Moving from On-the-job Training towards Organisational Learning

Tamara Sumner, John Domingue, Zdenek Zdrahal
Knowledge Media Institute
The Open University, UK
{T.Sumner, J.B.Domingue, Z.Zdrahal}@open.ac.uk

Alan Millican, Jayne Murray
Virtual University
British Aerospace, UK
{alan.millican, jayne.murray}@bae.co.uk

Abstract. The process of acquiring knowledge cannot be separated from the process of applying it. Integrating working and learning is not a desirable luxury – it is a fundamental requirement for businesses to remain competitive. In this article we describe the ENRICH approach to supporting organisational learning. The enrichment of workproducts with context is aided by suite of web based tools which integrate HTML documents, discussion threads and knowledge models. The approach is described using one of the three ENRICH industrial case studies - the British Aerospace Team Workbook.

1. INTRODUCTION

> "Businesses spend up to $100 billion each year to train workers. Yet estimates are that less than 10% of this training transfers to the job. So business wastes $90 billion each year...."
>
> [Review on US training effectiveness (Detterman 1993)]

Current theories of learning reveal why this is so: the process of acquiring knowledge cannot be separated from the process of applying it. Integrating working and learning is not a desirable luxury – it is a fundamental requirement for businesses to remain competitive. In fact, (Brown and Duguid 1991) argue that learning is the essential bridge between working and innovating and that the three processes are inextricably intertwined. They draw on the empirical work of Orr (Orr 1995) to argue for a fundamentally situated view of learning. In this view, problems with the traditional 'knowledge transfer' model of on-the-job training include the separation of simplified abstract principles from the rich detail of actual practice and the separation of learners from the workplace community.

Brown and Duguid use their arguments to critique the inadequacies of business process redesign and training programmes based on 'transfer model' assumptions of essentially deskilled, overly

simplified models of work. Instead, they advocate that technology and business processes should support the existing rich learning practices within the (typically co-located) workplace community by enabling indivduals within communities to somehow record and share their experiences. This influential article has inspired many projects in the area of knowledge management and organisatioanl learning, including ours. Indeed, the approach of capturing and sharing existing experiences to enhance future practices is the key goal of many organisational learning approaches, particularly in the area of software engineering (see (Landes, Schneider et al. 1998) for examples of 'experience factories'). In effect, experience factories view learning as a process of incrementally building on existing, valuable 'traditions' within the workplace.

However, a key challenge faced by experience factories (and other forms of 'best practices' memories) occurs when the desirable traditions or best practices aren't widespread in the first place. Landes, et. al report the paradox they have encountered in their project: workers are eager to use the experience factory to learn about quality management issues but there aren't enough people experienced with the quality management practices to contribute to the experience base in the first place (Landes, Schneider et al. 1998). Thus, for various reasons, an organisation may sometimes need to, or desire to, deliberately 'transcend' existing traditions, and promote widespread new practices within the workplace. This need can arise from external forces (for instance, new EU safety codes that effect current manufacturing processes in fundamental ways) or internal goals (the desire to adopt Quality Management practices). Given these tensions, it is perhaps more appropriate to view learning as a 'dialectic between tradition and transcendence' and the challenge is to support both, yet deny neither[1].

In the traditional model of on-the-job training, to promote the new practices, workers would typically receive a pre-prepared course in the new regulations, procedures, or processes – often at a different location from their place of work – and be expected to apply this abstracted knowledge later in their workplace. As indicated by the training effectiveness figures above, often with disappointing results. The goal of our research is to use technology, and advances in workplace learning theories, to rethink current training practices to: (1) move from one-shot training courses to continuous learning processes and (2) move from training focused on the individual to incorporate work group and organisational learning.

The vehicle we use to conduct this research is the CEDAR Toolkit. CEDAR is an experimental system supporting the creation and maintenance of Contextually-Enriched Document ARchives providing:

- a document-centred discourse space for structuring discussions around representations of work;

- client tools enabling users to articulate and refine domain concepts by incrementally enriching their representations of work with related discussions and underlying knowledge models;

- knowledge delivery and access mechanisms supporting individual and group learning;

- an intranet-based organisational memory server for knowledge capture, structuring and maintenance.

