Integrating working and learning: a document enrichment approach

Paul Mulholland, Zdenek Zdrahal, John Domingue and Marek Hatala

Knowledge Media Institute

The Open University

Milton Keynes, MK7 6AA, UK

http://kmi.open.ac.uk/

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

+44 1908 65-3800 (Fax: 3169)

Integrating working and learning is seen as a desirable alternative to traditional training regimes. An integrated approach to working and learning is more gradual, contextual, problem oriented and cost effective. Knowledge technology aims to catalyze workplace learning, but requires the right organisational culture and brings additional costs regarding the articulation, representation and transmission of knowledge. Our approach reduces these costs by making articulation a natural part of collaborative working, designing for both informal and formal knowledge, and facilitating the transition of socially situated knowledge through enriched documents. To be successful, our approach has certain prerequesites concerning organisational culture, and the nature of shared documents, organisational knowledge and work activities

1. Introduction

Companies and therefore their workers need to adapt to ever changing demands. A traditional top-down approach to this problem is for a company to identify 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. This model has a number of shortcomings. First, training has been shown to be ineffective. Detterman (1993) 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 (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 workers can play in determining their own development.

Many have recognized the potential to harness and support the learning that naturally occurs between colleagues within the workplace. This process, termed organisational learning, involves workers learning and sharing ideas as a normal part of the social and collaborative nature of their activities. It appears beneficial to nurture this process as, compared to conventional training, it is more responsive to immediate needs, being driven by problems occurring in the workplace, as well as being gradual, ongoing and conducted within the work context.

For organisational learning to occur, a number of social and cultural conditions have to be in place. First, this learning primarily occurs within a community of practice (Lave and Wegner 1991). These communities do not necessarily equate to the formal team structure of the organisation. Organisational learning emerges in collaborative activity, which requires an appropriate collaborative culture (Aryris and Schön 1996) and a process of socialization (Nonaka and Takeuchi 1995). Another common feature of organisational learning is the reuse and reworking of past experience and solutions (Argyris and Schön 1996, Levitt and March

1988). Organisational learning builds on, questions and modifies previous solutions and ideas. Although there are many examples of successful learning communities (Lave and Wegner 1991, Orr 1990), to be of large potential to the organisation, knowledge needs to be efficiently transferable, across space and time.

Knowledge technology aims to catalyze the organisational learning process for the service of the organisation, but there are two hurdles that knowledge technology needs to negotiate in order to be successful. First, knowledge technology requires the social and cultural conditions for organisational learning to be in place (McDermott 1999). Deploying the technology will not bootstrap an appropriate collaborative climate. There are a number of examples of how knowledge technology will otherwise fail (e.g. Orlikowski 1992). Second, knowledge technology imposes extra demands on the workforce. These demands are concerned with the articulation, representation and transmission of knowledge.

For knowledge to be stored and represented, it must first be articulated. This can be problematic even if the necessary cultural conditions are in place. Workers can find it difficult to say what they know. Articulated knowledge then needs to be represented in order that it can be stored and retrieved. This knowledge may be formal or informal. The advantage of formal knowledge is that it can be can be objectively indexed as part of the representation process, but much of the valuable knowledge of a company is tacit, hard to formalize, and therefore hard to classify. Represented knowledge then needs to be transmitted across community boundaries and appropriately interpreted by the recipients. This is hard to achieve as knowledge tends to be "sticky" (McDermott 1999) and difficult to extract from the social context in which it was created. Knowledge technology therefore aims to accelerate organisational learning but imposes costs of its own.

Our document enrichment approach to organisational learning (Sumner et al. 1999), like all knowledge technologies, cannot create the conditions of organisational learning where they do not already exist, but does aim to reduce the costs of articulation, representation and transmission. Articulation is supported by making it a natural consequence of collaboration within the workplace. Representation is supported by capturing both formal and informal knowledge and adopting a model of how the formalization of knowledge in a domain can be captured. The transmission of sticky knowledge is supported by sharing knowledge captured in key documents of the organisation, enriched by their context.

