CONTEXT: A REAL PROBLEM FOR LARGE AND SHAREABLE KNOWLEDGE BASES

Bob Jansen
CSIRO Division of Information Technology,
Locked Bag 17, North Ryde, NSW 2113, AUSTRALIA
Ph: +612 - 325 3100 Fax: +612 - 325 3101
email: jansen@syd.dit.csiro.au

ABSTRACT

Existing large and shareable repositories of knowledge, e.g., libraries, are dependent on the notion of context for their structure and efficient use. Books are classified and catalogued into subject areas, authors, keyword areas, etc. Each of these arrangements facilitate the appropriate access to, and hence use of, relevant parts of the repository. This paper argues that the notion of context, although complex and little understood in the AI community, must be one of the intrinsic parameters of any large and shareable knowledge-base system. It introduces some details of context as found in the literature to highlight its complexity and how little understood the notion of context is. Electronic document systems is suggested as a most promising area for researching into context applicable within large and shareable knowledge-base systems.

KEYWORDS

context, very large knowledge-base systems, shareable knowledge-base systems, Platonic principle of knowledge, epistemology, electronic document systems

1. INTRODUCTION

Knowledge-base systems (KBS) developed to date can generally be characterised by relatively small, well-bounded domains in comparison with human expertise. Some large domains have been addressed by KBS technology (e.g. MYCIN for bacterial infections (Buchanan and Shortliffe, 1984); Xcon for configuring computer systems, in excess of 6200 rules (Soloway *et al*, 1987); PEIRS for the interpretation of laboratory reports (Compton *et al*, 1992), in excess of 900 rules and growing at a rate of five rules per day) but even of those, the majority could be classified as research systems.

Extending KBS into real-world size, tackling real-world problems in a real-world fashion, requires some fundamental changes to the theories underlying the technology. Real-world applications are characterised by large sets of diverse knowledge; multiple, and possibly parallel, inferencing strategies; multiple representations, or at the very least the ability to translate required knowledge from one representation to another to suit the problem at hand; context-sensitive knowledge representations; context-sensitive knowledge processing (i.e., processing relevant knowledge in an appropriate fashion); the ability to deal intelligently with fuzzy or incomplete knowledge; discarding the Platonic principle for knowledge; etc. In general, the theories and technologies underlying KBS are deficient in catering for these types of characteristics. However, much research is underway to address some of these issues.

One area almost totally devoid of systematic research is that of context. Human intelligence will always be superior to machine intelligence whilst the latter has no notion or understanding of context, the relevance and applicability of knowledge to a

†The work reported in this paper has been supported in-part by grants from DITARD, Australia, under the FAIR/ISTAC program, and the European Software Factory (ESF) as contributions to a collaborative project involving CSIRO, Australia, and INRIA and CNRS, France; and CNRS, France, as a contribution to a collaborative project involving CSIRO and the CNRS Institut des Textes et Manuscrits Modernes (ITEM), France.

situation. Experience with the long-term maintenance of the Garvan ES1 expert system (Horn, 1990; Horn *et al*, 1985) has shown that experts process knowledge within specific contexts and taking such knowledge out of context by applying it in a KBS was the underlying *raison d'être* for the majority of maintenance (Compton *et al*, 1988; Compton and Jansen, 1988). Studies done as part of the Garvan ES1 project (e.g., Compton and Jansen, 1988) showed that for most KBS, the domain knowledge was represented as if it was all part of a single context, rather than being acquired from multiple contexts. This assumption, the single domain representation, leads to the requirement of the consistency of the knowledge. A KBS is said to be unable to produce an appropriate interpretation of the data if its knowledge is inconsistent.

Human expertise, however, is characterised by being inconsistent. After all, do most of us not remember phlogiston and surrounding theories from the early days of physics and chemistry. These theories are inconsistent with our beliefs today, yet we manage to be intelligent and produce intelligent interpretations of situations. Human intelligence is context-sensitive, it is represented in context-sensitive fashions and applied in context-sensitive ways. Researchers like Ogden, Richards and Russell (see below) claim that context is of such importance in epistemology that there can be no inference or knowledge without it. It appears to be a basic building block of knowledge and its processing.

