to determine the state of the technology in relation to knowledge processing during the present and coming generations of the technology.

It has also been noted in many disciplines that the qualitative phenomena during the growth of the logistic curve vary from stage to stage [43, 73, 104]. The era before the learning curve takes off, when too little is known for planned progress, is that of the inventor having very little chance of success. When an inventor makes a *breakthrough*, very rapidly his or her work is *replicated* at research institutions world-wide. The experience gained in this way leads to *empirical* design rules with very little foundation except previous successes and failures. However, as enough empirical experience is gained it becomes possible to inductively model the basis of success and failure and develop *theories*. This transition from empiricism to theory corresponds to the maximum slope of the logistic learning curve. The theoretical models make it possible to *automate* the scientific data gathering and analysis and associated manufacturing processes. Once automation has been put in place effort can focus on cost reduction and quality improvements in what has become a *mature* technology.

The fast, sustained, learning curve for electronic devices, and the scope for positive feedback in the information sciences, together result in a tiered infrastructure for the information sciences and technologies which is fundamental to their nature [61, 66]. It involves a succession of learning curves as rapid advances in one level of technology trigger off invention in others as shown in Figure 1. Advances in digital *electronic device technology* in the 1930s allowed the *virtual machine architecture* that detached computing from electronics to be developed in the late 1940s. This in turn triggered off developments in *problem oriented languages* in the 1950s, and the increasing reliability of computers enabled *interactive activity systems* to be developed in the 1960s. It is this last development that supports digital communications, interactive systems and networking and is most critical to the support of scholarly communities through the internet.

One important factor in the application of information technology that is apparent in Figure 1 is that there is a significant time interval between inventions and product innovation based on them. This seems to be remarkably constant at about sixteen to seventeen years across diverse areas of technology [102]. In terms of the learning curves this interval corresponds to the *replication/empirical* periods when the technology is first

being investigated. There is a corresponding interval during the *theory/automation* periods before products become mature, and automated mass-production at low cost is feasible. The resultant trajectories of *invention, research, product innovation, product lines, low-cost products* and *throw-away products* where replacement is cheaper than maintenance are shown superimposed on Figure 1. It becomes reasonable for pioneering applications to experiment with the new technology during the product innovation era, but it is not until low-cost products become available that widespread use will become common, and it is the mature technology of throw-away products that becomes used routinely as the expected mode of operation.

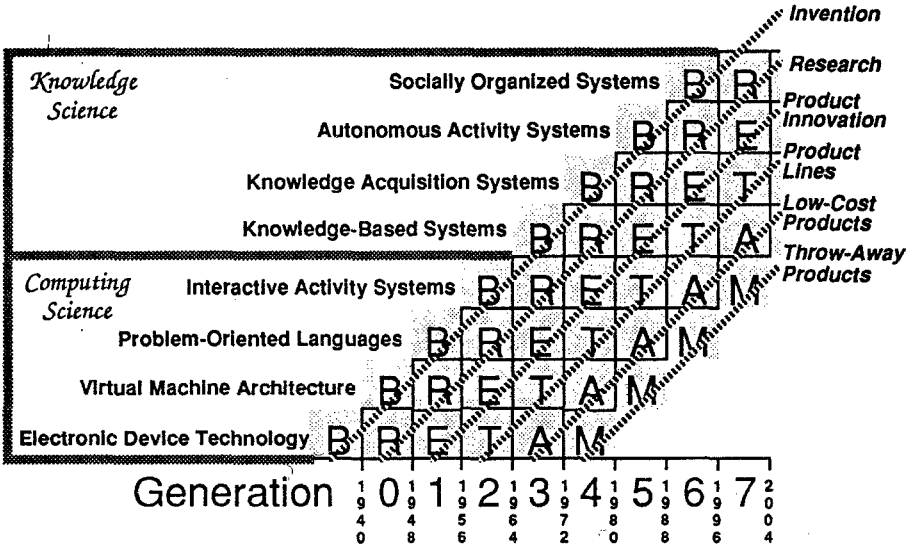

- **Breakthrough: creative advance made**
- **Replication period: experience gained by mimicing breakthrough**
- **Empirical period: design rules formulated from experience**
- **Theoretical period: underlying theories formulated and tested**
- **Automation period: theories predict experience & generate rules**
- **Maturity: theories become assimilated and used routinely**

Figure 1 Learning curves in the infrastructure of the information technology

2.2 The Fifth Generation Transition from Computer Science to Knowledge Science

A major transition from information technology to knowledge technology was triggered in the 1970s with the advent of *knowledge-based systems*, and this in its turn has stimulated research in *knowledge acquisition systems, autonomous activity systems,* and *socially organized systems*. Such developments belong to the new era of *knowledge science* associated with the fifth generation of computers that focuses as much on the content and intentions of computer-based activities as on the underlying technologies of hardware, software, communications and human-computer interfaces.

Figure 2 left shows the cross section of Figure 1 that is relevant to the state of the art in information technologies during the fifth generation of computers. The top three levels on the right of invention, research and innovation show why the fifth generation is generally recognized for its innovations in artificial intelligence [56, 108]. It was during this period

that knowledge-based system products such as expert system shells first became available. However, in terms of reliance upon proven technology, it is the lower levels of product lines and below that are significant. The fifth generation was that in which human-computer interaction was dramatically improved through *graphic users interfaces*, *object oriented languages* brought control of complex system development in software engineering, and *networking* became ubiquitous. All these innovations took for granted advances in the underlying device technology that offered very fast powerful and reliable processors and large high-speed memories at low-cost.

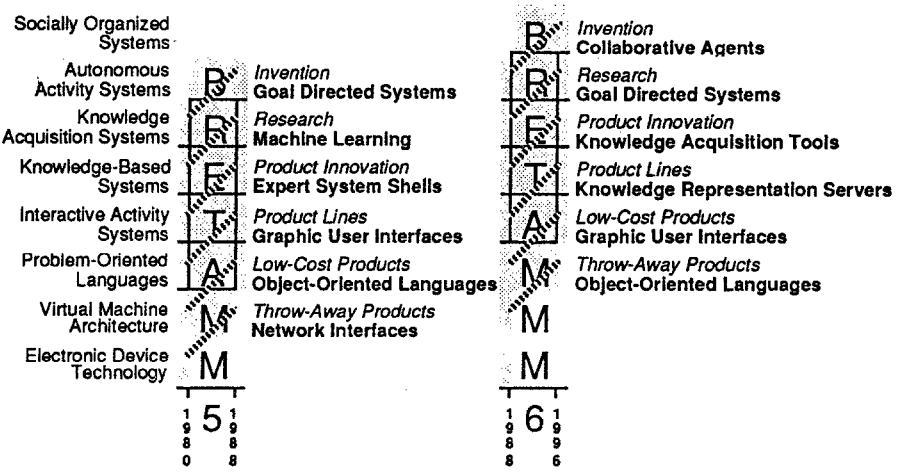

Figure 2 Significant technologies in the fifth and sixth generations

Figure 2 right shows the equivalent picture of what is happening now as we progress through the sixth generation of computers. Hardware and networking have become almost negligible in cost and almost indefinitely powerful. Large-scale distributed systems are becoming readily available in terms of equipment and hardware architectures. Object oriented languages, and their associated application programming support environments (APSEs) and class libraries, are becoming routinely available at very low-cost. Graphic user interfaces (GUIs) are become standardized and portable across platforms as a routinely available technology. By the end of this generation the lowest level of knowledge-based system technology will have become available as well-supported product lines. These will support large-scale conceptual modeling at the enterprise level, the integration of heterogeneous information and processes at lower levels, and the emulation of many aspects of human skilled behavior. Whether they are called "object oriented deductive databases," "second generation expert system shells," or "knowledge representation servers," or something completely different, is a matter for fashion, chaos theory, linguistics and marketing, to determine—we can already define their functionality and exhibit their application and that is enough.

3 Knowledge Representation Technology

The PEIRCE project fits into the infrastructure of information technology as shown in Figure 1 as an instance of one of a number of major knowledge representation and inference tools being developed as the *knowledge-based systems* learning curve goes

through its *theory* phase in the sixth generation. For knowledge representation technologies this phase marks the transition from product innovation to product lines, from exploratory system development to the availability of systems that offer high performance and scale to larger problems. It is interesting to track the complete learning curve to situate current developments in terms of their origins and prospects.

3.1 Breakthrough, Replication and Empirical Phases of Knowledge Representation

The primary breakthroughs in knowledge-based systems in the early 1970s were three parallel but unrelated developments: MYCIN [137] as the first successful demonstration of an effective knowledge-based system; Minsky's [106] Frames formalizing pre-breakthrough notions of semantic networks [125]; and the development of Prolog [89] as a principled, general-purpose knowledge representation and inference system. The MYCIN line of development led to experience with meta-representational tools such as TEIRESIAS [44] and Guidon [38], and replication of the EMYCIN shell combined with the concept of frames led to products such as KEE, ART and NEXPERT in the next phase of the learning curve in the 1980s. The Prolog line of development was given major impetus by its central role in the Japanese Fifth Generation computing research program commencing in 1982 [56, 108], and led to the development of deductive databases [105] as a parallel approach to knowledge representation alternative to those arising out of MYCIN and frames.

