A methodological approach to supporting organizational learning

Paul Mulholland, Zdenek Zdrahal, John Domingue and Marek Hatala

Knowledge Media Institute, The Open University, Walton Hall, Milton Keynes MK7 6AA, UK. email: {p.mulholland, j.b.domingue, z.zdrahal, m.hatala}@open.ac.uk

Ansgar Bernardi

German Research Center for Artificial Intelligence (DFKI) GmbH, Knowledge Management Group, Postfach 2080, D-67608 Kaiserslautern, Germany. email: ansgar.bernardi@dfki.de

Many organizations need to respond quickly to change and their workers need to regularly develop new knowledge and skills. The prevailing approach to meeting these demands is on-the-job training, but this is known to be highly ineffective, cause stress and devalue workplace autonomy. Conversely, organizational learning is a process through which workers learn gradually in the work context through experience, reflection on work practice and collaboration with colleagues. Our approach aims to support and enhance organizational learning around *enriched work representations*. Work representations are tools and documents used to support collaborative working and learning. These are enriched through associations with formal knowledge models and informal discourse. The work representations, informal discourse and associated knowledge models together form on organizational memory from which knowledge can be retrieved later. Our methodological approach to supporting organizational learning is drawn from three industrial case studies concerned with machine maintenance, team planning and hotline support. The methodology encompasses development and design activities, a description of the roles and duties required to sustain the long-term use of the tools, and applicability criteria outlining the kind of organizations that can benefit from this approach.

© 2001 Academic Press

KEYWORDS: ▮; ▮; ▮

1. Introduction

Learning as an integral part of working occurs naturally in the vast majority of organizations. Workers will, among other things, share stories, offer advice, adapt to new tools, and copy the behaviour of respected colleagues. This kind of learning, that is responsive to changes, generally social in nature, and conducted in the work context, can be described as organizational learning. Over the last few decades, a number of theories and definitions of organizational learning have been offered, though there are a number of common threads running across them. First, organizational learning primarily occurs within a community of practice (Lave & Wegner, 1991). These communities do not necessarily equate to the formal team structure of the organization. Second, organizational learning occurs in collaborative activity, which requires an appropriate collaborative culture (Argyris & Schön, 1996). A third common feature of organizational

learning is the reuse and reworking of past experience and solutions (Argyris & Schön, 1996; Levitt & March, 1988). Organizational learning builds on, questions and modifies previous solutions and ideas.

Organizational learning is often quite different from the processes employed by the organization when training staff. Traditional approaches to training involve strategists within the company identifying or predicting the skill gaps between the company they need to be and the current competencies of their staff. The gap is then bridged by conventional training methods that extend staff competencies to meet company requirements. Traditional approaches to training have a number of shortcomings. First, training has been shown to be ineffective. Detterman (1993) reported that 90% of the training was not transferred to the job, wasting approximately $90 billion per year in the US. This is believed to be due to training occurring outside the normal context of work, and therefore difficult to transfer. Second, this approach also leads to stress among employees, as skills have to be developed intermittently and abruptly (Ivergård, 1998). Third, the process is almost solely top-down, where senior personnel must identify what needs to be done and how, devaluing the role that workers can play in determining their own development.

Essentially, traditional approaches to training concentrate on providing theoretical knowledge, which is vital, but training often fails to support the worker in translating what they have learned into practical knowledge. There is a clear contrast between having theoretical knowledge about some subject, and knowing how to put that knowledge into practice. We are not claiming that organizational learning should replace all forms of training, but that learning through practice is undervalued and ill supported in many organizations. Although organizational learning can appear to be an attractive proposition, it is limited in its scope and pace if unsupported. There are a number of accounts of successfully integrated working and learning. A good example is Orr's (1990) account of photocopier engineers sharing "war stories" about faults in different kinds of machines and how they were solved. These stories were freely shared within a community of practice, and became a very effective unofficial learning resource, quite different to the official manuals provided by the company. Although this is an account of effective learning within a community, organizations require learning that can readily span space, time and community boundaries, providing knowledge on demand to those who need it.

The methodology and tools described in this paper were developed in the ENRICH project (Sumner, Domingue, Zdrahal, Millican & Murray, 1999). ENRICH was a European Union ESPRIT funded project investigating how organizational learning could be harnessed and catalysed for the service of the individual worker, workplace community, and the organization as a whole. Nurturing learning in an organization is a laudable aim in that it improves conditions for, and motivation of, workers. It is highly effective as lessons learned in context can be more easily applied to work. The ENRICH approach not only supports learning, but also captures newly evolving knowledge in the organization as a consequence of this support, to be stored in, and retrieved from an organizational memory. The organizational memory provides a knowledge repository that can exist independent of those who evolved and created that knowledge in the workplace. This knowledge therefore becomes a tangible asset of the organization. Knowledge only found in the heads of the people that work in the organization is a human asset of the company. As these workers become more proficient and knowledgeable, their value as

a company asset increases. But unlike some other assets of the company, workers are free to leave and take their knowledge with them. An example is Davenport and Prusak's (1998) account of the Taurus design team at Ford Motors. When the company later tried to learn lessons as to why the Taurus design had been so successful, they were unable to find out. The designers had all gone, and taken their unrecorded, uncaptured knowledge with them. Replacing these valuable and knowledgeable employees is expensive. Even after new employees have been hired and trained at great expense, the company will still often have failed to regain the knowledge that the workers took with them. However, capturing knowledge translates a transient human asset that can be lost, into an infrastructure asset of the company (Brooking, 1999).

Our approach to organizational learning has been derived from three case studies which we will refer to as the mining case, planning case and hotline case. The mining case was conducted within an industry facing increasing competition within a shrinking sector, requiring enormous productivity gains in order to survive. This led to a reliance on a small number of complex mining machines. As reserve capacity had been eliminated, the fast diagnosis and repair of machine faults was vital to the company. Successful maintenance of the machines relied heavily on experienced personnel. The retrieval and use of valuable experience was negatively affected by three factors. First, staff turnover was high, meaning valuable know-how was lost. Second, due to adverse conditions down the mine, each machine was worked on by up to seven shifts per day, in order to maintain 24 h activity. Knowledge sharing across shifts was not adequately supported. Third, the coal mine continuously employed three mining machines at distant geographical locations, and similar machines were used at other mines. The ability to share knowledge across locations would be of great benefit to the company. In view of these circumstances, the objective of the mining case study was to ensure that the continuously changing workforce can profit from long-time experiences and improve their maintenance performance. To this end, the mining case study integrates person-independent conservation of experience by structured documentation, dynamic evolution of knowledge by communication in the actual work context, and continuous quality assurance of the knowledge repository by established feedback and revision loops.

The commencement of the planning case coincided with a company-wide initiative to introduce a new approach to team planning. This was formed around a paper-based workbook being rolled out throughout the company that explained and supported a team planning methodology. This workbook was intended to guide teams at all levels within the organization. The key aims were to create consistency between plans in the same unit of the organization (for example the plan of a department being consistent with the plans of teams within it) and also use the plans as a vehicle to share knowledge across the organization. They also wished to combine team plans with their evolving best practice archive, permitting access to best practice during the planning process.

Within the hotline case, the business of the unit was to provide customer support. The product being supported was a complex controller to be customized and integrated within the customers' existing infrastructure to support the automation of certain tasks. Customer queries were initially dealt with by engineers providing front-line support. More complex queries that could not be resolved by front-line staff in a reasonable amount of time were passed to the back-office, who would then liase directly with the customer. Although the company had in place some support for tracking queries and

recording solutions, the organization wished to further support knowledge sharing in three ways: the sharing of ideas among front-line staff; front-line access to new solutions developed in the back-office; and direct customer access to solutions through an FAQ web site to decrease the burden on hotline staff.

Further aspects of the three case sites will be described later in the paper. Initially, an overview of the organizational learning theories on which the approach is based will be presented. This is followed by an overview of the ENRICH approach including a definition of important concepts and a generic description of the ENRICH lifecycle model. The applicability criteria will then be introduced, that describe the kind of organization that can be supported by the ENRICH approach. Further details of the approach will then be presented in the form of the developer and user methodologies.

2. The nature of organizational learning

Supporting organizational learning requires an understanding of the kinds of learning and knowledge found in the workplace. Learning can be characterized on the level of the individual, group and organization. For example, a new idea may be initially developed by an individual within a team. This idea may then be modified and elaborated through collaboration within the team, and become incorporated into revised work practices. This may become known and adopted by other teams and eventually even lead to changes in company policy. At each of the steps in this hypothetical example, the knowledge that started as a new idea takes different forms, plays different roles, and requires different kinds of support. Our approach focuses on four types of learning: reflection-in-action (at the individual level); domain construction and community of practice learning (at the group level) and perspective taking (at the organizational level). We will now describe our theoretical framework as to the nature of work-based learning and its relation to knowledge.

2.1. INDIVIDUAL LEVEL LEARNING

Our theoretical basis for identifying the occurrence of individual learning is Schön's (1983) theory of "reflection-in-action". Schön claims that a professional worker uses and applies a range of knowledge and skills during their duties (whether they be a doctor, engineer, architect, manager, etc.) in a seemingly effortless way, that would bewilder an untrained onlooker. He refers to this as "knowledge-in-action". Reflection-in-action can be triggered when the knowledge that a worker uses and applies during their professional duties produces an unexpected outcome. Often this will be because applying their professional knowledge has fallen short of dealing with the problem they were trying to solve. This breakdown in their work requires them to reframe the problem—view the problem from a different perspective. This allows them to reflect on their actions, question their assumptions and reshape their activity, whilst in the process of their everyday work.

Reflection-in-action highlights an important distinction between two kinds of knowledge applied within the organization: tacit and explicit. Tacit knowledge is used effortlessly and routinely as part of work, but cannot readily be described. Explicit knowledge is open to scrutiny and can be discussed among colleagues.

Reflection-in-action can be thought of as a process of generating and modifying explicit knowledge related to action that was guided by tacit knowledge. Individual learning and knowledge creation is therefore closely related to the interplay between tacit and explicit knowledge. Within the ENRICH approach, we aim to capture the creation of individual knowledge by focussing on reflection-in-action which sits on the boundary between tacit and explicit knowledge.