2. CONTEXT

Context is, by and large, an ignored problem in knowledge-based systems. Guha (1991) describes three aspects of the homogeneity of knowledge-based systems, namely: while many KBS structure the domain they rarely structure the knowledge about the domain; the KBS contains a single model of the world and to allow meaningful inferences, this model must be kept logically consistent; and the single model approach implies that the model should be kept independent of particular problems.

Yet, any practitioner in the KBS domain realises very quickly that human reasoning is not so circumscribed. Human experts reason using knowledge of the domain as well as knowledge within the domain; they use multiple models as part of their problem solving process; and their knowledge is not guaranteed to be logically consistent. Compton and Jansen (1988) describes results of a research project into the long-term maintenance of the Garvan ES1 knowledge-based system and how experts provide knowledge within well-defined chunks and contexts. More importantly, in context, these chunk of knowledge were consistent, but when taken out of context problems arise. This explains the problem associated with most knowledge-base systems to date, they require continual maintenance because they meet situations never before recognised and apply knowledge out-of-context to the situation. The knowledge is able to be applied out-of-context because all contextual information is lost by the process of massaging it into the sole model supported by the knowledge-based system.

The conclusion that can be drawn from this is that the Platonic principle, namely, that knowledge should be represented in isolation from its processing function, is erroneous and compliance with it eventually leads to problems[1]. Context has a major part to play in knowledge representation and processing, in fact as described below, it can be argued that without context there is no knowledge or knowledge processing. (Ogden and Richards, 1946; Russell, 1992)

The problem with context is that there is as yet no universally accepted definition, nor any comprehensive understanding of how to represent it or how it affects knowledge processing. The next section presents some descriptions of context as found in the literature, but only serves to highlight this complexity and lack of comprehension.

2.1. Nature of Context

The following presents some insight into the nature of context and implications of its use, as found in the literature.

2.1.1. The Meaning of Meaning

In *The Meaning of Meaning*, Ogden and Richards (1946) present a detailed account of the nature of context based primarily in linguistics, but also addressing other disciplines.

A context consists of a set of things with specific relationships among them. More importantly, in order to be defined as a

[1]Problems will only arise if the knowledge within the knowledge-base system requires change for any reason.

[2]This section is based on a more detailed analysis in Jansen (1994)

context, the set of things and relationships must have multiple similar instances. For example, if we perceived a match being struck in America and a small but finite time thereafter, a flame in China, then this combined instance would not be part of a context unless there was a definite linking relation between the two sub-events (ie. the match scrape and the subsequent flame) and there were other events of a similar nature (ie. two events occurring a significant distance apart but causally related, or one event in America causally related to another event in China). Given the little likelihood of causal relationship between such distant events, there seems little point of calling the condition a context or in fact attributing it to a context. One could however include it as an instance in a context of events which occur only once or are *unusual*, *unexpected*, or just plain *weird*. In this case, each set of objects in the context are a situation and the concept of weirdness with a relationship distinguishing each situation as *weird*.

The discussion then proceeds to identify various types of context including psychological, wherein the events are mental events; literary, wherein the context is a group of words, incidents or ideas that surrounds whatever is said to have to context; determinative, wherein they introduce the notion of a complex context comprised of simpler contexts; and the structure of a context as a set of static and dynamic components.

Ogden and Richards provide a number of interesting effects of context, effects that have a major impact for large and shareable knowledge-base systems.