The rest of this paper is structured as follows. The next two sections describe the design of the document enrichment architecture. Section two describes the identification and redesign of shared documents to capture informal knowledge. Section three describes the development of the underlying formal knowledge structure. Section four describes the evolution of the organisational memory. Section five discusses the applicability of the approach.

2. Characterizing shared documents

Two main kinds of design activity occur when deploying our approach. This section describes design activities concerned with identifying and redesigning key documents within the organisation, and the activities that occur around them. The next section describes construction of the knowledge model. Although these two tasks are described separately, they are in practice intertwined. Our tools and methodology aim to be an organisational learning solution for document centred organisations. By this we mean that knowledge crucial to the organisation is captured, stored and accessed from key documents within the organisation. Additionally, the documents are central to work activity, rather than functioning as receptacles of information. These documents facilitate the negotiation of shared meaning, enabling communities to come to some degree of consensus on a shared view of a particular domain (Brown and Duguid 1996). We term these documents, "working documents". Our first task when analyzing the organisation is to identify these working documents that function as community support tools and idenitify who they are incorporated with work practices, who uses them, and the nature of the collaborative activities in which they support.

Our integration of tools in a number of organisations has identified specific examples of working documents with a number of common characteristics:

- Each document supports a specific "case".
- The document is used to capture ongoing problem solving activity as well as the final solution.
- Individuals or team members have explicit responsibility for the working document and its associated tasks.
- Solving a problem captured in a document requires collaboration.

In our test site in the aerospace industry, the identified working document was a 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. For a second test site in the mining industry we identified fault reports which are initiated whenever one of the main mining machines presents a fault. This document is used to for communication, and record communications, between the technician responsible for investigating the fault, and other experts. This records all repair activities until the machine is back in service. A third test site, a customer support department, initiate a new case document when a customer contacts them with a new problem. This document, assigned to a particular technician, records all work carried out in dealing with the problem, including all interactions between the technician and the customer. The case is closed and archived when the customer query has been satisfied. A summary of the test sites is presented in table 1.

Test site	**Working document**	**Collaboration**	**Best practices**
Aerospace company	Team planning book for reflecting on performance and setting near term targets	Face-to-face planning meetings. Report and discussion of, ongoing team actions toward meeting objectives.	Growing archive of concise best practice descriptions grouped around company values.
Customer support centre	Electronic case files of customer problems from request to resolution.	Mainly asynchronous collaboration between colleagues focused around current customer problems.	On-line FAQs used as internal learning resource and publicly accessible by customers
Mining company	Fault logs pertaining to the main mining machines.	Synchronous communication between workers above and below ground concerning current machine fault.	Vast archive of machine faults logs for use by mining technicians

Table 1. Summary of identified working documents, collaboration patterns and use of best practices across the test sites.

When determining how these working documents could be redesigned, a crucial issue to investigate was the nature of the communication surrounding the use of the document. In each case, the working document was used to support some problem solving process (e.g. fixing a fault, constructing a plan). In each case, solving the problem involved certain kinds of collaboration around the document, and around the problem to which the document referred. Some of this collaboration was already captured within the document in some form.

It was considered an important part of the design process to investigate the nature of this collaboration. 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 (Lave and Wegner 1991). 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 (Brown and Duguid 1991). For a solution to be reusable it is necessary to view it in context and appreciate the reasoning behind that solution as found in the collaborative activities that led to its construction. Therefore, when designing around the working documents identified in an organisation, it is necessary to understand the collaboration and communication that surrounds them. 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.

In the initial analysis of the scenarios where our tools have been deployed, we investigated the nature of the communication surrounding their working documents and found some interesting differences. 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 distributed document and on the team notice board. Current performance against the target is also presented on the notice board. Similarly, synchronous communication formed an important role within our test site within the mining industry. Although all decisions and work concerning a machine fault were recorded on an electronic fault log, technicians used the telephone to discuss the fault. The fault log would then be updated later. This is because the technician usually identified the fault when underground, and the telephone was the only form of communication with workers above ground. In the customer helpline case-site, unlike the first two, a great deal of the collaboration around the problem tended to be asynchronous using email. The telephone is an important form of customer contact, therefore most collaboration between colleagues tended to use email to keep the telephone lines free.