2.2. GROUP LEVEL LEARNING

We interpret group level learning within organizations by using the notions of domain construction (Sumner, 1995), and community of practice learning, which draws on a number of sources (Schön, 1988; Lave & Wegner, 1991; Raelin, 1997; Cook & Brown, 1999). We intend the term "group" to encompass both officially sanctioned groups within the organization, often referred to as teams, and also unofficial groups emerging around shared interests, usually referred to as communities of practice (Lave & Wegner, 1991).

Domain construction describes the processes by which groups elaborate and evolve a shared knowledge of their domain over time. Domain construction takes many forms: negotiating and defining new domain concepts, making links between concepts, and making links between concepts and the tools of work. These new concepts will evolve over time. Initially, a new concept will start to be used in informal dialogue between workers. During these dialogues, debates will occur as to the utility of the concept, its precise meaning, and its relation to other domain concepts. If the concept is found to be useful, then a consensus will be reached, and the concept becomes a more fixed and formal part of the domain. During domain construction, concepts gradually evolve from being informal and ill defined, to being formal and well defined. Essentially, domain construction describes the process by which a group explicates and formalizes shared knowledge. Within ENRICH, we provide tools and methodology to support domain construction within groups.

Conversely, we use the term community of practice learning to describe the evolution and adoption of shared tacit knowledge at the group level. Group tacit knowledge refers to the emerging conventions by which a group works and shares ideas. For example, in Orr's (1990) account of engineers sharing war stories, even though the stories make explicit some knowledge of a machine and how to repair it, there are implicit factors governing what makes a good story and how a story should be told. These conventions are not governed by explicit rules but emerge through social contact within the group. Nonaka and Tackeuchi (1995) in their account of the knowledge creation spiral, emphasize the importance of making knowledge explicit so that it can be shared. Traditionally, in the knowledge management field, tacit group knowledge was accorded a lower status, as it is hard to capture, quantify and share. Cook and Brown (1999) suggest that this account underplays the importance of tacit group knowledge and argue for its equal status. They use the term "organizational genre" to describe the nature of group tacit knowledge. They offer the example of different communication mechanisms (e.g. email, memo) becoming used for particular kinds of purpose without any explicit rules being elaborated. Beyond the use of and choice of media, community of practice learning also encompasses the way different kinds of artefacts develop shared meaning. This is particularly noticeable in communities of designers. Schön (1988) describes how within

a design community, design classifications emerge, which he terms "design types". Designers use these as a communal method for talking about particular types of design, without ever defining precisely what they mean. A design type such as "end entrance building" serves as a tag to a set of design ideas. Similarly, Raelin (1997) highlights the importance of group tacit knowledge in his model of work-based learning, and describes the role it plays in intertwining knowledge of theory and practice. Following from these observations, within the ENRICH approach we aim to recognize the importance of group tacit knowledge and learning, which we refer to as the community of practice learning.

2.3. ORGANIZATIONAL LEVEL LEARNING

Perspective taking (Boland & Tenkasi, 1995) develops a model of how learning can occur across distinct communities in an organization. They argue that innovation in knowledge intensive companies, that comprise specialized communities, necessitates the ability to make perspectives within a community and be able to take the perspectives of other communities. Perspective making is the process by which a community develop their own domain knowledge and practices, which is closely related to our conceptualization above of group level learning. Perspective taking describes the process by which communities recognize, use and evaluate the perspectives of other communities as part of their work, and use these to reflect on their own work practices. An example of perspective taking would be for a community to read and reflect upon a document prepared by another community. Understanding the document and its potential relevance requires the reader to engage in the world of the community that prepared the document. This process of perspective taking permits a community to then view and evaluate themselves from another perspective. Additionally, perspective taking may be supported by forums that allow communities to develop a shared perspective, to support the sharing of knowledge across community boundaries.

Perspective taking relates to Argyris and Schön's (1996) account of double-loop learning. They distinguish single-loop learning which drives incremental organizational change, from double-loop learning which presents itself as a more radical strategic change. Double-loop learning requires an organization not only to question work practices and what they have learned, but also question how they have learned. By appreciating the perspectives of other communities, it becomes possible not only to reflect on the efficacy of ones current work practices, but also, on the way in which one actually evaluates ones success. Within ENRICH we develop support for taking other community perspectives and building shared cross-community perspectives.

3. The ENRICH approach

Two key features of the ENRICH approach are that organizational learning is supported around enriched work representations, and that these are incorporated within a lifecycle model that supports the continued evolution of the organizational memory. We will introduce the ENRICH approach by explaining what is meant by work representations and their enrichment, and then situate these within the generic lifecycle model.

3.1. WORK REPRESENTATIONS

Our approach focuses on *work representations*. First, we will describe what we mean by a "work representation" (abbreviated as WR). The nearest single word in English to the meaning we wish to express is "document", but our definition of WR differs from the generally accepted definition of document in two ways. First, the term document, can imply a paper centric or bureaucratic way of working, or imply typical document genres such as reports and memos. Brown and Duguid (1996) in their account of documents in the digital world, included artefacts such as radio and television programmes within their classification of what constituted a document. Similarly, we wish our definition of WR to include any artefact that can be contained in a web page, including audio, video and virtual reality. In order to adopt the ENRICH approach, these WRs need not necessarily exist in advance, but there must be clear potential for them, and a vision as to how they can be incorporated into the current work scenario.

Second, in terms of the role that it plays within the workplace we intend WR to have a narrower definition than document. Within our definition, a WR must be the following.

- A tool of knowledge work rather than just paperwork.
- Embedded in work activities that require (or would benefit from) communication.

Be a tool of knowledge work rather than just paperwork. This condition actually combines together two further conditions. The WR should be a tool, and that tool should support a knowledge intensive activity. These two aspects will be considered in turn. First, the WR needs to be a tool to undertake some work tasks. That is, the WR should be integral to the work, not merely a form to be filled in after this work has been completed. If a task within the organization is supported by a document, then consider how that task would be affected if the document was taken away. If the performance on the task would be unaffected (or even improved) by the absence of the document, then it is clearly not a WR. That document falls into the category of paperwork, which although possibly valuable to the organization as a whole, is a chore for the team, that takes them away from their core duties. We are not claiming that paperwork is a bad thing, but we are claiming that we cannot support organizational learning around paperwork. Second, the work with which the WR is tied, should be "knowledge work". This work has no simple solutions, and requires the worker to interpret information, reflect and learn. Examples of knowledge work are planning, designing, project specification and bid preparation. The WR should be used as a tool within this work. The WR should therefore structure, coordinate and record this work, capturing ongoing problem-solving activity as well as final solutions.

Be embedded in work activities that require communication. The kinds of activities that WRs solve are conducted within teams or communities. Communication and collaboration occurs around and through the WR. In work such as planning, the plan evolves over time, and the current plan acts as a focus for team members to present new ideas and negotiate a shared meaning. In our planning case, the WR is a tool used to support team value planning activities. During planning meetings, the tool is used to record potential objectives and search for relevant best practice descriptions. Between meetings, it is updated with ongoing work toward the plan and to display current targets. The associated team members who use the planning tool, do not spend all their working day
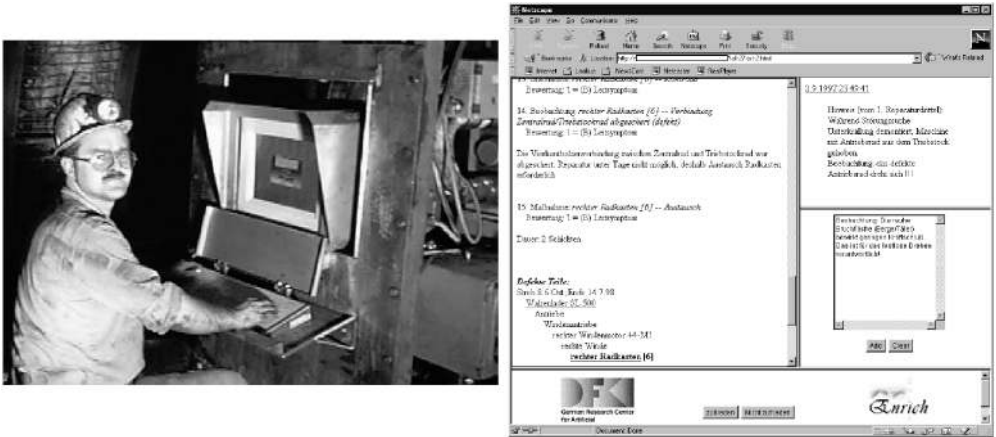
FIGURE 1. A PC in the environment of a black coal mine (left) and the WR showing a retrieved fault record with the discussion area.

planning or using this particular WR, but it supports a crucially important knowledge intensive aspect of their work.