This research is being conducted using a user-centred, case-driven approach focusing on the needs of specific workplace communities in three different industrial settings (see (Sumner, Domingue et al. 1998) for more info on these cases). In this article, we focus on our experiences from one of these cases, The Team Workbook. The Team Workbook is a product of BAe's Senior Management team and is designed to foster best practices through the intertwined teaching and use of planning methodologies such as Total Quality Management. In the near term, the goal of the Workbook programme is to transcend current practices by teaching practitioners how to apply quality

---

[1]   Some years ago, Pelle Ehn claimed that 'design is a dialectic between tradition and transcendance' (Ehn 1989). We think the same dialectic holds for learning.

management to their current work. In the long term, the programme's goals are to foster a learning organisation where individuals and groups continually reflect on and improve their work practices, and where distributed groups share and improve upon their 'traditions' and 'best practices'.

In the remainder of this article, we begin by analysing current theories in workplace learning to derive key learning processes that our approach must address. Next, we analyse the shortcomings of the current state-of-the-art from technical and cognitive perspectives. We then describe the CEDAR toolkit's technical and methodological objectives and illustrate the toolkit with a use scenario derived from the Team Workbook case. Finally, we reflect on the implications of our approach for training programmes in the workplace of the future and discuss future directions of our research.

2. NATURE OF WORKPLACE LEARNING

"Learning is the new form of labour" (Zuboff 1988) in a knowledge economy – it is absolutely vital that learning be effective and efficient. Yet, many industries rely on traditional 'school models' for most of their training needs even though workplace learning is fundamentally different from traditional school learning (Table 1). Empirical studies of professional practice, by ourselves and others (Lave 1991; Sachs 1995; Sumner 1995), show that while the focus is primarily on getting the job done, learning is inextricably intertwined with working. In order to do their job, professionals must continually learn to apply existing knowledge to routine or innovative situations and to construct new knowledge in response to changing workplace situations. Thus, learning is fundamentally embedded in ongoing work activities and these work activities, in turn, give rise to the problems driving the learning that must take place.

Table 1. School Learning versus Workplace Learning

	School	Workplace
Emphasis On:	Learning basic facts and skills	Getting the job done
Ultimate Goal:	"Knowing"	Developing "Best Practices"
Knowledge:	Static, Decontextualized, General	Dynamic, Situated, Practice-oriented
Topics / Problems:	Given by curriculum	Arise from and embedded in work situation
Scope of Learning:	Primarily Individual	Individual, Group, Organisation

Our research is concerned with building computational environments integrating both working and learning. Our previous work focused on supporting individual learning (Sumner and Stolze 1996; Sumner, Bonnardel et al. 1997), and knowledge modelling (Domingue, Motta et al. 1993; Motta and Zdrahal 1996; Motta 1998) in centralised work settings. In this project, we extend these previous efforts to the *organisational level*, and provide support for the integrated spectrum of learning activities (individual, group, and organisational) in *distributed work settings*. Before looking at the state-of-the-art, we will examine current theoretical perspectives and empirical findings of individual, group, and organisational learning to enumerate key challenges which an integrated approach must address.

2.1. Individual Learning: Becoming Reflective Practitioners

In his seminal book, "The Reflective Practitioner," Schön (Schön 1983) describes an action-breakdown-reflection cycle that underlies professional practice. In this cycle, practitioners engage in situated action until their expectations are not met and they experience a breakdown in the current work situation. At that moment, practitioners stop and reflect on how to overcome the breakdown before proceeding. These breakdowns in situated action present opportunities for learning because

there is an opportunity to construct new contextualized knowledge while solving a personally relevant problem (Fischer 1994).

However, detecting and overcoming breakdowns requires much skill and domain-specific knowledge. People newly hired into a workplace will often lack this necessary knowledge. Even "old-timers" can be challenged. In complex domains, no one person can possess all the knowledge necessary to complete a task (Rittel and Webber 1984). In dynamic domains, what constitutes "necessary" knowledge is continually changing. Practitioners, both old and new, need to continually learn and relearn how to: (1) recognise potential breakdowns, (2) identify knowledge relevant to the breakdown, and (3) apply this knowledge or construct new knowledge to overcome the breakdown.

– <u>Key challenges for supporting individual workplace learning are to support reflection-in-action processes</u> by helping practitioners analyse their work products and detect potential problems or opportunities. Systems need to deliver relevant information to practitioners that they may be unaware of to assist their reflections.

2.2. Group Learning: Becoming Communities of Practice

These days, teams and groups form the core work units in many industries. However, while teams themselves are widespread, examples of effective group practices are less so. The literature is filled with stories of dysfunctional group working practices (Curtis, Krasner et al. 1988; Grudin 1991). Studies reveal that effective groups are good at "perspective making" (Boland and Tenkasi 1995); i.e., they share customs, conventions and standard practices that help to get the job done more effectively (Brown, Collins et al. 1989; Lave 1991; Sachs 1995). It is precisely these often tacit customs and conventions that form "best practices."