These observations have a number of consequences for the design and integration of our tools. 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 working documents, and consider how they can be supported by, incorporated with, and improved by, our technology.

As we mentioned earlier, team planning meetings in our aerospace industry test site are conducted face-to-face. This of course continued with the new tools, 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 document is shown in figure 1, modified slightly for reasons of privacy. 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 is prepared using an extended version of the Digital Document Discourse Environment (D3E) (Sumner and Buckingham Shum 1998). In the coal mining case, asynchronous communication between the underground workers and the control centre was provided for the first time, along side the existing telephone link. Although asynchronous communication was not originally used, it may have advantages, for example due to the intermittent availability of staff during a shift. These trials are ongoing, and we await the outcome. If asynchronous communication proves successful then rationale capture is well supported and mining staff have a new useful tool. If the new link is not used, then a summary of the communication and rationale is still captured in the fault log. This would provide a more rationalized account but is far better than coercing staff into a suboptimal way of working, which in the end may be subverted and leave us with tools no one uses. The customer helpline case is more straightforward due to the existing preference for asynchronous communication, but careful integration with existing tools still needed to be addressed.

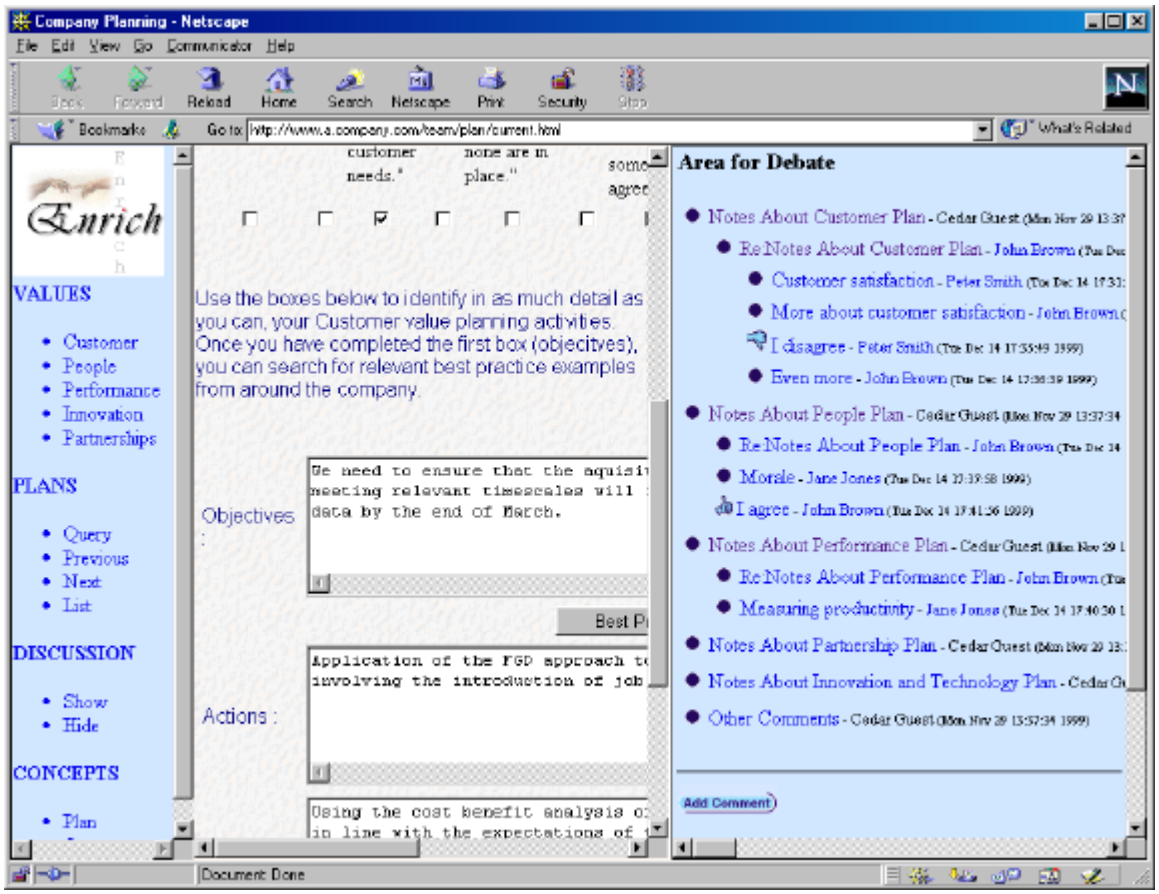

Figure 1. A redesigned planning document (in the centre), with a navigation bar (to the left) and an associated discussion and rationale space (to the right).

Another important design issue was the development of the initial seed structure of the discussion and rationale capture component of the new working document. This initial seed was based on the identified stereotypical steps or issues involved in the activity the working document 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.

We have described how we capture key activities in each organisation, occurring around working documents, including solutions to problems and the rationale behind them. We will now move on to consider how these form a reusable resource. Our document centred organisations have a further characteristic in common - the problem solving process draws on knowledge of previous solutions.

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. The three test sites mentioned already had in place some form of best practice archive, though these were not necessarily well developed and often captured little of the reasoning behind a solution, which is so important for reuse. Our aim in the long term is to extend the

best practice archive by selecting and editing important working documents. 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.

3. Constructing the formal knowledge model

The second part of the seeding process involves constructing the initial formal knowledge model. The role of the knowledge model is to establish viable connections between descriptions of work being entered in the current working document 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. An overview of how the knowledge model relates to the documents of the organisation is shown in figure 2.