The mining case illustrates that a suitable integration of the WR with ongoing work and communication processes can require additional effort. Within the mining case, the WR was based on existing paper-based hand-written fault records. Experienced miners working on the machines documented faults that had occurred and actions taken for their repair, filling paper books with written documentation. Whenever a complex fault had occurred requiring work across several shifts, experts would gather around the fault records to discuss the solution and their experiences, modifying and extending the written notes where appropriate. While the communication and documentation processes were thus well established, the results had serious shortcomings due to the characteristics of the employed material. First, observations were written down without pre-defined structure or form. As a consequence, many vital details required to make sense of and reuse the observation, were often omitted because nobody documented what seemed to be obvious in the actual fault situation but proved not-so-obvious in rétrospect—when it was to late. Second, information contained in fault records was not classified in a way that supported a worker in easily accessing useful past experience. Third, the information was not accessible where it was perhaps most needed—down the coal mine. The new tools solved the structure and classification problems by coupling the documentation with formal representations of the work context; details of this *enriching* will be illustrated below. Online access in the coal mine supported asynchronous communication between miners (see Figure 1). Using the tool, they could also gain easy access to useful support material (see Figure 2). What used to be sleeping documentation thus became a living part of everyday work and consequently, a well-suited and dynamic WR. The new tool required the use of a special computer with an explosion-safe casing and keyboard. Furthermore, the interface had to be simplified as a computer mouse cannot be used in such conditions, but this extra effort was worthwhile as it provided a tool for work, not paperwork to be completed afterward, and provided more
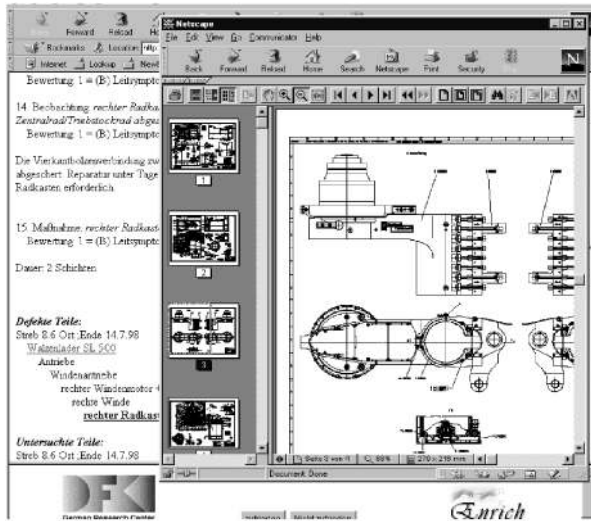
FIGURE 2. Additional information accessible through the WR.

opportunity for collaboration to occur around fault records. Further details can be found in Schmit, Bernardi and Becker (2000).

3.2. ENRICHED WORK REPRESENTATIONS

WRs support ongoing work toward some final product, such as a plan, design or bid. Many important ideas are developed and evaluated during that work. Numerous lessons are learned along the way. Once the work has been completed, these important ideas and lessons will reside in the team that carried out the work, but may not be so discernible in the final solution description. For example, a team constructing a plan will, along the road to a final solution, makes numerous approximations. These approximations will be evaluated. Some will be extended and form part of the final solution, others will be rejected. During the process, the way in which the team conceive, propose, evaluate and extend these solutions will itself be questioned and evolved. By the end of the process, the team will not only have a solution, but will know far more about the process by which a solution is created. As a knowledge asset of the company, this knowledge of the process, is at least, if not more, important than the solution itself (Kidd, 1994). A solution therefore becomes a far more valuable knowledge asset if it is contextualized with the process by which it was created. *Enriching a work representation means tightly coupling it with its context.*

This context takes two forms. First, the WR has an informal context in which it is used. Much of this informal context in the form of emails, sketches and further documents emerges from collaboration and communication among colleagues working on the task to which the WR relates. Second, a WR has a formal context in terms of how the knowledge it contains is related to the formal knowledge of the team that uses it, and the formal knowledge of the organization as a whole. These two aspects of context will be considered in turn.
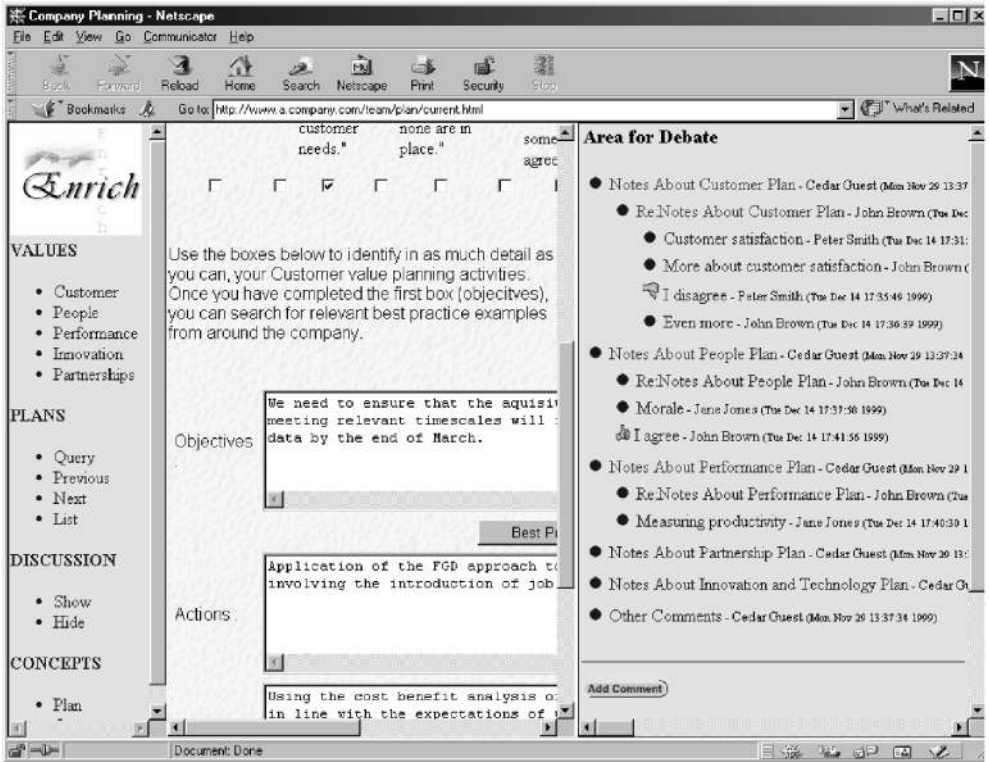
FIGURE 3. Work representation used to support team planning.

Informal context. WRs support ongoing work toward some task, but often the WR itself often only captures ideas that have reached some level of maturity and acceptance among the team. Many other activities and information sources surround the WR. Although these information sources play a vital role in the problem solving process, they become detached from the final solution. From a knowledge reuse perspective, the solution is therefore of less value. Additionally, during the lifecycle of the task in which the WR is being used, the infrastructure to support distributed collaboration between colleagues is often inadequate. Communications via, for example email and telephone are not conducted within the immediate context of the problem. The context has to be established by the collaborators. They have to explain what they wish to discuss and how that fits into the overall task. ENRICH provides facilities to couple collaboration and supporting information sources, with the WR to which they refer. Figure 3 shows a work representation from our planning case, attached to its informal context (modified for reasons of confidentiality). The centre of the window shows the team planning tool itself. This was based on an existing-paper-based planning tool. The right-hand part of the interface contains the associated informal context. Information in the form of asynchronous discussion and links to other online resources is situated on the right-hand side. This discussion space is structured around the planning objectives of the team, each of which is related to a part of the WR. This space is used to talk around, and link

resources to, the WR. The layout of the tool and how it is used was strongly influenced by the teams themselves, and provides an insight into their way of doing things. Clearly, the tool does not capture all informal context surrounding some task, but does make accessible an important part of the informal context.

Formal context. Every WR contains knowledge that relates to the knowledge of the team and organization, and the knowledge contained in other WRs and documents. Understanding the nature of the formal context involves an appreciation of the distinction between formal and informal knowledge. This distinction applies to knowledge that has been externalized into some form, whether as language, models, sketches, plans, etc. The vast majority of externalized knowledge falls into the category of informal knowledge, as it is contextual, contains (perhaps deliberate) ambiguities, and requires interpretation in order to derive its meaning. Conversely, formal knowledge is an explicit, unambiguous representation of semantics that can be captured in the form of a knowledge model. For example, in the mining case, the core of the knowledge model is a representation of the universally accepted structure of the mining machines. This part of the knowledge model relates the components of the machine in a hierarchy using the part-of relation. Each component is described by a set of attributes, among them the part number, a description, and a link to relevant background material. Based on this model, the computer can deduce for example that a fault in a sub-component influences its super-components (formal knowledge). For the human user, the model can serve as an index to the (non-formalized) background documentation.

Each knowledge model has two parts. First, the framework for the domain is described in an ontology (Gruber, 1993). The ontology makes concepts and relations of a domain explicit. Second, specific representations tied to particular WRs are placed in a knowledge base. The generic representation of mining machine parts is stored in an ontology. Specific machine parts, their individual characteristics and relations to other parts is represented in the knowledge base. An explanation of how the knowledge model is developed can be found in Section 5.3.

Within an ENRICH application, the role of the knowledge model is to establish viable connections between descriptions of work entered in the current WR and other information sources. The formal context serves to situate the WR and the knowledge it contains within the wider knowledge of the company. Our construction of knowledge models is supported by the WebOnto tool (Domingue, 1998), which allows the collaborative editing and viewing of models via a conventional web browser. Figure 4 shows a screen snapshot of WebOnto. The figure shows a small component of the ontology related to team planning. The left-hand part of the background window contains a list of all the classes defined within the model. The right-hand part of the window allows the class hierarchies and instances to be graphically browsed and edited. Each node within the ontology contains slots and values of the ontology.

