Int. J. Human-Computer Studies (1999) **51,** 615–641
Article No. ijhc.1999.0278
Available online at http://www.idealibrary.com on IDEAL®

Towards a knowledge technology for knowledge management

NICK MILTON, NIGEL SHADBOLT, HUGH COTTAM AND MARK HAMMERSLEY
Artificial Intelligence Group, School of Psychology, University of Nottingham, University Park, Nottingham NG7 2RD, UK. email: nrs@psychology.nottingham.ac.uk

Knowledge Management (KM) is crucial to organizational survival, yet is a difficult task requiring large expenditure of resources. Information Technology solutions, such as email, document management and intranets, are proving very useful in certain areas. However, many important problems still exist, providing opportunities for new techniques and tools more oriented towards knowledge. We refer to this as Knowledge Technology. A framework has been developed which has allowed opportunities for Knowledge Technology to be identified in support of five key KM activities: personalization, creation/innovation, codification, discovery and capture/monitor. In developing Knowledge Technology for these areas, methods from knowledge engineering are being explored. Our main work in this area has involved the application and evaluation of existing knowledge for a large intranet system. This, and other case studies, have provided important lessons and insights which have led to ongoing research in ontologies, generic models and process modelling methods. We believe that the evidence presented here shows that knowledge engineering has much to offer KM and can be the basis on which to move towards a Knowledge Technology.

1. Introduction

Large organizations are becoming increasingly aware of the importance of knowledge for efficiency and competitiveness. As such, Knowledge Management (KM) is becoming an ever more pervasive subject within the business community. Although there are a number of viewpoints and approaches to KM, they all centre around the notion that knowledge is a valuable asset that must be managed, and the essence of KM is to provide strategies to get the right knowledge to the right people at the right time and in the right format. However, as the prominence of KM grows, so does the realization that the systematic documentation, distribution and reuse of knowledge are difficult and time-consuming tasks.

To support KM, organizations are using a number of software technologies, including email, document management and intranets. As has been highlighted in the management literature, technology is not the complete answer to the difficulties of KM (Davenport & Prusak, 1998). The need to create the right organizational culture and infrastructure in which knowledge can be created and disseminated is important too. However, few would deny that technology is an important facilitator and can help in overcoming the

1071-5819/99/090615 + 27 $30.00/0

problems of KM. In our opinion, however, most current software tools have more to do with new ways of storing and communicating information than with the actual ways in which people create, acquire and use knowledge. It is no coincidence that these technologies are referred to as Information Technology, since they deal primarily with information, rather than with knowledge.

New methods and tools are needed that can supplement existing technologies, but which are specifically oriented towards knowledge. Such methods and tools could act as a bridge between people and current technologies. This would better support such key activities as knowledge creation, knowledge mapping, knowledge retrieval and knowledge use. We shall refer to this as Knowledge Technology (to distinguish it from Information Technology).

In this paper, we describe some of the research being undertaken to develop a Knowledge Technology for KM. The approach being adopted is based on the idea that a useful starting point for such a technology can be found in the techniques and tools used to perform and assist knowledge acquisition (KA) for the development of knowledge-based systems (KBSs).

Over recent years, KA tools and techniques have been applied and extended for use in broader business contexts than the development of KBSs (Shadbolt & Milton, 1999). Much of our work in this area has been part of a large collaborative project, called SPEDE, with Rolls-Royce, Rover and Parametric Technology as industrial partners. A major aim of the SPEDE project was to modify and validate existing KA techniques and tools for use in organizational initiatives, particularly dissemination of best practice, acquisition of knowledge for Business Process Reengineering and reuse of process-oriented knowledge (e.g. generic process models). As this, and allied work, have progressed it has become clear that ideas and methods from knowledge engineering do indeed show much promise in developing a Knowledge Technology for KM. This paper explores this promise.

The paper is organized into three main parts. The first part, comprising Section 2, covers the rationale for pursuing the use of KA methods for KM. It shows that there are important similarities between the fields of KA and KM. A framework is described with which to identify opportunities for Knowledge Technology and how KA techniques might fulfil these. The second part of the paper, comprising Sections 3–7, introduces the KA techniques and tools being used and describes three case studies. A key section here is Section 4, which is an extended case study on the acquisition of knowledge for an intranet system. Part two concludes with Section 7, which covers lessons learned from the case studies and part of a vision for Knowledge Technology. In the third part of the paper, we describe how parts of this vision are being realized in ongoing work, and briefly cover other approaches in the area. We finish with a summary and concluding remarks.

2. Applicability of KA to KM

There is a growing realization that many of the methods developed in knowledge engineering can be of benefit to KM (Wielinga, Sandberg & Schreiber, 1997). In this section, we review the arguments for pursuing a research programme that aims to evaluate and evolve the methodology, techniques and tools of KA for use in KM.

2.1. SIMILARITY OF ISSUES AND CONCERNS

The most obvious similarity between KA and KM is the importance of treating knowledge as an asset which, if used in the right way, provides for increased efficiency and effectiveness within the workplace. Subsequent similarities between the two domains stem from the nature of knowledge. Of particular relevance are the common problems found across the two domains. Three of the most prevalent problems found within the KM literature are as follows.

1. Organizations contain such a vast amount of knowledge that mapping all of it would be a futile endeavour (Davenport & Prusak, 1998).
2. Tacit knowledge is vital to an organization, yet is very difficult and time consuming to capture and codify (Nonaka & Takeuchi, 1995).
3. Ordinary language is the main form of communication, yet is so full of jargon, assumptions and ambiguities that people often fail to understand what others are trying to say (Allen, 1990).

Knowledge engineers have been dealing with such problems for 20 years. Many of the principles and techniques developed within KA aim to address such problems.

2.2. PRINCIPLES OF KA

Knowledge engineers have developed a number of principles which form the basis for the techniques and tools used for knowledge acquisition and modelling. Given the commonality of issues and concerns in KA and KM, we believe that there are strong reasons for adopting many of these principles in developing Knowledge Technology. A number of key principles are as follows.

- *Systematicity*: knowledge engineers have always placed great importance on systematic analysis. Alongside categories of knowledge and representational frameworks, other schemes have been developed such as categories of experts and expertise. These can be used to focus the selection of methods and the types of knowledge to be acquired (Shadbolt & Burton, 1995).
- *Reuse*: KA is a time-consuming activity. Reuse of previous knowledge models such as generic task models and ontologies has become a vital part of knowledge engineering practice (Schreiber *et al*., 1999).
- *Broad repertoire of techniques*: there is much evidence to suggest that different techniques can be more or less efficient in the types of knowledge they can elicit, the so-called differential access hypothesis (Hoffman, Shadbolt, Burton & Klein, 1995). Hence, to efficiently acquire the knowledge in a domain often requires a range of techniques.

Principles such as these offer potential advantages for use in KM. For example, current KM practitioners are generally unaware of the wide range of techniques available for acquiring knowledge. We believe that they would especially benefit from certain "contrived" techniques, such as concept sorting, shadowing and repertory grids, which can help uncover implicit knowledge (Shadbolt & Burton, 1995). In addition, the KM literature says little about the use of ontologies and generic models for the reuse of knowledge.

A number of other aspects of knowledge engineering are potentially beneficial to KM. These include methods and approaches for scoping, validation, maintenance and auditing. For instance, the modelling approach of CommonKADS includes models such as the Organization Model which have been developed to aid in the scoping activity (de Hoog, Benus, Vogler & Metselaar, 1996).

2.3. KNOWLEDGE TECHNOLOGY

We are using the term "Knowledge Technology", but what do we actually mean by this? To answer this necessitates identifying the requirements of KM and those key activities that might be supported by more knowledge-oriented techniques and tools. The growing literature on KM provides a number of alternative perspectives on this and various typologies of KM strategies and activities have been developed. Examples of these typologies include codification and personalization (Hansen, Nohria & Tierney, 1999); generation, codification and transfer (Ruggles, 1997); conceptualization, reflect, act and review (Wiig, de Hoog & van der Spek, 1997); create, identify, collect, organize, share, adapt and apply (O'Dell & Gayson, 1998). Confronted with these various typologies, we have developed a framework which integrates them. This framework currently comprises five key KM activities each of which could potentially be supported by Knowledge Technology. Figure 1 shows a diagrammatic representation of this framework.