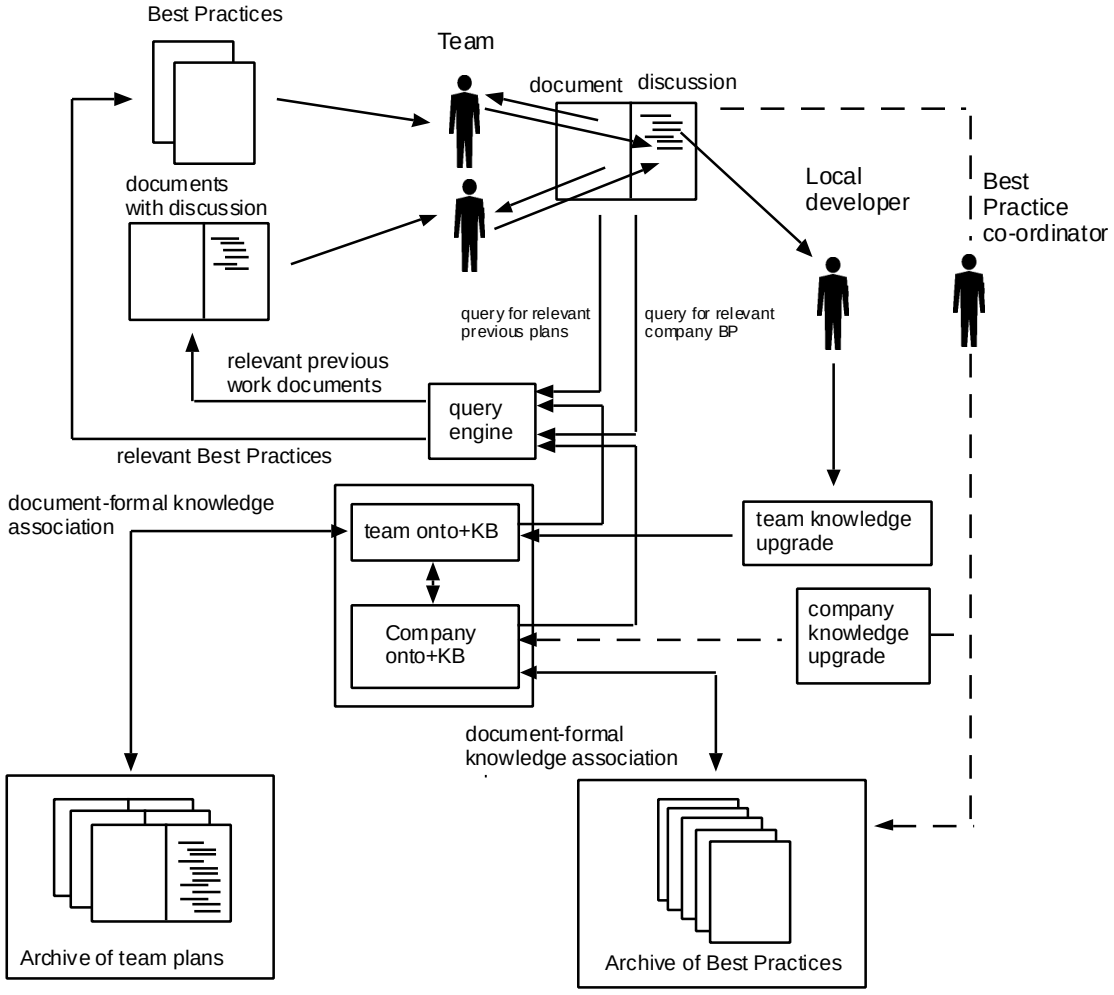

Figure 2. Overview of the document enrichment approach.

The search for relevant best practices could be done as a text search or keyword search. This approach is though 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. It is worth noting at this point that with our approach, there is not a separate search engine, distinct from the working document that is used to find the best practices. The working document effectively doubles as the search engine interface, therefore the search is based on what is found in the working document. Best practices are suggested to the team or individual as they go about their work. The required 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 (Domingue 1998), 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.