Our empirical studies of groups suggest that supporting communities to evolve their own vocabularies and to elaborate them to create a shared domain model is a critical step towards creating a common perspective (Sumner 1995) and an effective "community of practice." Shared vocabularies improve group communications; shared practices improve group coordination during complex tasks. Over time, as communities engage in negotiation and reflection about how to do their job better, their vocabularies and domain models become more elaborate and formal; i.e., their tacit understandings are articulated and refined towards more explicit knowledge forms. In many cases, they enrich their tools and work products with formal representations of their domain models in order to have better tool support for their work practices. The articulation and use of explicit domain models reinforce best practices by aiding the consistent reproduction and interpretation of work products. We refer to these processes of elaborating vocabularies, negotiating, and enriching as "domain construction".

– <u>A key challenge for supporting group learning is to support domain construction processes.</u> Systems need to enable practitioners to articulate their informal understandings, to view and discuss their emerging ideas and domain models and, importantly, to incrementally modify domain models as their understandings change.

2.3. Becoming a Learning Organisation

Typically, organisations are composed of multiple interacting communities, each with highly specialised knowledge, skills, and technologies. Important tasks like product design and innovation in knowledge-intensive firms require these diverse communities to bridge their differences and integrate their knowledge and skills to create a new, shared perspective (Boland and Tenkasi 1995). Some researchers argue that this social sharing is the crucial first step towards knowledge creation (Nonaka and Takeucji 1995).

This bridging process is not so much one of passive 'sharing' as one of active 'perspective taking' and it is often complicated by the fact that a community's shared vocabulary or domain model is

often tacit, making it uninspectable and difficult for another community to understand. Communities may share similar words and concepts at the surface level, but may actually be using them in entirely different ways (a phenomena dubbed 'ontological drift' (Robinson and Bannon 1991)) resulting in communication and coordination problems. Sometimes this bridging is further complicated by time – the critical experts are no longer with the company or are otherwise unavailable for collaboration. Supporting such long-term, asynchronous collaboration is particularly important in industries needing specialised expertise that also rely on mobile and flexible workforces. Ultimately for perspective-taking to be successful, shared objects and work products must be re-interpreted and assigned a shared meaning – a process that usually requires much debate and negotiation.

– <u>A key challenge for organisational learning is to support perspective taking.</u> Systems need to support knowledge sharing across workplace communities *and* across time. However, sharing knowledge is different from simply sharing information – people need support for interpreting each others' perspective and for negotiating a new, shared perspective.

2.4. Integrating Individual, Group, and Organisational Learning

Building on Argyris and Schön's definition (Argyris and Schön 1978), we define organisational learning as: (1) a process that takes place through the agency of the individual members, (2) where individual and group learning experiences become incrementally embedded in organisational memory. An integrated approach is necessary because one cannot support organisational learning without supporting individual agency. However, we extend their definition to take into account the different types of learning processes at each of the three levels; i.e. individual, group, and organisational. Table 2 summarises requirements for our integrated approach.

Table 2. Requirements for an Integrated Organisational Learning Approach

Scope	Learning Process	Requirements to Support Process
Individual: Within an individual work session	Reflection-in-action	• detecting potential problems or opportunities to improve work products • actively delivering new information to support reflection
Group: Within a community	Domain Construction	• articulation and elaboration of domain vocabularies and concepts • negotiation about emerging domain concepts • enriching of work products and tools with domain models
Organisational: Across Communities and Time	Perspective taking	• sharing of knowledge and work products • products linked with supporting context to aid interpretation and negotiation

3. STATE OF THE ART

Here, we examine the state-of-the-art with respect to the requirements outlined in Table 2. We begin by looking at an important type of structured discussion space, design rationale systems, which our approach builds on. Then we analyse related work in organisational memories, learning-on-demand systems, intranet-based document management systems and proposed extensions to the Web that CEDAR builds on.

