

Supporting organisational learning: an overview of the ENRICH approach

Paul Mulholland, John Domingue, Zdenek Zdrahal and Marek Hatala

Knowledge Media Institute

The Open University

Milton Keynes, MK7 6AA, UK

{P.Mulholland, J.B.Domingue, Z.Zdrahal, M.Hatala}@open.ac.uk

Traditional training separates learning from the work context in which the newly acquired knowledge is to be applied. This requires the worker themselves to apply imparted theoretical knowledge to knowledge in practice, a process that is grossly inefficient. The ENRICH approach builds on organisational learning theory to intertwine working and learning. The ENRICH methodology incorporates theories of learning at the individual, group and organisational level. Individual level learning is supported through the provision of semantically related resources to support problem reframing and to challenge assumptions. Group learning is supported through the evolution of domain concepts through work documents and representations linked to formal models of group knowledge, and the development of group practices and perspectives through enhanced sharing and collaboration. Organisational learning is supported through exposure to customs and conventions of other groups through shared best practices and knowledge models. The approach is being investigated in a range of industrial settings and applications.

1. Introduction

Learning as an integral part of working naturally occurs in the vast majority of organisations. 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 organisation learning. Over the last few decades a number of theories and definitions of organisational learning have been offered, though there are a number of common threads running across them. First, organisational learning primarily occurs within a community of practice [12]. These communities do not necessarily equate to the formal team structure of the organisation. Second, organisational learning occurs in collaborative activity, which requires an appropriate collaborative culture [1]. A third common feature of organisational learning is the reuse and reworking of past experience and solutions [1, 13]. Organisational learning builds on, questions and modifies previous solutions and ideas.

Organisational learning is often quite different from the processes employed by the organisation when training staff. Traditional approaches to training involve strategists within the company identifying or predicating the skills gap 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 [7] reported that 90% of training was not transferred to the job, wasting \$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 [11]. Third, the process is almost solely top-down,

where senior personnel must identify what needs to be done and how, devaluing the role 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 organisational learning should replace all forms of training, but that learning through practice is undervalued and ill supported in many organisations. Although organisational 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 [16] 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, organisations require learning that can readily span space, time and community boundaries, providing knowledge on demand to those who need it.

The overall aim of the ENRICH project [22] was to develop a methodology and technology that could harness and catalyze organisational learning processes for the service of the individual worker, workplace community, and the organisation as a whole. The ENRICH approach builds on knowledge and document management solutions by focussing primarily on support for the evolution and generation of knowledge rather than directly the management and capture of knowledge. Within ENRICH, knowledge is captured as a consequence of supporting its generation through the tools we provide. Individual level learning is supported through the provision of semantically related resources to support problem reframing and to challenge assumptions. Group learning is supported through the evolution of domain concepts through work documents and representations linked to formal models, and the development of shared practices and perspectives through enhanced sharing and collaboration. Organisational learning is supported through exposure to customs and conventions of other groups through shared best practices, and the development of explicit conceptualizations of common views. The ENRICH approach, like all knowledge management and organisational learning applications requires certain conditions within the host organisation, such as a sufficient culture of collaboration and cooperation. ENRICH, or any other application, cannot create these conditions where they do not exist [14].

The design and incorporation of an organisational learning application requires a theoretical position on how organisations can learn, and the nature of the knowledge created and used in the workplace. Toward this end, the next section will describe our perspective on the nature of organisational knowledge and learning. Sections three and four will then describe how working and learning are integrated around important representations used in the workplace, and the representation and evolution of formal knowledge. Section five will summarize how ENRICH provides ongoing support for organisational learning. Sections six and seven will outline the developer methodology and when the ENRICH approach can be used.

2. The nature of organisational learning

Supporting organisational 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 organisation. For example, a new idea may be initially created 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 organisational 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 [18] theory of "reflection in action". Schön claims that a professional worker (i.e. knowledge intensive worker) uses and applies a range of knowledge and skills during their professional 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 occurs 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 then 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. Once the breakdown has been resolved, they can return to working with knowledge in action, combined with explicit knowledge.

