

---

# Epistemological Approach to the Process of Practice

Richard Dazeley · Byeong Ho Kang

Received: 30 August 2007 / Accepted: 6 August 2008

**Abstract** Systems based on symbolic knowledge have performed extremely well in processing reason, yet, remain beset with problems of brittleness in many domains. Connectionist approaches do similarly well in emulating interactive domains, however, have struggled when modelling higher brain functions. Neither of these dichotomous approaches, however, have provided many inroads into the area of human reasoning that psychology and sociology refer to as the process of practice. This paper argues that the absence of a model for the process of practise in current approaches is a significant contributor to brittleness. This paper will investigate how the process of practise relates to deeper forms of contextual representations of knowledge. While researchers and developers of knowledge based systems have often incorporated the notion of context they treat context as a static entity, neglecting many connectionists' work in learning hidden and dynamic contexts. This paper argues that the omission of these higher forms of context is one of the fundamental problems in the application and interpretation of symbolic knowledge. Finally, these ideas for modelling context will lead to the reinterpretation of situation cognition which makes a significant step towards a philosophy of knowledge that could lead to the modelling of the process of practice.

**Keywords** Dynamic context · Epistemology · Hidden context · Knowledge based systems · Knowledge representation · Philosophy of knowledge · Process of practice · Situation cognition

---

R. Dazeley (✉)

School of Information Technology and Mathematical Sciences, University of Ballarat,  
P.O. Box 663, Ballarat, VIC 3353, Australia  
e mail: r.dazeley@ballarat.edu.au

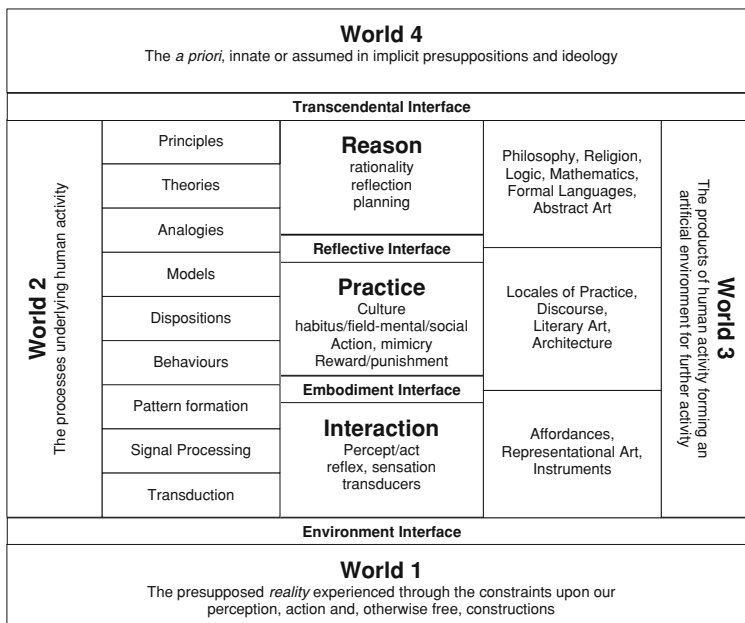
B. H. Kang

School of Computing and Information Systems, University of Tasmania, Private Bag 100,  
Hobart, Tasmania 7001, Australia  
e mail: Byeong.Kang@utas.edu.au

## Introduction

There have been a number of methods applied to the task of creating the thinking machine with varying degrees of success. For instance, artificial neural networks (ANNs) and other function fitting methods have shown great ability at being able to learn generalised solutions to problems. These systems can provide reasonable results to previously unseen situations. However, they generally require extensive and repeated training or calculation to form the required function. Alternatively, through encoding human knowledge directly, such as a knowledge based systems (KBSs), extremely accurate classification can be achieved without the need for repeated training. However, unlike ANNs, KBSs generally completely collapse when faced with a previously unseen situation. This is particularly clear when they require knowledge from outside the system's domain.

The differences between these dichotomous approaches were investigated by Gaines (2000). He looked at where within the overall framework for human activity, each group of methods have been applied with the greatest success. Figure 1 shows a reproduction of Gaines's (2000) *Levels and Worlds of Being* diagram, which attempts to capture the entire conceptual framework for human psychology, sociology, action and knowledge. The central column of this diagram presents a



**Fig. 1** Conceptual framework for human psychology, sociology, action and knowledge. The *central column* shows a three layer model of human entities. The four outer worlds set human entities within the context of Popper's (1968) worlds. The *boxes on the left of the central region* identify various models of knowledge transfer in relation to the three layers. The *boxes to the right of the central core* relate the three layers to various products of human entities. This model is a replication of the levels of worlds of being model (Gaines 2000, p. 115)

three layer model of human entities, such as people, companies or governments. Gaines (2000) claims that connectionist approaches have emulated the *interaction* level with increasing effectiveness, while digital computation has done similarly well with the top brain functions of *reasoning*.

Gaines's (2000) investigation into these different approaches found that the central area concerning the *processes of practice* is relatively untouched. This paper is interested in finding an applicable epistemological foundation that goes some way to filling this gap. This will be achieved by discussing the process of practice and how it relates to deeper forms of contextual representations of knowledge. This will include an investigation into current approaches to incorporating context. It is argued that the omission of these higher forms of context is one of the fundamental problems in the application and interpretation of symbolic knowledge. These ideas for modelling context will lead to the reinterpretation of situation cognition which makes a significant step towards a philosophy of knowledge that could lead to the modelling of the process of practice. The following sub-sections will define what is meant by the process of practice and the various types of context discussed. This will be followed by a brief walkthrough some historical philosophies of knowledge and the effects they had on the development of KBSs.

### The Process of Practise

In human psychology and sociology the process of practice is where the behaviour of people or organisations' is not governed by logic or a reflexive response to an event. Rather, they are ruled by their cultural, mental or social state at the time. For instance, a person's *habitus*, or predisposition to be effected by something (WordNet 2003), can significantly alter their response to a particular set of facts (Gaines 2000). One type of habitus, where behaviour is induced by normative rules, such as government codes or through business operations, has been successfully modelled by KBSs. However, these can still be problematic as normative rule sets are generally incomplete and require significant interpretation (Op cit). Gaines (2000) claims that the failure to adequately model the remaining process of practice is the "...greatest impediment to the development of expert systems" (Gaines 2000, p. 115).

Currently there are two general groups of approaches whittling away at the region of practice. Firstly, connectionist researchers have been investigating the emulation of higher brain functions. Generally, connectionist work has struggled to develop systems that model higher brain functions as they tend to be mostly reactive (Gaines 2000). Machine learning methods like temporal difference learning can be integrated with connectionist approaches providing the ability to spread rewards over temporally separated events in a generalised form (Sutton and Barto 1998). This offers the ability to learn some simple actions (Op cit). However, these systems struggle to learn the amount of knowledge required for complex higher brain functions, thereby, limiting their effectiveness. Simultaneously, approaches such as Lenat's (1995) Cyc project, have attempted to extend the application of reason downward through the codification of general knowledge. This approach is attempting to gain the connectionists' ability to generalise through brute force.