3.3. THE LIFECYCLE MODEL

Our generic lifecycle of how WRs are managed and associated with the formal knowledge model, and the required roles within the host organization is shown in Figure 5. A second important type of document is what we term best practice. In this context, best practices are archived WRs or other kinds of documents identified and
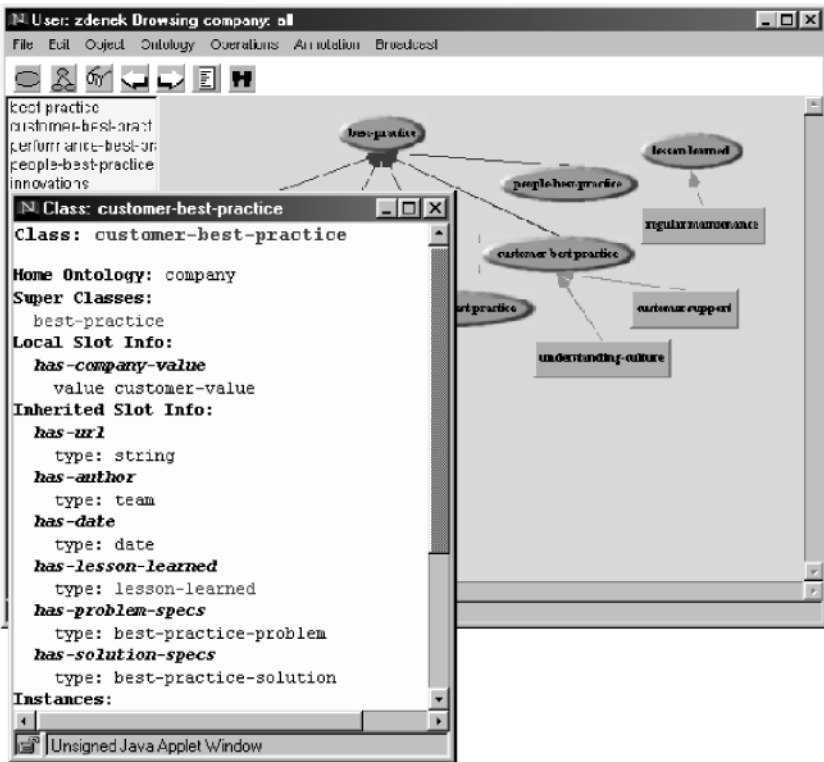
FIGURE 4. Formal context used to support team planning.

redesigned by the organization to be useful to a wider community. A WR is used by collaborating workers in the accomplishment of some task, such as team planning. A query engine proactively or reactively retrieves relevant best practices and previous team WRs with associated discussion, based on the information entered by the team. Best practices and previous WRs are retrieved from an archive according to semantic associations represented in the knowledge model. The knowledge model is comprised of an organization knowledge model and a team knowledge model. The organization knowledge model contains knowledge applicable across teams such as company structures and policies. The team knowledge model represents knowledge local to teams such as specialized knowledge of their particular domain. The archive of WRs and best practices, together with the organization and team knowledge models constitute the organizational memory.

Important issues related to the WR and ongoing progress are recorded in the discussion space. This is the process by which the WRs become enriched with informal context. Two new roles are required to support the semantic enrichment of WRs and the identification of best practice. The association of WRs to the team knowledge model, and modifications of the team knowledge model are undertaken by a *local developer*. Drawing on the work of Gantt and Nardi (1992), a local developer is generally a member of the team, who has acquired knowledge of, or been given some training in the ENRICH
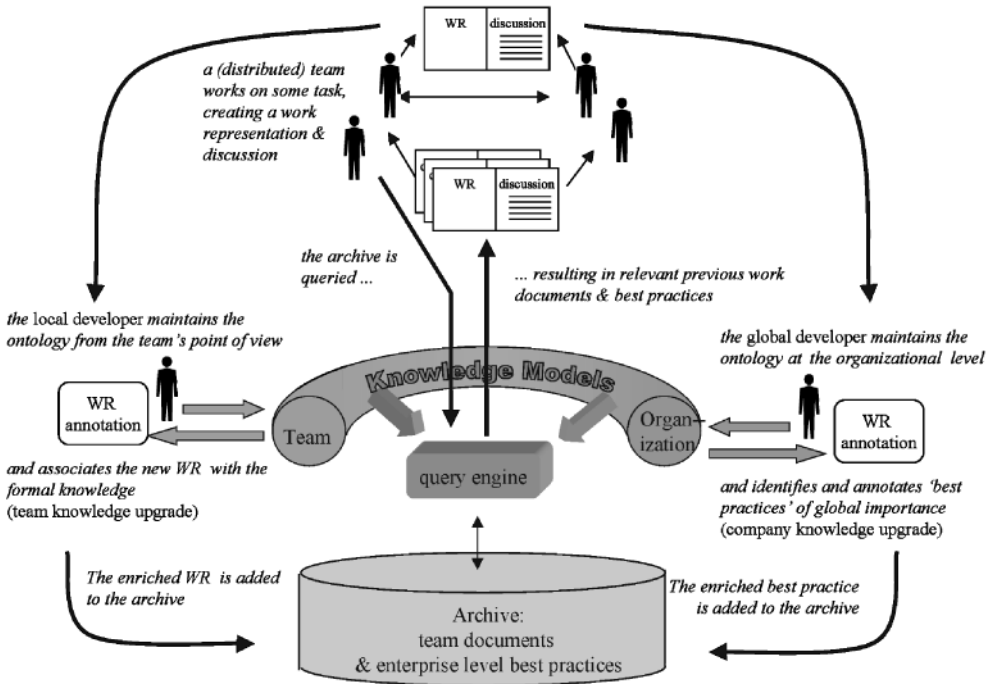
FIGURE 5. The ENRICH lifecycle model.

tools and methodology. They use their knowledge of the team and the domain in which they work to associate WRs with the team knowledge model, and make changes to it. The company knowledge model and best practice archive are the responsibility of the *global developer*. They detect best practice and then use the related WR to produce a new best practice document. New best practices or changes in the organization may lead to modifications of the company knowledge model.

4. Applicability criteria

The ENRICH approach is not applicable or desirable for all organizations. The applicability criteria below describe in general terms the characteristics of an organisation that are appropriate for the ENRICH approach. Even for organizations that can benefit from ENRICH, it is unlikely that every activity in the organization should be supported by the ENRICH tools. The applicability criteria therefore have a complementary function of helping to identify the tasks and members of the organization that the ENRICH development process should focus on. The applicability criteria are grouped around three facets of the organization.

- *Work characteristics*: The nature of the work, the required knowledge and the tools that support it.

- *Group characteristics*: The communities within which the work is conducted, patterns of collaboration and the use of tools.
- Organization characteristics: Learning and cooperation across communities and the tools that support it.

These are described below.

4.1. WORK CHARACTERISTICS

Four characteristics relate to the working and learning done on an individual level within the organization.

(1) *Work involves knowledge that is critical and valuable to the organization.* The adoption of ENRICH by an organization requires a significant allocation of resources, therefore the learning captured and supported has to be of high value to the organization. The ENRICH approach should therefore support knowledge or learning in an organization that is rare, valuable and without which the organization would be unable to function effectively. Brooking (1999) describes these as critical knowledge functions. What constitutes critical knowledge will depend on the context in which the organization operates. The critical knowledge of the company is not necessarily the most complex, but is the knowledge related to the core business of the organization.

(2) *This work involves changing, adapting and innovating as a part of the job.* A component of the ENRICH approach is the storage and retrieval of valuable knowledge, but one of the principle characteristics of ENRICH is the capture of newly evolving knowledge as a consequence of the support for learning. To benefit from the ENRICH approach, there must not only be knowledge critical to the organization, but this knowledge must also be evolving. In the workplace, the evolution of knowledge demonstrates itself through: workers having to change in response to external demands; adapting to new tools and technologies; or innovation being a necessary requirement of the work in order to stay ahead of competitors.

(3) *Past experience is valuable to performing this work.* Although work should be changing and throwing up new challenges, for ENRICH to be effective, there should be great value in past experience when dealing with new work tasks. There is little point in capturing new knowledge in the organization, if that knowledge is not going to be useful in the future. New knowledge captured by the organization can be thought of as having a life-span before becoming obsolete. During the life-span of the knowledge, it has to pay back with interest the expenditure required to capture and maintain it. The ENRICH approach is therefore designed to support knowledge that is evolving, but not in constant flux.

Many examples of work that fit all three of the above conditions may be described as knowledge work, or knowledge intensive work. The examples we have identified include planning, fault diagnosis and design. The final criterion in this section is concerned with the tools that support this work.

(4) *This critical knowledge work could be supported and captured around WRs.* The creation and capture of critical knowledge is focused around WRs. As we described earlier, WRs should be a tool of work, not just paperwork. New knowledge can then be captured naturalistically, during work. Forms filled in after the work is done are not

good WRs, as they capture only a rationalized account of the solution, not the often more valuable problem solving process that led to the solution. An important feature of this criteria is that the work need not be supported by a WR in advance, but any existing WRs, in the form of computer-based or paper-based tools or documents provide a good starting point for the design process.

4.2. GROUP CHARACTERISTICS

Four further characteristics relate to the groups within which the above work is carried out.

(5) *Workers are members of groups that share a common interest or goal*. The ENRICH approach captures knowledge through its support for learning and collaborating around WRs. To meaningfully collaborate around and discuss knowledge intensive tasks, the worker clearly needs to be associated with a group of colleagues that have similar expertise, and perform similar tasks. We use the term group to encompass both formal teams within the organization and also the informal communities of practice.

(6) *Group members have some degree of self-determination in how this work is carried out*. For new knowledge to emerge through collaboration, team members need some level of autonomy as to how their tasks are carried out. If work involves activities such as adapting, problem solving and innovating, then this criteria will almost certainly be met. To perform such roles, the workers need sufficient freedom to pursue solutions to novel problems. This self-determination will of course be bounded. To meet this criterion, workers do not need the freedom to rewrite their own job description, but do need to work in an environment where new ideas are permitted and encouraged.

(7) *Groups have a culture of cooperation*. A collaborative element to working and learning also requires a culture of cooperation. If the workers are put in competition with each other, then collaboration or offers to help colleagues will be a rare commodity. The ENRICH approach requires a culture of collaboration, in order that new ideas emerging from collaboration can be captured and reused. Attempts to incorporate knowledge-sharing support into a competitive environment lead to the technology being unused (Orlikowski, 1991). Enforcing its use in a competitive environment will most likely lead to sophisticated forms of subversion.

Identified examples that fit the above three criteria include members of a team collaboratively developing a team plan and teams of engineers freely offering help to colleagues on how to diagnose and fix faults. The final criterion in this section is concerned with how this collaborative activity relates to WRs.

(8) *Collaborative work could be supported and captured around WRs*. The criteria in this section do not require that extensive collaboration already occurs within groups, only that if it does not currently occur, it is for reasons that ENRICH can help to deal with (i.e. lack of adequate collaboration mechanisms) rather than socio-cultural factors. ENRICH aims to support collaboration within the context of work, by combining communication mechanisms with the WR.

4.3. ORGANIZATION CHARACTERISTICS

Three further criteria relate to the organizational level, which considers learning across group boundaries.