Reflection in action highlights an important distinction between two kinds of knowledge applied within the organisation: tacit and explicit. Tacit knowledge is the knowledge that is used effortlessly and routinely as part of work, but it is hard for the person who uses that knowledge to explicitly describe that knowledge and how it used. A designer may intuitively feel that a particular approach is right, or doomed to failure. They cannot explain why, but their intuition is nearly always correct. Their tacit knowledge is a result of experience and only becomes apparent through action. An important advantage of tacit knowledge is that it allows the worker to concentrate on other, less routine problems. Workers can rely on tacit knowledge until a breakdown occurs, when some assumption within which the tacit knowledge is grounded fails to hold, leading to reflection in action. This process of reflection brings explicit knowledge into play.

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. Explicit knowledge has the advantage that it is open to scrutiny and can be discussed among colleagues. 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 organisations using the notions of domain construction [20], and community of practice learning, which draws on a number of sources [6, 12, 17, 19]. We intend the term "group" to encompass both officially sanctioned groups within the organisation, often referred to as teams, and also unofficial groups emerging around shared interests, usually referred to as communities of practice [12].

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 [16] 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 Takeuchi [15] in their account of the knowledge creation spiral, emphasize the importance of making knowledge explicit

so that it can be shared. Traditionally, in the organisational learning field, tacit group knowledge was accorded a lower status, as it is hard to capture, quantify and share. Cook and Brown [6] suggest that this account underplays the importance of tacit group knowledge and argue for its equal importance. They use the term "organisational 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 artefact develop shared meaning. This is particularly noticeable in communities of designers. Schön [19] describes how within a design community, designs 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 [17] 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 appropriately recognize the importance of group tacit knowledge and leaning, which we refer to as community of practice learning.

2.3. Organisational level learning

Perspective taking [2] develops a model of how learning can occur across distinct communities in an organisation. 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 companies. 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 complex 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 making 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 [1] account of double loop learning. They distinguish single loop learning which drives incremental organisational change, from double loop learning which presents itself as more radical strategic change. Double loop learning requires an organisation 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 your current work practices, but also, on the way in which you actually evaluate your success. Within ENRICH we develop support for taking other community perspectives and building shared cross community perspectives.

3. Enriching work representations through organisational learning

The ENRICH approach aims to provide support for the four kinds of learning outlined above (reflection in action, domain construction, community of practice learning and perspective taking) and capture knowledge naturalistically as a consequence of this support. In particular, two kinds of support are provided to workers. 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. These work representations may be important documents or tools used as part of their work. Second, work representations are semantically linked to formal knowledge structures allowing workers to search for semantically related resources. These tools provide support for collaborating and learning in the context of work. As a consequence of these tools being used in the organisation, the representations of work become "enriched" in two ways. First, through collaboration using the asynchronous discussion tools, informal discussion evolves around the work representations enriching them with context. Second, over time further work representations become linked to the formal knowledge model, and the knowledge model itself is extended to reflect the incremental evolution and formalization of organisational knowledge, enriching the

work representations with formal semantics. These representations of work, enriched as a consequence of supporting work based learning, themselves become part of the organisational memory, and therefore become part of the learning resources of the organisation. ENRICH therefore supports organisational learning in two interrelated ways:

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

Two main kinds of design activity occur when deploying our approach. This section describes design activities concerned with identifying or designing representations of work within the organisation, and the collaborative activities that occur around them. The next section describes construction of the formal knowledge model to which the representations are linked. Although these two tasks are described separately, they are in practice intertwined.

Our approach centres around work representations. Before describing their identification or design 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 [5] 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. Second, in terms of the role that it plays within the workplace we intend WR to have a narrower definition than document. Knowledge crucial to the organisation is captured, stored and accessed from WRs within the organisation. Additionally, the WRs are central to work activity, rather than functioning as receptacles of information. These WRs facilitate the negotiation of shared meaning, enabling communities to come to some degree of consensus on a particular domain [5]. To reserve the use of the document in its more general form we adopt the term work representation (WR) to describe documents of work that may be instantiated in a range of media and serve an active rather than just recording role in the workplace.