They appear, maybe in deference to the mystical quality they have imputed as part of attempted definitions, to place context in the position of prime importance when it comes to the use of knowledge. They claim that no inductive generalisation, no knowledge or probable opinion as to what is not immediately given is possible without context. Their argument is based on the premise that these processes are in themselves, members of certain psychological contexts. Russell, see below, goes one further and argues that context forms the rationale for proper names and the very nature of knowledge itself. They hypothesise 'peculiar' relationships between psychological and external contexts as part of the processing of knowledge. If such linkage does not form, they claim 'we have been mistaken' (Page 57)

Ogden and Richards make a final association between situations and context by the statement that the relations and events that actually comprise a context can only be revealed by experience, and thus no amount of planning or modelling will elicit the complete structure. In terms of the processing of knowledge within contexts, they claim that the past history of an organism will direct its response to a situation. Some aspects of the organism's history will be more relevant and will guide its overall behaviour. This claim supports the notion that knowledge processing is context sensitive. More importantly, the recurrence of a part of a context will cause an organism to behave as it did before, but they do not explore what aspect of context causes this response.

2.1.2. Human Knowledge: Its Scope and Limitations

In *Human Knowledge: Its Scope and Limitations*, Russell (1992) attempts to bringing together a lifetime's work in striving to understand 'knowledge'. Unfortunately, much as Ogden and Richards, the language used in describing Russell's conclusions, leaves the reader more confused than before. He draws heavily on a presupposed correlation between psychology and physics in an attempt to provide a unified view.

Russell takes the stance of relating context to perceptions in space-time. He postulates a relation called *compresence* which holds between any 'qualities' when one person experiences them simultaneously[3] and then goes on to describe complexes of compresence and their existence based on the nature of their components. The argument then proceeds to establish the dualism of complexes of compresence and what would normally be called context. One possible effect of this dualism is that if true then it might be possible to use the laws of physics and mathematics to describe context. Furthermore, it is claimed that contexts exist not of themselves but only by nature of their

[3] "If I see something and at the same time hear something else, my visual and auditory experiences have a relation which I call 'compresence'. If at the same moment I am remembering something that happened yesterday and anticipating with dread a forthcoming visit to the dentist, my remembering and anticipating are also 'compresent' with my seeing and hearing. We can go on to form the whole group of my present experiences and of everything compresent with them." (Russell, 1992, Page 312)

components and there are reasons for supposing that at every place in physical space-time there is at every moment a multiplicity of occurrences, just as there is in mind, (i.e., multiple contexts.)

Two important results of Russell's work that affect knowledge-base systems are that one can never determine the completeness of a context; context obeys a Heisenberg Uncertainty Principle, i.e., context is context-sensitive (Page 323), and that knowledge acquisition and the concept of proper names are inexorably intertwined such that knowledge would cease if our knowledge was complete (Page 325). These claims, if true, state that a context is an open-ended, unbounded, thing, and as such has profound implications for its representation and use within knowledge-base systems.

2.1.3. What is Context

In *What is Context*, Shanon (1983) argues against three possible notions of context: linguistic; internal representation; and external state of affairs in the world. Each of these notions, whilst solving some particular problem associated with context, provides many more problems. For example, by pushing context out of the internal state and into the external world, we solve any contextual problems associated with internal states but do not address any problems associated with the external world. Shanon argues that context should be regarded as an interface element between the internal and external worlds, ie. between the psychological and linguistic notions as seen in Ogden and Richards and Russell. The argument in favour of this notion is based on proposed frameworks that have argued against the notion of mind/world separation. Shanon argues that these frameworks have introduced terminology which intertwined the cognitive agent and the environment in which it acts. One of these terms is 'situation' as coined in Barwise and Perry (1983).

2.1.4. Contexts: A Formalisation and Some Applications

In *Contexts: A Formalisation and Some Applications*, Guha (1991) describes the notion of context as a solution to the problem of writing a non-naïve theory. For example, a theory of commercial transactions includes not just the axioms describing buying and selling but also information about the assumptions made by the theory, when these assumptions are reasonable, when this theory is applicable, etc. Contexts are defined as rich objects within the domain which can not be completely described. Guha's experience is from within the CYC project (e.g., Guha and Lenat, 1990; Lenat and Guha, 1988). CYC began its life as an attempt at capturing all the common-sense knowledge as found in an encyclopaedia, but has since compressed its focus onto a smaller subset of an encyclopaedia. It is interesting to note the emergence of context into the later papers about CYC as if it was realised at quite a late stage that there was a contextual component and that it was important. Guha (1991) describes several operations based on a context-calculus, whereby contexts are manipulated as objects in their own right. This follows naturally from their initial definition of context as an object in the domain. The structure of a context as defined is akin to the notion of a world or Herbrand Universe (Herbrand, 1967).