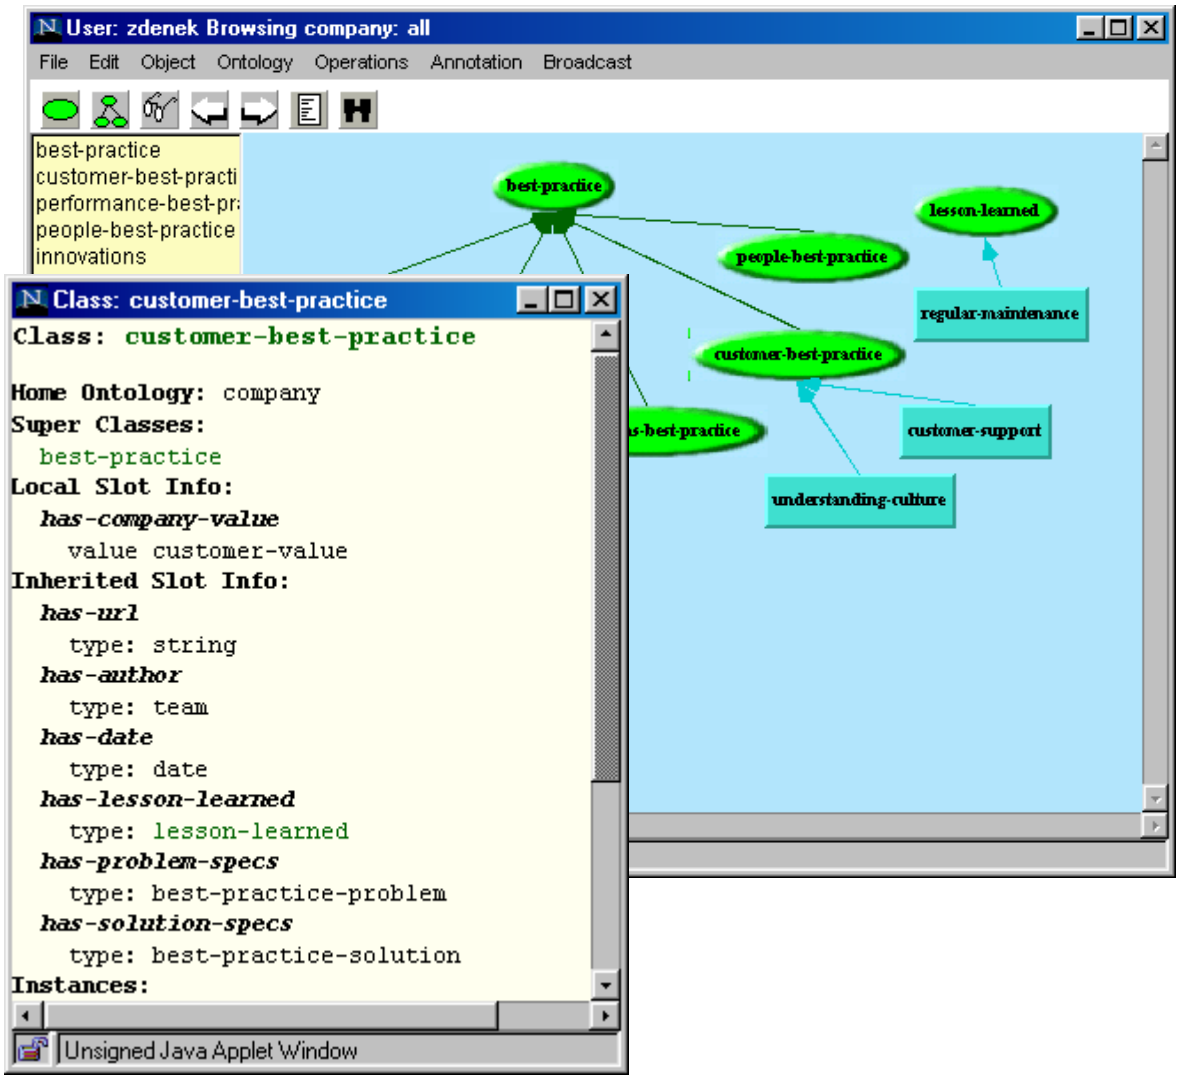

Figure 3. A screen snapshot showing part of the ontology underlying the selection and representation of documents in our team planning scenario.

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 (Uschold and Gruninger 1996). 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 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` class). These included for example the benefits ofa 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 far more important in the mining and customer support test sites. Within these scenarios the extent to which knowledge could be shared was far more dependent on domain characteristics such as the type of machine used. Domain knowledge within these formal models was used to scope or limit the relevance of best practices to the current problem. These test sites had elaborate problem solving methodologies, but in use, the methodology was instantiated within particular machine components and faults, and opportunities for abstraction were more limited. Both test sites had legacy domain knowledge models that could be incorporated within the archive seed.

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.

4. Evolution of the organisational memory

The SER (Seeding, Evolutionary growth, Reseeding) model (Fischer et al. 1994) 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. The last two sections have described how we seed the organisational memory in terms of document redesign, structuring the discussion and rationale capture space, and construction of the seed knowledge model. We now move onto the second part of the process: the evolution of the organisational memory and how it is supported. The evolution of the organisational memory tends to be left to the host organisation, with limited support. The organisational memory evolves in four ways:

- Additions to documents and their discussion space.
- Additions and deletions to the best practice archive.
- Modifications to the knowledge model.
- Indexing new documents.

The most prevalent form of evolution is contributions made, as part of the normal working process, to documents and their associated discussion and rationale capture space. This provides teams with a record of their ideas, work and progress over time. This evolving historical collection of working documents provides a useful reference for teams or communities in the organisation, and is archived for this purpose. Although of direct value to the workers concerned, these documents do not make a very sharable resource in their raw form.