(9) *Groups can learn from each other, and benefit from each others' perspective.* ENRICH provides methods for organizational level learning, across teams. These will only be useful if different teams have some shared interest. If there is no relationship between the concerns, issues and ideas of different teams, they will be unable to learn from each other. This criterion does necessarily mean that there are different teams carrying out almost identical jobs. Sharing across teams could be concerned with higher level, non-day-to-day, issues such as how to build up an effective relationship with a customer. As sharing between groups is intended to support teams in viewing problems from new perspectives, differences in work and approach between groups, could actually be beneficial to the learning process.

(10) *There is a culture of cooperation across groups.* ENRICH provides support for sharing between groups. This will only occur if there is a culture of cooperation between them. It is important that official company policies and rewards should be consistent with this aim. This criterion though does not require groups to have open access to each others' WRs, documents and ideas. Many workers will wish, quite understandably, to initially voice and discuss new ideas within their immediate group, not broadcast them to the entire organization. Additionally, many day-to-day work activities as captured by the group would also fail to "travel", because other teams would fail to appreciate the context in which the work is situated.

(11) *Cooperation and sharing across groups can be supported and linked to WRs.* ENRICH supports team members in gaining access to relevant resources originating outside the team. For this to occur, there is a need to develop documents and other resources that can travel across team boundaries. This is the role we intend for best practices. There is no requirement for best practices to be in place before the adoption of ENRICH, but may be developed during either the seeding or evolution phase.

5. The ENRICH developer methodology

The development model, based on the Seeding, Evolutionary growth, Reseeding (SER) model (Fischer, McCall, Ostwald, Reeves & Shipman, 1994) is shown in Figure 6. This representation of the development methodology, based on Bergmann, Breen, Göker, Manago and Wess (1999), comprises two kinds of boxer: product boxes and process boxes. A product box represents some input to, or output from a process, and is represented as a rectangle with tram lines down each side. A process box indicates some activity to be carried out, and is shown as a lozenge shape. The development model has two main parts. First, there is a relatively linear process from the initial project "Vision" to the "Organizational memory seed". This part of the methodology involves two processes called "Preliminary analysis" and "Seeding", plus an intermediate product labelled "Scenario". The development process from vision to seed does not follow a conventional waterfall model of software development. Two main feedback points are identified: the preliminary analysis may lead to modification of the project vision; and the seeding process may lead to modification of the project scenario. The second part of the development model is cyclical. The seed organizational memory will evolve though use. The evolved organizational memory will then be reseeded to create a new seed. Most of section five will focus on the first part of the development model, that turns a "Vision"
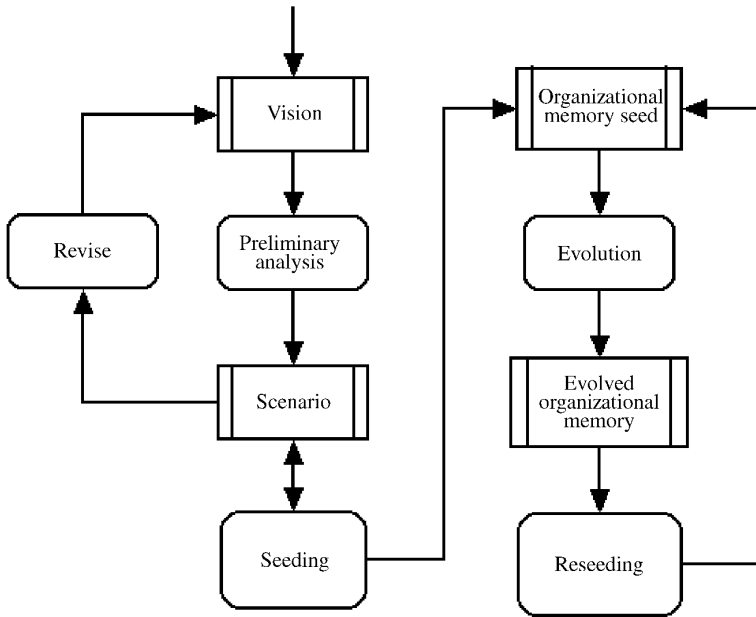
FIGURE 6. Overview of the ENRICH development model.

into an "Organizational memory seed". The main products and process will now be explained in a little more detail.

In order to start, a project needs a *vision*. The vision articulates some need or goal for the organization concerned with how learning could be conducted in the workplace and how knowledge could be captured and shared. The vision may originate from one or more places in the organization, for example as a goal of senior management, or an identified problem in the workplace. The vision should, though, be more than a quotation from a mission statement. It should ideally focus on a practical problem of the organization, the solution of which can be seen to benefit both senior management and the company's employees. From a senior management perspective, potential benefits from a financial (e.g. cost saving) and competitive (e.g. stealing a lead on our rivals) perspective should be visible. From an employee perspective, enhancements to work practices should be foreseeable. The aim of the *preliminary analysis* is to ascertain the feasibility of the vision, and carry out ground work required for the design process. The preliminary analysis initially involves comparing the organization against the applicability criteria described earlier. Further tasks include selecting the work domain for the application, undertaking a requirements analysis, and appointing a project team. The preliminary analysis stage is described in Section 5.1. The preliminary analysis process may lead to the project team to *revise* the vision. It is useful to update the vision to reflect the lessons learned, as the vision document can serve as a public account of the ongoing aims of the project. The output from the preliminary analysis is what we term a *scenario*. The scenario builds on the vision by formulating practical proposals as to how the vision could be achieved. The scenario will contain an initial formulation of what will constitute

the WR, what kind of knowledge model will be required, the work supported, and the workers involved. The *seeding* process is the largest part of the development methodology, covering the design of the WR, knowledge model, technical infrastructure, and the selection of metrics against which the new application can be assessed. Two key parts of this process, WR development and knowledge model development, are described in Sections 5.2 and 5.3. Section 5.5 describes some of the tools developed to support these processes. The seeding process also involves training, and the appointment of personnel to sustain the organizational memory, described in Section 6.

The *organizational memory seed*, being the output of the seeding process, marks the point of hand-over to the recipient organization. The organization should then be able to use and routinely maintain the organizational memory with no or limited external support. This also marks the point of entry into the seed-evolution-reseeding loop. *Evolution* of the organizational memory seed occurs with the support of local developers from the teams, and the global developer of the organization. Use of the *evolving organizational memory seed* is regularly monitored to identify when it no longer supports current work practice, or is no longer able to clearly represent currently evolving knowledge, in the context of previous knowledge of the organization. This is the cue for the reseeding process. The *reseeding* process has some similarities to the initial seeding process, but as the existing knowledge model, work representations, documents and installed organizational memory support tools can be taken into account, the amount of development work required should be relatively small. Advice on the reseeding process is presented in Section 5.4.

5.1. PRELIMINARY ANALYSIS

The aim of the preliminary analysis is to ascertain the feasibility of the vision, and carry out ground work required for the design process. The potential for using the ENRICH application should be assessed using the applicability criteria presented in Section 4, covering work, group and organizational issues. The complete ENRICH methodology incorporates the applicability criteria into an applicability checklist that suggests preparatory work to be undertaken if the criteria are not met. The applicability checklist helps to ascertain the suitability of the ENRICH approach, and should also identify the kind of work in the organization that could be supported, the teams involved in this work, and the kinds of tools and documents they use. This forms what is termed the *application scope* and serves to focus the development work on areas of the organization meeting the applicability criteria.

A project team responsible for the development process needs to be appointed. If the team responsible for the development process have been hired as consultants for the project, the full team will also need to contain advisors from the recipient organization. These advisors will include one or more representatives from the work areas for which the ENRICH application is being developed, plus representatives from IT and knowledge management groups within the company. The full project team will include members with skills in web server technology, knowledge modelling, software design and evaluation and knowledge management. Even if the ENRICH application is being delivered by an external consultancy, the presence of members of the recipient organization with any of the above skills on the project team will aid the hand-over of the

organizational memory. All of these skills are required within the recipient organization if the ENRICH development model is being self-applied. Experts in user-interface design and graphic design will also be required to ensure the usability and aesthetics of the tools.

The application scope identifies the area of work to which ENRICH can be applied. This needs to be narrowed down further to identify specific teams within the organization. This will involve visiting candidate teams and presenting the project vision to them. Although many teams may meet the applicability criteria, teams whose workers have an understanding and motivation for the project are identified. There must also be managerial support for the project within those teams. It is increasingly accepted within knowledge management that the most effective way to proceed is to initially identify promising teams, and design for them. The scope of the application can then be extended to a team at a time, once early success has been demonstrated. Lessons learned from, and solutions developed for, the initial teams can be extended and modified to support these further teams. Gradually, a generalized solution for a sector of the organization may emerge. This approach carries less risks than attempting to initially design a generic solution for a wide range of teams.

Once teams have been identified, a requirements analysis should be undertaken, comprising three components.

- Identify technical infrastructure requirements of the project.
- Identify intended role of the work representation and its relation to existing tools and documents.
- Identify intended role of the knowledge model and its relation to existing knowledge models if present.

A particular methodology for requirements analysis may be adopted to guide the identification of requirements in the above four areas, such as contextual inquiry (Beyer & Holtzblatt, 1998).

The final task within the preliminary analysis phase is the formulation of clear objectives for the project. To ensure continued support for the project there needs to be a demonstration of its success. Success cannot be demonstrated without clear objectives. Objectives should be as specific as possible, so that performance against them can be measured. Selected specific objectives, depending on the application area may include the following.

- Reduce time taken to diagnose a machine fault.
- Increase number of teams adopting the planning methodology.
- Increase percentage of customer queries resolved in a single contact.

At least some of the selected objectives should be realisable in the short term. Long-term (financial) support for the project cannot be expected without some demonstrable short-term benefits. The project team should also formulate long-term benefits such as predicted competitive advantage for the organization. These objectives will feed into the assessment strategy formulated as part of the seeding process.

5.2. WORK REPRESENTATION DEVELOPMENT

The first WR development task is to establish a WR development group to support the design and testing of the WR and associated documents. These should on the whole be

members of the team who are to use the designed application as part of their work. Members of this group should have a strong motivation for the aims of the project and be skilled members of their own work area. As the final application is web-based, then ideally at least a few members of the development group will be familiar with internet technologies, such as web browsers and on-line discussion spaces. This will allow them to act as local experts and explain the new technology to other members of the team.