2.1.5. Context: Identifying its Elements from the Communication Point of View

In *Context: Identifying its Elements from the Communication Point of View*, Mittal and Paris (1993) attempt to characterise various aspects of context and to bring them together into a single global picture. They identify the following five aspects of context, each having some effect on communication: the problem solving situation; the participants involved; the mode of interaction in which the communication is occurring; the discourse taking pace; and the external world. Problem solving is broken down into the tasks being achieved; methods for achieving the tasks; the results expected; the content of the knowledge base and its representation; the system; and whether the system is single- or multi-agent. The participants is broken down into the expertise of the user in the problem domain; the specific beliefs a user may have (or misconceptions); and the goals and plans of the user. The mode of interaction is broken down into single vs multiple participants; the medium of interaction; and the use by the system of the user's feedback. The discourse is broken down into communicative goals (i.e. what a speaker intends with each utterance); how the information is communicated (i.e. rhetorical devices used); and ideational information or what is actually communicated to a hearer. The external world is broken down into the genre of the text and the

environment in which the communication takes place. Despite a relatively detailed decomposition, this paper highlights the complex nature of context. Some of the detailed characteristics appear to be repeated, indicating that there are still many things to be learned about context.

2.1.6. Knowledge in Context: A Strategy for Expert Systems Maintenance

In *Knowledge in Context: A Strategy for Expert Systems Maintenance*, Compton and Jansen (1988) introduce Ripple-down Rules as one of the first exemplars of context-sensitive knowledge representation and processing. Ripple-down rules aims to represent, and apply, the rules acquired from an expert in the context in which they are given. Similar to case-based reasoning, ripple-down rules represents the context of the knowledge acquisition within the notion of a situation. In ripple-down rules each situation is represented by one branch of the inference tree.

The *raison d'être* for ripple-down rules was the realisation that knowledge engineering, obtaining knowledge from experts and incorporating it into expert systems, is difficult and time consuming. It was hypothesised that these difficulties arise because experts never report on how they reach a decision, rather they justify why the decision is correct. These justifications vary markedly with the context in which they are required, but in context they are accurate and adequate. The difficulties in knowledge engineering arise from taking the justification out of context. Thus knowledge engineering may be obviated, particularly in the long term maintenance of expert systems, if the rules experts provide are used in the context in which they are given.

The advantage of ripple-down rules over case-based reasoning is that context is explicitly represented within the decision tree and knowledge chunks, i.e., production rules, are only activated in the appropriate context for the current problem. The implementation of ripple-down rules, as a pointer-based production rule structure, is not exceptional in its own right, but does provide, at the same time, the possibility of detailed views of particular contexts and their interaction, as well as more broader visualisation of the decision making process.

2.2. Conclusions on Context

The above discussion highlights the various ways of looking at and processing context and the lack of any consensus between disciplines. The one aspect that seems to be common to all is the notion of a situation and this seems the most likely place where any sign of consensus is likely to emerge.

The types of contexts discussed involve several areas of philosophy, linguistics and logics. The external and internal, or mental and physical, contexts are related to the mind-body problem so frequently discussed in philosophy. Literary and linguistic contexts are associated with the written word and communication and the primary domain of linguistic researchers but increasingly addressed by the electronic publishing community.

As for the processing of context, Guha (1991) presented the notion of a context-calculus defining several primitive operations over contexts as whole objects, but all of the above papers recognised that knowledge processing is context sensitive. In fact, Ogden and Richards (1946) and Russell (1992) characterised context with some religious aspects wherein it became a central notion to the nature of knowledge itself.