Design Rationale. Design rationale systems provide structured discussion spaces based on various notations such as issues, pros, and cons (Conklin and Begeman 1988), and questions, options, and criteria (Buckingham Shum 1996). These systems assume that as practitioners work, they also add their reasoning and justifications to the system using the provided notation and thus create an

'organisational memory' as they work. Experiences indicate that design rationale systems can be very useful for supporting long-term asynchronous negotiation and collaboration across time (Fischer, Grudin et al. 1992). However, one weakness of this approach is the lack of support for context: often there is little integration between work products and discussions about the products. This separation results in a loss of necessary context for understanding and interpreting both the design rationale and the products (Ruhleder 1994). Integrating work products with their related discussion is also important for sustainability of the rationale repository; integration makes the repository easier to access during actual work, which in turn promotes making further additions to it (Fischer, Lemke et al. 1991).

Organisational Memories. Organisational memories are motivated by the desire to preserve and share the knowledge and experiences that reside in an organisation. As such, most systems focus on capturing the knowledge, storing it, and making it accessible, rather than explicitly supporting the creation of new knowledge. By themselves, organisational memories are a necessary but insufficient step towards organisational learning. Several analyses based on case studies (Ackerman 1993) and critiques from social and psychological perspectives (Bannon and Kutti 1996) conclude that maintenance of contextuality is of crucial importance for supporting learning or 'active remembering'. Likewise, the generic nature of memories that attempt to serve all needs across large organisations often inhibits successful location and interpretation of relevant information. Recent efforts are instead targeting smaller, more focused approaches such as task-based memories (Ackerman and Mandel 1997), methodology-based memories (Hidding 1997) and community memories that support the incremental evolution of both structure and content (Marshall, Shipman et al. 1994). These approaches are a positive step towards enabling organisational memories to support the learning needs of individuals and groups.

Learning on Demand. Many approaches to supporting learning-on-demand or just-in-time learning are divorced from actual work contexts and, instead, are embedded within traditional curriculum-driven educational models. Their emphasis is on using the Internet to support distance learning. Other approaches intertwining working and learning in design domains, such as critiquing systems, have proven effective in supporting workplace learning for both newly hired and experienced designers (Sumner, Bonnardel et al. 1997). The challenge for this type of learning-on-demand is to have a rich shared context between the user and the system in order to determine the user's potential information needs (Fischer, Nakakoji et al. 1993). Another crucial feature, particularly with regard to supporting reflection-in-action, is proactivity: systems need to point out potential problems at the 'right time' when practitioners are best able to take advantage of the provided information (Lemke 1990).

Intranet-based Document Management Systems. Intranets are the fastest growing segment of the Internet market. Intranets use World Wide Web technologies to support an organisation's internal information needs, which are often document-centred since documents permeate much of organisational practice (Brown and Duguid 1996). Such uses were heralded by pre-Web groupware products such as Lotus Notes™ (Orlikowski 1992). However, as recently noted by Xerox™ (with more than 200 servers and 20,000 users), "this internal Web, as an environment for supporting organisational work, is falling short of our expectations and hopes in significant ways" (pg, 81) (Rein, McCue et al. 1997). They found off-the-shelf intranet software to be most successful when deployed to support simple, well-understood work practices where the information was factual and did not involve interpretation. They advocated creating a new style of intranet, with support for document management roles. Similarly, we view systems such as Intranets and Lotus Notes as important generic enabling technologies that we can build on to create organisational learning systems, rather than end-points in themselves.

Extensions to the Web. A number of extensions for placing semantic representations within Web documents have recently been put forward. The Resource Description Framework (RDF) (W3C, 1998b) is a draft proposal for representing metadata - data about the data - within a document. Encoding metadata into Web pages would have benefits in a number of areas. For example,

semantic search engines and site catalogues could be created. At the core of RDF is a model for representing named properties and their values. The properties can be used to encode the attributes of and relationships between resources within a document. Because RDF has been designed to represent *data* and not *knowledge* its representational power is relatively weak when compared with standard knowledge modelling languages (see for example, (Farquhar, Fikes et al. 1995; Motta 1998)). At a syntactic level RDF models would be represented in XML (W3C, 1998a) a proposed extension to HTML which is at the recommendation stage.

Analysis of these Approaches. Common shortcomings across many of the above approaches are passivity, lack of extensibility, and lack of context. The passivity of many approaches requiring users to search for relevant information in large memory stores limits their ability to support reflection-in-action. Several approaches lack extensibility which inhibits their evolution in the workplace to take into account new knowledge or work practices; i.e., organisational memories that are difficult for practitioners to extend. Such non-extensible systems cannot support domain construction processes underlying group learning. Finally, common among many approaches was lack of support for context, either to assist human interpretation and learning (i.e., the separation of design rationale from work product) or to support intelligent system interpretation (e.g., the necessity of shared context for learning-on-demand mechanisms). In the next section, we will describe how our approach towards supporting organisational learning addresses these shortcomings.