Best practices describe knowledge that can be shared within the organisation, and in some cases, even across organisational boundaries. Most of these evolve from working documents, though some process of selection and editing needs to be applied. Many working documents will not contain lessons of utility to a wider audience. Those that do, may have limited impact in their current form. The process of selecting best practice from the growing array of documents produced in the organisation may take a number of forms. Documents may be submitted by teams or individuals when identifying one of their own work episodes as potentially useful best practice. This process often requires some form of incentive, even in a sharing culture. One of our test sites had an extensive incentive and reward scheme to promote the submission of best practice. Within our test sites, separate best practice teams were also common, whose purpose was to identify best practices within the organisation, so they could be given a wider audience.

Once a best practice has been identified it needs to be edited by the best practice coordinator to form a reusable resource. The nature of the editing process needs to be informed by an understanding of how and to what extent context is bound up with the lesson, technique or insight to be shared. The approach taken to editing a best practice lies on a context continuum. At one extreme there is the position that the identified knowledge can be extracted from the episode in which it occurred and shared in isolation of its context. At the other extreme is the position that all context, not matter how seemingly irrelevant, must be preserved as context is inextricable. The best practice coordinator may also make changes to the company ontology, as a results of knowledge contained in new best practices.

From our own perspective, we reject out of hand the first of the two positions, subscribing to the view that knowledge is intertwined with the context in which it was created (Brown and Duguid 1991). The opposing

position, though persuasive, is often neither desirable nor practical. The particular stance taken to this issue is dependent on organisational characteristics, and draws on similar issues to those that motivate the design of the formal knowledge model. Best practices within our team planning scenario maintain less context. The best practice is intended to have immediate impact, be thought provoking and describe techniques or insights that can have wide applicability across the organisation. In the other test sites, far more context is preserved as context is a crucial component of the knowledge. In each case, preserving the discussion space in some form is an important part of the design. The full discussion or rationale that accompanied the working document submitted or identified as best practice is pruned to leave some main issues of debate to act as a discussion seed related to how the lessons learned can be applied in varying situations. Evolution of the best practice archive will also periodically involve the deletion of existing best practices as they become surpassed by newly evolving best practices.

Contributions to documents and discussions sometimes lead to modifications of the knowledge model. Contributions may offer new insights into relations between concepts in the model, or newly developed concepts. Over time, groups establish and formalize new concepts and terminology relating to their work. Once established within the community they can be incorporated into the formal representation of their work. This evolving knowledge may be part of the methodology of how they perform their work (process knowledge) or an elaboration of their shared areas of expertise (domain knowledge). Also, new documents, particularly best practice documents need to be indexed in terms of the knowledge model. This work is carried out, drawing on the work of Gannt and Nardi (1992), by a local developer, who is a member of the work team, but has a greater familiarity with the technology and methodology. This formalization of their knowledge is recorded in their personal team ontology which forms a part of the overall company ontology.

Reseeding involves more radical changes to central concepts in the knowledge model, overhauling the methodology, and redesigning (or possibly even differently identifying) the key documents 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 the evaluation process.

5. Applicability of the approach

Our document enrichment approach aims, as do other knowledge technologies, to support the organisational learning process, but also reduce the additional costs imposed on the articulation, representation and transmission of knowledge. Our approach is though not universally applicable. From our experiences deploying the tools and methodology, a number of requirements have been identified. These encompass the nature of key documents within the organisation, the activities that surround them, and the characteristics of knowledge within the organisation. As with all organisational learning support tools, the social and cultural conditions (e.g. the existence of communities, collaborative working, etc.) have to already be in place.

In terms of document characteristics, the organisation needs to utilize working documents that capture, coordinate and guide ongoing work activity. These documents, though having a canonical, prescribed structure, must describe experience of some event covering some period of time, such as a planning or diagnostic process. There also needs to be a separate set of documents that serve as best practices, describing some past lesson or success.