3. CONTEXT AND LARGE KBS

Jansen and Compton (1993) show that KBS increase in size for two main reasons; knowledge refinement and knowledge enhancement. In knowledge refinement, the KBS grows more expert at its task without expanding the scope of its domain. In knowledge enhancement, the scope of the KBS increases and it is applicable to more tasks.

Knowledge refinement is the process of improving the resolution of the KBS, i.e. improving its capability of resolving finer and finer situations within the same domain. In this scenario, context must play a crucial role. As the knowledge chunks become more specialised they become less applicable. Thus the KBS must be capable of determining the relevance of a chunk of knowledge to the problem and this is essentially context. Thus the knowledge and hence its processing becomes more and more context sensitive. Ripple-down rules was developed as a direct result of this process within the Garvan ES1 system.

Knowledge enhancement is the process of adding more and more knowledge to the KBS, or enabling knowledge re-use, such that the KBS can be applied to other domains. The problem here is that the KBS must not apply irrelevant knowledge in any situation, and this also is essentially a problem of context. Thus the structure of the KBS becomes context sensitive.

4. CONTEXT AND SHAREABLE KBS

There are essentially two different scenarios for shareable KBS: multi-user and multi-domain. Multi-user enables many users to use the same KBS, either by real-time sharing (i.e., on-line transaction processing) or using an identical copy of the KBS. Multi-domain enables the KBS to be used with different domains.

Multi-user scenarios have been extensively developed in the information systems community, with the development of on-line transaction processing and database systems. The basic technology is essentially understood and mature. Context is implemented by the use of transactions and by physically and logically subsetting the storage medium (i.e., areas and subschemas/views).

Multi-domain scenarios are more difficult for non-trivial problems mainly because of the amount of knowledge potentially involved. Each domain could consist of domain-dependent and domain-independent knowledge. The domain-independent knowledge could be represented once and applied to each domain as appropriate. Context would be implemented as domains, thus logically, if not physically, subsetting the knowledge.

The real problem in this scenario is the structuring and representation of the knowledge. Different domains will require different aspects to their knowledge. For example, what might be seen a common sense in one domain could be crucial in another domain. This problem has already been realised by the CYC project (Lenat and Guha 1988; Guha and Lenat, 1990) wherein context has lately featured larger in their published papers. The complexity of this problem is indicated by a down-grading of CYC's goal, namely from that of representing all the knowledge in an encyclopaedia to that of a small part of an encyclopaedia[4]. Context will feature a major part in this process but dependence on semantics will make this an unsolved problem for some time yet.

Hence, whilst all KBS's are shareable in that, for example, we can apply MYCIN to the problem of configuring computers and it will sensibly conclude that it does not know anything, the notion of shareability conveys notions of relevance and applicability which are aspects of context.

5. CONCLUSION

The above discussion has highlighted a number of issues with regard to large and shareable knowledge bases and the representation and use of context. One easy conclusion that can be drawn is that context is an indispensable part of any KBS, yet as has been shown in the history of KBS technology, we can build very useful KBS without any notion of context. It is the expansion of a KBS into a large and shareable KBS that will inevitably bring problems if the notion of context is not included as part of its intrinsic parameters.

We will require an extensive international effort into the nature, representation and use of context. This problem is too big to be conquered by any one research team. It will require multi-disciplinary skills and, in practise, real applications to act as testbeds for theories.

We should take advantage of existing large and shareable knowledge repositories to study the effect of context, as well as apply context within all new developments. It is only through studying existing scenarios to produce theories regarding context and the application of those theories that we will solve this problem.