The missing link between these approaches is the ability to store or represent complex knowledge in a generalised manner. This has been partially addressed in a number of symbolic based methodologies by attempting to represent the context of knowledge. For instance, Repertory Grids, Formal Concept Analysis (FCA) and Ripple-Down Rules (RDR) all integrate either implicit or explicit contextual information. However, these methodologies treat context as a static entity, neglecting many connectionists' work in learning hidden and dynamic contexts and are still significantly limited in their ability to represent knowledge in a fully generalised manner.

One of the most promising methods for modelling the central region of practice though, as claimed by Gaines (2000), is through a system using multiple rule sets based on exceptions within a prioritised hierarchy. Furthermore, Gaines (2000) claims that Multiple Classification RDR (Kang and Compton 1992, 1994; Kang et al. 1995) goes some way to providing such a knowledge base (KB) structure. The multiple conclusions allow for a less rigid outcome, which permits the resolution of conflicting constraints by providing a number of alternate possibilities. However, such a system still only provides specific alternate classifications rather than a general constraint, as Gaines (2000) claims is needed. Thus, a more general system of selecting admissible actions that satisfy the identified constraints, but are otherwise indeterminate, is more capable of realistically modelling human practice (Gaines 2000).

## Context

The strength of RDR based methods in partially addressing the problem of modelling habitus, is based on their incremental nature, exception based knowledge structure and the ability to capture static contexts. Context can be thought of as the starting state or assumptions prior to inferencing. In the case of RDR, context is determined during inferencing. This allows simple maintenance within the context determined, through rule exceptions. This contextual based maintenance allows RDR based methodologies to capture more transient and tacit knowledge (Richards and Busch 2000, 2001). This is without the need to globalise knowledge to a general context, where no prior knowledge or assumptions are required before codifying, unlike many other KBS methodologies, such as Knowledge Acquisition and Design Structuring (KADS).

This failure of most methodologies to consider context has significantly hampered their knowledge acquisition process. However, methods, including RDR based techniques, that do incorporate either implicit or explicit context, assume that it is *a priori*, and therefore, deductive. This assumption leads to *static* representations of contextual based knowledge. However, context in certain situations could be considered *a posteriori*, and therefore, inductive (Brezillon 1999). For instance, during a conversation the context or assumptions of the people involved is constantly altered as the interaction meanders through different topics. Each new topic requires a different set of assumptions. Situations such as this require a dynamic and changeable representation. Dynamic contexts require the ability to

have multiple start states that result in the same or similar contexts. Alternatively, a number of similar starting assumptions can result in many different contexts.

Brezillon (1999) suggests that dynamic context can only be modelled during the solving of a problem or through interaction, and thus not directly coded in a static structure. A number of authors, such as Widmer and Kubat (1996) go further, suggesting that these new contexts are often differentiated by variables or facts that are hidden, even to those altering their contextual position. Such contexts are sometimes referred to as *hidden contexts*.

In this paper the idea of a hidden context is one that is only identifiable through influences on a system that cannot be captured, such as human emotion or intuition. These situations arise where aspects of a contextual situation are hidden from the current state and the path taken to reach that position. A hidden context differs from a dynamic context due to the contextual position being unknown, whereas dynamic contexts are changing over time. Furthermore, the cause of a dynamically changing context can potentially be due to the influence of a hidden context.

The dynamic, and generally hidden, contextual nature of habitus, potentially prevents contextually static approaches from being able to encroach further into the central region of practice. It is common knowledge that different people react differently to particular events. Furthermore, the same person may react differently to the same event at different times. For instance, people will often wait to tell someone some news because they are “*waiting for the right time*.” This hidden and dynamic context of habitus means facts about a situation will be interpreted differently, resulting in a different response. Thus, the contextually static system’s inability to incorporate this contextual information hinders their ability to accurately model the *process of practice*.

## **Traditional View of Knowledge**

Symbolic reasoning is built around the storage and codification of human knowledge. Yet the question “*what is knowledge?*” is still very much an open question. This section will investigate this question and what has resulted from the different answers to this question. The traditional expert systems’ view of knowledge was founded on the *physical symbol hypothesis* (Newell and Simon 1976), which takes the view that knowledge is comprised of symbols, and the connections between those symbols, representing pieces of reality. Furthermore, that *intelligence* comes from the appropriate manipulation of these symbols and relationships. This AI perspective of knowledge is not new and closely draws from the philosophies of Wittgenstein, Descartes, and ultimately Plato’s archetypes (Compton and Jansen 1990; Compton 1992).

The result of the physical symbol hypothesis was that Expert System researchers assumed that such symbols and relationships should be extractable and usable without any further need of the expert. This fundamentally reductionist strategy was the logical extension and led to the *knowledge principle* (Feigenbaum 1977; Lenat and Feigenbaum 1988, 1991). The knowledge principle essentially suggested that the success of an expert system is dependent on the amount of information about

symbols and their relationships in the KB and not the inferencing or reasoning strategy applied. Therefore, when difficulties arose, the common misconception was that this came about from ineffective knowledge acquisition, and furthermore, that experts do not “...communicate the underlying knowledge very well” (Compton and Jansen 1990, p. 242). A closer investigation of these problems, however, reveals three primary areas of difficulty:

- *Knowledge base consistency.*  
Each additional symbol extracted from an expert, must be integrated into the KB, such that no existing conclusions in the system, that are already correct, are changed. Tools have been developed for validating and verifying knowledge that can provide some assistance in a few systems. However, once identified, an inconsistency still must be corrected and hand tailored by the Knowledge Engineer (KE) (Compton and Jansen 1988).
- *Inconsistent use of symbols by experts.*  
Shaw (1988) identified that experts frequently disagree on what is “*correct knowledge*” when presented with it in the form of a KB. However, generally these same experts have no problem discussing ideas with each other, even though they apparently have conflicting views on the individual pieces (Compton and Jansen 1988). This disagreement, however, can significantly hamper knowledge acquisition.
- *Experts’ justifications can be affected by the KE and the tool used.*  
People tend to describe concepts relative to a listener’s level of knowledge. Furthermore, the contextual knowledge provided by an expert can often depend on the framework of the questioner (Compton et al. 1991). Thus, a KE may ask a question with certain expectations, which can be subtly conveyed to the expert, who will tend to respond within that context.

The last two points challenge the original physical symbol hypothesis idea of knowledge as being a static entity that can be extracted. These two points suggest that the individual symbols can vary in meaning between experts and the connections between those symbols vary according to the context in which the knowledge is being used. This suggests that any form of static extraction of symbols and connections ultimately is doomed to fail in all but the most rigid of knowledge domains. The problems in knowledge acquisition originate not so much from the extraction of the knowledge from the expert, but from the KE’s attempts to twist and warp the acquired knowledge into a form suitable for the KBS in which they were attempting to fit the knowledge.

### Traditional Knowledge Representation

There have been numerous methodologies developed attempting to either, acquire already globalised knowledge, or, of representing knowledge so that it facilitates the easier inclusion of the gathered knowledge. These methods are primarily from, what has been referred to as, *second-generation ES*. Second-generation ES moved away from the more ad hoc KB structuring used in the earlier shell based systems, called *knowledge transfer* (Newell and Simon 1976), towards a *knowledge-level modelling*

(Siewiorek et al. 1982) approach. These new model-based ES approaches were believed to provide a richer knowledge representation (KR) and, thus, the ability of capturing deeper and different types of knowledge, such as tacit knowledge (Richards 2001).