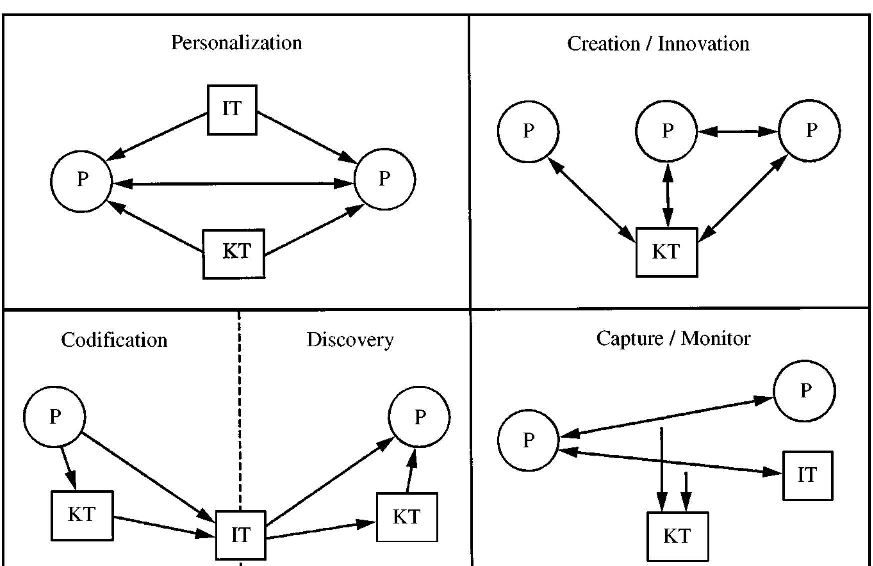

FIGURE 1. Framework for identifying knowledge technology opportunities.
(*Key*: *P = Person*, *KT = Knowledge Technology*, *IT = Information Technology*)

The five key KM activities are described below.

- *Personalisation* is the activity of sharing knowledge mainly through person-to-person contacts. This can be facilitated by investment in current IT systems (Hansen *et al.*, 1999). There is also an opportunity for Knowledge Technology to enhance this process by providing tools to allow employees to communicate more effectively, for example by ensuring that they are clear in their terminology and the ways in which they conceptualize a domain. Promising research here, based in knowledge engineering, includes the APECKS system which allows domain experts to create personal ontologies which are used to identify discrepancies and prompt discussion (Tennison & Shadbolt, 1998).
- *Codification* is the activity of capturing existing knowledge and placing this in repositories in a structured manner. This is the most likely area where a Knowledge Technology based on KA techniques might be applied, the aim being to make the process more efficient, for instance by using generic models, and more effective, by using a range of specialized techniques. This has been the main focus of the extended case study described in Section 4.
- *Discovery* is the activity of searching and retrieving knowledge from repositories and databases, such as using internet and intranet systems. There is potential here for Knowledge Technology to aid in search procedures such and automatic construction of ontologies. Promising research based on knowledge engineering, particularly CommonKADS, is ongoing in this area (Crow & Shadbolt, 1999).
- *Creation/innovation* is the activity of generating new knowledge, vital if an organization is to remain competitive (Nonaka & Takeuchi, 1995). Present technologies have fallen short of providing any significant impact on knowledge creation (Bond & Otterson, 1998), and there seems little doubt that in the foreseeable future this is likely to remain a primarily human endeavour. There is, however, an opportunity for Knowledge Technology to be of assistance, if only in providing sophisticated brainstorming tools. The Unilever case study described in Section 5.2 provides promising evidence for this idea.
- *Capture/Monitor* is the activity of capturing knowledge as people carry on their normal tasks such as interacting with people and computer systems. This is an attractive notion as it does not have the overhead of taking people off-line in order to capture their knowledge. One promising opportunity is to provide knowledge tools that both aid people in their activities and in so doing capture important knowledge, such as providing an audit trail of decision-making. The Andersen Consulting case study described in Section 5.1 hints at the possibilities of this.

2.4. KNOWLEDGE TECHNOLOGY AND KNOWLEDGE ENGINEERING

Using this framework, it is clear that Knowledge Technology should have two important features: (1) it should encompass a wide range of methods, and (2) for most applications, it should be useable by relative novices. The latter point is particularly important, and is perhaps the most problematic for the ideas we are promoting since current knowledge engineering practices are often difficult to learn and apply effectively. Because of this, the training and evaluation of novices in knowledge engineering practices has become one of the major areas of our research.

From a knowledge engineering perspective, we see Knowledge Technology as primarily comprising techniques and associated tools which exist at the "knowledge level", enabling the creation, use and reuse of knowledge models. Using the framework above, it is clear that the most promising area for evaluating this is in support of codification. Based on experience of both the research and practice of KA, and a provisional understanding of the requirements of knowledge codification, we propose certain requirements for technological support.

- Codification requires a number of methods and representations, on which advice should be given using a flexible meta-methodology.
- Codification requires tools that will support software-assisted knowledge acquisition (SAKA). Hence, the basic principles of KA should be easily learned and applied by users of the technology, and enforced in the tools e.g. ensuring unambiguous terminology and consistent semantic categories.
- Best practices should be established by reusing knowledge in the form of ontologies, process libraries and generic process models.

The main research described in this paper takes these three requirements as initial premises on which to evaluate existing KA techniques and tools for the population of an intranet system. The next section introduces the KA techniques and tools being used and the following section describes an extended case study in their application and evaluation.

3. The PC-PACK knowledge acquisition toolkit

As a background to the empirical work and case studies presented later, this section describes the particular KA tools used in the research.

3.1. PC-PACK

PC-PACK is a portable PC-based knowledge engineering workbench comprising an integrated set of software tools and representations which have been found useful in a range of knowledge engineering projects (Zanconato & Davies, 1997; Montero & Scott, 1998). Its inspiration lies in research performed on the ACKnowledge project (Shadbolt, Motta & Rouge, 1993) and later in the VITAL project (Motta, O'Hara, Shadbolt, Stutt & Zdrahal, 1996). PC-PACK was designed and built by Epistemics Ltd (www.epistemics.co.uk). It comprises 14 tools. Description of these and the associated model-driven methodology can be found in O'Hara, Shadbolt and Van Heist (1998).

The subset of tools found to be most useful in our current work is follows.

— *A protocol editor*: this allows the annotation of interview transcripts, notes and documentation in a variety of formats.

— *A hypertext tool*: this allows more sophisticated annotation, and the creation of a linked set of topics or concepts to provide documentation of knowledge fragments acquired by the other tools.

— *A laddering tool*: this facilitates the creation of hierarchies of knowledge elements such as concepts, attributes, processes and requirements.

— *A control editor*: this facilitates the building of process control diagrams.
— *A card sort tool*: this facilitates the grouping of objects or concepts into classes, using the metaphor of sorting cards into piles.
— *A repertory grid tool*: this uses personal construct theory to identify attributes in domain, and to group objects or concepts according to how similar they are with respect to these attributes.
— *An entity-relationship tool*: this allows the user to construct compact networks of relations between concepts.
— *HTML generator*: this publishes hypertext pages in HTML form.

Further descriptions of the use of these tools can be found in Shadbolt and Milton (1999).

3.2. TYPICAL USE OF PC-PACK

The PC-PACK tool set can be used in a variety of ways, depending on the domain and goals in question. Its primary use is by experienced knowledge engineers when producing the knowledge document (conceptual model of expertise) for a KBS.

A typical scenrio is shown in Figure 2. This includes a series of KA sessions between the knowledge engineer and domain experts, progressing from natural to more contrived techniques. Many of the tools in PC-PACK support the contrived KA techniques typically used by experienced knowledge engineers (Hoffman *et al*., 1995).

Each KA session is tape recorded and transcribed verbatim. The transcripts are analysed using the Protocol Editor tool to identify knowledge elements such as concepts, processes, attributes and values. These elements are then classified into hierarchical graphs using the Laddering tool. The Control Editor permits data flow and control flow representations to be constructed. Successive KA sessions become increasingly structured and focus upon appending, modifying and validating knowledge previously elicited. Annotations, explanations and notes are recorded in hypertext pages associated with each knowledge element.

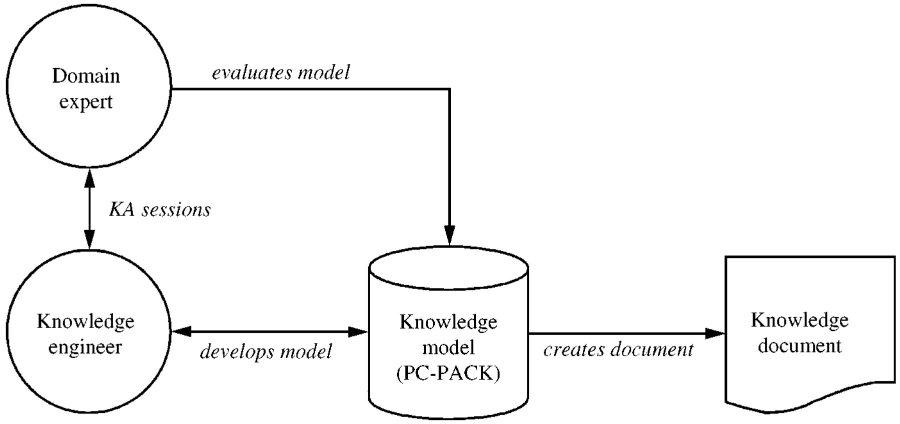

FIGURE 2. Typical use of PC-PACK.

PC-PACK is designed to be used during KA sessions and validation plenary sessions to elicit feedback on previously constructed knowledge models. Such tools as the Matrix tool, Card Sort tool and Repertory Grid tool help elicit new concepts, attributes and classes if required.