Sowa's [140] book on *conceptual structures* written in the early 1980s draws on these breakthroughs in knowledge-based systems, and emphasizes three particular requirements which together have characterized ongoing research arising out of his seminal work. First, *natural language* is seen as the primary form of human knowledge representation, and adequate treatment of natural language as an essential requirement of a knowledge representation schema. The primacy of language in human knowledge processes may seem obvious to common sense, but the lack of a formal linguistic model of the semantics of natural language has led to its significance as a knowledge representation schema being down-played in artificial intelligence research. Natural language processing has been presented as one area of application of artificial intelligence research, but it has not been treated as the generative, central theme of knowledge representation research. Second, *formal logic*, which *has* been accepted as the generative, central theme of knowledge representation research, was also seen as essential in Sowa's analysis of conceptual structures. All conceptual structures should have a translation in formal logic, and all operations on structures should be logically well-defined. Third, a *visual language* representation of conceptual structures is also seen as an essential requirement at a formal level ensuring complete intra-translatability between natural language, visual language and logical language. Visual languages are attractive for their perspicuity in presenting complex knowledge structures [113], and have become recognized as a valid way of presenting formal proofs in mathematical disciplines [17].

3.2 The Current Theory Phase of Knowledge Representation

We are now in the *theory* phase of the knowledge-based systems learning curve in which the experience of the 1970s and 1980s is being used to develop foundational models of knowledge representation. One significant line of theoretical development has been the analysis of the *computational complexity* of reasoning in knowledge representation schema relative to the features available. The precursor of this line of research was Brachman and Levesque's [28] studies of the complexity of subsumption in frame-based

languages, and the major results were developed in the late 1980s and early 1990s such as the undecidability of KL-ONE [132], intractability through the growth in the number of anonymous concepts [110], and the detailed relations between particular features and complexity results [48, 93]. Another significant line of theoretical development has been the application of abstract date type theory to knowledge representation schema. The precursor of this line of research was Cardelli's [32] and Aït-Kaci's [2] proposals for the typing of record structures and frames, and the major results were developed in the late 1980s and early 1990s such as semantics of types for database objects [115], efficient algorithms for type computation [7] and the design of LIFE [6] as fully typed logic and functional programming language.

The significance of the theory phase in a technological learning curve is that it begins to provide secure foundations for significant applications and the basis for product lines with life times sufficiently long to support the growth of user communities. A major new factor is also emerging in the 1990s in that the massive decline in costs of computing and networking technologies is changing the structure of the industry. Commercial ventures now have to compete with grass-roots community ventures developing freeware and shareware, and many product lines are being developed and supported without conventional commercial frameworks. This is consistent with fundamental socio-economic changes in which differing functions of industry, such as the provision of manufactured products and the provision of employment leading to the redistribution of wealth, are becoming separated in the post-modern economy [68]. It is also consistent with trends in commercial software pricing and support in which, as shown in Figure 2, major software products have become throw-away products whose cost is so low that they can be supported only indirectly through dissemination of updates over the internet. Much unix software is provided as freeware through organizations such as the Free Software Foundation, and the internet itself has evolved as a public domain activity [98, 124] with an elaborate infrastructure of RFC's defining standards that become socially adopted *de facto* rather *de jure*.

3.3 The KL-ONE Family, CLASSIC, LOOM, BACK, KRIS and KRS

The dominant form of implemented and available knowledge representation system in the late 1980s and 1990s has been those systems deriving from Brachman's KL-ONE [29] development, and variously called *term subsumption systems* or *terminological logics*. The breakthrough made in KL-ONE compared with earlier developments of semantic networks is that it incorporated Woods [150] and Brachman's [26] theoretical insights into the essential intensional nature of concepts. This led to a clear distinction between *concepts* and *individuals*, between *T-box* representation and inference, and *A-box* representation and inference. This is captured in logical theory by Zalta's [152] statement that abstract objects (concepts) *encode* properties whereas concrete objects (individuals) *exemplify* them. This distinction is also foundational for deductive databases and leads to a central role for *object identities* and their surrogates [42].

The original KL-ONE implementations were by no means deductively complete, and an interesting stage in the evolution of term subsumption systems was Brachman and Levesque's development of KRYPTON [27] as a system with similar functionality but embedded in a theorem prover. KRYPTON was too slow to be useful and this triggered off the theoretical studies of the relation between representational power and the worse-case computational complexity of inference that were discussed above. Complexity considerations have determined research directions in KL-ONE-like systems in recent

years. Brachman decided to undertake the development of a knowledge representation schema with reduced features such that deductive completeness could be achieved in polynomial time, and this resulted in the implementation of CLASSIC [23], which has been applied in AT&T to the management of large software engineering projects [45]. Other widely used KL-ONE derivatives such as LOOM [96] and BACK [109] have continued to offer rich representational capabilities and accept incomplete inference. KRIS [12] is an interesting attempt to offer full logical representation and deductive completeness in such a way that there is no penalty for generality in inference with knowledge structures that do not utilize the representational features known to lead to intractability.

The current state of the art is that CLASSIC, despite its reduced representational capabilities, has been shown [64] to exhibit the intractable growth of anonymous concepts suggested by Nebel [110], and to be deductively incomplete in that it is able to represent the intractable problem of SAT with positive literals [93] while not being able to solve it. This casts doubt on the value of reduced representational power, and this doubt has been addressed in practical terms by Doyle and Patil [49] who conclude that, while there are sound arguments for such limitations, the capabilities of the resultant systems will often fail to satisfy reasonable application requirements. One impact of this is that system designers may add functionality that provides the missing capabilities in ways that are less principled than those of a general server. A second is that problems may have to be represented in unnatural ways that are conducive to poor performance. A third is that the effect of the limitations on deductive capabilities may not be apparent to users, leading to errors. Empirical studies [83] of a number of the major implemented KL-ONE-like systems in terms of speed and deductive capabilities on a range of problems, which originally showed KRIS to be generally very much slower than CLASSIC, also resulted in implementation insights which have allowed KRIS to be reimplemented to be similar in speed to CLASSIC for those problems which CLASSIC can solve [13]. Thus, a 'pay as you go' strategy of supporting a full set of features and deductive completeness but solving in polynomial time problems that do not exhibit worse case intractability now seems both attractive and feasible.

A major problem with term-subsumption systems developed with standard logical foundations is that the theory does not readily extend to encompass the range of data types required in practical applications, such as numbers, dates, strings and associated constraints. The knowledge representation server, KRS [64], originated as a fast and light-weight reimplementation of CLASSIC as a open architecture class library in C++. However, to allow new data types to be added in a principled fashion the theoretical model adopted was one of constraint algebras [65, 71] rather than predicate calculus. KRS is modeled through an algebra of *constraints* as composable primitives for which composition is idempotent, associative and commutative. The computational units are: for the TBox, *concepts*, identified with collections of *constraints* indexed by roles; and for the ABox, *individuals* represented as unique identifiers, each of which has associated conceptual variable whose value in any particular state of the knowledge base is a concept. This has the normal semantics for KL-ONE systems that concepts are precisely a composition of constraints—changing the constraints changes the identity of the concept, whereas individuals are precisely identities of objects—changing the associated constraints changes the *state* of the individual, not its identity. The class library implementation of KRS allows new families of constraints to be added as classes accepting a small number of well-defined messages, such that inference can continue to

be carried out by the general engine and remains sound and complete. The semantics of CLASSIC can be modeled as based on a single but unconventional data type, of sets with lower and upper set and cardinality bounds.

The significance of the constraint algebra model is that it bridges between the theoretical foundations of KL-ONE systems and those of deductive databases. Aït-Kaci's [2, 3] ψ-calculus model of deductive data bases with operational semantics through term rewriting is an example of a constraint algebra. Ohori's [115] notion of inference in deductive databases as a series of improving approximations may be modeled as a trajectory of constraints of increasing strength in the A-box. That is, monotonic deductive inference is a process that adds additional constraints to individuals, and terminates on a particular dimension when the constraint on that dimension has become an atom, or *value*, in that the addition of further constraints necessarily leads to a fixed point or a contradiction. These links between KL-ONE systems and deductive databases have been further strengthened with the development of semantics for recursive definitions in the T-box [110] and associated algorithms for the computation of subsumption between recursively defined concepts [46]. Maximum fixed-point semantics for subsumption between recursive definitions are simply computed by an optimistic algorithm that assumes subsumption unless a contradiction is reached, and corresponds to the *solution lemma* of Aczel's [1] model of non-well-founded sets.

Thus, overall, there are now a number of implemented, available knowledge representation systems in the KL-ONE family that are being used in a variety of large-scale applications. The theoretical foundations of these systems continue to develop and are leading to algorithms that are efficient in space and time, and automatically exploit the trade-off between representational power and tractability of inference without reducing the functionality available to users. An important development in the usability and integratibility of these systems is the agreement of an emerging standard for application programming interfaces and knowledge interchange, KRSS [118], and the availability of translators from emerging knowledge interchange standards [112] such as KIF [78] and Ontolingua [81]. There is also convergence between the features and capabilities of term subsumption knowledge representation systems and those of deductive databases.