The majority of these methodologies, however, still view knowledge as static, which once captured remains correct. They are in fact just new approaches to facilitating the conversion of experts' knowledge into a form that can be stored in the KB. They generally do this by forcing the expert to provide their knowledge to the engineer in a more structured approach, causing them to globalise their knowledge at the point of extraction. This reduces the need for the engineer to extensively convert the knowledge. While this is a significant improvement over the ad-hoc approach, the fundamental problem is that it transfers part of the difficulty of knowledge acquisition from the engineer to the expert, rather than actually solving the root issue of KR. It would be preferable to simply store the knowledge in the contextual form naturally gathered than requiring anyone having to globalise it first.

### Traditional Knowledge Maintenance

While knowledge maintenance is presented here separately from KA and KR it is in fact a microcosm of the same issues already described. It requires both the ability to acquire new knowledge, as well as being able to insert it into the existing KB. However, it usually presents as a much more difficult problem for ES developers. There are three primary reasons for this added difficulty. First, the KE usually only wants to make the minimal amount of changes when performing maintenance and does not wish to restructure the entire KB. Second, often the original experts and engineers are no longer as intimately involved with the project, making updating problematic. Finally, the ES's domain can often shift, including new areas not covered by the existing system, after its deployment. Without a complete redevelopment this can significantly increase maintenance problems, as found in the Xcon system (Bachant and McDermott 1984). However, as the previous section briefly argued, KA, KR and KM are not necessarily separate tasks and if a more holistic approach is used then the difficulty created by these separate views can be resolved.

### Knowledge Based Systems

There are numerous methodologies and tools that have been developed to aid in the creation of KBSs, which attempt to address the KA problem. There are many books, papers and reviews (such as Lenat and Guha 1990; Schreiber et al. 1993; Schreiber 1993; Leake 1996; Gennari et al. 2002) that investigate the issues associated with the problem in detail. This section will only very briefly discuss two approaches developed: KADS and Cyc.

### *Knowledge Acquisition and Design Structuring*

Knowledge Acquisition and Design Structuring is the outcome of the European research project ESPRIT-I P1098 initiated in 1983 with the aim of developing a

comprehensive, commercially viable methodology for knowledge-based systems (Wielinga et al. 1992). KADS is recognised as one of the first true methodologies developed specifically for the development of KBSs and is still widely used today. It was developed in particular to address the problems found in the *knowledge acquisition bottleneck*. The researchers' view of the problem was that knowledge acquisition failed in traditional systems through an inability to get to the *deep-knowledge* of the experts, primarily due to a lack of structural constraints on the experts during knowledge extraction.

Rather than viewing KA as filling a container of knowledge, the KADS perspective is an operational model that displays a form of observed behaviour which is "...specified in terms of real world phenomena" (Wielinga et al. 1992, p. 6). Basically, KADS has an array of modelling techniques, where the expert and engineer work together to build up a set of tasks so that knowledge acquisition can be approached in a systematic way. This ensures, as best as possible, that nothing gets left out and the people concerned are aware of where they are up to in the process (Compton et al. 1993). This divide and conquer approach to KA is the basic underpinning of all *task* orientated methodologies.

The reason approaches, such as KADS, work reasonably effectively during development is because the expert themselves model a particular individual task. This forces the expert to step out of their current context into the global knowledge domain. This process of the expert now providing more globalised knowledge prevents the engineer from being required to extensively convert the knowledge further. However, the resulting KB is still global in context and static in representation. Wielinga et al. (1992) even defines domain knowledge in KADS as being static knowledge.

This static representation, however, leads to one of KADS greatest shortcomings: knowledge maintenance. Nowhere in the KADS methodology, is the knowledge maintenance issue addressed directly. Instead it assumes that the spiral life cycle model will continue infinitum. This omission has resulted in alterations to the basic methodology, such as *structure preserving design* (Schreiber 1993), which preserves the *information content and structure* in the knowledge-level model, within the final artefact. Therefore, the system not only provides the static domain knowledge but also the relationships between the artefact and the original knowledge sources and/or meta-classes. This makes development a process of adding implementation detail to a knowledge-level model, which makes it more possible to trace omissions or inconsistencies in an artefact back to the relevant part of the model, considerably simplifying maintenance (Killin 1993; Schreiber 1993). This essentially, although only in a limited form, can be seen as an attempt to include some form of context within the symbols.

## Cyc

Cyc, short for encyclopaedia, did not start out as an attempt to develop new methodologies for knowledge acquisition. Rather, its intention was primarily to solve the problem of *brittleness*, where a small amount of missing knowledge can be



significantly detrimental to the system. However, as a result it has had implications for both KA and KR methods. A KB's brittleness comes from its concentration of specialised knowledge within a single narrow domain. Therefore, when knowledge is required from just beyond this domain the KBS collapses. Fundamentally, this is part of the same problem the previous methodology was attempting to fix. KADS's unarticulated view had been if we can extract deeper knowledge then the system will be less brittle. Lenat et al. (1990) however, argues that brittleness is the result of insufficient commonsense knowledge within the KB. It is this lack of commonsense knowledge that Cyc has been developed to address. The Cyc system is a universal schema with millions of directly entered and inferred commonsense axioms that make up hundreds of thousands of general concepts (Lenat 1995).

While methodological development was not the driving force behind Cyc, it was one of the results. During the system's development it was obvious that existing methodologies were woefully inadequate at scaling to the required size or at representing particular concepts. Cyc incrementally developed its own representation language then, to address repetition issues that eventuated, and periodically smoothed out the resulting structure (Lenat et al. 1990). It uses a frame-based language embedded in a first order predicate calculus framework with a series of second-order extensions that allow the representation of defaults, reification, and for reflection (Lenat et al. 1990; Pittman and Lenat 1993; Guha and Lenat 1994; Lenat 1995). The inference engine used for Cyc was also incrementally constructed using more traditional computer science data-structures and algorithms.

Lenat (1995) argues that the majority of *assertions* could not be made correctly without the use of some form of context. For instance, the statement "*you cannot see a persons heart*" assumes that the person is not currently undergoing open heart surgery. Alternatively, the assertion could also represent a metaphorical meaning. Cyc's solution was to place each assertion into one or more explicit contexts, through the use of *microtheories*. Within each context, the assertion is then given its default conclusion. Each context is itself an individual KB. Additionally, the provision for being able to import assertions from other contexts was also included in Cyc allowing the combination of contexts and contradictory assertions to be resolved (Lenat et al. 1990; Lenat 1995).