A cycle of analysis, model building, acquisition and validation sessions is repeated until the experts and knowledge engineer are satisfied that the ontology constructed in PC-PACK represents that required to fulfil the goals of the project.

This is a brief coverage of one use of PC-PACK. It neither assumes any previous knowledge has been gathered, nor that any generic knowledge can be applied. In reality, the aim would be to reuse as much previously acquired knowledge as possible in the form of ontologies, generic models and process libraries.

3.3. APPLICATION OF PC-PACK FOR KM

PC-PACK has been used extensively within the SPEDE project to assist in the acquisition of best practice and business process knowledge. Elsewhere, it is being used in a number of organizations to support such activities as brainstorming, knowledge auditing, knowledge sharing, group elicitation and decision-making. The following two sections describe some of this work. In particular, the next section contains an extended case study in which the typical use of PC-PACK described above was evaluated for its use by novice knowledge engineers and its efficacy in knowledge codification.

4. The application of KA technology for KM at Rolls-Royce

Rolls-Royce plc is one of the world's leading organizations in the design, development and manufacture of jet engines and is also a leading industrial partner in SPEDE. The company identified the relevance of early findings from SPEDE to its own KM programme and embarked upon a bilateral programme with the University of Nottingham to exploit KA techniques for the rapid development of components of the Rolls-Royce Capability Intranet.

Rolls-Royce's Capability Intranet is intended to become the company's quality system, providing quick and easy access to all the latest information needed by staff in order to complete tasks accurately and reliably, including the capture of lessons learned and evolving best practices. The scope of the Capability Intranet spans business processes, manufacturing processes, product definitions, technical skills and training. It includes quality manuals, working practices, information about technologies and capabilities and specific examples of good (and bad) practice based on real case examples.

A technology transfer programme was conceived in early 1998 whereby established KA tools and techniques could be applied and evaluated within the context of developing knowledge-rich web sites for the Rolls-Royce Capability Intranet. The programme has involved a series of coached projects, typically comprising two company employees on secondment to a special facility based at the University of Nottingham for a period of 12 weeks. Sixteen groups have passed through the facility over the past year amounting to 38 Rolls-Royce employees. At the moment, a total of 12 Rolls-Royce employees are on the programme, though this is expected to rise significantly over the next year as personnel and facilities are expanded.

4.1. AIMS AND OBJECTIVES

4.1.1. Technology development

The primary purpose of the programme is to observe and analyse the use of KA techniques and tools within a real KM project environment, guiding development of methods and technologies. The short repeatable cycle allows variations in approach to be rapidly evolved and deployed.

4.1.2. Skills transfer

The programme aims to equip graduate-calibre participants with no prior experience with the necessary skills to deliver a useful Capability Intranet site containing relevant company knowledge. At the end of each project, the participants take away with them a real appreciation of the relevance of KM and a firm foundation in its practice that can be applied in many situations in their future careers.

4.1.3. Business Benefits

The main aim for Rolls-Royce is the rapid capture of internal best practice for dissemination throughout the organization at minimum cost and in particular without taking valuable experts away from their normal work to write detailed technical papers. By developing an internal pool of knowledge engineering skills, all additional knowledge gained stays within the company rather than leaving with external consultants.

4.2. APPROACH

At the time of writing, this programme operates on a revolving basis with three new projects (six participants) commencing at intervals of approximately six weeks. We are concentrating on the development of content for knowledge-rich web sites, with preceding stages of the overall lifecycle (e.g. project selection) and subsequent stages (e.g. maintenance) remaining substantially outside the scope of this paper.

4.2.1. Project selection and organization

Each project addresses a specific business process which the company's management has prioritized as being in need of knowledge capture and dissemination. Typical projects are in the fields of design, engineering and manufacturing.

The overall business goal for a project is established by the project *Sponsor* (Capability Owner), a senior business manager who allocates resources and eventually approves the finished work for release on the Capability Intranet. The target user community is represented by the *Customer* who plays a hands-on role in managing the project, meeting with the team at least once in a week to guide their work and ensure relevant and focused deliverables. Programme *Participants* include graduate entrants still within their induction period and long-service employees on secondment from their everyday jobs.

Each project typically involves up to 20 other staff either for one-off KA sessions or in ongoing evaluation roles. Three classes of domain roles are distinguished. *Experienced Practitioners* are those currently employed in the target process and in which they are demonstrably proficient. These provide the breadth of initial knowledge content and will benefit from the web site as an aid to following best practice. *Experts* are often found outside the directly targeted activities, for example as managers or in laboratories, and include the authors of existing written manuals. Their depth of knowledge and experience make them valuable critics and reviewers of site content. *Inexperienced Practitioners* are likely to be the immediate beneficiaries of the web site. Their involvement is essential to ensure that the knowledge collected is presented in a way that will be useful and usable.

4.2.2. Training

The first two weeks of each programme comprise an intensive training course in which participants with no prior understanding of knowledge engineering are given the basic skills and knowledge needed to complete the programme. Thereafter, each project is supported by a dedicated *coach* who is an experienced knowledge engineer and/or web designer with a good understanding of the necessary skills for knowledge web development. The coaches provide guidance and advice where necessary and at the same time monitor how the tools and techniques are being used (and misused).

The first week of training covers the core principles of KM together with relevant experiences and lessons learned from the fields of business process analysis, knowledge engineering and web development. The basics of knowledge representation and modelling are covered. Presentations from previous project participants are included to demonstrate knowledge web sites that have already been deployed.

The second week of training reinforces the practical skills of knowledge acquisition, web design and project management which the participants will need to apply in order to realize their goals. We run an extended exercise covering the full lifecycle of a KM project including knowledge elicitation, analysis, modelling, web site design and web page authoring. We also start to address specific project requirements through initial meetings with the customer and group problem-solving workshops including new participants and those who are already half way through their projects.

Throughout the programme, significant value is obtained through peer learning and a regular exchange between teams at different stages of their projects. This has created and perpetuates the programme's own organizational memory of best practice experience about delivering successful KM projects.

4.2.3. Understand and confirm project requirements

Following the initial training fortnight, each team must rapidly understand its sponsor's business goal and evolve a specific set of project aims and objectives. Given the limited time available for the project, much of the preparation work will have been done up front by the customer. Participants spend most of their third week with the customer employing KA techniques to fill in any gaps that might still remain. During this week they will do the following.

- Obtain an overview of the target process, focusing on the business success factors and the important points in the process where actions must be swift and effective.
- Clarify who will use the content of the web site, meeting some of these people to gain an informal appreciation of their roles and how the web site will help their work.
- Identify the important content that will be needed by the target users at these key points in their business cycle, obtain relevant documentation and start to list the local jargon used.

By the end of their third week, the team should have evolved a "Knowledge Storyboard" (Applehans, Globe & Laugero, 1999) forming the basis of a specification for the web site that they will deliver. This is normally a one-page diagram linking the target business cycle to roles, information leverage points (Davenport, 1996) and topics to be included as web content.

The project proposal is normally a formal meeting with the sponsor, customer and other stakeholders at which the team present their plan and deliverables and obtain approval to proceed with the remainder of the project. It is essential that the requirements are stabilized at an early stage in the project if the team is to deliver a good knowledge web at the end.

4.2.4. Web design

The web design arises out of the content (Parker, 1997). This includes a logical web structure for ease of navigation and clear, consistent page layout for effective knowledge transfer to the reader (Siegel, 1997; Tufte, 1997). The design must accommodate many different needs: the web will be used by experts and novices; it will become a how-to guide for employees carrying out the target activities as well as a reference source for those outside the process; it will be delivered over a range of IT platforms and at least for the foreseeable future some of the users will be wholly unfamiliar with web technology.

A draft web design is produced at an early stage in the project, based on the content that users say they need. This will mature over several weeks and reduce the risk of time-consuming changes after the structure has been populated with knowledge content. It also helps to guide the capture of content and keep the team focused on the priorities established by their customer. Target users enjoy reviewing prototype web pages and this helps them to focus on the knowledge content when the time comes for validation.

Delivered web sites must support multiple viewpoints; usually a process viewpoint, a product viewpoint and others as appropriate to the needs of the particular project. This is provided by multiple navigation routes through the site, and also by careful structuring of the text on each page. A set of standard page layouts is evolving for different types of knowledge element, e.g. each activity will be described on a separate page and each of those "Activity Pages" will conform to a standard structure. In the future, it might be possible to dynamically construct pages for different classes of user.

4.2.5. Collate and validate knowledge content

The main section of the project involves the application of KA techniques and tools to develop knowledge content for the web. This begins by mapping the target processes, identifying and exploring tasks/task-groups with an emphasis on the previously prioritized leverage points at which the project aims to make the biggest impact.

Our approach combines the standard model-driven and software-assisted KA techniques embodied in PC-PACK and described in Section 3 together with elements of an object-oriented business analysis approach (SOMA) whose roots also lie in knowledge engineering (Graham, 1998). In particular, we are evolving a series of standard reusable forms ("cards") for describing different types of business objects. This has proven particularly effective in enabling novice knowledge engineers to achieve rapid productivity.