The process of designing the work representation can be broken down into six separate stages. The design process, conducted with members of the organization, may adopt methods or techniques developed to support such a participatory approach to design. Useful methods and guidelines can be gleaned from a number of sources, including participatory design (Schuler & Namioka, 1993), contextual design (Beyer & Holtzblatt, 1998), and informant design (Scaife, Rogers, Aldrich & Davies, 1997). The six stages of the design process will now be explained in turn.

(1) *Design the tool/document that will form the WR*. From the requirements analysis, there should be a reasonably firm idea as to what tasks the WR will support and what existing tools and documents, if any, it should be based upon. As stated earlier, the WR should directly support some knowledge intensive task. This will almost certainly be some sort of problem-solving task such as planning, diagnosis, design, analysis, etc. We cannot give a precise definition as to what a potentially successful WR will be like, creativity and the participation of the WR development group are required to help ensure success. The project team will find it helpful to make rough sketches of the proposed WR, and talk these through with the WR development group. Usability issues should be considered, but of prime importance is the scenario of use for the WR—who is going to use it, how, and when? The project team cannot expect the workers from the teams to establish this for themselves.

This design process may interact with knowledge model development, as the new WR may have embedded within it a method for searching for related documents. For example, within our team planning case, text entered into the body of the plan is used to motivate an ontology supported search of the best practice archive. Such a facility would require some text analysis functions to be developed as part of the work on the knowledge model.

(2) *Select scope of each WR*. Each WR will have a life-span within the workplace. Each WR supports some problem solving processes, and will be linked to specific problems. Once a problem has been solved, that WR becomes archived, and the team or individual worker will use a new WR to support a new problem-solving process. For example, in our mining case, a separate WR was used for each recorded fault. The relationship could be more complex, for example a WR may support a number of individual cases, or a certain type of case. The design team need to ensure that the transition between WRs is consistent with natural transitions within the work process.

(3) *Determine interrelations between WRs across time and tasks*. Depending on the nature of work, individual WRs may be independent, or may be linked to each other, each forming a subtask of some larger ongoing problem. In our hotline and mining cases, WRs produced over time were independent of each other. This will be the common pattern where work is divided into discrete cases, such as in a fault diagnosis scenario. Conversely, within the planning case, WRs had a temporal relationship to each other. Such a scenario places extra demands on the design process, as it may be necessary to

provide easy ways of carrying over data from one plan to its successor. The knowledge model will also need to represent this temporal relationship so that the evolution of a plan for a particular team can be traversed. In some scenarios, more than one WR may be used concurrently to support related tasks. For example, each task may be divided into two WRs, focussing on different aspects. The relation between WRs may be more complex, if for example, they are supporting stages in a design process. The design process often takes the form of a breadth first search of a space of solutions, where a number of candidate solutions are evaluated in parallel and then specialized. In this kind of domain, even greater support would be required for traversing related WRs, and moving data between them.

(4) *Decide how each individual WR is to be divided.* Some WRs such as a team plan may be quite large, covering a number of pages. These will need to be divided into separate sections so that the user can easily navigate to the appropriate part. Each individual section, like sections of a book, should be conceptually meaningful, and relate to some topic, issue or task. The breakdown of the WR should also relate to how the work it supports is broken down into separate activities. The designer needs to decide whether each section should be displayed on a separate web page, or as subsections of a single web page. The divisions of the WR may also have ramifications for the knowledge model and structure of the asynchronous discussion space. For example, each section may be related to a concept in the knowledge model, and may have its own thread within the asynchronous discussion. The designers need to determine the degree of freedom that workers will need when producing new WRs. For example, in some domains the WR may have a highly stereotypical structure, each problem having predictable stages. For example, a WR to support fault diagnosis may be divided into problem description, solution, testing, etc. For some domains, the structure of the WR may be more ad hoc, and the worker may need more control over how individual WRs are broken up.

(5) *Decide how the asynchronous discussion tool is to be used.* The asynchronous discussion component may play functionally different roles in different work settings. First, it may be used primarily to support distributed collaboration within a work team, allowing field workers to collaborate in the context of the problem with a centrally located staff. For example, within our mining case, the tool is used for asynchronous discussion between workers situated down the mine and those in the control centre. Second, it may often be used to communicate and present ideas to some target group. For example, in our hotline case, lessons concerning how a particular product can be used in different situations are added to the discussion space concerned with that product. In this scenario, the discussion space does not support asynchronous collaboration, but rather the presentation of an idea in context to a target audience. Third, the asynchronous discussion space can be used primarily to link the WR to other resources. Within the team planning case, the asynchronous discussion space is commonly used to generate hyperlinks from a discussion thread relating to a part of the plan to external resources, serving as evidence of how each goal of the plan is being achieved.

(6) *Decide how discussion threads are divided and related to the WR.* Once the role of the asynchronous discussion space has been established, the designer needs to formulate the discussion seed, in terms of the primary subject areas in which a message can be placed. Determining the seed structure requires consideration of two separate views of the WR: the externally imposed perspective; and the perspective acquired from the WR content.

Each WR describes an individual problem (e.g. develop the August 2000 plan in light of organizational demands). That problem will be of a type (e.g. planning by teams in division X of company Y), and therefore may have an expected set of sub-problems and issues that need to be considered (e.g. financial performance, customer satisfaction, etc.). These expected problems and issues can provide whole or a part of the structure of the discussion seed. Alternatively, some or all of the discussion seed structure can be automatically derived from the content of the individual WR itself. For example, each subsection of the WR may have its own discussion thread.

5.3. KNOWLEDGE MODEL DEVELOPMENT

The design of the knowledge model requires members of the organization to form a knowledge acquisition group. Additionally, the knowledge acquisition group may include members from outside the organization, acting as consultants contributing specialist knowledge of value to the organization. There may be a large overlap between the knowledge acquisition group and the WR development group. Unlike the WR development group, the knowledge acquisition group requires greater input from members of the wider organization, not just the team. This is particularly important if, for example, the knowledge model relates to best practices, and other resources originating outside the team. In such a situation, members of the organization who provide this resource, would be included in the knowledge development group. In addition, if the organization already has a knowledge model, then personnel working on this will be included within the group, to support integration with, and reuse of, the existing resource. For the knowledge model to be effective, it needs to be consistent with the perspective of the team who will use it, but the knowledge model should also be compatible and interpretable within a wider organizational context.

When designing a knowledge model of an area on work within an organization, the temptation can be to model whatever is possible, and particularly aspects that seem easy to model, such as organizational hierarchies. This is a poor approach, that will inevitably lead to a great deal of wasted effort. Knowledge should be acquired and represented for a clear purpose. The knowledge model is used to describe the knowledge of the teams and organization, and knowledge contained in individual WRs and documents in order to facilitate the following.

- The description of a WR currently being used by a team in order to promote its reuse.
- The capture of knowledge evolving in the teams and organization as a knowledge asset of the organization.
- Describe Best Practices and other resources to enable their appropriate selection to support workers in undertaking their current tasks.
- Support other required knowledge services such as the detection of critical states.

The knowledge acquisition process should therefore focus on knowledge that is *relevant* (i.e. knowledge used in WRs), *valuable* (i.e. knowledge that can be an asset to the organization), *discriminating* (i.e. knowledge that can determine which resources are useful for a particular problem) and *practical* (i.e. supporting useful knowledge services). Some specifications of the knowledge required and the role it will play should have been

TABLE 1
The generic work-representation class

Slot name	Documentation	Slot type
Has-associated-task	WR supports some task	Instance of the `task` class
Has-associated-group	Team in which it belongs	Instance of the `group` class
Has-associated-readership	Readership may be team or wider	Instance of the `group` class
Has-related-work-representation	Next, previous, subset of, etc. Depending on the nature of the work	Instance of the `work-representation` class
Has-associated-best-practice Has-associated-supporting-document	Best practices used by the WR Other resources that relate to the WR	Instance of best-practice class Instance of `document` class

identified from the preliminary analysis. The design team may adopt an established methodology to support this process such as CommonKADS (Schreiber *et al.*, 2000).

The knowledge model developed for an ENRICH application has three main parts: a generic ENRICH ontology; specialized ontology for the application; and a knowledge base describing the specialized application. The generic ontology is common across our current cases, and provides a framework for developing knowledge models for further cases. The specialized ontology is developed according to the purpose for which the knowledge model will be used within a specific organization. The knowledge base represents individual WRs and other concepts identified in the ontologies.

The knowledge represented in the generic and specialized ontologies is divided into three types: process knowledge; domain knowledge and media knowledge. Process knowledge describes the process by which the knowledge work is carried out and what resources are needed to undertake it. Domain knowledge describes the domain of work, such as work products and machinery. Media knowledge describes where resources can be found. Below, we describe the main generic classes used within our existing case studies.

Media knowledge describes the range of resources available including WRs, best practices and supporting documents. Their representation in the ontology should include reference information and a description of content. Reference information may include the URL of the resource, author, publication date, title, etc. In terms of content information, each WR will probably have an associated task and team, which are represented within domain knowledge. The ontology should be specialized to represent how a particular WR relates to other WRs. For example, WRs may be related linearly or hierarchically. WRs may also be formally related to best practices and other resources. A generic description of a work representation class is shown in Table 1.

A separate class will be required to represent any documents held in a best practice archive (see Table 2). Two important aspects of a best practice to be Represented are the lesson that can be gleaned from the best practice, and the conditions that one needs in

TABLE 2

The generic `best-practice` *class*

Slot name	Documentation	Slot type
Has-interesting-lesson	The reason it is a best practice. What you can learn from it	Instance of class `lesson-learned`
Has-associated-condition	They are not universally applicable. What is needed	Instance of class `associated-condition`

place in order to operationalize the best practice. The best-practice class may be specialized, and have further subclasses to represent the kinds of best practice found in the organization.