3.4 Deductive Databases

The development of deductive databases has paralleled that of knowledge representation systems and, although the two technologies are rapidly and completely converging, there has been little overlap in the research communities involved. The deductive database approach is important in having two different derivations, each of which is highly significant in its own right. From one perspective, deductive databases may be seen as deriving from the imposition of integrity constraints on relational databases [77] and having strong links with developments of conventional databases, particularly the computational availability of a *semantic model* of the database at run time [87]. From another, they can be seen as deriving from the representation of *complex objects* through many-sorted extensions to logic programming languages, resulting in record structures and typing in Prolog and its derivatives [16, 34]

There are now a number of implemented and available deductive databases that are use by numbers of researchers such as O_2 [15] and CORAL [126]. CORAL is available for ftp (ftp.cs.wisc.edu) as an extensible class library in C++ with links to the EXODUS

storage manager to support persistent relations, and is documented as an extensible, embeddable system. Aït-kaci's LOGIN [8] development has also been highly influential as a natural embedding of complex record structures and inheritance in Prolog. This is now subsumed by LIFE and discussed separately below. The theoretical relations between term subsumption knowledge representation systems and deductive databases should lead to an increasing exchange of technologies. Knowledge representation research could have benefited some time ago from incorporation of results on recursive definitions and set structures from deductive databases, and deductive database research would have had more wide-ranging impact if the technology had been used from the beginning to tackle some of the significant problem domains arising out of artificial intelligence research.

3.5 Theorem-Provers, Boyer-Moore, Otter, HOL and Isabelle

An important source of implemented and available inference technology is the range of theorem provers and proof generators for first- and higher-order logics that have been developed by various groups and made available through ftp. Boyer and Moore's [24] first-order system that was developed at SRI in the mid-1970s has been widely used and extended with input from many researchers [25] as has Otter [101] from Argonne National Laboratory. Higher-order logic theorem provers are also available and in widespread use in research on hardware and software verification, notably Gordon's HOL [79] and Paulson's Isabelle [119]. These systems are a rich source of both deductive inference techniques and test cases for knowledge representation systems. Their main weakness in a knowledge representation context is that the inheritance relations of subsumption in semantic networks have to be expressed in terms of general implication, and the lack of special proof techniques for inheritance causes the systems to be slow through excessive search when dealing with quite simple problems [8]. An illustration of this is Schubert's *steamroller* problem which was designed to be intractable for resolution theorem provers, but which can be solved simply and rapidly if inference over explicit sorts is incorporated [147].

Another important advance in inference technology in recent years has been the development of symbolic *model-checking* techniques for very large state spaces [30] using binary decision trees [95]. The case by case reasoning involved in model checking has tended to be regarded in knowledge representation research as a source of intractability to be avoided by choice of representational schema [53], whereas there are strong arguments for its being more effective than general theorem proving in many situations [82]. One reason that model checking has been neglected in knowledge representation research is that the subsumption relation tends to be defined extensionally which confounds theorem proving and model checking. It would clarify the issues involved if *subsumption* were taken to be a term for intensional implication, or entailment, such that "concept A subsumes concept B" is defined intensionally in terms of all the properties of A being implied by those of B. Intensional subsumption is most naturally modeled as an algebraic relation between constraints, such that the composition of the constraints of A with those of B results in the constraints of B as a fixed point. The case by case reasoning of model checking corresponds to extensional subsumption and is most naturally modeled as a modal inference involving search over possible worlds. The two techniques of constraint propagation and world search combine naturally as separate, but interacting services in a knowledge representation server [71].

3.6 LIFE

LIFE is a unique development combining types, functions and relations in a hybrid, yet highly integrated, knowledge representation, deductive database, theorem proving and functional programming system. Its theoretical foundations and current implementation in C as Wild-LIFE (available for ftp on gatekeeper.dec.com) are the culmination of Aït-Kaci's research over the past decade, commencing with his thesis on a lattice-theoretic approach to computation based on a calculus of partially-ordered types [2], and continuing with a wide range of theoretical and applied papers on the features and effective implementation of systems for reasoning with complex object definitions [3-9, 138, 139]. Over the years, Aït-Kaci has developed a number of hybrid systems combining different methodologies such as LOGIN incorporating types in Prolog, Le Fun a relational and functional programming language, and FOOL a pattern-oriented functional language, and these are now all subsumed in LIFE as shown in Figure 3.

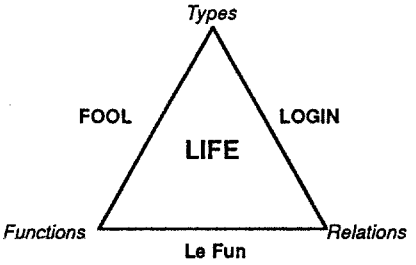

Figure 3 The architecture of LIFE and its predecessors

Aït-Kaci's co-workers in LIFE have included those from the term subsumption knowledge representation community such as Smolka who has worked with Nebel on representation and reasoning with attributive descriptions [111], and Ellis who together with Levinson initiated the PEIRCE project. LIFE is one of the most eclectic developments in knowledge representation and inference systems, and the open availability of a working system running on many platforms with access to source code and many related theoretical papers may be expected to have a major impact on other developments.

3.7 PEIRCE

The PEIRCE project has much in common with those described above in that it is making available a principled implementation of a general knowledge representation system in source form to the research community. It differs from them in taking conceptual graphs as its focus and providing graph-theoretic operations which go beyond deductive inference in providing, for example, the generalizations underlying analogical reasoning [92]. This makes the name particularly appropriate since Peirce treated abduction, induction and deduction as patterns of human reasoning having equal status within a unified framework [127], and hence gave logical foundations for the philosophy of science and for many developments in artificial intelligence that cannot be reduced to deductive theorem proving. The emphasis on visual languages which has been central to conceptual graph development is also reflected in the PEIRCE project, and Sowa's [140] original proposal for conceptual graphs stems directly from Peirce's own research on *existential graphs* as a proof technique in logical reasoning [128].

The central role of natural language in the foundations of conceptual graphs has been reflected in a regular track of natural language research at the annual conceptual graphs workshops over the years, and is being incorporated in PEIRCE through a working group on natural language parsers and generators [51]. There has been a rich vein of research on semantic processing of natural language that draws on KL-ONE knowledge representation schema for its implementations, but seems to have surprisingly little impact on the foundations and underlying theory of those implementations. This is documented in recent proceedings on natural language and logic [143] and in the final report of the LILOG project [85] managed by IBM in Germany. Given the many deep relations between term subsumption systems and conceptual graphs, it will be interesting to see how well techniques developed on these projects carry over to PEIRCE.

The current status of PEIRCE in March 1993 is that C++ source code, a manual and test data sets are available for ftp (from ftp.cs.uq.oz.au). The most significant feature of the initial implementation is its incorporation of the algorithms for efficient search of large databases of conceptual graphs developed by Levinson and Ellis [94]. These utilize the same types of techniques that have allowed KRIS to be brought up to speed with CLASSIC described above [13], and provide a powerful central engine for the other components of PEIRCE currently under development.

3.8 Summary

The current sixth generation of computing is one in which the technological foundations of knowledge-based systems are being secured by the widespread availability of principled and efficient implementations of knowledge representation and inference servers. The first generation of such implemented systems arose out of research in diverse communities with differing origins and objectives and little overlap in personnel and professional discourse. A second generation is now arising based on a synthesis of theory, features, and experience across these differing sub-disciplines of information technology. It is through such synthesis that major advances will be made.

4 Changing Patterns of Scholarly Communication

The conceptual graphs list server fits into the infrastructure of information technology as shown in Figure 1 as an instance of one of a number network-based community support services that became possible as the *interactive activity systems* learning curve went through its theory phase in the fifth generation. Personal computers communicating through wide area networks at low cost have made possible new modes of discourse and knowledge dissemination in scholarly communities. A change in the pattern of scholarly communication away from paper-based journals is a truly historic event since it is three hundred and twenty five years since the first two scholarly journals came into being: the *Journal des Sçavans* in January 1665 in France, and the *Philosophical Transactions* of the Royal Society in March 1665 in England. Now, as we prepare to enter the next millennium, the two have grown to some fifty thousand, and what was the blessing of improved scholarly communication has become the curse of information overload. The growth of knowledge that has its seeds in the seventeenth century enlightenment has become a prolific jungle, the understanding and control of which through the information sciences and technologies is an important knowledge objective in its own right.

4.1 Electronic Journals

One major question which has to be answered is whether the technology to replace paper is yet ready. The problems created by the use of an obsolescent medium have long been recognized, but there has been little progress in overcoming them. In 1937, H G Wells promoted a "permanent world encyclopaedia" giving access to all of human knowledge, and remarks:

> "our contemporary encyclopaedias are still in the coach-and-horses phase of development, rather than in the phase of the automobile and the aeroplane. Encyclopaedic enterprise has not kept pace with material progress. These observers realize that the modern facilities of transport, radio, photographic reproduction and so forth are rendering practicable a much more fully succinct and accessible assembly of facts and ideas than was ever possible before." [148]

In 1939, Churchill's science advisor J D Bernal echoed the same issues in relation to journals:

> "In the old ideal of science, communications were the only link between scientists. Now the very quantity of scientific information has made its diffusion an enormous problem, with which the existing machinery has utterly failed to cope. The present mode of scientific publication is predominantly through the 33,000 odd scientific journals. It is, as we have already shown, incredibly cumbersome and wasteful and is in danger of breaking down on account of expense." [21]

Some six years later Roosevelt's science advisor Vannevar Bush was also echoing these issues in his article in *Atlantic Monthly* [31] which is often cited as first expressing the need for hypertext.