The Cyc project's brute force approach shows potential as a basis for domain specific KBs to be built upon. However, Cyc's ad-hoc development methodology requires explicitly stated knowledge that can often be dated or invalidated by the time it is eventually used in a real world system. This paper asserts that: commonsense knowledge is one of the most a posteriori-contextually dependent forms of knowledge, due to its high dependence on culture and time. For instance, individual axioms not only change conclusions between contexts, but they also can change within the same context between different culturally independent minority groups. For instance, the common sense knowledge in Cyc is representative of "*...TheWorldAsAGroupofWhiteWesternMiddleClassProfessionalMenBelieveItToBe*" which may not be the same as *TheWorldAsSomeOtherPeopleBelieveItToBe*" (Adam 2000, p. 243). Furthermore, the contexts themselves within these groups also changes over time. Adam (2000) exposes how "...culturally and temporally variable common sense can be" (Adam 2000, p. 243), by pointing out it was

“...common sense in Victorian times to assume that higher education shrivelled women’s reproduction organs” (Adam 2000, p. 243). These issues can be seen if we look at one of Lenat et al’s (1990) own examples of knowledge in the Cyc system:

Payments of less than ten dollars are usually made with cash; those over 50 dollars are usually made via cheque or credit card (Lenat et al. 1990, p. 43).

Such an assertion is highly cultural, location and time specific and, therefore, very susceptible to failing in a contextually-dynamic environment. For example, given such knowledge, one must ask “*how relevant is this?*” to the following: a peasant farmer in central china, who has no concept of how much a dollar is worth; to people in a war zone that have no access to secure financial institutions; or, to someone living in 2015 where all transactions are made with smart cards. Therefore, even though various static-contexts are identified in Cyc, the absence of *dynamic-contexts* renders its approach as far too simplistic and highly susceptible to obsolescence, for the development of a (near) complete KB of common sense knowledge. For instance, Clancey (1991) likens common sense knowledge to that of chaos theory, where projects such as Cyc attempt to collect it like “...so many butterflies” (Clancey 1991, p. 245).

## Situated Cognition Applied to Artificial Intelligence

Problems such as inconsistent use of symbols and experts’ justifications being affected by the context of the listener, oppose the original *physical symbol hypothesis* idea of knowledge as being a static entity that can be extracted. This returns us to the original question of “*what is knowledge?*” Situation Cognition moves away from these traditional definitions of knowledge and instead argues that knowledge is generated at the time of use and that context-independent assertions cannot accurately model human cognition (Menzies 1996, 1998). The proponents of Situated Cognition (SC) tend to fall into one of two camps.

1. *Weak SC* argues that when the human agent uses a particular description of knowledge that they use the context of the current problem or situation to continually reinterpret that description.
2. *Strong SC* takes a significant step further, claiming that context has such a potent influence on the human agent that systems should be purely reactive and that we should discard symbolic representations altogether (Menzies 1996, 1998).

Therefore, SC, in its *weak* and most common interpretation, views knowledge instead as being mostly context based (Menzies 1996). SC and its subfields, situated automata (Maes 1990; Waldrop 1990) and situated action (Agre 1990; Suchman 1987), are philosophically justifiable through work such as Bartlett (1932), Piaget (1970), Jenkins (1974) and Bransford et al. (1977). SC attacks the *Platonic* view of knowledge and memory and instead claims that knowledge is re-constructed each time it is needed (Agre 1990; Compton and Kang 1993; Maes 1990; Clancey 1991).

This dynamic re-construction of knowledge involves the combination of two factors: firstly, where the agent has come from to reach this particular point, and the

location of the agent's eventual goal. For instance, Clancey provides the example that "...at a base level a person is always like a dancer balancing [his/]her next steps against the inertia of past movements and [his/]her view of where s[he] is going" (Clancey 1991, p. 244).

### Knowledge Acquisition in Context

A new perspective on knowledge does not necessarily imply that the method by which knowledge is acquired needs to be reconsidered. Essentially, all that SC does is allow a KE to view the knowledge they gather differently. The problems identified for knowledge acquisition in Sect. "Traditional View of Knowledge" did not originate from the extraction of knowledge from the expert, but from the KE attempts to globalise the acquired knowledge into a form suitable for the KBS. Therefore, the key to improved KBS development is not so much in *how knowledge is acquired*, but *how that knowledge should be represented*.

### Representing Knowledge in Context

As discussed earlier the traditional *physical symbol hypothesis* views extracted knowledge as unchanging. Therefore, ES methodologies can only build KBs with a static KR. The SC view of knowledge, however, indicates that knowledge is relative to the context it was re-constructed, and therefore, changeable. Thus, to truly reduce the problem of KA, a representation that incorporates context should be used. It follows then that a KR methodology that intends to incorporate context, must be able to assimilate change. This does not negate the possibility of specifying a KB prior to its application, but it does mean that the previously specified store of knowledge should be dynamically alterable.

Ideally a system should have a KR technique that is sufficiently flexible that it could change knowledge or include the context of the supplied knowledge. This would allow knowledge extracted from an expert to be significantly more easily included within the KB and, thereby, preventing the need for either the KE or expert to interweave knowledge into a global perspective. This would significantly simplify the problems associated with knowledge acquisition, as the conversion from contextualised knowledge to globalised knowledge would no longer be required.

### Maintaining Knowledge in Context

Menzies (1998) identifies one of the primary changes in knowledge engineering brought about due to the SC view as a new emphasis on maintenance rather than design. This is due to the context in which knowledge can be located is often only identifiable in an online environment during the maintenance phase. This directly challenges the approach of second generation KBS methodologies like KADS, which focus on design and do not even discuss maintenance in the core set of tasks. Menzies (1998) suggests that one approach to situated knowledge engineering is to

build two KBs. The first captures system knowledge like any other KB, while the second models the impact of the first KB on its environment and vice versa. For example, the BRAHMS system by Clancey et al. (1998) and REMAP by Ramesh and Dhar (1992) both attempt to model the environment around a KB. Another approach is to directly represent the context of the knowledge being stored within the KB itself. This context can either be stored implicitly or explicitly. For instance, the Cyc ad-hoc approach uses microtheories, which allows for a form of explicit contextual information to be included in the KB.

## Contextual Knowledge-Based Systems

There are only a limited number of tools and methodologies that have been developed to aid in the creation of KBS through the incorporation of context. This section will briefly introduce some widely used approaches; namely, FCA, Repertory Grids and RDR.

### *Formal Concept Analysis in Contextual Knowledge Acquisition*

Formal Concept Analysis, originally developed by Wille (1981), is based on lattice theory. FCA is not a KA methodology, rather it is mathematical tool used for the discovery of concepts. It was designed to provide a basic answer to two fundamental questions in relation to a concept; how to appropriately classify objects within a particular context, and what the dependencies are between attributes (Wille 1981). Its relevance to KBS comes from its representation theory and how that can be used in KR.

Formal Concept Analysis is rooted in the philosophical idea of a concept as a single unit of thought that is made up of two components: *extent*, which covers all objects (entities) belonging to a particular concept; and *intent*, which comprises the attributes (properties) that are valid for all those objects (Wille 1981). Put another way, a formal context,  $K$ , is a triple  $K(G, M, I)$  containing the set of objects  $G^1$  forming the *extension* of the concept, and a set of attributes  $M^2$  represent the *intention* of the concept, which are linked together by the binary relation  $I$  (Richards 1998a; Wille 1981).