Domain knowledge is used to represent the domain of expertise within which the members of the organization work. Domain knowledge is concerned with the officially prescribed tools and techniques of the professional domain to which the work teams belong, such as engineering, aircraft design or mining. This is the kind of knowledge that a worker would accumulate during a formal training or apprenticeship regime. The kind and level of domain knowledge required will largely depend on its discriminating power in determining the usefulness of a resource to the current problem. Only domain knowledge that has discriminating power, or is viewed as being worthy of capture as an infrastructure asset, should be elicited and modelled. As domain knowledge is officially prescribed, then once represented, it will usually be reasonably stable, and require little or no modification during the evolution of the knowledge model. As domain knowledge is concerned with the tasks, methods, tools and materials of the work being supported, these will be represented by the specialized ontology. Domain knowledge will be connected to both media knowledge and process knowledge. Media knowledge represents documents that support some task. If these tasks are represented in the ontology, they will form a part of the domain knowledge. Teams and departments of the organization and their roles may also form a part of the domain knowledge. The final part of the ontology, process knowledge will also be connected to domain knowledge.

Process knowledge has two main components: knowledge of the problem-solving process by which work is undertaken, and knowledge about the resources needed to tackle this problem. Unlike domain knowledge, process knowledge evolves from lessons learned in the workplace and their context. Process knowledge is therefore the most dynamic part of the ontology during the evolution process. Process knowledge includes the lessons learned and associated conditions that can be used to index best practices. The interesting characteristics of problems described in WRs are also represented in process knowledge. Lessons and problems should only be represented if they are interesting, i.e. they constitute relevant and valuable knowledge.

A generic form of the lesson learned class is shown in Table 3. Minimally, each lesson will have an associated task, problem and solution. A class hierarchy may be used to categorize problems and solutions. An associated condition class will represent what needs to be in place for a lesson to be applied. The associated conditions could include financial, skill and time requirements.

TABLE 3
The generic `lesson-learned` *class*

Slot name	Documentation	Slot type
Has-associated-task	Best practice relates to a task	Instance of the task ontology component of domain knowledge
Has-problem	Describes a problem	Instance of problem
Has-solution	Describes a solution	Instance of solution

5.4. THE RESEEDING PROCESS

After some period of evolutionary growth, the organizational memory will need to be reseeded. The timing of the reseeding process will depend on the extent to which knowledge related to the work area changes. Reseeding may be initiated by a particular event, such as the issuing of new guidelines and methodology of how work should be conducted, requiring extensive changes to the knowledge model. Similarly, workplace changes effecting how WRs need to be used may lead to reseeding. Even in the absence of large-scale changes to work knowledge or practice, reseeding should occur periodically, to formalize evolving knowledge into the knowledge model, a process that users often find difficult to undertake without external support. Requirements for the reseeding process will be gathered through monitoring performance against the assessment measures, and the observations of the global developer and local developers. Reseeding focuses on three aspects of the organizational memory: the knowledge model; the discussion space; and documents including the WR.

Three core tasks for the reseeding of the knowledge model are: identification of redundancy; pruning and abstraction; and formalization. First, identification of redundancy is required as the knowledge model should be no larger than is required to support the available services such as searching for resources. Areas of the ontology that are found unnecessary or are no longer relevant should be removed. Second, pruning and abstraction should be used to remove duplications and unnecessary details within the ontology and generalize parts of the ontology into higher level categories. Abstraction will apply in particular to process knowledge, where over time, it may have become possible to abstract lessons learned from a number of cases. Knowledge may also be abstracted into a higher level ontology, for example knowledge from a team ontology may be generalized and moved to the company ontology. Third, new knowledge used in the workplace, some of which may be identifiable from the discussion space, may be formalized in the ontology. Although this may occur during the evolution phase, being conducted by the local and global developers, formalization is a difficult process to carry out without specialized assistance, and will therefore need to be considered during the reseeding phase. Other reseeding tasks may include partitioning or joining ontologies, or extending the scope of the ontology to other kinds of knowledge within the organization.

Similar subtasks such as identification of redundancy and pruning are required for the discussion space. This process should apply to messages contained in discussion spaces

and the thread structure of individual discussion spaces. The reseeding process should also involve reassessment of the WR in terms of its consistency with work practice, interface and layout. The structure of the initial best practices should also be reconsidered, to identify whether their form and layout adequately supports reuse. If the initial organizational memory seed did not contain a best practice archive, then after a period of evolutionary growth, the potential for best practices may be more clearly understood, and a best practice archive can be developed from selected WRs created during evolutionary growth.

5.5. THE ENRICH TOOLKIT

The tools developed during the ENRICH project provide support for the development of WRs, knowledge models and their integration. The CEDAR† (Contextually-Enriched Document Archive) toolkit (Hatala, 2000) was developed to support the publishing and management of WRs and other documents within the organization. CEDAR takes as input a raw HTML file that can then be broken up, associated with a discussion space and published according to a user defined template. The CEDAR tool also has document management functions for representing document groups and user profiles. CEDAR currently integrates with Hypernews‡ and aGNeS§ discussion forums. CEDAR is available in executable form for Windows 95, 98 and NT.

Within our existing case sites, the OCML¶ (Operational Conceptual Modelling Language) (Motta, 1999) knowledge modelling language has been used. The browsing and editing of ontologies is supported by the graphical WebOnto†† tool (Domingue, 1998). The WebOnto interface is a Java applet, that communicates with the WebOnto server. Our approach of using OCML and WebOnto allowed the ontology plus knowledge base to be stored centrally on a server, referring to target on-line documents via URLs. This allows multiple ontological views of a set of documents to be developed, without having to edit the documents themselves. Association of web-based documents with the ontology is supported by the Annotation Tool (Hatala & Hreno, 1999). This tool is used to view a published document, that can then be annotated with a selected ontology held on the WebOnto server. The tool allows an annotation to refer a whole document, a document section such as a chapter, or any part of a section, down to a single character.

6. Training and role development methodology

To support the continued evolution of the organizational memory, three main roles are required.

- Users trained in how to use the new tools as part of their work.
- Local developers working within the teams, and supporting the work of their colleagues.

---

† CEDAR toolkit: http://enrich.open.ac.uk/cedar
‡ Hypernews: http://www.hypernews.org
§ aGNeS: http://www.vestris.com/agnesbbs
¶ OCML: http://kmi.open.ac.uk/knowledge-modelling
†† WebOnto: http://webonto.open.ac.uk

- Global developer coordinating and monitoring use of the ENRICH application from the perspective of the organization.

The users are the members of work teams or communities that will use the new tools as part of their work. It is intended that they will use WRs, the discussion space, and search the organizational memory as an integral part of their work. For this to be realized, the users should be trained in how to use the tools. The users should also be given an account of the background and rationale for the ENRICH application, and a clear scenario of how the new tools will impact upon their work.

Overall, the duties of the local developer involve supporting team members and the local evolution of the organizational memory. Their main duties are maintaining team WRs and documents, editing of the discussion space, and association of WRs and documents with the ontology. First, the local developer will manage and modify WRs and other documents as required. For this, the local developer will need training in publishing WRs, and making minor changes to the layout. Second, the local developer is responsible for editing the discussion spaces. This will include modifying messages or deleting messages sent in error. The local developer may also be required to prune the discussion spaces of archived WRs in order to facilitate their reuse. Alternatively, this may be left exclusively as a task to be undertaken during the reseeding phase. Third, the local developer will also be responsible for indexing WRs generated and completed during work. The local developer will not require a full understanding of the ontology to perform this task, as the job will generally involve adding representations of WRs to the knowledge base. This task may be performed using a version of the Annotation Tool. Alternatively, an HTML form may be prepared allowing the local developer to describe a WR by selecting from predefined values for a list of slots.

The global developer is the champion of the ENRICH application on the organizational level. Their main duties will include management of the ontology, WRs, best practices and other documents from an organizational perspective. First, the global developer will have to inspect, debug and update the company knowledge base and possibly modify team knowledge bases if changes are required beyond the remit of the local developer. Second, the global developer will need to be skilled in the publishing tools used to present and organize WRs and other documents. If the CEDAR toolkit is adopted, then the global developer will need to know how to create new templates and manage the group and document structures. Third, the global developer will be involved in detecting best practice from the work of teams and adding these to the best practice archive. This activity will need to be integrated with any existing best practice initiative of the organization. If the archive is not in place at the start of the evolutionary growth phase, then the global developer will devise a methodology of how best practices should be prepared in order to facilitate their effective reuse.

7. Summary and comparison with other approaches

The ENRICH approach is designed to support the four kinds of learning outlined in section two (reflection-in-action, domain construction, community of practice learning and perspective taking) and capture knowledge naturalistically as a consequence of this support. Learning is supported around work representations. As a consequence of

learning, the WRs become enriched in two ways. First, workers are given the facility to collaborate and negotiate within the context of their work, via asynchronous discussion spaces linked to important work representations. This builds the supporting context for the WR. Second, WRs become linked to the formal knowledge model, and the knowledge model itself is extended to reflect the incremental evolution and formalization of organizational knowledge. This enriches the WRs with a knowledge context. These WRs, enriched as a consequence of supporting work-based learning, themselves become part of the organizational memory, and therefore become part of the learning resources of the organization. ENRICH therefore supports organizational learning in two interrelated ways.

- Tools support organizational learning around representations of work, which become enriched as a consequence.
- The enriched representations form an evolving organizational memory that supports further learning in the organization.

The ENRICH approach has three main benefits. Unlike many knowledge and document management systems, the ENRICH tools and methodology provide the following.

- Support for integrating working and learning around enriched representations of work within the organization.
- Knowledge services, including annotation and reasoning mechanisms, based on the semantics of the work process.
- Explicit support for the ongoing capture and evolution of knowledge, according to the ENRICH lifecycle model.

Integrating working and learning around WRs. An important component of the ENRICH methodology is concerned with the identification and development of WRs, and their incorporation into work activity. This process is definitely not a case of adapting work practices to fit a new tool, but neither is it adapting a tool to fit work practice. The suite of ENRICH tools offer new methods for collaborative work, tracking progress and presenting work outcomes. The design process is conducted in participation with the target population to explore these new opportunities.