An overall map of the target process, naming and grouping the tasks and indicating their goals, inputs and outputs is summarized in a hierarchical decomposition referred to as the *Process Ladder*. All other relevant concepts are entered in a classification hierarchy referred to as the *Concept Ladder*. There is also a *Glossary* of all the key terminology together with a short definition and synonyms. The *Yellow Pages* is a list of who–knows–what linked to contact information. These will eventually become pages in the delivered web site.

The KA process progresses from identifying the names of relevant knowledge elements and classifying them within the appropriate knowledge ladders through to completing the necessary details about them. The *Activity Card* lists the attributes that are needed to usefully define each element in the Process Ladder, such as goal, controls and resources. A number of other standard knowledge templates are used for elements in the Concept Ladder.

The primary sources of knowledge are experienced practitioners who explain how the job is really done. There is additional reference to a vast array of existing documentation including quality documents, textbooks and training material. At this stage, alternative approaches are evaluated and consensus on best practice is sought. Finally, the experts review the completed cards and confirm that each is factually correct.

4.2.6. Documentation and presentation

In the final period of the project, the knowledge content gathered is converted to web pages using the web structure and page design principles already established. Target users are consulted once again to confirm that the content will indeed be relevant and useful to their work. Associated documentation is written and relevant findings are presented to company stakeholders.

4.3. RESULTS

Participants with no prior understanding of knowledge engineering or web design are now consistently delivering useful webs at the end of 12-week coached projects. The programme has reached this stage by evolution through two main phases involving significant exchange of ideas and perspectives amongst a multidisciplinary research and training support team involving specialists in KA, knowledge modelling, web design, business analysis and project management.

4.3.1. Business benefits

Estimates by Rolls–Royce suggest that web development productivity is around five times higher than other more conventional methods after training time is removed. Not only that, but the knowledge gathered is seen to be better structured, better validated and better focused on user requirements, such that knowledge retrieval and use are enhanced. Furthermore, there is something like a 10 times reduction in the time needed from experts.

In addition, Rolls-Royce managers have noted that the trainee knowledge engineers also acquire significant expertise in the domains on which they have been producing knowledge webs. Several departments now consider that this training programme is the most effective way to train new engineering graduates for roles as effective knowledge workers, quite aside from anything else they might learn about knowledge technology itself.

4.3.2. Trade-off between usability and reusability

During the first six months of the programme, we applied relatively pure KA techniques, producing impressive knowledge models and numerous web pages. However, none of the web sites so created were immediately usable. Since then we have incorporated broader principles from object-oriented business analysis and web design with the result that webs are now in use. We are getting significant structural reuse from knowledge templates (e.g. activity cards) but overall ontological integrity was probably compromised in the process. Our next objective is to obtain better reuse of knowledge content across domains without compromising the quality of delivered knowledge webs.

4.3.3. KA approach

Many of the conventional KA techniques used for KBS development are also beneficial for the development of knowledge web content. However, some of the approaches developed by KA research groups have had to be adapted in order to deliver the tactical needs of our KM projects. In particular, classic KA knowledge models are built with the intention of implementing run-time systems. A knowledge web project is complete when it gets the right information to the right people at the right time so that they can take action and create value (O'Dell & Grayson, 1999). Hence, classic KA knowledge models tend to contain many more entities and relationships than are necessary to achieve this goal.

We have found that it is unproductive to go to experts first or to focus on getting all the knowledge from one or two experts. This caused the webs to get too big and too deep to be immediately useful. By starting with experienced practitioners and developing a consensus view of the domain, we have been able to focus better on satisfying the business requirements for a particular project.

4.3.4. Knowledge structures

The biggest realization leading to quality web sites in that an ontologically pure knowledge model does not automatically deliver usable knowledge webs, even when the contents of the PC-PACK knowledge base can be converted into HTML pages. Furthermore, there is no simple mapping between knowledge elements (processes, concepts, attributes, values) in PC-PACK and web pages. The granularity and structure needed in a usable web page is very different from that produced in the hypertext annotation of a formally structured KBS knowledge document.

4.3.5. Tool utilization

The most commonly used PC-PACK tools are the Protocol Editor, Laddering tool, Control Editor and Hypertext tool. A number of lessons have been learned in the use of these tools by relative novices in the context of knowledge codification. Since these have relevance to our ongoing research, they will be covered later in Section 6.

4.4. CONCLUSION

We believe this case study clearly demonstrates that knowledge engineering techniques and tools do form the basis of a Knowledge Technology that is becoming increasingly easy to learn and use by novices within real KM situations. However, just as KA alone does not result in successful KBSs, nor does it make a complete knowledge web. The knowledge model helps assure completeness and reusability (integrity). Web design assures usability and usefulness, the key elements necessary to ensure user acceptance (Davis, 1993).

5. Two further case studies

The extended case study described in the previous section illustrates one important use of KA techniques and tools in support of KM activities. This is not the only use of Knowledge Technology but is the main focus of this paper since we have achieved more work to date within this area. To further illustrate the potential of KA techniques to be the basis of a Knowledge Technology, two further case studies are briefly covered in this section.

5.1. ANDERSEN CONSULTING

Andersen Consulting is a global management and technology consulting organization whose mission is to help its clients change to be more successful. Use is made of KA tools for both KBS development and for KM activities. Montero and Scott (1998) give a full description of the use of PC-PACK in systems development, so we will focus on its other role: as an aid to knowledge auditing, sharing and decision-making.

During such processes as planning and design, numerous decisions are made to arrive at the best solution. However, much of the knowledge of why a particular decision was made and what alternatives were discarded (i.e. the "audit trail") is often lost. This is because those involved either forgot what happened or failed to keep adequate records. Hence, stakeholders in such decisions such as managers, decision implementers and future decision makers have no access to knowledge which could be vital for them. A tool that can capture the rationale behind decisions would, therefore, be an important aid to this aspect of KM.

Andersen Consulting make particular use of the PC-PACK Requirements Laddering tool for this purpose. It allows them to construct representations showing various requirements and associated issues, positions, arguments and assumptions. This also allows them to define and visualize alternative decisions as sets of positions on issues. This tool is used in combination with another PC-PACK tool, the Protocol Editor. According to Andersen Consulting, the result is not only useful for modelling but is a "powerful brainstorming tool for the interaction between expert and analyst" (Montero & Scott, 1998).

5.2. UNILEVER

Unilever is an international organization with over 300 000 people world wide and having over 1000 brands. The importance of knowledge and Knowledge Technology to

Unilever's continued success is recognized at the highest levels of management. Unilever currently use PC-PACK in both in its "traditional" role during the development of KBSs, and in more non-technological areas of KM. It is the latter which we briefly describe below.

The main use of PC-PACK to support KM at Unilever is in "knowledge workshops". By projecting what would normally appear on the VDU onto a large screen, a group of experts can interact and represent the current state of their thinking. In this way, they can build and critically analyse an ontology of the domain in question. Tools, such as the Laddering and Matrix tools, are used dynamically to add, delete and modify elements. In this way, the group can reach a consensus on the main knowledge categories and relationships in the domain. Since this is not known beforehand, these sessions are as much about "brainstorming" as they are about knowledge capture. One can see this as an example of what Leonard-Barton (1995) refers to as "creative abrasion" leading to creativity and innovation.

One use of such a method is to identify certain product ranges and their associated attributes. This provides knowledge for such activities as production restructuring or devising new marketing strategies. A more elaborate use is to build ladders of typical problems, their causes and possible solutions. Attributes of these are then identified using the Repertory Grid and Matrix tools. Automatic translation of ladders and matrices produce various graphical representations. These can be used to summarize what causes and solutions apply to a particular problem.

In summary, key aspects of the Unilever approach include the following.

- Support of Knowledge Auditing using reporting and analysis tools.
- Creation of multiple views of the same domain.
- Identification of what is known, and (perhaps more importantly) what is not known.

6. Lessons learned

The most important lesson to come from the case studies is that KA techniques and tools are proving successful for a number of KM applications. However, much development is still required. For codification, in particular, our initially naïve hopes that ideas from KA could be exported relatively unchanged to satisfy the requirements of KM have been tempered by the reality of the situation. The lesson is that KA is a good starting point, not a ready-made answer.

At Unilever and Andersen Consulting, experienced knowledge engineers and consultants are involved in running sessions and operating tools and techniques. The nature of the Rolls-Royce case study, however, involves novice knowledge engineers performing these activities. Two main lessons have resulted from this. First, with careful training and coaching, novices can become competent practitioners in knowledge acquisition and modelling within a relatively short period of time. Second, there exist a number of problems commonly encountered by novice users of KA techniques as supported by PC-PACK. Such problems highlight important gaps that exist between current KA tools and those that may be required for certain aspects of Knowledge Technology. Ideas to fill these gaps form part of a vision for Knowledge Technology. Common problems and parts of this vision are described below.