These comments of fifty years ago still apply today since we have made remarkably little progress in applying modern information technology to mainstream knowledge dissemination. In 1978 Lancaster documented the problems of existing paper media in great detail, and saw "paperless information systems" as providing a solution to the problem [91]. Following an economic analysis [133], an NSF-funded experiment started in 1978 in which a psychological journal on mental workload was developed and operated on the EIES computer-based conferencing system [86]. The final report [136] states that "the results fell short of initial aspirations" and "many of the participants, not themselves experienced with computers, found EIES inconvenient or difficult to use." However, dissatisfaction with the technology is not surprising at a time before the advent of UUCP and USENET, when access to the system was at 110 or 300 baud, largely through teleprinters, and with no possibility of transmitting anything other than textual material in a discipline that traditionally makes heavy use of graphs and figures. Such technical limitations were also apparent in the larger British Library BLEND experiment in the UK from 1981 through 1983 [114]. Surveys of digital journals in the late 1980s indicated that none had succeeded in establishing a viable publication [54], and that some providing a useful service, such as *Clinical Notes On-Line*, had been discontinued [47].

Experimental publication of digital journals continues, and there have been positive reports on some activities in the 1990s such as the *Public-Access Computer Systems Review* on BITNET [14]. The most important of recent developments has been the publication by the American Association for the Advancement of Science in the Spring of 1992 of a new journal, *The Online Journal of Current Clinical Trials*, as an online, peer-reviewed journal [117] making full use of new electronic documentation technologies [146]. There are also many innovative uses of the network to support scholarship that replace or complement some of the functions of journals [124, 129].

4.2 Multimedia Digital Publication

While the technologies for emulating current paper publications have only become available in recent years, it seems that as we attain them we are becoming capable of offering functionality that goes way beyond them, for example, in terms of incorporating sound and video, hypertextual links, datasets, analysis and simulation tools, and overt, operational knowledge structures [20, 33, 41, 121]. There have been many experimental multimedia document systems developed in the past decade that show that it is possible to deliver full typographic documents with embedded figures, that have the same quality and functionality as a paper document, and offer much greater potential functionality for indexing, retrieval, analysis, reuse, and active components [116, 130, 131, 149, 151]. The availability of low-cost CD-ROM production and readers makes it possible to publish such documents economically [84, 90]. The MIME standard [22] supports their distribution on Internet, and evolving standards for document structure and products for document format conversion support interchange between formats [70].

These possibilities enhance the attractiveness of electronic publication as a more versatile medium, and authors are already beginning to take advantage of them. For example, Claerbout has published a research monograph both as a book [36] and as a CD-ROM [37] in which data and programs are linked to the plots in the book as "interactive figures" such that the results plotted may be re-analyzed as the reader wishes. One has to be careful that active document features do not limit the possibility of parallel publication of paper and electronic versions of the same material, but this is not a major limitation. Gaines and Shaw [74] have published a paper on a knowledge-based system for room allocation as camera-ready copy in a book such that the parallel electronic version runs as an expert system to solve specific problems.

4.3 Social Factors in Knowledge Dissemination

The new technology has also to take into account the social factors involved in the acceptance of the new medium, for example, that academic publication involves complex reinforcement mechanisms in which society rewards the dissemination of knowledge through status symbols such as publication in 'high-quality' journals that then translate into direct economic benefits [103]. It is also important to note that the critical processes of a scholarly community rely on the 'frozen' status of a published work so that commentaries apply to well-defined and widely available arguments. Electronic publications have to support mechanisms for refereeing, authenticating and time-stamping publications so that quality and priority of publication may be evaluated in ways that can be trusted, and editing of versions of electronic documents can be detected [69].

Currently electronic journals seem to have continuing difficulties in competing with paper journals, largely because social reward mechanisms are still linked to well-established paper journals. One major impact of the internet has been on preprint distribution in some disciplines where rapid access to current publications is very important. A theoretical physics archive on internet receives some 600 new preprints a month and has some 8,000 users [144]. TeX is used to format most of the papers so that full typography, diagrams and mathematical formulae are supported.

One of the most interesting phenomena of internet has been the growth of scholarly discourse that completely bypasses conventional publication media by using electronic archives and electronic mailing list servers that carry material in electronic mail format

not intended for paper publication. It is impossible to keep track of the number of list servers now operating on the internet because they are simple to install and manage at any site and serve specialist communities who do not generally advertise their existence. This informality has raised concerns about the growing use of digital networks for discourse between scholars. A survey undertaken by Meadows in Britain showed many correspondents to be concerned about the development of electronic cliques of senior scholars that were invisible to those not invited to join [10]. Concern was also expressed about the circulation of data and theories that has not been adequately assessed and yet were being used by other scholars as if they were published results.

4.4 List Servers Supporting Scientific Communities

The KSI tracks a number of list servers in order to model their impact of scholarly activities, and also to develop requirements specifications and technologies for enhancing the utility of the servers, such as text analysis, information filtering, and graphic presentation of discourse structures. An interesting list server that has many parallels to cg is csg-l (on LISTSERV@VMD.CSO.UIUC.EDU) which is a list server for those interested in Powers [122] theory of behavior as the control of perception. This list has some 130 members, commenced in August 1990, carries about 130 Kbytes of mail a week and has some 6 Mbytes of mail in its archives for 1992. Bill Powers is a frequent contributor and his original book is still the primary intellectual source for the group discussions.

The community involved does not publish extensively in standard journals, but the quality of research and scholarly discourse carried on the network is very high. The community also has its own software activity similar in status to PEIRCE in that source code has been made available to the community and is being used for a variety of purposes. The software is called "Little Man" and demonstrates the principles of the behavioral control of perception in action through a simulation of visual and kinesthetic control.

By way of comparison the cg list server has some 220 members, commenced operation in April 1990, has some 3.5 Mbytes of mail in its archives for 1992. The scholarly value of the archives of both of these groups is very high. They contain cutting edge discourse between active researchers with common concerns and differing perspectives.

It is interesting to attempt to estimate the impact on research productivity of the existence of such mail servers. One objective of the Japanese sixth generation research proposal was the "systematic acceleration of scientific research" [57, 141], and internet mailing list servers already seem to be contributing to the achievement of that objective. For example, the speed and breadth of response to Gerard Ellis and Robert Levinson's call in April 1992 for an international collaborative effort to develop PEIRCE, a conceptual graph work bench, was quite remarkable, and contrasts strongly with similar proposals for collaborative development in other disciplines that generally originate at annual meetings and do not result in any significant coordination between meetings. The proceedings of the first international workshop on PEIRCE [51] from Las Cruces in July 1992 details a major project structure involving some ten different areas of activity, some 60 researchers, and some 50 research institutions in 10 different countries. There are already deliverables, and it will be very interesting in the PEIRCE workshop at this meeting in Quebec to see how much development has been achieved over the first eighteen months of operation of the project.

5 Representation, Discourse, Logic and Truth

It is reasonable to become somewhat euphoric about the progress in knowledge science surveyed in the preceding sections of this paper. We do seem to be moving towards levels of support of human knowledge processes that are a dramatic advance on past practice, in which information technology provides a new medium both for knowledge dissemination and for the overt, operational representation of knowledge. However, it is salutary at this point to look beyond the achievement of effective implementations of the n'th generation of systems such as LIFE and PEIRCE and ask how much will remain to be done. How much of human knowledge processes will we have encompassed and what will provide the challenges for future development?

5.1 Representation and Knowledge

One of the major presuppositions to be challenged in artificial intelligence research in recent years is the primacy and significance of overt knowledge representation. Clancey, One of the MYCIN pioneers, has criticized approaches to expert system development based the assumption that expertise can be captured in overt knowledge:

> "The new perspective, often called *situated cognition*, claims that all processes of behaving, including speech, problem-solving, and physical skills, are generated on the spot, not by mechanical application of scripts or rules previously stored in the brain. Knowledge can be represented, but it cannot be exhaustively inventoried by statements of belief or scripts for behaving. Knowledge is a capacity to behave adaptively within an environment; it cannot be reduced to representations of behavior or the environment." [39]

He argues that overt representations of knowledge are only partial models of the knowledge processes underlying human behavior:

> *"A representation is not equivalent to knowledge*
>
> A representation of what a person knows is just a model of his or her knowledge, a representation of a capacity. Knowledge cannot be reduced to (fully captured by) a body of representations. Knowledge cannot be inventoried.
>
> *The meaning of a representation cannot be made explicit*
>
> Meaning can be represented, but it cannot be defined once and for all, captured fully by representations. The meaning of a representation is open, though there are culturally stable representations of meaning (e.g., word senses).
>
> *The context in which a program is used cannot be made explicit*
>
> Context can be represented, but the world cannot be objectively and exhaustively described; cultural or social circumstances cannot be reduced to a set of facts and procedures." [40]

In the artificial intelligence literature *connectionist* systems are proposed as providing the capacity to behave adaptively within an environment [19], and the reductionist link from mind to brain is a significant alternative perspective to those which emphasize overt knowledge. One problem with this perspective is that it may be presented dogmatically as excluding mental phenomena:

> "I am willing to infer that folk psychology is false, and that its ontology is chimerical. Beliefs and desires are of a piece with phlogiston, caloric, and the alchemical essences." [35]

However, there is no need to adopt such a unitary reductionist position, and the more interesting research question is how to bridge the gap between the physical processes of neural networks and the phenomena of mind. The primacy of action and the behavioral basis of adaptation leading to skilled performance without imputation of a knowledge

level is being addressed by research on situated agents [97]. A more fundamental problem is that connectionist reduction focuses on the individual brain and does not address the social context of expertise.