The advantage in the FCA approach is that lattice theory provides a simple mathematical vocabulary for discussing order, and especially when used in systems where there is a natural sense of hierarchies (Birkhoff 1938), due to its theory of substructures (Wille 1981). This advantage results in the ability of a complete lattice, called a *concept lattice* or *semantic net*, providing a conceptual clustering of objects via the *extents* as well as a representation of the implications between the attributes via the *intents* (Wille 1992). It is this ability to express all relationships between attributes, describe objects in terms of the concepts contained and to show the relationships between those concepts, which makes FCA a powerful representation for knowledge (Richards 1998a).

<sup>1</sup> G stands for *Gegenstande* in German.

<sup>2</sup> M stands for *Merkmale* in German.

Generally, FCA has not been used as a KR methodology by itself. Instead it has been used for knowledge acquisition (1992; Wille 1989) in many KBS based areas, such as CBR (Diaz-Agudo and Gonzalez-Calero 2001) and ontology construction and maintenance (Stumme et al. 2000), as a means for discovering knowledge embedded in cases through these identified relationships. Within these areas FCA has been highly effective at extracting contextual knowledge.

Clancey (1991) mentions, however, that although semantic network based approaches can embody a cognitive model that show human-like behaviour, they are limited to a "...grammatical model of cognition" (Clancey 1991, p. 251). Therefore, they fail to capture non-verbal concepts or to "...model the perceptual-conceptual learning..." (Richards 1998a, p. 155) evident in humans' attachment of meaning to a grammar (Richards 1998a). Clancey (1991) continues by pointing out that concepts are being regarded as *things* rather than the "...processes of perceiving and processes of behaving" (Clancey 1991, p. 252). Additionally, Wille (1981) himself identifies the method's lack of scalability for the inclusion of all relevant attributes and relationships with large contexts, potentially limiting its use.

### *Personal Construct Psychology and Repertory Grids*

Kelly's (1955) book "*Psychology of Personal Constructs*" introduces his *personal construct theory* (also referred to as *Personal Construct Psychology*, PCP), where he postulates that "A person's processes are psychologically channelled by the way in which he anticipates events" (Kelly 1955, p. 46). Basically, he viewed all people as their own *personal scientist at anticipation*, where an anticipatory model of the environment was reflexively applied to the self (Shaw and Gaines 1992). Shaw and Gaines saw PCP as a "...constructivist position in modelling human knowledge...that characterises conceptual structures in axiomatic terms" (Shaw and Gaines 1992, p. 31).

Kelly's (1955) emphasis is on the space created by the process of making distinctions rather than being defined by the *elements* identified. Within this *psychological space* a *construct* is like a single plane that slices through a large collection of events (Kelly 1970). Fundamentally, a construct contains a triple of two disjoint distinctions that are mutually subsumed by the third, referred to as the *range of convenience* or the *subsuming concept* (Richards 1998a, b). Essentially, the disjoint pair bound the extremes of the range of convenience or concepts. Therefore, this provides a means for bounding concepts, allowing knowledge acquisition to start with a more encompassing view and acquire knowledge inwardly; rather than the traditional approach, of starting at the centre and acquiring knowledge outwardly.

In order to actually acquire the relevant constructs and elements for a domain, Kelly also developed the *repertory grid technique* (Richards 1998a, b). It is designed to bypass a person's cognitive defence and provide an avenue to their underlying construction system by asking them to compare and contrast various examples. Thus, repertory grids use the idea that people are more able to offer context based examples than defining globally based rules. Essentially, the repertory

grid is a method of finding concepts, the structures within those concepts and the relationships between those structures without eliciting them directly. This can be more effective than the conventionally based methodologies that directly extracted models (Richards 1998a).

Kelly's repertory grid technique has had widespread success in a number of domains and continues to be applied in both manual and computerised applications. Some systems that use the repertory grid are the Knowledge Support System Zero (KSS0), Expertise Transfer System and AQUINAS (Boose et al. 1989). Gaines and Shaw (1993) also have provided a general framework to assist users in the elicitation of the conceptual structures when using repertory grids (Richards 1998a).

Personal Construct Psychology is similarly simple in its acquisition of knowledge to RDR. Both view the concept of model derivation as being difficult and problematic, preferring the simpler technique of eliciting knowledge directly from an expert without the need of a KE. However, the captured knowledge is different, in that RDR develops an assertional KBS, while PCP captures a conceptual model from which a terminological KBS can be extrapolated (Richards 1998a). As discussed in Sect. 3.4.1, Richards (1998a, b) developed a combination of RDR and FCA in order to create a terminological KBS with RDR.

### *Ripple-Down Rules*

Compton et al. (1988, 1989) extended Popper's (1963) theory of hypothetico-deductive reasoning to the application of knowledge engineering. He suggested that experts do not provide information on their insight or how they reached a particular conclusion; but instead, they provide a justification for excluding the other possible hypotheses from within a particular context (Compton and Jansen 1988, 1990; Compton et al. 1988, 1989). It was suggested that these dichotomies between insight and justification (Compton et al. 1988, 1989) arise from a traditional misinterpretation of the form of knowledge provided by experts. This new hypothesis for knowledge engineering has resulted in the development and deployment of a new collection of methodologies and applications based on the knowledge acquisition and inferencing philosophy of RDR.

The original impetus for the development of RDR came from Compton et al's (1988, 1989) work on the GARVAN-ES1 expert system and the maintenance issues that arose. Through this maintenance experience, Compton et al. (1988, 1989) found that in order to create a maintainable ES through dialogue with domain experts, then the system should build into its methodology a means for incorporating an expert's justification of a hypothesis within a particular context.

The resulting methodology developed, RDR, uses a simple exception structure aimed at partially capturing the context that knowledge is obtained from an expert. It was assumed that the context was the sequence of rules that had evaluated to provide the given conclusion (Beydoun 2000; Compton and Jansen 1988; Compton et al. 1991; Preston et al. 1993, 1994). Therefore, if the expert disagrees with a conclusion made by the system they can change it by adding a new rule with whichever level of generality was required by the expert. However,

the new rule will only fire if the same path of rules is evaluated with the same outcomes as previously (Compton et al. 1993). Since its development it has accurately been described as knowledge acquisition through *fault patching* (Menzies and Debenham 2000). RDR was so named to capture this notion of appending *fix-up* rules to the bottom of a KB, thereby, rippling the formula for a conclusion down over many rules, rather than attempting to fix existing rules (Jansen and Compton 1989).

The context of a new rule is its position within the RDR structure. A rule can only fire when the rules that led to it, before its creation, have fired. Therefore, the rules context is the previous path of rules and their outcomes. This is effectively capturing the rules context implicitly within the structure itself. This corresponds to how a human expert, explaining the justification of their conclusion to a trainee, will start from where they believe the trainee deviated from the correct solution, providing the justifications required from that point to the expert's correct conclusion. Furthermore, it also corresponds to a person's method of debate or argument. When debating an issue each person presents their view but they start their justifications from the current context of their opponents.