Our first task when analyzing the organisation is to identify these WRs that function as community support tools and identify how they are incorporated with work practices, who uses them, and the nature of the collaborative activities they support. Our integration of tools in a number of organisations has identified specific examples of WRs with a number of common characteristics:

- WRs support specific work cases or episodes.
- WRs are used to capture ongoing problem solving activity as well as final solutions.
- Individuals or team members have explicit responsibility for a WR and its associated tasks.
- Solving a problem captured in a WR requires collaboration.

We will now illustrate the identification and redesign of a WR using one of the case sites from our project. In our case site in the aerospace industry, the identified WR was an existing paper based team planning tool. A team plan is completed periodically to support reflection on current team performance, set near term targets, actions for delivering them and measurements of success. Preparing the plan is a collaborative activity involving the whole team, and a facilitator, who is the member of the team responsible for recording contributions in the plan. Each new plan builds on its predecessors, and the ongoing experiences of the team.

When designing the digital version of the team planning tool, a crucial issue to investigate was the nature of the collaboration and communication surrounding the use of the plan. This is because collaboration reveals the context and reasoning behind a solution, and a solution in context is a far more effective learning resource. There is a great deal of support for the idea that solving real world problems should be thought of as situated and collaborative [12]. Important lessons from an earlier problem solving episode are inextricably linked to context and much is lost if they are decontextualized into a canonical form [4]. As we

wish to capture the reasoning behind a solution we need to know how collaboration currently occurs (e.g. via email, face-to-face, written memos, etc.) and how it can form part of the captured solution. Our aim to capture important events within the collaboration also has the dual role of serving as a design rationale for the life of the current problem, supporting those involved in reaching an optimal solution. With our case within the aerospace industry, face-to-face collaboration forms an important role in constructing a team plan. Contributions to the meeting are recorded by a facilitator. Results of the meeting, including targets are publicly displayed and periodically updated, often as a document for distribution, and on the team notice board. Current performance against the target is also presented on the notice board.

Although one of our aims is to capture and share the collaboration and rationale behind a solution, enforcing new patterns of work to make the capture process easier would not be appropriate. For example, if synchronous communication is the current way of working, then trying to force workers to communicate asynchronously would be doomed to failure. We need to understand the role of communication in the activities around the WRs, and consider how they can be supported by, incorporated with, and improved by, our technology.

Face-to-face collaboration at planning meetings continued with the new tool, with contributions being entered by a facilitator. An added benefit of the new approach is that capturing and laying out the rationale in a more structured way helps to ensure that alternative suggestions are adequately explored in the meeting. The redesigned planning tool is shown in figure 1, modified slightly for reasons of confidentiality. The centre of the window shows the new version of the previously paper-based planning document. The discussion and rationale space is shown to the right. This document was prepared using an extended version of the Digital Document Discourse Environment (D3E) [21].




Figure 1. A redesigned planning document (in the centre), with a navigation bar (to the left) and an associated discussion and rationale space (to the right).

Another important design issue was the development of the initial seed structure of the discussion and rationale capture component of the WR. This initial seed was based on the identified stereotypical steps or issues involved in the activity the WR supported. These were identified from interviews with workers and an examination of the original document structure. Discussion seeds in this initial version covered issues or steps such as problem description, objectives and actions. In our team planning scenario (see figure 1), the discussion is seeded around the five company values (customer, people, performance, partnership and innovation). During the planning meeting this is used to record the main points raised, concerning the justification of the plan. During the following months, this area is used to record progress against, and discussions about the measurable objectives outlined in the plan. At the next planning meeting, the discussion area for the previous plan gives a record of how outcomes related to objectives, which is then taken into account when developing the new plan.