6.1. TERMINOLOGY

Use of jargon is an aspect of any field, and none more so than knowledge engineering. Since PC-PACK was designed to be used by expert knowledge engineers, it uses many terms that are unfamiliar, abstract and non-obvious to novice users. In addition, certain terms are used synonymously in different KA techniques, such as attribute, construct and dimension. Such terminological difficulties are both daunting and confusing to novices. What Knowledge Technology requires is a better vocabulary. One possible solution we find helpful is to use the grammatical categories of linguistics. For instance, when performing transcript analysis, novices find it easier to think of concepts as nouns, values as adjectives and processes as comprising active verbs (i.e. "doing words").

6.2. CATEGORIZING KNOWLEDGE

A major problem found by novice users is in deciding what knowledge category a knowledge element should be classified within. Novices spend far too long agonizing over such decisions, and then often produce inconsistent results, such as classifying processes and attributes as concepts, and classifying concepts as attributes. Of course, some of these problems involve deep epistemological issues, such a when is a concept a concept and not an attribute. However, when a novice user assigns "car" to attribute and "colour" to concept, this is usually a mistake rather than a sophisticated ontological decision. To a certain extent, the use of grammatical categories mentioned above can help alleviate such problems. Another aid is to give users a predefined generic ontology (described in Section 7). A further idea we find useful is to give novices template sentences to check their classifications (e.g. "the ⟨attribute⟩ of ⟨concept⟩ is ⟨value⟩"). In the longer term, automatic categorization based on ontologies is an attractive solution. To operate effectively, this may require methods to handle the apparent fact that categories can vary their ontological status across contexts and domains.

6.3. INCONSISTENT RELATIONSHIPS

Another very common problem is found when novice users build hierarchies. They often use inconsistent relationships in the same hierarchy, such as using "part-of" relationships in a taxonomy and "is-a" relationships in a compositional hierarchy. Since the importance of inheritance in a taxonomic hierarchy is lost on most novices, they also fail to see the importance of maintaining a consistent semantic structure. As before, the use of a predefined generic ontology can greatly ease this problem. Another solution is to provide users with hierarchies that can accommodate mixed relationships. This is to be incorporated in the next version of PC-PACK.

6.4. CHOICE OF TOOL

Novice users find it difficult to know when and how best to use the many tools in PC-PACK. Having a suite of tools with many representational formats is a major benefit to the expert user, but tends to be a drawback for the novice user overwhelmed by the choice. The common response is to persist with 2 or 3 tools and ignore the rest. One solution to this problem is to use a model-assisted approach known as Generalized Directive Models (GDM). This method is supported in PC-PACK with a tool called the

GDM tool, the theory and operation of which has been described in detail elsewhere (O'Hara, Shadbolt & van Heijst, 1998). Though useful to expert knowledge engineers, novices users find the GDM concept difficult to grasp. Part of the problem is terminological difficulties. Another problem is to do with the nature of the generic models that are supported in the GDM tool. This is a pervasive problem with PC-PACK that is described next.

6.5. REPRESENTING THE PHYSICAL DOMAIN

A problem with PC-PACK, and KA tools in general, is the emphasis on problem-solving structures. For example, the GDM tool provides a library of problem-solving models, but does not support or advise on the modelling of physically oriented knowledge (e.g. how physical objects and phenomena are transformed by physical processes). This has been found to be a drawback with current PC-PACK representations since many of the applications of KM in industrial contexts involve physical objects and operations. We have found that an ontological commitment that divides the world into a "knowledge domain" and a "physical domain" allows one to continue using models and approaches developed in knowledge engineering, whilst supplementing these with a new set of models and approaches. The use of this categorization scheme in developing PC-PACK is discussed further in Section 7.

6.6 KNOWLEDGE STRUCTURES

An important use of PC-PACK for KM is to facilitate the population of company intranets. The rigour of knowledge engineering methodologies provides knowledge that is well-structured, focused to the customer requirements and more easily maintained. A problem, however, occurs if users directly transfer the knowledge structures of PC-PACK to the web page structures of an intranet. Since the knowledge structures of PC-PACK have been designed for the development of expert systems they are based on certain fundamental approaches useful in rule-based systems (e.g. frame-based description models). The resulting knowledge structures, however, are rarely the way in which web pages should be organized to satisfy user requirements. Indeed, having an intranet that is structured in a strictly hierarchial fashion is often too rigid for the more flexible hypertext-based approach. Having said that, the representations commonly used in hypermedia systems are compatible with those used in AI (Woodhead, 1991). Hence, what users require is advice and support in mapping the knowledge structures built in PC-PACK to the structures of company intranets. We find that users develop their own methods of doing this. For example, they use the Laddering tool to map out a hierarchy for their web pages, then use the Hypertext tool to provide hypertext links to the knowledge structures. This is an example of how Knowledge Technology must allow flexible use of tools within certain constraints, even if this means users can violate the correct semantic or syntactic nature of the tools.

7. Ongoing research work

In many respects, the work with Rolls-Royce is much more than a case study demonstrating that tools and techniques from knowledge engineering can be of benefit to KM.

Its real importance is as an ongoing research environment to facilitate the development and evolution of prototype ideas into established tools and techniques. In this section, we describe the ongoing research work and how it might address the lessons learned.

7.1. ONTOLOGIES AND GENERIC MODELS

Ontologies are becoming a major research area in knowledge engineering. Our main work here has been in developing ontologies and generic models for use in business environments (Cottam, Milton & Shadbolt, 1998). There is ongoing research to assess and adapt these in order to be of more benefit to KM, especially for more efficient and effective knowledge codification. As has been described, a number of the problems found in the Rolls-Royce case study may be reduced by providing users with a predefined ontology.

From initial research, it became clear that the ontology and associated generic models should be consistent with a number of popular management approaches: process-oriented approaches (Hammer & Champy, 1993), systems viewpoints (Senge, 1990), and cyclical models of business activities (Deming, 1993). It should also be consistent with the knowledge-intensive modelling approaches developed within knowledge engineering, such as CommonKADS.

7.1.1. Knowledge-Physical Ontology (KPO)

To satisfy these objectives, the knowledge-physical ontology (KPO) was developed. This was originally based on a process-oriented approach, which took as its foundation a taxonomy of over 250 generic business processes (APQC, 1998). Analysis comprised both a bottom-up and top-down approach, the result being a taxonomy of business processes described in Cottam *et al*., (1998). Further developments to this produced the present ontology.

A fundamental notion of the KPO is that at the highest levels, a business can be viewed as comprising two process domains, the knowledge domain and the physical domain. This dichotomy of the world into a knowledge (or information) domain and a physical (or material) domain is long established in philosophy and has a number of advantages over other high-level ontological categories such as those used in CYC (Sowa, 1995). Within these two high-level domains are a number of sub-domains. For instance, the knowledge domain can be seen as comprising a mental domain (i.e. that of human agents) and a simulation domain (e.g. computer models, scientific laws, mathematics), as well as a number of meta-domains. The physical domain comprises such sub-domains as the physical-model domain and the pilot-process domain.

A business process either resides within a single domain, or crosses domain boundaries. For example, processes such as IMPORT and EXPORT cross a single domain boundary. Processes that cross the boundary between the knowledge domain and physical domain include IMPLEMENT and EXECUTE (knowledge to physical) and OBSERVE and MEASURE (physical to knowledge). Figure 3 shows a number of generic processes within and between two process domains.

An important class of processes represented in the KPO are those that combine simple generic processes to form loops or spirals. Examples of process loops include the well-known feedback loop of systems theory and Karl Wiig's "knowledge management

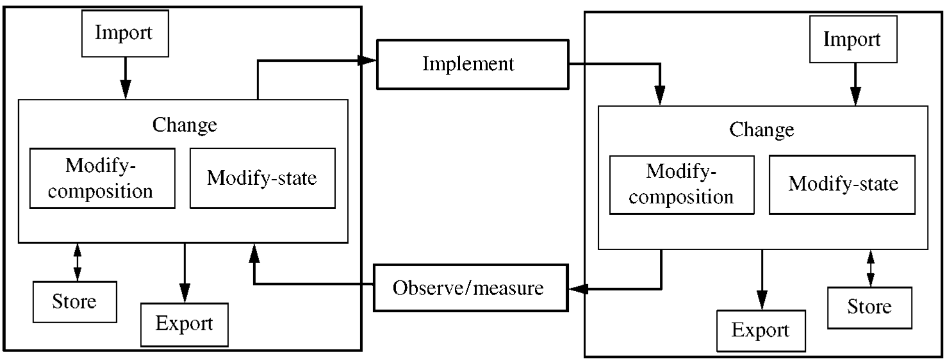

FIGURE 3. Example of generic business processes.

life cycle" model (Wiig *et al*., 1997). Examples of process spirals include the Deming cycle of quality management (Deming, 1993) and the risk-driven spiral model of software development (Boehm, 1988).