Ripple-Down Rules' very simple no-model approach to KA has revealed many advantages. For instance, because a rule can only fire within a particular context, then when it is being created it only requires validation within that context (Dazeley 2007). Therefore, new knowledge can easily be validated during the moment of acquisition. Furthermore, the expert does not need to know anything about the actual underlying structure of the knowledge they are constructing. This is a marked difference from model based approaches where the KE needs to consider and validate each new piece of knowledge across the entire KB. In this system the normal KB development phases of requirements definition, analysis, design, implementation and maintenance can be eliminated. RDR omits the design stage completely and merges the implementation and maintenance phases seamlessly (Beydoun 2000), effectively providing a holistic approach to KA and KM where there is little or no difference between the tasks.

## **Hidden and Dynamic Context Applied to Artificial Intelligence**

The application of ideas from weak situation cognition has already shown great potential to significantly improve Knowledge Acquisition and maintenance in many domains. However, in its weak form SC assumes context is static and once found can be codified. This is once again falling into the same trap that KBS developers did when using the *physical symbol hypothesis*. This assumption simplifies the problem domain allowing KBS developers to ignore more difficult unsolved situations. This is fine when used in situations where context is for the most part static, such as well defined areas like medical diagnostics. However, in many real world situations, such as general knowledge, this abstraction reduces the potential of such methodologies being successfully applied.

The more radical view of strong SC agrees that the process of abstraction and representation simplifies the environment too much. Brooks (1991) claims that



giving a computer a human *Merkwelt*<sup>3</sup> is flawed because it is based on our own introspection. There is no evidence that what we perceive as our *Merkwelt* is in fact what occurs internally and "...it could just as easily be an output coding for communication purposes" (Brooks 1991, p. 141). Brooks (1991) identifies two primary problems with the use of the human *Merkwelt*. The computer, robot, or software agent does not perceive the world in the same way a human does. Furthermore, even if it did, there is no guarantee the human introspectively derived *Merkwelt* is correct.

While these arguments for strong SC are legitimate, the conclusion to reject all symbolic representation is flawed. As noted by various authors (Anderson 1990; Menzies 1996, 1998; Patel and Ramoni 1997; Vera and Simon 1993a, b, c) there are numerous examples showing that the symbolic based descriptions of experts are richer and more abstract than those of novices (Menzies 1998). Therefore, these descriptions must represent some interpretable meaning. However, what the strong SC argument shows is that weak SC based KBSs still suffer from an over simplification of the environment being represented.

It can, however, be argued that a middle ground could be taken. Given the weak SC argument that knowledge is constantly reinterpreted in the context of the problem, then it is possible to expand upon this notion by elaborating on "*what is context?*" Context in the Cyc, PCP, FCA and RDR have all assumed context to be a priori known and, therefore, identifiable, codifiable and static. For instance, when driving a car in Australia you should give way at an intersection to cars coming from the right. Likewise, cars to your left will give-way to you. However, if one of the cars to your left is an emergency vehicle with its siren blaring then you also give-way. Therefore, the different context in this case is the presence of the emergency vehicle. These two contexts are static, because the knowledge used in each is constant and does not change.

*Strong SC*, however, claims that the effect of context is overpowering, rendering such work as futile. For instance, long term weather prediction is inaccurate due to the existence of so many factors making up the forecast. Each combination of these factors represents a different context. *Strong SC* does not, however, alter the underlying notion that context is an influence; just that it has a greater affect than weak SC claims. Nor does it adequately define context or detail why some form of symbolic representation does not aid learning to some degree. Thus, a middle ground reinterpretation of SC, referred to as *Intermediate SC*, should provide an elaboration on the types of context possible. Building on from the *weak SC* claim:

*Intermediate SC* suggests that humans continually reinterpret a particular description of knowledge based on their current situation, which is dynamically alterable due to hidden contexts from within their *Merkwelt*.

An interpretation of SC such as Intermediate Situation Cognition still allows for symbolic reasoning but indicates that it should be able to handle different

<sup>3</sup> *Merkwelt* is the term used by Jakob von Uexkull in his 1934 paper 'A Stroll through the Worlds of Animals and Men: A Picture Book of Invisible Worlds', to refer to the complete set of environmental factors that have an affect on a species regardless of whether they are perceptible or not.



reinterpretations of the same knowledge in the same context. For instance, a blackjack player's decision to draw a card when on seventeen is only partly based on the mathematical odds. Numerous factors from within their *Merkwelt* could also be affecting their choice. For example, they may think there are many smaller cards than usual left in the deck or they may simply have a *feeling*. In this situation the *Merkwelt* is being influenced by hidden emotional or subconscious states or contexts, which alters their response from previous times they have had the same hand. This outwardly produces what appears to be a dynamic contextual response, because a different action was taken for the same context, the hand dealt.

Some may claim, however, that dynamic context is represented in the contextually-static methodologies, through breaking down a drifting context into a series of discrete static representations. However, this is abstracting a level of complexity from a domain that could contain important information. It is also turning a domain into a "*blocks world*" type of toy problem. It is this over simplification that Strong SC, and therefore intermediate SC, truly opposes.

Furthermore, recognition of hidden and dynamic contexts is certainly not new. For instance, Arbib (1993) alludes to the existence of hidden contexts in schema theory, in that "...no single, central, logical representation of the world need link perception and action" (Arbib 1993, p. 273). Widmer and Kubat's (1996) FLORA system was designed with the idea that contexts may be hidden and unrealised.

One of the most widely researched areas of dynamic context is in the Information Filtering research stream, dealing with the issue of concept drift. Fundamentally, there is little difference between a user's concept and the context of their behaviour. The concept is the idea behind a user's action, whereas, the context represents the circumstances of an event. Thus, the user's concept governs the circumstances of the event, and therefore, the current context. Consequently, as the user's concept drifts, often dynamically, so does the context.

## Implications for Knowledge Engineering and Maintenance

The idea of introducing hidden and dynamic contexts into KBS research could potentially aid in understanding many of the historical and current problems in knowledge engineering and the maintenance of such systems. The fundamental issue, however, is how does a symbolic representation capture such information. Once a symbol is captured as a piece of knowledge, we not only are required to store a contextual relationship in parallel with that knowledge, we must also be able to alter it dynamically. Additionally, we should be able to find new contexts not expressed by an expert.

## Conclusion

This paper discussed how existing approaches to AI have failed to adequately model the area of human reasoning that psychology and sociology refer to as the process of practise. It is this failure that has significantly contributed to brittleness in symbolic

reasoning. This paper investigated how the process of practise relates to deeper forms of contextual representations of knowledge.

This paper briefly outlined two philosophies of knowledge. The first, *physical symbol hypothesis* and its later, less ad hoc, form of *knowledge-level modelling* were discussed. The primary difficulty in these views was the conversion from contextually based to globally based knowledge and that these methodologies are inherently static in nature. Also discussed were a number of methodologies that use this interpretation of knowledge and how this presented further difficulties during knowledge acquisition and especially maintenance.

The emergence of a SC view of knowledge addressed these issues, which was the second philosophy discussed. It was identified that KR methodologies need to be used that allow for dynamically changing knowledge. RDR is one methodology that has attempted to meet this SC view. Other methodologies that also use SC as their underlying philosophy, such as PCP and FCA were discussed.