The organisation needs to have the right activities surrounding these documents. Collaborative work must occur around the documents so that knowledge and practices can be naturalistically articulated. To capture and represent ongoing activities, not just solutions to problems, work must be integrated with the documents. This ensures that both formal and informal knowledge can be captured.

There are further conditions relating to the nature of the knowledge within the organisation, and how it is used, questioned and created. Some of the knowledge within the company needs to be formalizable in order that a seed ontology can be developed. Overall, the knowledge of the company needs to be changing, but relatively stable. The company needs to be open to new knowledge so that the organisation can learn, through document enrichment, but also there needs to be some stability, characterized by the gradual evolution, rather than constant revolution of company knowledge.

A wide range of organisations meet these broad document, activity and knowledge criteria, as have been found within the initial test sites for the approach.

6. Conclusions

Organisational learning, unlike conventional training is gradual, contextual and problem driven. Knowledge technology aims to facilitate organisational learning, but it needs the right social and cultural conditions in place, and incurs extra costs concerning the articulation, representation and transmission of knowledge. Our document enrichment approach aims to deliver the benefits of organisational learning while minimizing increased cost. Our approach has been deployed in three industrial settings, and although lowering the ongoing costs of using knowledge technology, has certain requirements in terms of the nature of documents in the organisation, the kinds of knowledge they are used to express, and the activities surrounding them.

Acknowledgements

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

References

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

Brown, J. S. and Duguid, P. 1991, Organizational Learning and Communities of Practice: Towards a Unified View of Working, Learning, and Innovation, *Organization Science*, **2,** (1), 40-57.

Brown, J. S. and Duguid, P. 1998, Organizing knowledge, *California Management Review*, **40** (3), 90-111.

Detterman, D. K. 1993, The Case for the Prosecution: Transfer as an Epiphenomenon, in D. K. Detterman and R. J. Sternberg (eds.) *Transfer on Trial: Intelligence, Cognition, and Instruction.* (Ablex Publishing Corporation, Norwood, New Jersey).

Domingue, J. 1998, Tadzebao and WebOnto: Discussing, Browsing, and Editing Ontologies on the Web. *Proceedings of the 11th Knowledge Acquisition for Knowledge-Based Systems Workshop* (April 18th-23r). 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)*, (April 24 - 28), Boston, MA.

Gantt, M. and Nardi, B. 1992, Gardeners and Gurus: Patterns of Cooperation Among CAD Users. *Human Factors in Computing Systems (CHI '92)*, (May 3-7), Monterey, CA.

Ivergård, T. B. K. (1998). 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 P. Vink, E. A. P. Koningsveld and S. Dhondt (eds) *Human Factors in Organizational Design and Management - VI*, (Elsevier, Amsterdam).

Lave, J. and Wegner, E. 1991, *Situated Learning: Legitimate Peripheral Participation*. (Cambridge University Press, Cambridge, UK).

Levitt, B. and March, J. 1988, Organisational learning, *Annual Review of Sociology*, 14, 319-340.

McDermott, R. 1999, Why information technology inspired but cannot deliver knowledge management, *California Management Review*, 41(4), 103-117.

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

Orlikowski, W. 1992, Learning from Notes: Organizational Issues in Groupware Implementation. Computer Supported Collaborative Work, *CSCW '92*, Toronto, Canada.

Orr, J. 1990, Sharing knowledge, celebrating identity: War stories and community in a service culture. In D. S. Middleton and D. Edwards (eds) *Collective Remembering: Memory in Society*. (Sage Publications, Beverly Hills, CA).

Sumner, T. and Buckingham Shum, S. 1998, From Documents to Discourse: Shifting Conceptions of Scholarly Publishing. *In the Proceedings of ACM CHI'91Human Factors in Computing Systems* (April 18-23) (ACM, Los Angeles).

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,* (October, 16-22) Banff, Canada.

Uschold, M. and Gruninger, M. 1996, Ontologies: principles, methods and applications, *The Knowledge Engineering Review*, 11(2), 93-136.