7.1.2. Generic Ladder

Based on the KPO, a Generic Ladder has been created that aims to provide a high-level structure for use with PC-PACK. Although the highest-level classes in the ladder are aimed at being general across all domains, the lower-level classes are aimed primarily at the business domain.

Figure 4 shows the upper levels of the generic ladder. At the highest level, the taxonomy comprises four main classes: concepts, attributes, tasks and relationships. The concept, attribute and task sections of the ladder have sub-classes associated with the physical domain, the knowledge domain or combinations of the two.

7.1.3. Current and future use of the KPO

The KPO is currently being taught in the work with Rolls-Royce as one perspective with which to view the business domain. Participants are then free to use the KPO as and when they feel appropriate.

A future use of the KPO would be to provide the basis for a grammar of business processes, which could supplement the grammer of problem-solving methods already used in the PC-PACK GDM tool. As discussed by Pentland (1995), the modelling of business processes can greatly benefit from a grammatical approach. The grammar, once established, would provide a model-driven approach to process capture which may involve improvements to the Control Editor tool in order to aid the novice user. For example, generic template processes could be displayed, prompting the user to enter domain-specific terminology.

Another future use of the KPO could be to facilitate knowledge creation. There are a number of ways in which this may take place. One is to support analogical reasoning, which appears to play key role in creative design (Goel, 1997; Bond & Otterson, 1998). It is believed that this can be aided by the symmetry that exists across the knowledge and

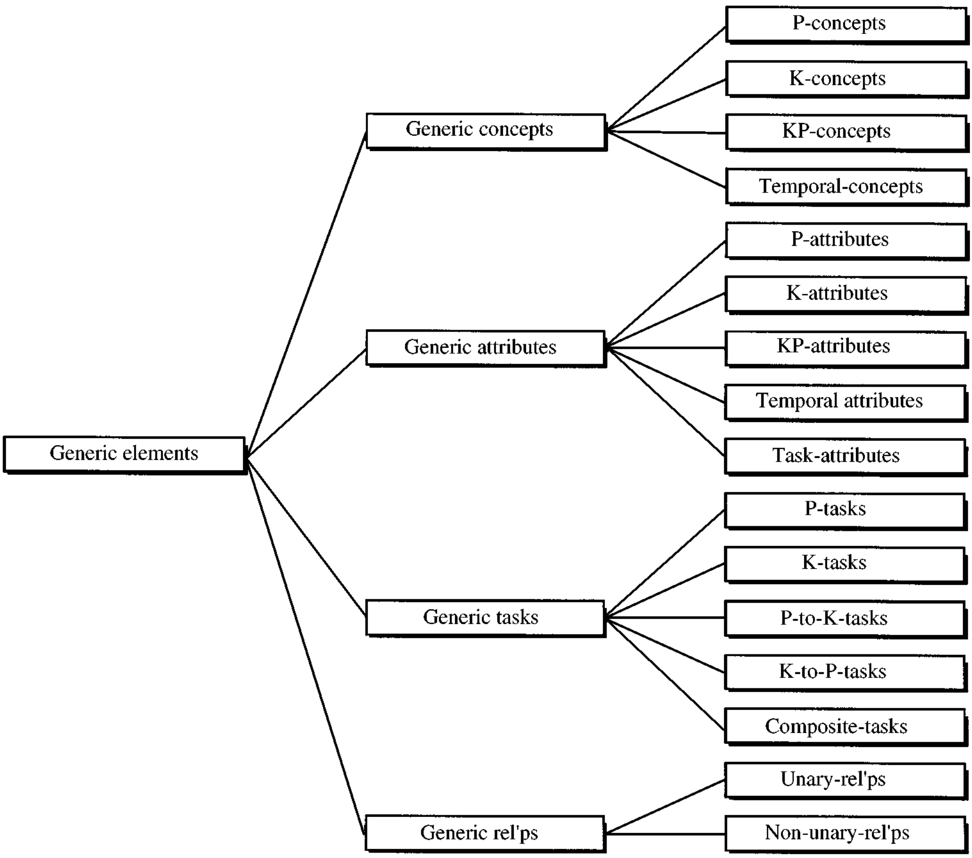

FIGURE 4. Upper levels of the Generic Ladder.
(*Key*: *P = Physical, K = knowledge, KP = Knowledge & Physical*).

physical domains, i.e. the same basic processes occur in both domains (e.g. import, modify-composition, store). Describing the processes that occur in the two domains in a more general language may prompt the user to notice where good practices in one domain might apply in the other domain. Another way in which the KPO may aid creative thinking is by embodying it in brainstorming tools, in a similar way to the Unilever case study presented earlier.

7.1.4. Current and future use of the Generic Ladder

The Generic Ladder is currently used in the Rolls-Royce work to aid in classifying and categorizing concepts, attributes and tasks, i.e. as a paper-based adjunct to the use of the Protocol Editor and Laddering tool. When developed further, the Generic Ladder will form the basis for a structured thesaurus allowing the next version of the Protocol Editor to automatically identify knowledge elements. That is, all terms already contained in the thesaurus would be automatically highlighted in the Protocol Editor, saving a lot of time and effort of novice users. Following this, the next version of the Protocol Editor may incorporate a simple semantic analyser to infer the class of new knowledge elements by linking thesaurus entries with key terminology. This would be a step towards producing a fully automatic knowledge classifier.

7.1.5. Goal-directed templates

The Generic Ladder is used to provide taxonomic information for all the knowledge elements in a domain. In order to represent the other relationships between knowledge elements, each node on the Generic Ladder has an associated template. This is a frame-based representation showing what relationships hold between that node and other nodes in the ladder, for example the "Dataflow" relation would link a task in the physical domain with a concept in the knowledge domain. A number of template types, called goal-directed templates, have been developed as described in Cottam, Shadbolt and Milton (1998).

7.2. GENERAL PROCESS ONTOLOGY AND THE PROCESS KNOWLEDGE EDITOR

Since SPEDE and the work with Rolls-Royce both embrace a process perspective, a second ontology has been developed that aims to identify the important distinctions between objects in process models. The general process ontology (GPO) has been developed that lists a set of basic process concept and relation types. These can then be used for the definition of knowledge relevant to process description. This is a base ontology that is designed to support process KA. It forces the distinction between basic process concepts, such as activity, data, role and organizational group. It also defines the relation types that can hold between these entities, such as sequence, dataflow, performs role and requires skill. These basic types have been determined based on an extensive analysis of existing ontologies, enterprise and process modelling formats, and the knowledge requirements of the analysis activities conducted on process descriptions. Additional basic types are being considered for inclusion in the GPO and it is envisaged that it will undergo modifications following further evaluation.

The Process Knowledge Editor is a software tool that explicitly supports the use of the GPO for defining the structure of process knowledge prior to acquisition (as well as the acquisition task itself). Basic types are selected from the GPO and act as the building blocks for the construction of well-structured process models. A core subset of the GPO is hardcoded into the Process Knowledge Editor, and the user may select additional basic types from the GPO dependent on the current knowledge acquisition task. The core subset of the GPO includes the basic process concept types of *Activity*, *Data*, *Role*, *Skill*, *Resource*, *Result* and *Decision*. It also includes the basic relation types of *Sequence*, *Dataflow*, *Performed by Role*, *Uses Resource*, *Requires Skill*, *Refers to Data* and *Acts as Support Role*.

The user may enter concepts of any basic concept types and then arrange them in a hierarchy. An example of an *Activity* hierarchy can be seen in Figure 5, and is based on knowledge acquisition that was carried out at Rolls-Royce on the design process for jet engine compressor blades. The leaf nodes of the hierarchy represent the activities that are actually performed, whilst the higher-level nodes represent categorizations of the lower-level concepts.

The Process Knowledge Editor uses process modules as a means of decomposing the overall process description. The module view is used to enter relations between