These systems have performed exceptionally well in many areas where knowledge-level systems have repeatedly failed. Nevertheless, strong SC literature argues that such systems will not achieve true robustness and intelligence as context is too problematic to be represented symbolically. Finally, this paper has suggested a clarification to what forms context can appear and suggest that for symbolic systems to succeed they must embrace all these contextual forms and that this meets the concerns of Strong SC. This has resulted in the presentation of a new interpretation of SC for AI referred to as Intermediate Situation Cognition which makes a significant step towards an epistemology that could lead to the modelling of the process of practise.

**Acknowledgements** The majority of this paper is based on research carried out while affiliated with the Smart Internet Technology Cooperative Research Centre (SITCRC), Bay 8, Suite 9/G12, Australian Technology Park, Eveleigh, NSW 1430 and the School of Computing, University of Tasmania, Locked Bag 100, Hobart, Tasmania.

## References

- Adam, A. (2000). Deleting the subject: A feminist reading of epistemology in artificial intelligence. *Minds and Machines*, 10, 231–253.
- Agre, P. E. (1990). Book review of plans and situated: Actions: The problem of human machine communication. *Artificial Intelligence*, 43(3), 369–384.
- Anderson, J. (1990). *Cognitive psychology and its implications* (3rd ed.). New York: W. H. Freeman and Company.
- Arbib, M. (1993). Book review Allen Newell, unified theories of cognition. *Artificial Intelligence*, 59I have reprint in *contemplating minds*, 265.
- Bachant, J., & McDermott, J. (1984). R1 revisited: Four years in the trenches. *AI Magazine*, 5(3), 21–32.
- Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. London: Cambridge University Press.
- Beydoun, G. (2000). *Incremental knowledge acquisition for search control heuristics*. Sydney: University of New South Wales.
- Birkhoff, G. (1938). Lattices and their applications. *Bulletin of the American Mathematics Society*, 44, 793–800.
- Boose, J. H., Bradshaw, J. M., Kitto, C. M., & Shema, D. B. (1989). From ETS to aquinas: Six years of knowledge acquisition tools. In *Proceedings of the Fourth AAAI sponsored Knowledge Acquisition for Knowledge based Systems Workshop, Banff* (pp. 5.1–5.17).

- Bransford, J. D., McCarrell, N. S., Franks, J. J., & Nitsch, K. E. (1977). Toward unexplaining memory. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting, and knowing: Toward an ecological psychology* (pp. 431–466). Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Brezillon, P. (1999). Context in artificial intelligence: II. Key elements of contexts. *Computer and Artificial Intelligence*, 18(5), 425–446.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159.
- Clancey, W. J. (1991). Book review Israel Rosenfield, the invention of memory: A new view of the brain. *Artificial Intelligence*, 50(2), 241–284.
- Clancey, W. J., Sachs, P., Sierhuis, M., & van Hoof, R. (1998). BRAHMS: Simulating practice for work system design. *International Journal of Human Computer Studies*, 49, 831.
- Compton, P. (1992). Insight and knowledge. In *AAAI spring symposium: Cognitive aspect of knowledge acquisition* (pp. 57–63). Stanford: AAAI Press.
- Compton, P., Horn, R., Quinlan, R., & Lazarus, L. (1989). Maintaining an expert system. In J. R. Quinlan (Ed.), *Applications of expert systems* (pp. 366–385). London: Addison Wesley.
- Compton, P., & Jansen, R. (1990). A philosophical basis for knowledge acquisition. *Knowledge Acquisition*, 2, 241–257.
- Compton, P., & Jansen, R. (1988). Knowledge in context: A strategy for expert system maintenance. In *Proceedings of Second Australian Joint Artificial Intelligence Conference (AI88)*, Sydney (Vol. 1, pp. 292–306).
- Compton, P., & Kang, B. H. (1993). Knowledge histories. In *Proceedings of 2nd Australian Cognitive Science Conference* (pp. 139–141). Melbourne, Australia.
- Compton, P., Kang, B. H., Preston, P., & Mulholland, M. (1993). Knowledge acquisition without analysis. In G. Boy & B. Gaines (Eds.), *Knowledge acquisition for knowledge based systems* (Vol. 723, pp. 278–299). Berlin: Springer.
- Compton, P., et al. (1988). Maintaining an expert system. In *Proceedings of Fourth Australian Conference on Applications of Expert Systems*, Sydney (pp. 110–129).
- Compton, P., et al. (1991). Ripple down rules: Possibilities and limitations. In *Proceedings of Sixth Banff Knowledge Acquisition for Knowledge based Systems Workshop (KAW91)*, Banff (Vol. 1, pp. 6.1–6.18).
- Dazeley, R. (2007). *To the knowledge frontier and beyond: A hybrid system for incremental contextual learning and prudence analysis*. Ph.D. dissertation, University of Tasmania, 281 p.
- Diaz Agudo, B., & Gonzalez Calero, P. A. (2001). Classification based retrieval using formal concept analysis. *ICCBR 2001. Lecture Notes in Computer Science*, 2080, 173–188.
- Feigenbaum, E. A. (1977). The art of artificial intelligence: Themes and case studies in knowledge engineering. In *International joint conference on artificial intelligence* (pp. 1014–1029). Stanford: Stanford University.
- Gaines, B. (2000). Knowledge science and technology: Operationalizing the enlightenment. In *Proceedings of the 6th Pacific Knowledge Acquisition Workshop* (pp. 97–124). Sydney: University of New South Wales.
- Gaines, B., & Shaw, M. L. G. (1993). Knowledge acquisition tools based on personal construct psychology. *The Knowledge Engineering Review*, 8(1), 49–85.
- Gennari, J. H., Musen, M. A., Fergerson, R. W., Grosso, W. E., Crubezy, M., Eriksson, H., et al. (2002). *The evolution of protégé: An environment for Knowledge Based Systems Development* (Technical Report No. SMI 2002 0943): Stanford Medical Informatics (SMI).
- Guha, R. V., & Lenat, D. B. (1994). Enabling agents to work together. *Communications of the ACM*, 37(7), 127–142.
- Jansen, R., & Compton, P. (1989). The knowledge dictionary: Storing different knowledge representations. In *Proceedings of the 5th Australian Conference on Applications of Expert Systems* (pp. 143–162). Sydney: CSIRO Division of Information Technology.
- Jenkins, J. J. (1974). Remember that old theory of memory? Well, forget it! *The American Psychologist*, 29(11), 785–795.
- Kang, B. H., & Compton, P. (1992). Knowledge acquisition in context: The multiple classification problem. In *Proceedings of the 2nd Pacific Rim International Conference on Artificial Intelligence*. Seoul, Korea.
- Kang, B. H., & Compton, P. (1994). Multiple classification ripple down rules. In *Third Japanese knowledge acquisition for knowledge based systems workshop*.