One area that provides a rich training ground is that of libraries and electronic documentation. Libraries are those existing knowledge repositories that already utilise context in their structure and books, and documents provide bounded domains for the application for context. The impetus for this research will be the electronification of our paper-based knowledge repositories, a process that is gathering much momentum and can provide needed financial backing for research projects. The publishing community sees a

[4]Doug Lenat in his address to FGCS'88 stated the initial goal, namely the representation of all of the encyclopaedia. Lately (i.e., 5 years later), they have said they are focussing on volume 1 only. It is a really difficult problem.

major change in the method of publishing and disseminating information, a change made possible by the advances in multimedia technology. This awareness, coupled with the possibilities offered by electronic multimedia libraries, offers a rich picking ground for research into context.

ACKNOWLEDGMENTS

We acknowledge the help of Bob Colomb, Department of Computer Science, University of Queensland and Patrick Brézillon, CNRS LAFORIA Université Paris VI, in reviewing a draft of this paper.

REFERENCES

Barwise J & Perry J, *Situations and Attitudes*, Harvard University Press, 1983

Buchanan BG & Shortliffe EH (eds), *Rule Based Expert Systems, The MYCIN Experiments of the Stanford Heuristic Programming Project*, The Addison-Wesley Series in Artificial Intelligence, reprinted 1984

Compton P & Jansen B, *Knowledge in Context: A Strategy for Expert Systems Maintenance*, Lecture Notes in Artificial Intelligence, J. Siekmann (ed), Subseries of Lecture Notes in Computer Science 406, AI'88 Proceedings of 2nd Australian Joint Artificial Intelligence Conference, Adelaide, November 1988

Compton P, Edwards G, Srinivasan A, Malor R, Preston P, Kang B & Lazarus L, *Ripple Down Rules: Turning Knowledge Acquisition into Knowledge Maintenance*, Artificial Intelligence in Medicine, 1992, in press

Compton P, Horn K, Quinlan JR, Lazarus L & Ho K, *Maintaining an Expert System*, Proceedings of the Fourth Australian Conference on Applications of Expert Systems, Sydney, May 1988

Guha RV & Lenat DB, *Cyc: A Midterm Report*, AI Magazine, Vol. 11, No. 3, Fall 1990, pp32-59

Guha RV, *Contexts: A Formalization and some Applications*, MCC Technical Report No. ACT-CYC-423-91, 1991

Herbrand J, *Investigations in Proof Theory: The Properties of True Propositions*, in From Frege to Godel: a Source Book in Mathematical Logic 1879-1931, J van Heijenoort (ed), Harvard University Press, 1967, pp525-581

Horn K, *Garvan ES1: An Expert System for the Interpretation of Thyroid Laboratory Tests*, M.App.Sc. Thesis, University of Technology, Sydney, 1990

Horn KA, Compton P, Lazarus L & Quinlan R, *An Expert System for the Interpretation of Thyroid Assays in a Clinical Laboratory* , The Australian Computer Journal, Volume 17, No 1, February 1985

Jansen B, *Context in Context*, submitted to the Knowledge Acquisition Workshop KAW'94, Banff, Canada, January 1994

Jansen B & Compton P, The *Knowledge Dictionary: Maintaining Knowledge-base Systems*, Academic Press, 1993, in press

Lenat D & Guha RV, *The World According to CYC*, MCC Technical Report ACA-AI-300-88, 1988

Mittal VO & Paris CL, *Context: Identifying its Elements from the Communication Point of View*, in Proceedings of the IJCAI'93 Workshop on "Using Knowledge in its Context", France 1993, pp87-97

Ogden CK & Richards IA, *The Meaning of Meaning: A Study of the Influence of Language upon Thought and of the Science of Symbolism*, Routledge & Kegan Paul Ltd.(pub) London, 1946

Russell B, *Human Knowledge: Its Scope and Limitations*, Routledge (pub), 1992 edition

Shanon B, *What is Context*, Journal for the Theory of Social Behaviour: 20:2 0021-8308 82:50, 1983

Soloway E, Bachant J & Jensen K, *Assessing the Maintainability of XCON-in-Rime: Coping with the Problems of a VERY Large Rule Base*, Proceedings of AAAI 87 Sixth National Conference on Artificial Intelligence, Seattle, Washington, 1987