Provision of knowledge services. The underlying ontology that is used to describe the work domain and the position of WRs and other documents within this, supports a range of reasoning mechanisms to provide intelligent analysis and search mechanisms to inform workers of resources related to their current activities. This form of reasoning can be realized in other knowledge management systems, but ENRICH is explicitly designed to support the development of these reasoning methods. The underlying language in which the ontology is represented (i.e. OCML) incorporates facilities to support these reasoning mechanisms, and a number of tools are provided to support the use, exploration and modification of the ontology (e.g. WebOnto).

Capture and evolution of knowledge. The ENRICH approach builds upon knowledge and document management solutions by focussing primarily on support for the evolution and generation of knowledge rather than a direct management and capture of knowledge. Within ENRICH, knowledge is captured as a consequence of supporting its generation through the tools provided. Specifically, the ENRICH approach supports the four kinds of learning outlined earlier and captures new knowledge as a consequence of

TABLE 4
Types of learning supported and captured in ENRICH

	Learning type	How supported	How captured
Individual	Reflection-in-action (Schön, 1983)	Provision of semantically related resources to support problem reframing or to challenge assumptions.	Ongoing discourse around WRs related to current problems and issues.
Group	Domain construction (Sumner, 1995)	Elaboration and negotiation of domain concepts around work represen-tations linked to formal knowledge models.	Evolution of group knowledge models and associations between knowledge models and WRs.
	Community of practice learning (Schön, 1988; Lave & Wegner, 1991; Raelin, 1997; Cook & Brown, 1999)	Development of collective practices and perspectives through enhanced sharing and collaboration around work representations.	Expressed in the customization and use of WRs. Implicit group knowledge contained in artefacts shared seam-lessly within the group.
Organization	Perspective taking (Boland & Tenkasi, 1995)	Exposure to customs and conventions of other groups through shared best practices and explicit conceptualizations of common views.	Cross-community discourse. Adoption and modification of best practices from other communities. Development of shared knowledge models.

this support. This is summarized in Table 4. At the individual level, ontology driven search agents provide resources related to the current task as expressed in the WR. This supports individuals in reframing current problems and challenging assumptions on which the current approach is based. Some of the new knowledge created during reflection-in-action is captured in the discourse area as problems with, and changes to, the approach are reported, and discussed among colleagues.

Domain construction, the process by which teams negotiate and extend their team knowledge is supported by tools based on WebOnto described earlier, generally with the support of the local developer. As a consequence, the evolution of team knowledge is captured, enhancing the performance of the ontology-driven search agents. Additionally, at the group level, community of practice learning is enhanced through support provided for collaboration and the sharing of work products, problems and solutions. The development of community of practice knowledge becomes apparent, as the local developer customizes the WR and its integration with other work resources, illustrating the group's work genre. Changes to the WRs and their use expresses group knowledge that may later become accepted and formalized through domain construction.

Perspective taking is supported through the retrieval of resources from other groups within the organization, particularly through the best practice archive. Once again, tools

based on WebOnto allow connections to be made across knowledge models, or an inter-team model of shared expertise to be developed, which is captured for further use in the knowledge model. Team members, as part of their work, comment on their use of a best practice in their own situation. These captured discussions of best practice are available to the global developer when modifying the best practice archive or company knowledge model.

8. Conclusions

The ENRICH tools and methodology are designed to support organizational learning around work representations. As a consequence of this, the work representations become enriched with context and semantic associations, thereby evolving the organizational memory. The four types of learning identified in the methodology are reflection-in-action, domain construction, community of practice learning, and perspective taking. These cover all levels of learning from the individual worker to the organization as a whole. Our approach to enriching representations of work to support organizational learning has been tested and refined through three industrial case studies.

REFERENCES

ARGYRIS, C. & SCHÖN, D. A. (1996). *Organisational Learning II*: *Theory*, *Method and Practice*. Reading, MA: Addison Wesley.

BERGMANN, R., BREEN, S., GÖKER, M., MANAGO, M. & WESS, S. (1999). *Developing Industrial Case-Based Reasoning Applications*: *The INRECA Methodology*. Lecture Notes in Artificial Intelligence, Vol. 1612, Subseries of Lecture Notes in Computer Science. Berlin: Springer.

BEYER, H. & HOLTZBLATT, K. (1998). *Contextual Design*: *Defining Customer Centered Systems*. San Francisco, CA: Morgan Kaufmann.

BOLAND, R. J. & TENKASI, R. V. (1995). Perspective making and perspective taking in communities of knowing. *Organization Science*, **6,** 350–372.

BROOKING, A. (1999). *Corporate Memory*: *Strategies for Knowledge Management*. London: Thomson.

BROWN, J. S. & DUIGUID, P. (1996). The social life of documents. *First Monday*, **1,** http://www.firstmonday.com.

COOK, S. D. N. & BROWN, J. S. (1999). Bridging epistemologies: The generative dance between organizational knowledge and organizational knowing. *Organization Science*, **10,** 381–400.

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

DETTERMAN, D. K. (1993). The case for the prosecution: transfer as an epiphenomenon, In D. K. DETTERMAN & R. J. STERNBERG Eds. *Transfer on Trial*: *Intelligence*, *Cognition*, *and Instruction*. Norwood, NJ: Ablex Publishing Corporation.

DOMINGUE, J. (1998). Tadzebao and WebOnto: discussing, browsing, and editing ontologies on the web, in *Proceedings of the 11th Knowledge Acquisition for Knowledge-Based Systems Workshop*, 18–23 April, Banff, Canada.

FISCHER, G., MCCALL, R., OSTWALD, J., REEVES, B. & SHIPMAN, F. (1994). Seeding, evolutionary growth and reseeding: supporting the incremental development of design environments. *Human Factors in Computing Systems (CHI'94)*, 24–28 *April*, *pp.* 292–298, *Boston*, MA, New York: ACM Press.

GANTT, M. & NARDI, B. (1992). Gardeners and Gurus: patterns of cooperation among CAD users. *Human Factors in Computing Systems (CHI'92)*, pp. 107–117, Monterey, CA, 3–7 May, New York: ACM Press.

GRUBER, T. R. (1993). A translational approach to portable ontology specification. *Knowledge Acquisition*, **5**, 119–220.

HATALA, M. (2000). Contextually-enriched documents: publishing for organizational learning. KMi Technical Report 85. Available from http://kmi.open.ac.uk/techreports.

HATALA, M. & HRENO, J. (1999). Annotation of documents with knowledge modelling concepts. ENRICH Technical Report. Available from http://enrich.open.ac.uk.

IVERGÅRD, T. B. K. (1998). Measures to balance the demand for skill in a fast changing technological business environment and the supply of skill available in the existing work force in P. VINK, E. A. P. KONINGSVELD & S. DHONDT, Eds. *Human Factors in Organizational Design and Management—VI*, Amsterdam: Elsevier.

KIDD, A. (1994). The marks are on the knowledge worker. *Human Factors in Computing Systems (CHI'94)*, Boston, MA 24–28 April, New York: ACM Press.

LAVE, J. & WEGNER, E. (1991). *Situated Learning*: *Legitimate Peripheral Participation*. Cambridge, UK: Cambridge University Press.

LEVITT, B. & MARCH, J. (1988). Organisational learning. *Annual Review of Sociology*, **14**, 319–340.

MOTTA, E. (1999). *Reusable Components for Knowledge Modelling*: *Case Studies in Parametric Design Problem Solving*. Amsterdam: IOS Press.

NONAKA, I. & TAKEUCHI, H. (1995). *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford, UK: Oxford University Press.

ORLIKOWSKI, W. (1991). Learning from notes: organisational issues in groupware implementation. *Computer Supported Cooperative Work (CSCW'92)*, Toronto, New York: ACM Press.

ORR, J. (1990). Sharing knowledge, celebrating identity: war stories and community in a service culture. In D. S. MIDDLETON & D. EDWARDS, Eds. *Collective Remembering*: *Memory in Society*. Beverly Hills, CA: Sage Publications.

RAELIN, J. A. (1997). A model of work-based learning. *Organization Science*, **8**, 563–578.

SCAIFE, M., ROGERS, Y., ALDRICH, F. & DAVIES, M. (1997). Designing for or designing with? Informant design for interactive learning environments. *Human Factors in Computing Systems (CHI'97)*, Atlanta GA, New York: ACM Press.

SCHMIT, P., BERNARDI, A. & BECKER, W. (2000). Elektronisches Störungsbuch auf dem Bergwerk Ensdorf. *Glückauf*, **136**, 258–264.

SCHÖN, D. A. (1983). *The Reflective Practitioner*: *How Professionals Think in Action*. New York: Basic Books.

SCHÖN, D. A. (1988). Designing: rules, types and worlds. *Design Studies*, **9**, 181–190.

SCHREIBER, G., AKKERMANS, H., ANJEWIERDEN, A., DE HOOG, R., SHADBOLT, N., VAN DE VELDE, W. & WIELINGA, B. (2000). *Knowledge Engineering and Management*. Cambridge, MA: MIT Press.

SCHULER, D. & NAMIOKA, A. (1993). *Participatory Design*: *Principles and Practices*. Hillsdale, NJ: Lawrence Erlbaum Associates.

SUMNER, T. (1995). The high-tech toolbelt: a study of designers in the workplace. *Human Factors in Computing Systems (CHI'95)*, Denver, CO 7–11 May, New York: ACM Press.

SUMNER, T., DOMINGUE, J., ZDRAHAL, Z., MILLICAN, A. & MURRAY, J. (1999) Moving from on-the-job training towards organisational learning. *Proceedings of the 12th Banff Knowledge Acquisition Workshop*, Banff, Alberta, Canada, 16–22 October.

AUTHOR QUERY FORM

HARCOURT PUBLISHERS

Queries and / or remarks

Manuscript Page/line	Details required	Author's response
1	Please furnish Keywords.	
4/18	Orr (1990) — Please check year (1990 & 1991).	
19	Brown and Duiguid (1996) — Please furnish page numbers.	