5.2 The Primacy of Social Action

These issues relating to the fundamental nature of knowledge rapidly become of practical and technological importance in studies of knowledge acquisition and knowledge engineering [60]. One finds it impossible to treat human experts in isolation from the situation in which their expertise is exhibited as problem-solving behavior, and studies of multiple experts cast doubt on the uniformity and consistency of the knowledge apparently being used by a coherent group of agents to solve similar problems [135]. As Touraine has argued, human action does not seem to derive from represented knowledge and logical inference:

> "the sense of an action can neither be reduced to the adaptation of the actor to a more or less institutionalized normative system, nor to the expression of the mind through a social activity" [145]

and one begins to see expert agents in Gadamer's terms:

> "The nature of experience is conceived in terms of that which goes beyond it; for experience can never be science. It is in absolute antithesis to knowledge and to that kind of instruction that follows from general or theoretical knowledge. The truth of experience always contains an orientation towards new experience. That is why a person who is called 'expert' has become such not only through experiences, but is also open to new experiences. The perfection of his experience, the perfect form of what we call 'expert', does not consist in the fact that someone already knows everything and knows better than anyone else. Rather, the expert person proves to be, on the contrary, someone who is radically undogmatic; who, because of the many experiences he has had and the knowledge he draws from them is particularly equipped to have new experiences and learn from them." [55 p.412]

5.3 The Collective Stance—Humanity as an Autopoietic Unity

A critical question arises as to the relation between the overt, represented, formal knowledge structures so valued in science and artificial intelligence research, and the behavior of human experts so valued in the pragmatic survival processes of human society. The model developed in the KSI to address these issues draws on a wide range of biological, psychological and social studies and theories to generate a framework for the study of knowledge processes that differs radically from those of current artificial intelligence research. It takes a *collective stance* [63, 72] and models the human species as a single organism, an autopoietic entity that is distributed in time and space by recursive partitioning into parts similar to the whole. The parts into which the human organism is recursively partitioned include societies, organizations, groups, individuals, roles, and neurological functions. Many concepts that apply to individuals may be applied to social systems, not as metaphors or analogies, but because, from a systemic perspective, it is the same concept that is being applied to different partitions of the system.

Notions of expertise arise because the organism adapts as a whole through adaptation of its interacting parts. The behavioral mechanism is one of exchange of reinforcement through some parts allocating tasks to others. The preferential allocation of tasks to those parts which have performed well previously also gives those parts access to experience enabling them to adapt and perform better in the future. This positive feedback leads to functional differentiation of the parts and the distribution of activities [58]. From a

functionalist perspective, the coordination of the activities among the parts leads to phenomena of communication, discourse and language. The short lives of individual parts would lead to loss of knowledge by the organism as a whole unless compensated by social interactions supporting knowledge transfer, including the generation and storage of overtly represented knowledge.

The improved performance resulting from adaptation may be modeled as the part involved having acquired a model of its task. Reflective processes in which parts model the behavior of other parts including themselves leads to a hierarchy of models of increasing abstraction, and detachment from direct experience. The modeling of human activity in terms of behavioral contingencies and its modeling in terms of symbolic interactions are complementary analyses of the same phenomena. There will be some situations which are more richly represented by one of the analyses and poorly by the other. Many of the phenomena of human action and expertise are behavioral and do not involve significant symbolic representations. The Copernican revolution greatly simplified our model of the solar system by dropping an earth-centered stance and adopting a sun-centered stance—we ceased to place ourselves at the center of the physical universe. The call for a collective stance proposes a similar revolution that simplifies our model of the human system by dropping an individual-centered stance and adopting a collective stance—we cease to place ourselves at the center of the cognitive universe.

How humanity came to be regardable as a single organism is not a major issue for the theory. One could assume that it has always been so, and that we tend not to see it so because egocentric perspectives emphasizing individual humans are more useful in our everyday lives and cultures. One could assume that it was not always so but that humanity has evolved to be more naturally and usefully regarded as a single organism. Again, such considerations involve no deep issues since no such assumptions are necessary to the argument, and both these perspectives and their associated vocabularies will be treated as legitimate models. In particular, notions such as fitness will not be taken as explanatory variables but rather as post-hoc rationalizations. A chaotic universe has a probability of producing any system including autopoietic systems. Once such systems exist and are modeled properties emerge [134]. As Peirce remarks:

> "Law begets law; and chance begets chance...the first germ of law was an entity which itself arose by chance, that is as a First." [120]

Jantsch [88] and Prigogine [123] have developed detailed models of how organization emerges from chaos. Gould has analyzed the fossil record and modeled the genesis and extinction of a wide variety of species as low probability random events [80]. Monod has given a biochemical model of life as an improbable phenomena that, once it exists, follows deterministic laws [107]. When a living system comes into existence it acts to persist, but, from the systemic perspective advanced by Maturana [100], this is the definitional property by which we recognize its existence as a living system, not an additional property going beyond active persistence.

5.4 Implications for Conceptual Graphs Research

The collective stance model proposed in this section does not denigrate the overt representation of knowledge, but rather *situates* it as one mechanism by which the human species mediates its coordination processes. This is consistent with the philosophical issues that Sowa [140] raises in developing a theory of conceptual structures, and also

with Peirce's underlying *evolutionary metaphysics* from which his own theories of logic and knowledge derive [52]. In particular, it seems appropriate to raise these very general issues in the context of this paper because the primacy afforded natural language in knowledge representation through conceptual graphs seems more amenable to extensions which take into account the anarchy of human knowledge processes than do those approaches that give that primacy to formal logic. Technically, one can see conceptual graphs as providing a translation schema from natural language phrases to graphs and hence to logical expressions. However, this translation schema is semantically incomplete because no interpretation of many of the terms involved is specified. It is not purely syntactic because closed-class words are used to determine the logical form and force of the phrase. The lack of interpretation of open-class words is not a weakness of the translation but rather its strength because it corresponds to those aspects of meaning brought to the phrase by the interpreter, including those aspects which are inferred from the situation in which the phrase has arisen.

Sowa has expressed this very clearly:

> "The semantic base depends on what the speaker knows about the topic. The way in which the speaker presents the topic depends on pragmatics—context, external circumstances, and the listener's expectations. There is no reason to believe that all these aspects of meaning originate in a single base structure. Instead, a sentence is derived from six different kinds of information:
>
> - *Conceptual graphs* are the logical forms that state relationships between persons, things, attributes and events.
> - *Tense* and *modality* describe how conceptual graphs relate to the real world. They state whether something has happened, can happen, will happen, or should happen.
> - *Presupposition* is the background information that the speaker and listener tacitly assume.
> - *Focus* is the new point that a speaker is trying to make.
> - *Coreference links* show which concepts refer to the same entities. In a sentence, these links are expressed as pronouns and other *anaphoric references*.
> - *Emotional connotations* are determined by associations in the minds of the speaker and listener" [140, p.10]

It is very important that these considerations remain central to the development of a conceptual graphs work bench. They are vital to the wide ranging role that such a work bench might play in the modeling and support of human knowledge processes. It is clearly relevant to compare the work bench with other knowledge representation tools as was done in Section 3, but it is also very important that the work bench does not become subsumed under the considerations that have motivated the development of those other tools. The issues involved in conceptual graphs theory and technology are of very much wider scope.

As an illustration of the interaction between technical features and deep knowledge-theoretic issues consider the notion of a *context* in conceptual graphs. A context box may be seen as a simple aggregation operator allowing a sub-graph to be itself treated as a proposition. However, this simple mechanism has the power to deal with many difficult issues, including the anarchy of knowledge processes. For example, it gives a representational schema for the linguistically and philosophically difficult analysis of *propositional attitudes* that Barwise and Perry analyze in terms of situation theory [18]. What it makes possible is the recursive treatment of statements *about* a complex knowledge structure as statements in the same language as that knowledge structure.

The boundary imposed by the context box enables one to express, for example, that a term used by one person is being used in the same sense as a different term by another person, or that the same lexical term being used by two different people is being used with completely different meanings, both situations being common in knowledge acquisition [135]. The ability to define and manipulate contexts is fundamentally important in a wide variety of applications and it is important that a conceptual graphs work bench support it. It will be also interesting to see whether it possible, since there is no universal mechanism for dealing with deductions about contexts, to include abstract operations on contexts of a high degree of generality and applicability. In the same way that one may gain insights into *inheritance* in its full generality as a simple graph rewriting operation, one may gain related insights into *context* in its full generality as a source of various graph-theoretic transformations.

6 Conclusions

The cumulative impact of developments in information technology during the past fifty years is resulting in major qualitative changes in applications of the technology, from information processing to knowledge processing. These changes not only involve new technologies, such as those for knowledge representation and acquisition, but also new social impacts such as the establishment of major scientific communities whose primary mode of discourse is through the internet. The conceptual graphs cg list server is a major example of the effective use of the internet to coordinate the intellectual activities of a widely dispersed research community. The PEIRCE project is an outstanding example of the effective use of the internet to initiate, design and coordinate development of a major software system. This paper has examined these activities from three different perspectives, placing the cg list server and PEIRCE project in the context of the development of the infrastructure of information technology, comparing the project with other developments in knowledge representation, comparing the cg list server with related activities, and looking beyond PEIRCE at what will have been achieved, and what issues will remain to be addressed, once this project, and related ones, have been completed.