Many of the lessons learned when solving earlier problems, can be captured in working documents and their associated discussion and be a valuable resource to other teams and individuals when solving similar problems. Documents describing some past successful work episode, selected and approved by the organisation are often referred to as best practices. Our aerospace industry case had in place some form of best practice archive, though these were not structured with any particular reuse methodology in mind. Our aim in the long term is to extend the best practice archive through the selection and editing of important team plans. The discussion associated with the working document has two important roles at this point. First, it supports the best practice co-ordinator in ensuring the rationale and context behind a solution are presented. Second it can be used to seed important points for discussion, to be considered by a team when considering the best practice in the context of their own problem or plan. For the initial seed, existing best practices were redesigned in our new environment and a discussion seed was developed, based on problem solving steps identified in the use of best practices.




Figure 2. Overview of how work representations, best practices and the formal knowledge are interrelated.

Our generic model of how WRs and best practices are managed and associated with the formal knowledge model, and the required roles within the host organisation is shown in figure 2. 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 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 a company knowledge model and a team knowledge model. The company 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.

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 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 Gannt and Nardi [10], a local developer is generally a member of the team, but who has acquired knowledge of, or been given some training in the ENRICH 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 best practice coordinator. They detect best practice and then use the related WR to produce the best practice document. New best practices or changes in the organisation may lead to modifications of the company knowledge model. In the next section we will describe how the knowledge model is formulated.

4. Constructing and evolving formal knowledge

Within the organisational learning and knowledge management communities, an important distinction is made 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. The classification system used in a library could be thought of as an example of formal knowledge. The classification of a particular book in a library gives an explicit unambiguous description as to the content of the book. The advantage of formal knowledge is that the semantics, that would normally only be derived from human interpretation, are accessible to a computer. Within the organisational learning, and other related communities, these formal descriptions of some domain are referred to as knowledge models. Knowledge models contain taxonomies of concepts and rules specifying relations between them. Typically, this part of the knowledge model is known as an ontology. Representation of important aspects of a domain as an ontology allows a computer to perform a semantically driven search to find related resources. This approach, unlike electronic databases that primarily group documents, emphasizes developing a representation of knowledge to which sophisticated reasoning mechanisms can be applied. 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 potentially useful best practices. The role of the knowledge model is therefore to support effective reuse of knowledge contained in best practice documents, and previous working documents of the team.

The search for relevant best practices could be done as a text search or keyword search. This approach is inadequate for our purposes. A conventional search proceeds by looking for documents containing the requested text. In our case, the words on which the search is based (i.e. the words entered in the working document) are not necessarily words appearing in the most appropriate best practice document. Some process of inference, some knowledge, is required to make the connection between the working document and the best practice. The search process is therefore far more than matching text and can draw on quite abstract concepts. The process of constructing the knowledge model can also have a beneficial side effect of mapping out the learning resources of the organisation, allowing them to identify their strengths and areas of their work that need more support.

Our construction of knowledge models is supported the WebOnto tool [8], which allows the collaborative editing and viewing of models via a conventional web browser. Figure 3 shows a screen snapshot of WebOnto. The represented ontology relates to our team planning scenario. The left part of the larger window contains a list of all the classes defined within the model. The right part of the window allows the classes and instances to be graphically browsed and edited. Each node within the ontology contains slots and values of the ontology. The smaller window to the front of figure 3, provides a detailed view of the customer-best-practice node.




Figure 3. A screen snapshot showing part of the ontology underlying the selection and representation of documents in our team planning scenario.

Within our approach, the formal model represents three kinds of knowledge: process, domain and media. Process knowledge describes how a best practice relates to a working document. Domain knowledge describes the domain of work, such as work products and machinery. Media knowledge describes where resources can be found. Our approach to building the knowledge models always kept in sight the way in which the knowledge would be used: to successfully discriminate between and select best practices based on information in the working document. Some characteristics of the organisation, its structure and products, although important in other situations gave little leverage in the selection of best practices and therefore played a lesser role in the knowledge model. This aspect of our approach can be described as

minimal ontological commitment [23]. Also, as will be described below, the relative importance of the process knowledge and domain knowledge differed considerably between the test sites. Determining the significance of each of the three kinds of knowledge, and determining their role in finding best practices involved a great deal of preliminary work observing work practices, interviewing staff and assessing the documents and how they are used. We will now briefly describe the three kinds of knowledge and the role they play.