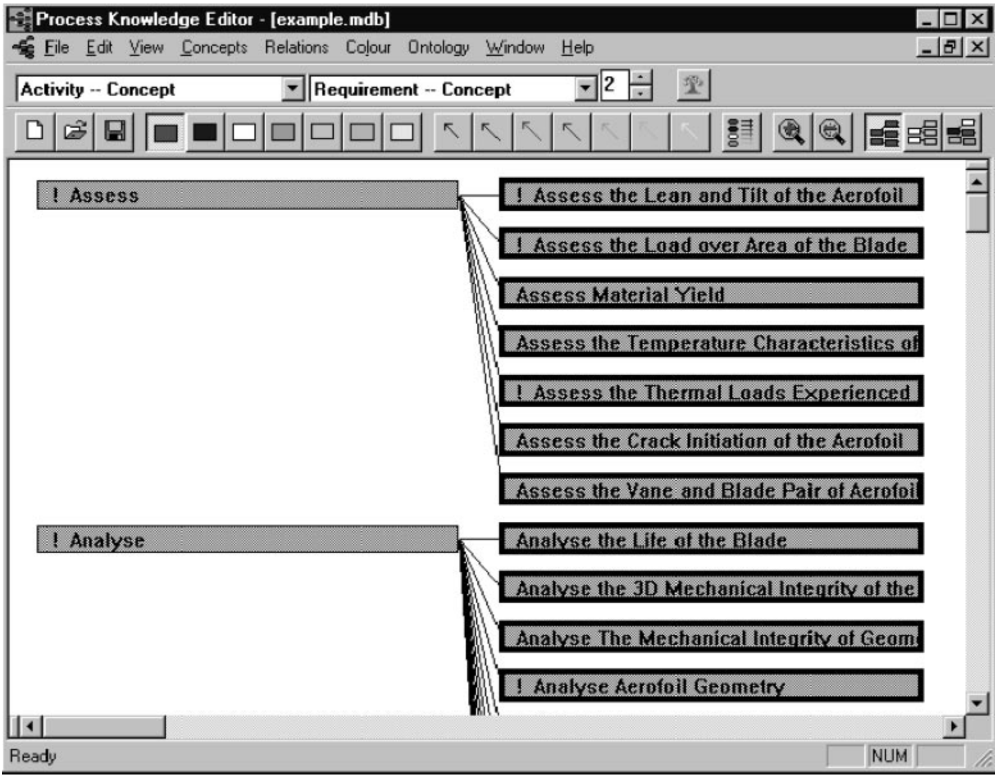

FIGURE 5. Activity Hierarchy.

concepts. Figure 5 shows the leaf node activities highlighted in bold. This indicates that a module had been defined for that activity. A process module from the compressor blade design example can be seen in Figure 6. The basic concept types are indicated in double-angled brackets. The use of modules allows the user to focus on a particular section of interest. In the example given, the modules are indexed by *Activity*, though they could be indexed in other ways if required. Modules facilitate information hiding and also support reuse by decomposing the process description.

The overall behaviour of a relation type can also be viewed. In the compressor blade design example, the network of *Dataflow* relations is of particular interest. The network is constructed merely by entering relations via the separate modules. This results in a very complex and heavily interconnected *Dataflow* network, that can be viewed in the Process Knowledge Editor. Viewing the overall *Dataflow* network gives a strong feel for the data interfaces to the process as a whole and also where the major data bottlenecks exist.

The Process Knowledge Editor also supports the hypertext documentation of the concepts comprising the process description. Automated generation of hypertext documentation can also be performed via interrogation of the process description.

The combination of GPO and Process Knowledge Editor allows the adoption of a goal-directed approach to process knowledge acquisition. The content and structure of

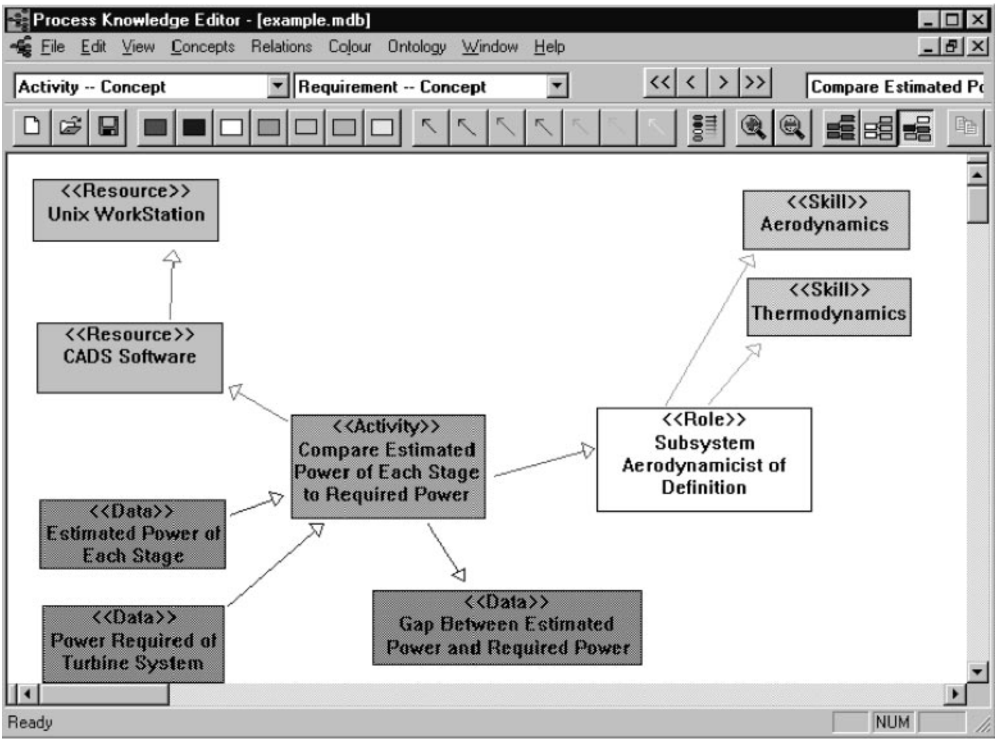

FIGURE 6. A process description module.

a process model depends on the intended use of the model, such as the analyses that will be performed on the model or the knowledge that the model must contain for implementation of a re-engineered process. The intended use of process knowledge thus defines high-level knowledge requirements. The linking of high-level requirements to the type of knowledge to be acquired can be explicitly supported within the Process Knowledge Editor using goal-directed templates.

The Process Knowledge Editor and accompanying ontologies are currently being evaluated in the work with Rolls-Royce.

8. Other approaches

In this section, we briefly cover other work that has investigated the application of methods and tools from knowledge engineering to KM.

The closest work to ours is research based at the University of Amsterdam. For instance, Weilinga, Sandberg and Schreiber (1997) have argued that many of the methods and models built for knowledge engineering purposes can be of benefit to KM. In particular, they propose that problem-solving models from the CommonKADS library can be linked to specific models of KM activities to provide a model-driven approach to KM. They illustrate these points by references to examples of modelling then redesigning knowledge records held by museums. We believe this is an important step forward towards a Knowledge Technology, particularly the model-driven approach that exploits existing problem-solving models and ontologies. However, for the industrial domains which have been our main concern, we believe the traditional knowledge engineering emphasis on problem-solving models needs to be supplemented with physically oriented models. Hence, the use in the KPO and generic ladder of both the knowledge domain, which can reuse models and ontologies developed in knowledge engineering, and the physical domain. We believe the addition of the physical domain is an important step towards a broader approach to KM.

An approach that included models of the physical domain was the KACTUS project (Laresgoiti, Anjewierden, Bernaras, Corera, Schreiber & Weilinga, 1996), which focused on methods and tools to enable reuse and sharing of technical knowledge. In particular, KACTUS aimed to provide a link between standard product models (i.e. STEP) and models from knowledge engineering, particularly CommonKADS. A major principle underlying this approach was that the representation of real-world objects always depends on the context in which the object is used. KACTUS referred to this context-dependent representation as a "viewpoint", and aimed to describe explicitly the "viewpoints" associated with a particular knowledge representation. This was seen as being the key to enabling knowledge reuse.

This KACTUS emphasis upon "viewpoints" has strong similarities with the goal-directed templates of our ontology work. The basic process object types used within our ontologies also have a number of similarities with the basic types found within the TOVE enterprise ontology (Gruninger & Fox, 1995). However, the TOVE ontology was designed for quite a different purpose: to support deductive reasoning over enterprise databases.

A large-scale process-oriented initiative that aims to promote enterprise knowledge reuse is the Process Interchange Format (PIF). PIF is a neutral interchange format that has been developed to enable the exchange of process descriptions across a wide variety of software tools such as "process modelers, workflow software, flow charting tools, planners, process simulation systems and process repositories" (Lee, Gruninger, Jin, Malone, Tate & Yost, 1998). The creation of a neutral interchange format for process descriptions reflects the general desire for enterprise knowledge reuse. However, it addresses the issue of reuse more from the perspective of software tools exchanging existing knowledge rather than the perspective of knowledge acquisition.

9. Summary and conclusions

The paper has presented work that is extending and modifying existing knowledge engineering tools and techniques to make steps towards a Knowledge Technology for KM. For organizations, we see the main purpose of Knowledge Technology as providing solutions to certain key problems associated with KM, such as scoping what knowledge is to be captured and disseminated, dealing with tacit knowledge, and facilitating better communication. Based on a review of the growing literature in KM, a framework has been developed that has identified five key KM activities, each of which provides opportunities for Knowledge Technology. Techniques and associated tools from knowledge engineering are being developed and evaluated for a number of these opportunities. Results from case studies indicate that techniques and tools developed for KA purposes can be of real benefit to such KM activities as codification, creation/innovation and capture/monitor.

Due to the vast amounts of knowledge in a organization, a prime requirement for Knowledge Technology is that it should be capable of wide use throughout the organization by relatively novice users. Enabling this has provided an important aspect of the research agenda being investigated as part of the SPEDE project. The work with Rolls-Royce provides a unique opportunity to test and evolve tools and techniques to become more easily learned and applied by novices. In parallel with this is the development of a training programme that aims to transform employees with little or no experience of knowledge engineering and KM to become effective knowledge managers within a relatively short period of time. We believe that the importance of the training aspects of a Knowledge Technology based on knowledge engineering should not be underestimated due to the novelty and complexity of the principles and techniques involved. Hence, we see this as a key aspect of our research which feeds into the development of the technology itself.