Acknowledgments

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada. I am grateful to John Aldwinckle, Debbie Leishman, Mengchi Liu, Mildred Shaw, and Kwang Sim for many stimulating discussions that have influenced this paper. I am grateful to Bosco Tjan for making the cg list archives available.

References

[1] P. Aczel, *Non-Well-Founded Sets*, Menlo Park, California: CLSI. 1988.

[2] H. Aït-Kaci, A lattice-theoretic approach to computation based on a calculus of partially-ordered types. University of Pennsylvania, Philadelphia. 1984.

[3] H. Aït-Kaci, "An algebraic semantics approach to the effective resolution of type equations," *Theoretical Computer Science*, vol. 45, pp. 293-351, 1986.

[4] H. Aït-Kaci, "Solving type equations by graph rewriting," in *Rewriting Techniques and Applications*, J.P. Jovannaud, Editor. Springer: Berlin. p. 158-179, 1986.

[5] H. Aït-Kaci, "A glimpse of paradise," in *Next Generation Information Technology*, J.W. Schmidt and A.A. Stogny, Editor. Springer: Berlin. p. 17-25, 1991.

[6] H. Aït-Kaci, "An overview of LIFE," in *Next Generation Information Technology*, J.W. Schmidt and A.A. Stogny, Editor. Springer: Berlin. p. 42-58, 1991.

[7] H. Aït-Kaci, R. Boyer, P. Lincoln, and R. Nasr, "Efficient implementation of lattice operations," *ACM Transactions Programming Languages Systems*, vol. 11, no. 1 pp. 115-146, 1989.

[8] H. Aït-Kaci and R. Nasr, "LOGIN: A logic programming language with built-in inheritance," *Journal of Logic Programming*, vol. 3, no. 3 pp. 185-215, 1986.

[9] H. Aït-Kaci and A. Podelski, Towards a meaning of LIFE. Paris Research Laboratory, Digital Equipment Corporation. pp. 1-42, 1991.

[10] A. Anderson, "Networks for thinking in cliques?," *Science*, vol. 253, no. 5019 pp. 506, 1991.

[11] R.U. Ayres, "Envelope curve forecasting," in *Technological Forecasting for Industry and Government: Methods and Applications*. Prentice-Hall: Englewood-Cliffs, New Jersey. p. 77-94, 1968.

[12] F. Baader and B. Hollunder, "KRIS: Knowledge representation and inference system," *ACM SIGART Bulletin*, vol. 2, no. 3 pp. 8-14, 1991.

[13] F. Baader, B. Hollunder, B. Nebel, and E. Franconi, "An empirical analysis of optimization techniques for terminological representation systems or: Making KRIS get a move on," in *Proceedings of KR'92: Third International Conference on Principles of Knowledge Representation and Reasoning*. Morgan Kauffman: San Mateo, California. p. 270-281, 1992.

[14] C.W. Bailey, "Electronic (online) publishing in action: the *Public-Access Computer Systems Review* and other electronic serials," *ONLINE*, vol. 15, no. 1 pp. 28-25, 1991.

[15] F. Bancilhon, C. Delobel, and P. Kanellakis, ed. *Building an Object-Oriented Database System: The Story of O2*. Morgan Kaufmann: San Mateo. 1992.

[16] F. Bancilhon and S. Khoshafian, "A calculus for complex objects," *Journal Computer System Sciences*, vol. 38, no. 2 pp. 326-340, 1989.

[17] J. Barwise and J. Etchemendy, "Valid inference and visual representation," in *Visualization in Mathematics*, Zimmerman and Cunningham, Editor. American Mathematical Association: 1990.

[18] J. Barwise and J. Perry, *Situations and Attitudes*, Cambridge, Massachusetts: MIT Press. 1983.

[19] W. Bechtel and A. Abrahamsen, *Connectionism and the Mind*, Cambridge, Massachusetts: Blackwell. 1991.

[20] I.D. Benest, "A hypertext system with controlled hype," in *Hypertext: State of the Art*, R. McAleese and C. Green, Editor. Intellect: Oxford, UK. p. 52-63, 1990.

[21] J.D. Bernal, *The Social Function of Science*, London: Routledge. 1939.

[22] N. Borenstein and N. Freed, MIME (Multipurpose Internet Mail Extensions): Mechanisms for specifying and describing the format of internet message bodies. Internet. 1992.

[23] A. Borgida, R.J. Brachman, D.L. McGuiness, and L.A. Resnick, "CLASSIC: a structural data model for objects," in *Proceedings of 1989 SIGMOD Conference on the Management of Data*. ACM Press: New York. p. 58-67, 1989.

[24] R.S. Boyer and J.S. Moore, *A Computational Logic*, New York: Academic Press. 1979.

[25] R.S. Boyer and J.S. Moore, *A Computational Logic Handbook*, Boston: Academic Press. 1988.

[26] R.J. Brachman, "What's in a concept: structural foundations for semantic nets," *International Journal of Man-Machine Studies*, vol. 9, pp. 127-152, 1977.

[27] R.J. Brachman, V.P. Gilbert, and H.J. Levesque, "An essential hybrid reasoning system: knowledge and symbol level accounts of KRYPTON," in *Proceedings of IJCAI85*. Morgan Kauffman: San Mateo, California. p. 547-551, 1985.

[28] R.J. Brachman and H.J. Levesque, "The tractability of subsumption in frame-based description languages," in *Proceedings of AAAI-84*. Morgan Kauffman: San Mateo, California. p. 34-37, 1984.

[29] R.J. Brachman and J. Schmolze, "An overview of the KL-ONE knowledge representation system," *Cognitive Science*, vol. 9, no. 2 pp. 171-216, 1985.

[30] J.R. Burch, E.M. Clarke, K.L. McMillan, D.L. Dill, and L.J. Hwang, "Symbolic model checking: 10 to the power of 20 states and beyond," in *Proceedings of Logic in Computer Science*. IEEE: New York. p. 428-439, 1990.

[31] V. Bush, "As we may think," *Atlantic Monthly*, vol. 176, pp. 101-108, 1945.

[32] L. Cardelli, "A semantics of multiple inheritance," in *Semantics of Data Types*, G. Kahn, D.B. MacQueen, andG. Plotkin, Editor. Springer: Berlin. p. 51-67, 1984.

[33] T. Catlin, P. Bush, and N. Yankelovich, "InterNote: Extending a hypermedia framework to support annotative collaboration," in *Hypertext'89 Proceedings*. Association for Computing Machinery: New York. p. 365-378, 1989.

[34] S. Ceri, G. Gottlob, and L. Tanca, *Logic Programming and Databases*, Berlin: Springer. 1990.

[35] P.M. Churchland, *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*, Cambridge, Massachusetts: MIT Press. 1989.

[36] J.F. Claerbout, *Earth Soundings Analysis: Processing versus Inversion*, Cambridge, Massachusetts: Blackwell. 1991.

[37] J.F. Claerbout and J.A. Dellinger, *Stanford Exploration Project: SEP-CD-1: (1) Earth Soundings Analysis—Processing versus Inversion; (2) Anisotropic Seismic Wave Propagation; (3) Progress Reports SEP-72 and SEP-70*, Palo Alto, California: Stanford University Press. 1991.

[38] W.J. Clancey, "From GUIDON to NEOMYCIN and HERACLES in twenty short lessons," in *Current Issues in Expert Systems*, A.V. Lamsweerde and P. Dufour, Editor. Academic Press: London. p. 79-123, 1987.

[39] W.J. Clancey, "Viewing knowledge bases as qualitative models," *IEEE Expert*, vol. 4, no. 2 pp. 9-23, 1989.

[40] W.J. Clancey, "The knowledge level reinterpreted: modeling socio-technical systems," *International Journal of Intelligent Systems*, vol. 8, no. 1 pp. 33-49, 1992.

[41] D.C. Compton, K. White, and S. DeWine, "Techno-sense: making sense of computer-mediated communication systems," *Journal of Business Communication*, vol. 28, no. 1 pp. 23-43, 1991.

[42] G.P. Copeland and S.N. Khoshafian, *Identity and versions for complex objects. MCC Technical Report DB-138-86*, Austin, Texas: MCC. 1986.

[43] D. Crane, *Invisible Colleges: Diffusion of Knowledge in Scientific Communities*, Chicago: University of Chicago Press. 1972.

[44] R. Davis, "TEIRESIAS: Applications of meta-level knowledge," in *Knowledge-Based Systems in Artificial Intelligence*, R. Davis and D.B. Lenat, Editor. McGraw-Hill: New York. p. 229-461, 1982.

[45] P. Devanbu, P.G. Selfridge, B.W. Ballard, and R.J. Brachman, "A knowledge-based software information system," in *IJCAI'89: Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*. Morgan Kaufmann: San Mateo. p. 110-115, 1989.

[46] R. Dionne, E. Mays, and F.J. Oles, "A non-well-founded approach to terminological cycles," in *Proceedings of KR'92: Third International Conference on Principles of Knowledge Representation and Reasoning*. Morgan Kauffman: San Mateo, California. p. 761-766, 1992.

[47] B. Dixon, "Science and the information society," *Scholarly Publishing*, vol. 20, no. 1 pp. 3-12, 1988.