Process knowledge describes the inferential process of how a best practice can be selected as appropriate for the current problem. The nature of this part of the knowledge model was strongly influenced by the structure of the planning or problem solving process as reflected in the working documents. These provide concrete evidence as to how the process works, where within the process best practice knowledge can be useful, and the reasoning behind the existing use of prior knowledge when solving problems, either from an existing archive or personal experience. This was used to identify abstractions of the problem solving domain utilized by workers and their application of heuristics. The process of selecting best practices could then be embedded within an understanding of the problem solving or planning process as a whole. Then, turning specifically to the selection of best practices, development of the process knowledge model focussed on how information in the working document could be used to rule in or rule out a potential best practice. This drew on formal methodologies or classifications embedded in the current use of the working document, and newly identified discriminating characteristics of best practices lying outside the formal methodology.

On analysis, the structure and terminology of each kind of document reflects a formal methodology of its use, which is also reflected in the structure of the best practice archive. In our team planning case, the planning process is structured around five values, serving as dimensions along which their current and target performance can be measured. One value is “Customer” which measures knowledge of the customer and the quality of procedures in place to meet customer needs. The team’s current and target performance on these values is scored on a scale from one to nine. Most best practices are classified as to the value they address, and some have “before and after” scores giving an indication of the benefit the best practice can provide for the team.

Our analysis augmented this formal methodology with further attributes that helped to improve discrimination between best practices. These focused on two issues. First, we represented the range of lessons embedded within best practice (represented in the lesson-learned concept). These included for example the benefits of a regular maintenance agreement. Second, we represented the problems this lesson could solve and the nature of the solution state (represented as *best-practice-problem* and *best-practice-solution*). The planning process and the utility of lessons learned tended to be fairly independent of the domain of expertise of the workers. A marketing team could learn from a manufacturing team and vice versa. This contrasted sharply with the other two test sites where domain knowledge played a far greater discriminating role.

Domain knowledge is used to represent the domain of expertise within which the members of the organisation worked. This kind of knowledge was found to be relatively less important in our aerospace case site, as many of the lessons and concepts related to team planning were independent of the particular domain of work. The opposite was true in one of our other cases in which the knowledge of how to fix faults on various kinds of mining machine was represented in the ontology. In this case the applicability of many of the lessons to be learned tended to be related to background domain knowledge concerning the type of machine and how it was being used.

Media knowledge describes where knowledge and learning resources can be found. This included not only the URLs of documents, but also external pointers to books, manuals, training resources and contact details for people able to offer particular kinds of consultancy or support.

5. Supporting individual, group and organisational learning: a summary

The ENRICH framework provides support for the four kinds of learning identified earlier. Crucially the ENRICH approach supports these four kinds of learning and captures new knowledge as a consequence of this support. This is summarized in table 1.

On 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.

| | Learning type | How supported | How captured |
|---------------------|--|--|--|
| Individual | Reflection in action [18] | 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 [20] | Elaboration and negotiation of domain concepts around work representations linked to formal knowledge models. | Evolution of group knowledge models and associations between knowledge models and WRs. |
| | Community of practice learning [6, 12, 17, 19] | 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 seamlessly within the group. |
| Organisation | Perspective taking [2] | 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. |

Table 1. Types of learning supported and captured in ENRICH.

Domain construction, the process by which teams negotiate and extend their team knowledge is supported by tools based on WebOnto [8] described above, 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, on 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 customises the WR and its integration with other work resources, illustrating the group's work genre. The shared WRs may also contain important artefacts that support the sharing of group tacit knowledge. One of our case sites embeds engineering models into the WRs as a method for sharing ideas that could otherwise be difficult to articulate. Changes to the WRs and their use as a method to share complex artefacts 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 organisation, particularly through the best practice archive. Once again, tools based on WebOnto [8] 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 best practice coordinator when modifying the best practice archive or company knowledge model.