- Kang, B. H., Compton, P., & Preston, P. (1995). Multiple classification ripple down rules: Evaluation and possibilities. In *Proceedings of the 9th Knowledge Acquisition for Knowledge based Systems Workshop*. Banff: University of Calgary.
- Kelly, G. A. (1955). *The psychology of personal constructs*. New York: Norton.
- Kelly, G. A. (1970). A brief introduction to personal construct theory. In D. Bannister (Ed.), *Perspectives in personal construct theory* (pp. 1–29). London: Academic Press.
- Killin, J. (1993). Managing and maintaining an operational KADS system. In G. Schreiber, B. J. Wielinga & J. Breuker (Eds.), *A principled approach to knowledge based development* (p. 457). London: Academic Press.
- Leake, D. B. (1996). *Case based reasoning: Experiences, lessons, and future directions*. Menlo Park, CA: AAAI Press/MIT Press.
- Lenat, D. B. (1995). CYC: A large scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11), 33–38.
- Lenat, D. B., & Feigenbaum, E. A. (1988). On the threshold of knowledge. In *Proceedings of the Fourth Australian Conference on Applications of Expert Systems* (pp. 31–56).
- Lenat, D. B., & Feigenbaum, E. A. (1991). On the threshold of knowledge. *Artificial Intelligence*, 47(1), 185–250.
- Lenat, D. B., & Guha, R. V. (1990). *Building large knowledge based systems: Representation and inference in the cyc project*. Reading, MA: Addison Wesley Publishing Company, Inc.
- Lenat, D. B., et al. (1990). Cyc: Toward programs with common sense. *Communications of the ACM*, 33(8), 30–49.
- Maes, P. (1990). Designing autonomous agents. *Robotics and Autonomous Systems*, 6(1), 1–196.
- Menzies, T. (1996). Assessing responses to situated cognition. In *Proceedings of the 10th Banff Knowledge Acquisition for Knowledge based Systems Workshop*. Spain: Catalonia.
- Menzies, T. (1998). Towards situated knowledge acquisition. *International Journal of Human Computer Studies*, 49, 867–893.
- Menzies, T., & Debenham, J. (2000). Expert system maintenance. In A. Kent & J. G. Williams (Eds.), *Encyclopaedia of computer science and technology* (pp. 35–54). New York, NY: Marcel Dekker Inc.
- Newell, A., & Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3), 113–126.
- Patel, V., & Ramoni, M. (1997). Cognitive models of directional inference in expert medical reasoning. In P. F. Feltyovich & R. Hoffman (Eds.), *Expertise in context* (pp. 67–99). Cambridge, MA: MIT Press.
- Piaget, J. (1970). *Genetic epistemology*. New York, NY: Columbia University Press.
- Pittman, K., & Lenat, D. B. (1993). *Representing knowledge in CYC 9*, MCC Technical Report CYC 175 93P.
- Popper, K. R. (1963). *Conjectures and refutations*. London: Routledge and Kegan Paul Ltd.
- Popper, K. R. (1968). Epistemology without a knowing subject. In B. V. Rootselaar (Ed.), *Logic, methodology and philosophy of science (Vol. III)* (pp. 333–373). Amsterdam: North Holland.
- Preston, P., Edwards, G., & Compton, P. (1993). A 1600 rule expert system without knowledge engineers. In *Moving Towards Expert Systems Globally in the 21st Century* (pp. 220–228). (*Proceedings of the Second World Congress on Expert Systems 1993*)
- Preston, P., Edwards, G., & Compton, P. (1994). A 2000 rule expert system without a knowledge engineer. In *Proceedings of the 8th AAAI sponsored Banff Knowledge Acquisition for Knowledge based Systems Workshop* (pp. 17.1–17.10).
- Ramesh, B., & Dhar, V. (1992). Supporting systems development by capturing deliberations during requirements engineering. *IEEE Transactions on Software Engineering*, 18, 498–510.
- Richards, D. (1998a). *The reuse of knowledge in ripple down rule knowledge based systems* (p. 336). University of New South Wales: Sydney.
- Richards, D. (1998b). Ripple down rules with formal concept analysis: A comparison to personal construct psychology. In *Proceedings of the 11th Workshop on Knowledge Acquisition, 1: KAT 4*. Banff: SRDG Publications.
- Richards, D. (2001). Knowledge based system explanation: The ripple down rules alternative. *International Journal of Knowledge and Information Systems*, 5(1), 2–25.
- Richards, D., & Busch, P. A. (2000). Measuring, formalising and modelling tacit knowledge. In *International congress on intelligent systems and applications (ISA'2000)*.
- Richards, D., & Busch, P. A. (2001). Acquiring and applying contextualised tacit knowledge. In *Australian conference for knowledge management & intelligent decision support (ACKMIDS' 2001)*.

- Schreiber, G. (1993). Operationalizing models of expertise. In G. Schreiber, B. J. Wielinga & J. Breuker (Eds.), *A principled approach to knowledge-based development* (p. 457). London: Academic Press.
- Schreiber, G., Wielinga, B., & Breuker, J. (1993). *KADS: A principled approach to knowledge-based system development*. Amsterdam: Elsevier.
- Shaw, M. L. G. (1988). Validation in a knowledge acquisition system with multiple experts. In *International conference on fifth generation computer systems* (pp. 1259–1266). Tokyo: Springer.
- Shaw, M. L. G., & Gaines, B. (1992). Kelly's "geometry of psychological space" and its significance for cognitive modeling. *The New Psychologist*, 23–31.
- Siewiorek, D. P., Bell, C. G., & Newell, A. (1982). *Computer structures: Principles and examples*. New York: McGraw-Hill.
- Stumme, G., Studer, R., & Sure, Y. (2000). Towards an order-theoretical foundation for maintaining and merging ontologies. In *Proceedings of Referenzmodellierung 2000* (pp. 136–149). Aachen: Shaker.
- Suchman, L. A. (1987). *Plans and situated actions: The problem of human-machine communication*. New York, NY: Cambridge University Press.
- Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: A Bradford Book, The MIT Press.
- Vera, A., & Simon, H. A. (1993a). Situated action: A response to reviewers. *Cognitive Science*, 17, 77–86.
- Vera, A., & Simon, H. A. (1993b). Situated action: A symbolic interpretation. *Cognitive Science*, 17, 7–48.
- Vera, A., & Simon, H. A. (1993c). Situated action: Reply to William Clancey. *Cognitive Science*, 17, 117–133.
- Waldrop, M. M. (1990). *Fast, cheap, and out of control*. *Science*, 248, 959–961.
- Widmer, G., & Kubat, M. (1996). *Learning in the presence of concept drift and hidden contexts*. *Machine Learning*, 23(1), 69.
- Wielinga, B. J., Schreiber, T. A., & Breuker, J. A. (1992). KADS: A modeling approach to knowledge engineering. *Knowledge Acquisition*, 4(1), 5–54.
- Wille, R. (1981). Restructuring lattice theory: An approach based on hierarchies of concepts. In I. Rival (Ed.), *Ordered sets: Proceedings of the NATO Advanced Study Institute held at Banff, Canada* (pp. 445–472). D. Dordrecht: Reidel Publishing.
- Wille, R. (1989). Knowledge acquisition by methods of formal concept analysis. In E. Diday (Ed.), *Data Analysis, Learning Symbolic and Numeric Knowledge* (pp. 365–380). New York, NY: Nova Science Publishers.
- Wille, R. (1992). Concept lattices and conceptual knowledge systems. *Computers Mathematical Applications*, 23(6), 493–515.
- WordNet. (2003). *WordNet*. Retrieved January 12, 2006, from <http://dictionary.reference.com/>