[48] F.M. Donini, M. Lenzerini, D. Nardi, and W. Nutt. Tractable concept languages. in *IJCAI'91: Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*. 1991. Sydney, Australia: Morgan Kaufmann.

[49] J. Doyle and R.S. Patil, "Two theses of knowledge representation: language restrictions, taxonomic classification, and the utility of representation services," *Artificial Intelligence*, vol. 48, no. 3 pp. 261-297, 1991.

[50] J.J.V. Dujin, *The Long Wave in Economic Life*, London: George Allen & Unwin. 1983.

[51] G. Ellis and R. Levinson, ed. *Proceedings of the First International Workshop on PEIRCE: A Conceptual Graphs Workbench.* ftp.cs.uq.oz.au /pub /TECHREPORTS /department /TR0241.ps.Z: University of Queensland, Australia. 1992.

[52] J.L. Esposito, *Evolutionary Metaphysics: The Development of Peirce's Theory of Categories*, Athens, Ohio: Ohio University Press. 1980.

[53] D.W. Etherington, A. Borgida, R.J. Brachman, and H. Kautz, "Vivid knowledge and tractable reasoning: preliminary report," in *IJCAI'89: Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*. Morgan Kaufmann: San Mateo. p. 1146-1152, 1989.

[54] D.T. Freedman, "The false start of the electrical journal," in *ASIS'87: Proceedings 50th Annual Meeting American Society Information Science*. Learned Information: Medford, New Jersey. p. 79-82, 1987.

[55] H.G. Gadamer, *Wahrheit und Methode*, Tübingen: Mohr. 1972.

[56] B.R. Gaines, "Perspectives on fifth generation computing," *Oxford Surveys in Information Technology*, vol. 1, pp. 1-53, 1984.

[57] B.R. Gaines, "Sixth generation computing: a conspectus of the Japanese proposals," *ACM SIGART Newsletter*, vol. 95, pp. 39-44, 1986.

[58] B.R. Gaines, "Positive feedback processes underlying the formation of expertise," *IEEE Transactions on Systems, Man & Cybernetics*, vol. SMC-18, no. 6 pp. 1016-1020, 1988.

[59] B.R. Gaines, "An ounce of knowledge is worth a ton of data: quantitative studies of the trade-off between expertise and data based on statistically well-founded

empirical induction," in *Proceedings of the Sixth International Workshop on Machine Learning*. Morgan Kaufmann: San Mateo, California. p. 156-159, 1989.

[60] B.R. Gaines, "Social and cognitive processes in knowledge acquisition," *Knowledge Acquisition*, vol. 1, no. 1 pp. 251-280, 1989.

[61] B.R. Gaines, "Intelligent systems as a stage in the evolution of information technology," in *Intelligent Systems: State of the Art and Future Directions*, Z.W. Ras and M. Zemankova, Editor. Ellis Horwood: Chichester, UK. 1990.

[62] B.R. Gaines, "Knowledge support systems," *Knowledge-Based Systems*, vol. 3, no. 3 pp. 192-203, 1990.

[63] B.R. Gaines, "Between neuron, culture and logic: explicating the cognitive nexus," *ICO: Intelligence Artificielle et Sciences Cognitives au Québec*, vol. 3, no. 2 pp. 47-61, 1991.

[64] B.R. Gaines, "Empirical investigations of knowledge representation servers: Design issues and applications experience with KRS," *ACM SIGART Bulletin*, vol. 2, no. 3 pp. 45-56, 1991.

[65] B.R. Gaines, "Knowledge representation servers: a generic technology for knowledge acquisition and knowledge-based systems," in *Knowledge Acquisition for Knowledge-Based Systems*, H. Motoda, R. Mizoguchi, J. Boose, andB. Gaines, Editor. Ohmsha: Tokyo. p. 413-430, 1991.

[66] B.R. Gaines, "Modeling and forecasting the information sciences," *Information Sciences*, vol. 57-58, pp. 3-22, 1991.

[67] B.R. Gaines, "The tradeoff between knowledge and data in data acquisition," in *Knowledge Discovery in Databases*, G. Piatetsky-Shapiro and W. Frawley, Editor. MIT Press: Cambridge, Massachusetts. p. 491-505, 1991.

[68] B.R. Gaines, "Manufacturing in the knowledge economy," in *Proceedings of ICOOMS'92: International Conference on Object-Oriented Manufacturing*. University of Calgary: Calgary. p. 19-36, 1992.

[69] B.R. Gaines, "Social dynamics of scholarly publication," in *AAAI-92 Workshop on Communicating Scientific and Technical Knowledge*. AAAI: Menlo Park, California. p. 1-8, 1992.

[70] B.R. Gaines, "An agenda for digital journals: the socio-technical infrastructure of knowledge dissemination," *Journal of Organizational Computing*, vol. , pp. to appear, 1993.

[71] B.R. Gaines, "A class library implementation of a principled open architecture knowledge representation server with plug-in data types," in *IJCAI'93: Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*. Morgan Kaufmann: San Mateo, California. p. to appear, 1993.

[72] B.R. Gaines, "Modeling expertise in individuals and organizations," *International Journal of Expert Systems Research and Applications*, vol. , pp. to appear, 1993.

[73] B.R. Gaines and M.L.G. Shaw, "A learning model for forecasting the future of information technology," *Future Computing Systems*, vol. 1, no. 1 pp. 31-69, 1986.

[74] B.R. Gaines and M.L.G. Shaw, "Documents as expert systems," in *Research and Development in Expert Systems IX. Proceedings of British Computer Society Expert Systems Conference*, M.A. Bramer and R.W. Milne, Editor. Cambridge University Press: Cambridge, UK. p. 331-349, 1992.

[75] B.R. Gaines and M.L.G. Shaw, "Integrated knowledge acquisition architectures," *Journal for Intelligent Information Systems*, vol. 1, no. 1 pp. 9-34, 1992.

[76] B.R. Gaines and M.L.G. Shaw, "Basing knowledge acquisition tools in personal construct psychology," *Knowledge Engineering Review*, vol. 8, no. 1 pp. to appear, 1993.

[77] H. Gallaire, J. Minker, and J.M. Nicolas, "Logic and databases: a deductive approach," *ACM Computing Surveys*, vol. 16, no. 2 pp. 153-185, 1984.

[78] M.R. Genesereth and R.E. Fikes, Knowledge Interchange Format, Version 3.0 Reference Manual. Computer Science Department, Stanford University. 1992.

[79] M.J.C. Gordon, "HOL: A proof generating system for higher-order logic," in *VLSI Specification, Verification and Synthesis*, G. Birtwistle and P.A. Subrahmanyam, Editor. Kluwer: Dordrecht. p. 73-128, 1988.

[80] S.J. Gould, *Wonderful Life: The Burgess Shale and the Nature of History*, New York: Norton. 1989.

[81] T.R. Gruber, "A translation approach to portable ontology specifications," *Knowledge Acquisition*, vol. , pp. to appear, 1993.

[82] J.Y. Halpern and M.Y. Vardi, "Model checking vs. theorem proving: a manifesto," in *Artificial Intelligence and the Mathematical Theory of Computation: papers in Honor of John McCarthy*, V. Lifschitz, Editor. Academic Press: New York. p. 151-176, 1991.

[83] J. Heinsohn, D. Kudenko, B. Nebel, and H.J. Profitlich, "An empirical analysis of terminological representation systems," in *Proceedings of the Tenth National Conference of the American Association for Artificial Intelligence*. MIT Press: San Jose. p. 767-773, 1992.

[84] T. Hendley, *CD-ROM and Optical Publishing Systems*, Westport, Connecticut: Meckler. 1987.

[85] O. Herzog and C.R. Rolllinger, ed. *Text Understanding in LILOG*. Vol. LNCS-546. Springer: Berlin. 1991.

[86] S.R. Hiltz, *Online Communities: A Case Study of the Office of the Future*, Norwood, New Jersey: Ablex. 1984.

[87] R. Hull and R. King, "Semantic database modeling: survey, applications and research issues," *ACM Computing Surveys*, vol. 19, no. 3 pp. 201-260, 1987.

[88] E. Jantsch, *The Self-Organizing Universe*, Oxford: Pergamon Press. 1980.

[89] R. Kowalski, *Logic for Problem Solving*, North Holland: New York. 1979.

[90] S. Lambert and S. Ropiequet, ed. *CD ROM: The New Papyrus*. Microsoft Press: Redmond, Washington. 1986.

[91] E.W. Lancaster, *Toward Paperless Information Systems*, New York: Academic Press. 1978.

[92] D. Leishman, "Analogy as constrained partial correspondence over conceptual graphs," in *Proceedings of KR'89: First International Conference on Principles of Knowledge Representation and Reasoning*. Morgan Kauffman: San Mateo, California. p. 223-234, 1989.

[93] M. Lenzerini and A. Schaerf, "Concept languages as query languages," in *AAAI'91: Proceedings of the Ninth National Conference on Artificial Intelligence*. AAAI Press/MIT Press: Menlo Park, California. p. 471-476, 1991.

[94] R. Levinson and G. Ellis, "Multilevel hierarchical retrieval," *Knowledge-Based Systems*, vol. 5, no. 3 pp. 233-244, 1992.

[95] H.T. Liaw and C.S. Ling, "On the OBDD-representation of general boolean functions," *IEEE Transactions on Computers*, vol. 41, no. 6 pp. 661-664, 1992.