Results from the Rolls-Royce case study strongly suggest that KA techniques have much to offer to those performing knowledge codification. Use of the PC-PACK knowledge toolkit and associated KA techniques has proved a more efficient and effective way of populating intranet pages than conventional approaches. In addition, PC-PACK is being used by organizations for other activities such as brainstorming and knowledge auditing.

Most research work in the area of Knowledge Technology takes a primarily top-down approach by focusing on the development of strategies, ontologies and generic models. In comparison, our research has taken a more pragmatic approach that combines both top-down and bottom-up approaches. This has enabled a close collaboration between the identification of practical problems, the development of immediate solutions and the evolution of research ideas. Thus, the ontologies and modelling approaches being developed aim to be both useful for, and usable by, relative novices.

It is still early days in the development of Knowledge Technology, yet there is every indication that knowledge engineering is well placed to develop such a technology provided it pays close attention to the KM requirements of organizations.

Most of the work described has been supported by an IMI grant (GR/K78737) from the Engineering and Physical Sciences Research Council of Great Britain. We thank all those members of the SPEDE project who contributed to the work presented here. In particular, we thank Paul Riley and the Rolls-Royce employees who have been associated with the SPEDE project for their ideas and assistance.

References

Allen, T. (1990). People and technology transfer. *The International Centre for Research on the Management of Technology*, August 1990, 5.

Applehans, W., Globe, A. & Laugero, G. (1999). *Managing Knowledge: A Practical Web-Based Approach*, Reading, MA: Addison-Wesley.

APQC (1998). Website at http://www.apqc.org/.

Boehm, B. W. (1988). A spiral model of software development and enhancement. *Computer*, **21,** 61–72.

BOND, P. & OTTERSON, P. (1998). Creativity enhancement software: a systematic approach. *International Journal of Technology Management*, **15,** 173–191.

COTTAM, H. D., MILTON, N. R. & SHADBOLT, N. R. (1998). The use of ontologies in a decision support system for business process re-engineering. *Proceedings of IT&KNOWS Conference of the 15th IFIP World Computer Congress*. Vienna/Budapest.

COTTAM, H. D., SHADBOLT, N. R. & MILTON, N. R. (1998). Acquiring knowledge for business process re-engineering. *Proceedings of AAAI-98 Workshop*. Menlo Park, CA: AAAI Press.

CROW, L. R. & SHADBOLT, N. R. (1999). Knowledge Engineering and Software Agents. *Proceedings of the AAAI Spring Symposium on Intelligent Agents in Cyberspace*. Stanford, CA: AAAI Press.

DAVENPORT, T. H. (1996). Finding the information that matters. *CIO Magazine*, 1 June 1996.

DAVENPORT, T. H. & PRUSAK, L. (1998). *Working Knowledge: How Organizations Manage What They Know*. Boston, MA: Harvard Business School Press.

DAVIS, F. D. (1993). User acceptance of information technology — system characteristics, user perceptions and behavioral impacts. *International Journal of Man-Machine Studies*, **38,** 475–487.

DE HOOG, R., BENUS, B., VOGLER, M. & METSELAAR, C. (1996). The CommonKADS organization model: content, usage and computer support. *Expert Systems with Applications*, **11,** 29–40.

DEMING, W. E. (1993). *The New Economics*. Cambridge, MA: MIT Centre for Advanced Engineering Study.

GOEL, A. K. (1997). Design, analogy and creativity. *IEEE Expert*, **12,** 62–70.

GRAHAM, I. (1998). *Requirements Engineering and Rapid Development: An Object-Oriented Approach*. Harlow, UK: Addison-Wesley.

GRUNINGER, M. & FOX, M. S. (1995). The logic of enterprise modelling. In J. BROWN & D. O. SULLIVAN, Eds. *Reengineering the Enterprise*, pp. 83–98. London: Chapman & Hall.

HAMMER, M. & CHAMPY, J. (1993). *Reengineeering the Corporation*. New York: Harper Collins.

HANSEN, M. T., NOHRIA, N. & TIERNEY, T. (1999). What's the strategy for managing knowledge. *Havard Business Review*, 106–116.

HOFFMAN, R., SHADBOLT, N. R., BURTON, A. M. & KLEIN, G. (1995). Eliciting knowledge from experts: a methodological analysis. *Organizational Behavior and Decision Processes*, **62,** 129–158.

LARESGOITI, I., ANJEWIERDEN, A., BERNARAS, A., CORERA, J., SCHREIBER, A. Th & WIELINGA, B. J. (1996). Ontologies as vehicles for reuse: a mini-experiment. *Proceedings of the 10th Knowledge Acquisition for Knowledge-Based Systems Workshop*. Banff, Alberta, Canada.

LEE, J., GRUNINGER, M., JIN, Y., MALONE, T., TATE, A. and YOST, G. (1998). The process interchange format and framework. *Knowledge Engineering Review*, **13,** 91–120.

LEONARD-BARTON, D. (1995). *Wellsprings of Knowledge*. Boston, MA: Harvard Business School Press.

MONTERO, L. & SCOTT, C. T. (1998). Improving the quality of component business systems with knowledge engineering. *Proceedings of the 11th Knowledge Acquisition for Knowledge-Based Systems Workshop*. Banff, Alberta, Canada.

MOTTA, E., O'HARA, K., SHADBOLT, N., STUTT, A. & ZDRAHAL, Z. (1996). Solving VT in VITAL: a study in model construction and knowledge reuse. *International Journal of Human-Computer Studies*, **44,** 333–371.

NONAKA, I. & TAKEUCHI, H. (1995). *The Knowledge-Creating Company*. New York: Oxford University Press.

O'DELL, C. & GRAYSON, C. J. (1998). *If Only We knew What We Know: The Transfer of Internal Knowledge and Best Practice*. New York: The Free Press.

O'HARA, K. & SHADBOLT, N. R. & VAN HEIJST, G. (1998). Generalised directive models: integrating model development and knowledge acquisition. *International Journal of Human-Computer Studies*, **49,** 497–522.

PARKER, R. C. (1997). *Guide to Web Content and Design*. New York: MIT Press.

PENTLAND, B. T. (1995). Grammatical models of organizational processes. *Organisation Science*, **6,** 541–556.

Ruggles, R. (1997). Knowledge tools: using technology to manage knowledge better. Website at http://www.businessinnovation.ey.com/mko/html/toolsrr.html.

Schreiber, A. Th., Akkermans, J., Anjewierden, A., De Hoog, r., shadbolt, N., VanDe Velde, W. & Wielinga, B. (1999). *Engineering and Managing Knowledge: The CommonKADS Methodology*. Cambridge, MA: MIT Press (forthcoming).

Senge, P. M. (1990). *The Fifth Discipline*. New York: Doubleday.

Siegel, D. (1997). *Creating Killer Web Sites*. Indianapolis, In: Hayden Books.

Shadbolt, N., Motta, E. & Rouge, A. (1993). Constructing knowledge-based systems. *IEEE Software*, 34–39.

Shadbolt, N. R. & Burton, A. M. (1995). Knowledge eliciation: a systematic approach. In J. R. Wilson & E. Corlett, Eds. *Evaluation of Human Work: A Practical Ergonomics Methodology* (2nd revised edn). London: Taylor and Francis.

Shadbolt, N. & Milton, N. (1999). From knowledge engineering to knowledge management. *British Journal of Management*, (forthcoming).

Sowa, J. F. (1995). Top-level ontological categories. *International Journal of Human-Computer Studies*, **43,** 669–685.

Stewart, T. H. (1998). *Intellectual Capital*. London: Nicholas Brealey Publishing.

Tennison, J. & Shadbolt, N. R. (1998). APECKS: a tool to support living ontologies. *Proceedings of 11th Knowledge Acquisition for Knowledge-Based Systems Workshop*. Banff, Alberta, Canada.

Tufte, E. R. (1997). *Visual Explanations: Images and Quantities, Evidence and Narrative*. Cheshire, CN, USA: Graphics Press.

Wielinga, B., Sandberg, J. & Schreiber, G. (1997). Methods and techniques for knowledge management: what has knowledge engineering to offer? *Expert Systems with Applications*, **13,** 73–84.

Wiig, K. M., De Hoog, R. & Van Der Spek, R. (1997). Supporting knowledge management: a selection of methods and techniques. *Expert Systems with Applications*, **13,** 15–27.

Woodhead, N. (1991). *Hypertext and Hypermedia: Theory and Application*. Wilmslow, UK: Sigma Press.

Zanconato, R. & Davies, A. (1997). Design for a knowledge-based decision support system to assist near Littoral Airborne Early Warning (AEW) operations. *Proceedings of the 4th Joint GAF/RAF/USAF Workshop in Human Computer Teamwork*. Kreuth, Germany.