6. The ENRICH development model

The ENRICH development model involves three steps, based on the Seeding, Evolutionary growth, Reseeding (SER) model [9]. The SER model was proposed to describe the importance of appreciating the

necessary evolution of systems that are embedded in organisations. The designer can never fully envisage how the system will come to be used. The users themselves must have a stake in the evolution of the system, and the overall design must be flexible and amenable to evolution.

Seeding creates through a process of participatory design, an initial state encompassing tools and an archive. Evolutionary growth occurs through use. Emerging scenarios of use lead to modifications of tools and the development of the archive as further knowledge is articulated. The reseeding process occurs after some period of use and can be thought of as a process of revolution rather than evolution. For example, key concepts in the formal knowledge repository may need to be reorganized, or extensive changes made to tools in response to emergent work patterns. Sections 3 and 4 described how we seed the organisational memory in terms of the identification or development of WRs, structuring the discussion and rationale capture space, and construction of the initial knowledge model. The evolution of the organisational memory tends to be left to the host organisation, with limited developer support. The organisational memory evolves through the way knowledge is captured as a consequence of supporting its construction, and with the support of local developers and the best practice coordinator. The evolution process was described in the previous section.

Reseeding involves more radical changes to central concepts in the knowledge model, overhauling the methodology, and redesigning (or possibly even differently identifying) the WRs of the organisation. Reseeding would generally be conducted by the organisation in collaboration with the original designers. The nature of the reseeding process would be motivated by the results of the evolution process and the ascertained success of the tools as indicated by an evaluation process.

The adoption of ENRICH by an organisation requires a significant allocation of resources, therefore the learning captured and supported has to be of high value to the organisation. The kind of knowledge or learning in an organisation that it would be appropriate to support using ENRICH would be learning or knowledge that is rare and without which the organisation would be unable to function effectively. Brooking [3] describes these as critical knowledge functions. What constitutes critical knowledge will depend on the context in which the organisation operates. If the organisation is a company that produces some product, then knowledge related to the development, manufacture, and marketing of that product will probably constitute many of its critical knowledge functions. Before starting to deploy ENRICH within an organisation, using the SER approach, the first step would be to carry out a preliminary analysis, which would identify critical knowledge functions. It is then necessary to determine whether these critical knowledge functions can be supported by ENRICH. This knowledge must be supported, or have the potential to be supported by WRs that can be used collaboratively. Some other important criteria relate to the nature of the critical knowledge. At least some of the critical knowledge should be formalizable in order that a seed ontology can be developed. This knowledge should also evolve through work, as ENRICH is specifically designed to support learning. If knowledge is static and well-organized, and even though the knowledge is critical, there is no need to learn, then a document or information retrieval solution can be adopted instead. Finally, the knowledge, though evolving, should be relatively stable. Knowledge cannot be built upon if there is radical disagreement between workers, and the knowledge is in a constant state of flux. The ENRICH methodology assumes that new knowledge is generated on the boundaries of current knowledge.

7. Conclusions

The ENRICH tools and methodology are designed to support organisational learning around representations of work. As a consequence of this, the work representations become enriched with context and semantic associations, thereby evolving the organisational memory. The four types of learning 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 organisation as a whole. Our approach of enriching representations of work through organisational learning is currently being tested in a range of industrial applications.

Acknowledgements

This research is funded by the ESPRIT programme for IT for learning and training in industry, project P29015.