[96] R.M. MacGregor, "A deductive pattern matcher," in *Proceedings of AAAI88*. Morgan Kaufmann: San Mateo, California. p. 403-408, 1988.

[97] P. Maes, ed. *Designing Autonomous Agents*. MIT Press: Cambridge, Massachusetts. 1990.

[98] C. Malamud, *Exploring the Internet*, Englewood Cliffs, New Jersey: Prentice-Hall. 1992.

[99] C. Marchetti, "Society as a learning system: discovery, invention and innovation cycles revisited," *Technological Forecasting & Social Change*, vol. 18, pp. 267-282, 1980.

[100] H.R. Maturana, "The organization of the living: A theory of the living organization," *International Journal Man-Machine Studies*, vol. 7, no. 3 pp. 313-332, 1975.

[101] W.W. McCune, *OTTER 2.0 Users Guide*, Argonne, Illinois: Argonne National Laboratory. 1990.

[102] G. Mensch, *Stalemate in Technology: Innovations Overcome the Depression*, Cambridge, Massachusetts: Ballinger. 1975.

[103] R.K. Merton, "Priorities in scientific discovery," *American Sosciological Review*, vol. 22, no. 6 pp. 635-659, 1957.

[104] M.D. Mey, *The Cognitive Paradigm*, Dordrecht, Holland: Reidel. 1982.

[105] J. Minker, ed. *Foundations of Deductive Databases and Logic Programming*. Morgan-Kaufman: San Mateo, California. 1988.

[106] M. Minsky, "A framework for representing knowledge," in *The Psychology of Computer Vision*, P.H. Winston, Editor. Simon and Schuster: New York. p. 211-280, 1975.

[107] J. Monod, *Chance and Necessity*, London: Collins. 1972.

[108] T. Moto-oka, ed. *Fifth Generation Computer Systems*. North-Holland: Amsterdam. 1982.

[109] B. Nebel, "Computational complexity of terminological reasoning in BACK," *Artificial Intelligence*, vol. 34, pp. 371-383, 1988.

[110] B. Nebel, *Reasoning and Revision in Hybrid Representation Systems*, Vol. LNCS-422. Berlin: Springer. 1990.

[111] B. Nebel and G. Smolka, "Representation and reasoning with attributive descriptions," in *Sorts and Types in Artificial Intelligence*, K.H. Bläsius, U. Hedstück, andC.-R. Rollinger, Editor. Springer: Berlin. p. 112-139, 1990.

[112] R. Neches, R. Fikes, T. Finin, T. Gruber, R. Patil, T. Senator, and W.R. Swartout, "Enabling technology for knowledge sharing," *AI Magazine*, vol. 12, no. 3 pp. 36-56, 1991.

[113] J.T. Nosek and I. Roth, "A comparison of formal knowledge representations as communication tools: predicate logic vs semantic network," *International Journal of Man-Machine Studies*, vol. 33, no. 227-239, 1990.

[114] P. Oakeshott, "The 'BLEND' experiment in electronic publishing," *Scholarly Publishing*, vol. 17, no. 1 pp. 25-36, 1985.

[115] A. Ohori, "Semantics of types for database objects," *Theoretical Computer Science*, vol. 76, no. 1 pp. 53-91, 1990.

[116] M. Palaniappan and G. Fitzmaurice, "InternetExpress: an inter-desktop multimedia data-transfer service," *IEEE Computer*, vol. 24, no. 10 pp. 58-67, 1991.

[117] J. Palca, "New journal will publish without paper," *Science*, vol. 253, no. 5027 pp. 1480, 1991.

[118] P.F. Patel-Schneider, Description Logic Specification from the KRSS Effort. The DARPA Knowledge Sharing Initiative Knowledge Representation Working Group. 1992.

[119] L.C. Paulson, "Isabelle: The next 700 theorem provers," in *Logic and Computer Science*, P. Odifreddi, Editor. Academic Press: New York. p. 361-386, 1990.

[120] C.S. Peirce, *Reasoning and the Logic of Things*, Cambridge, Massachusetts: Harvard University Press. 1898.

[121] R.L. Philips, "MediaView: a general multimedia digital publication system," *Communications of the ACM*, vol. 34, no. 7 pp. 75-83, 1991.

[122] W.T. Powers, *Behavior: The Control of Perception*, New York: Aldine. 1973.

[123] I. Prigogine, *Order out of Chaos*, Toronto: Bantam. 1984.

[124] J.S. Quarterman, *The Matrix: Computer Networks and Conferencing Systems Worldwide*, Reading, Massachusetts: Digital Press. 1990.

[125] M.R. Quillian, "Semantic memory," in *Semantic Information Processing*, M. Minsky, Editor. MIT Press: Cambridge, Massachusetts. p. 216-270, 1968.

[126] R. Ramakrishnan, D. Srivastava, and S. Sudarshan, "CORAL—Control, relations and logic," in *Proceedings of the 18th Very Large Database Conference*. Morgan Kaufmann: San Mateo. 1992.

[127] N. Rescher, *Peirce's Philosophy of Science*, Notre Dame, Indiana: University of Notre Dame Press. 1978.

[128] D.D. Roberts, *The Existential Graphs of Charles S. Peirce*, The Hague: Mouton. 1973.

[129] H. Rosenbaum and G.B. Newby, "An emerging form of human communication: computer networking," in *ASIS'90: Proceedings 53rd Annual Meeting American Society Information Science*. Learned Information: Medford, New Jersey. p. 300-325, 1990.

[130] J. Rosenberg, M. Sherman, A. Marks, and J. Akkerhuis, *Multi-media Document Translation: ODA and the EXPRES Project*, New York: Springer. 1991.

[131] J. Savoy, "The electronic book Ebook3," *International Journal Man-Machine Studies*, vol. 30, pp. 505-523, 1989.

[132] M. Schmidt-Schauß, "Subsumption in KL-ONE is undecidable," in *Proceedings of KR'89: First International Conference on Principles of Knowledge Representation and Reasoning*. Morgan Kauffman: San Mateo, California. p. 421-431, 1989.

[133] J.W. Senders, "An on-line scientific journal," *The Information Scientist*, vol. 11, pp. 3-9, 1977.

[134] N. Sharif, "An introduction to modeling," in *Systems Modeling for Decision Making*, N. Sharif and P. Adulbhan, Editor. Asian Institute of Technology (distributed by Pergamon Press): Bangkok. p. 1-21, 1978.

[135] M.L.G. Shaw and B.R. Gaines, "A methodology for recognizing conflict, correspondence, consensus and contrast in a knowledge acquisition system," *Knowledge Acquisition*, vol. 1, no. 4 pp. 341-363, 1989.

[136] T. Sheridan, J. Senders, N. Moray, J. Stoklosa, J. Guillaume, and D. Makepeace, Experimentation with a multi-disciplinary teleconference and electronic journal on mental workload. Massachusetts Institute of Technology. pp. 1-329, 1981.

[137] E.H. Shortliffe, *Computer-Based Medical Consultations: MYCIN*, New York: Elsevier. 1976.

[138] G. Smolka, Residuation and guarded rules for constraint logic programming. Paris Research Laboratory, Digital Equipment Corporation. pp. 1-14, 1991.

[139] G. Smolka and H. Aït-Kaci, "Inheritance hierarchies: semantics and unification," *Journal Symbolic Computation*, vol. 7, pp. 343-370, 1989.

[140] J.F. Sowa, *Conceptual Structures: Information Processing in Mind and Machine*, Reading, Massachusetts: Adison-Wesley. 1984.

[141] STA, *Promotion of R&D on Electronics and Information Systems That May Complement or Substitute for Human Intelligence*, Tokyo: Science and Technology Agency. 1985.

[142] P. Stoneman, *The Economic Analysis of Technological Change*, Oxford: Oxford University Press. 1983.

[143] R. Studer, ed. *Natural Language and Logic*. Vol. LNCS-459. Springer: Berlin. 1990.

[144] G. Taubes, "Publication by electronic mail takes physics by storm," *Science*, vol. 259, pp. 1246-1248, 1993.

[145] A. Touraine, *La Sociologie de l'Action*, Paris: Éditions du Seiul. 1965.

[146] M. Walter, "Scholarly journals: case studies of SGML for electronic publishing," *Seybold Report on Publishing Systems*, vol. 21, no. 18 pp. 3-20, 1992.

[147] C. Walther, "A mechanical solution of Schubert's steamroller by many-sorted resolution," in *Proceedings of AAAI-84*. Morgan Kauffman: San Mateo, California. p. 330-334, 1984.

[148] H.G. Wells, *World Brain*, New York: Doubleday. 1938.

[149] S.A. Weyer and A.H. Borning, "A prototype electronic encyclopedia," *ACM Transactions Office Information Systems*, vol. 3, pp. 63-88, 1985.

[150] W.A. Woods, "What's in a link: Foundations for semantic networks," in *Representation and Understanding: Studies in Cognitive Science*, D.G. Bobrow and A.M. Collins, Editor. Academic Press: New York. p. 35-82, 1975.

[151] N. Yankelovich, N. Meyrowitz, and A. van Dam, "Reading and writing the electronic book," *Computer*, vol. 18, no. 10 pp. 15-30, 1985.

[152] E.N. Zalta, *Intensional Logic and the Metaphysics of Intentionality*, Cambridge, Massachusetts: MIT Press. 1988.