References

- [1] C. Argyris and D.A. Schön, *Organisational Learning II: Theory, Method and Practice*, Addison Wesley, Reading, MA, 1996.
- [2] R.J. Boland, and R.V. Tenkasi, Perspective making and perspective taking in communities of knowing, *Organization Science* 6 (1998), 350-372.
- [3] A. Brooking, *Corporate memory: Strategies for knowledge management*, Thomson, London, 1999.
- [4] J.S. Brown and P. Duguid, Organizational Learning and Communities of Practice: Towards a Unified View of Working, Learning, and Innovation, *Organization Science* 2 (1991), 40-57.
- [5] J.S. Brown and P. Duiguid, The social life of documents, *First Monday* 1 (1996), <http://www.firstmonday.com>.
- [6] S.D.N. Cook and J.S. Brown, Bridging epistemologies: The generative dance between organizational knowledge and organizational knowing, *Organization Science* 10 (1999), 381-400.
- [7] D.K. Detterman, The Case for the Prosecution: Transfer as an Epiphenomenon, in *Transfer on Trial: Intelligence, Cognition, and Instruction*, D. K. Detterman and R. J. Sternberg, eds., Ablex Publishing Corporation, Norwood, New Jersey, 1993.
- [8] J. Domingue, Tadzebao and WebOnto: Discussing, Browsing, and Editing Ontologies on the Web, in *Proceedings of the 11th Knowledge Acquisition for Knowledge-Based Systems Workshop*, April 18-23, Banff, Canada, 1998.
- [9] G. Fischer, R. McCall, J. Ostwald, B. Reeves and F. Shipman, Seeding, evolutionary growth and reseeded: Supporting the incremental development of design environments. in *Human Factors in Computing Systems (CHI '94)*, April 24 - 28, Boston, MA, 1994.
- [10] M. Gantt and B. Nardi, Gardeners and Gurus: Patterns of Cooperation Among CAD Users, in *Human Factors in Computing Systems (CHI '92)*, May 3-7, Monterey, CA, 1992.
- [11] T.B.K. Ivergård, Measures to balance the demand for skill in a fact changing technological business environment and the supply of skill available in the existing work force, in *Human Factors in Organizational Design and Management - VI*, P. Vink, E. A. P. Koningsveld and S. Dhondt, eds, Elsevier, Amsterdam, 1998.
- [12] J. Lave and E. Wegner, *Situated Learning: Legitimate Peripheral Participation*, Cambridge University Press, Cambridge, UK, 1991.
- [13] B. Levitt and J. March, Organisational learning, *Annual Review of Sociology* 14 (1988), 319-340.
- [14] R. McDermott, Why information technology inspired but cannot deliver knowledge management, *California Management Review* 41(1999), 103-117.
- [15] I. Nonaka and H. Takeuchi, *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*, Oxford University Press, Oxford, UK, 1995.
- [16] J. Orr, Sharing knowledge, celebrating identity: War stories and community in a service culture. in *Collective Remembering: Memory in Society*, D. S. Middleton and D. Edwards, eds., Sage Publications, Beverly Hills, CA, 1990.
- [17] J.A. Raelin, A model of work-based learning, *Organization Science* 8 (1997), 563-578.
- [18] D.A. Schön, *The Reflective Practitioner: How Professionals Think in Action*, New York, Basic Books, 1983.
- [19] D.A. Schön, Designing: rules, types and worlds, *Design Studies* 9 (1988), 181-190.

- [20] T. Sumner, The high-tech toolbelt: A study of designers in the workplace, in *Human Factors in Computing Systems (CHI'95)*, May 7-11, Denver, CO, 1995.
- [21] T. Sumner and S. Buckingham Shum, From Documents to Discourse: Shifting Conceptions of Scholarly Publishing, in *Human Factors in Computing Systems (CHI '98)*, April 18-23, ACM, Los Angeles, 1998.
- [22] T. Sumner, J. Domingue, Z. Zdrahal, A. Millican and J. Murray, Moving from On-the-job Training towards Organisational Learning, in *Proceedings of the 12th Banff Knowledge Acquisition Workshop*, October 16-22, Banff, Canada, 1999.
- [23] M. Uschold and M. Gruninger, Ontologies: principles, methods and applications, *The Knowledge Engineering Review* 11 (1996), 93-136.