4. OUR APPROACH: CONTEXTUALLY-ENRICHING DOCUMENTS

When discussing the critical role of learning in the new knowledge economies, Zuboff described 'smart machines' that could assist in actively 'informating' practitioners as they work (Zuboff 1988). In our view, it is not smart machines per se that serve to informate, but contextually-enriched documents since documents form the core of many business practices and are the objects that many practitioners work with daily.

In a recent study looking at document use by knowledge workers, Kidd found that the *process of articulating and refining work products* was often more important for informing practitioners than the products themselves (Kidd 1994). In a broader context, Brown and Duguid contrast two basic models of document use: 'documents as darts' where documents serve as a means of objectively transmitting knowledge and 'documents as a means for supporting social processes' (Brown and Duguid 1996). They argue that *a key role of documents is to support negotiation and interpretation as communities struggle to reach a shared understanding*. Likewise, several leaders in collaborative working and learning have argued for a re-thinking of the status of 'representations of work' (Bannon 1995; Suchman 1995). They argue that rather than regarding these work products as static accounts or descriptions of the way things are, *representations should be regarded as starting points for discussion about the way things ought to be*.

In our view, the key to integrating working and learning is to support these socially-based, process-oriented views of representations of work. As indeed, these processes derived from analyses of document use are the same as the processes underlying individual, group, and organisational learning: reflection, articulation, elaboration, negotiation, interpretation, and sharing.

The core of our approach centres on enabling practitioners to progressively enrich their 'representations of work' with important contextual cues and information arising from these social processes. Representations of work take on many forms, including project specifications, design solutions, project bids, planning documents, etc. In most organisations, these representations take the form of documents, usually in paper form and increasingly in digital form as they are published on organisational intranets. Contextual cues and information takes on many forms including discussions surrounding the document, shared vocabularies or practices underlying the document, relationships to organisational competencies, and other related or dependent work products.

While this contextual information is vital for supporting key processes such as interpretation, we argue that simply capturing this information is insufficient. A key contribution of our 'enriching' approach is to capture these important contextual cues and information in such a way as to *tightly couple* them with the representations of work. This enriching approach has three direct benefits:

- The <u>richer context</u> supports improved human-human communication and collaboration by keeping the context for interpreting a document coupled to the document itself.

- Enriching is <u>a user-centred form of extensibility</u>, encouraging practitioners to articulate their tacit understandings and incrementally refine them towards more explicit knowledge representations by reifying the context and its interconnections.

- The richer context makes possible <u>active forms of computer support</u> (i.e., learning-on-demand) by providing a richer machine-interpretable context.

5. CEDAR TOOLKIT AND METHODOLOGY

Earlier, we discussed how organisational memory approaches based on the 'one-size-fits-all model' were problematic, resulting in memories that were too general to be useful or understood. Following our belief in the centrality of interpretation, we are focusing on creating task-specific instances of organisational memories based on our contextually-enriched document approach. Specifically, we are developing a toolkit supporting the cost effective construction and customisation of task-specific memories, and a methodology guiding others in incorporating the memories into work practices to achieve sustained use and growth. In this section, we briefly describe the generalities of the CEDAR toolkit and methodology before turning to a detailed description of how this approach is being applied to create a specific system – the Team Workbook.

Broadly speaking, the CEDAR approach enables (1) developers (information providers and knowledge engineers) to construct an initial organisational memory 'seed', (2) end-users to view and extend the memory contents, and (3) computational agents to deliver critical information at the right time to support reflection-in-action. The CEDAR toolkit is based on a server/client architecture communicating over a corporate intranet. At the heart of the system, the CEDAR archive is supported by a suite of CEDAR tools supporting knowledge modelling, document publication, archive maintenance, and context-based information retrieval.

The CEDAR archive consists of an organisational memory containing contextually-enriched documents. These hypermedia documents are work products enriched with related communications (e.g., structured on-line discussions), task-specific knowledge models (e.g., knowledge about quality management tools and techniques), and broader organisation-level knowledge models (e.g., models of customer-supplier relationships, corporate structure, or corporate values). These hypermedia documents are based on standard web-protocols and formats.

From the perspectives of <u>end-users</u>, task-specific memories created using the CEDAR toolkit will have several common characteristics, including:

- a document-centred virtual discussion space enabling distributed communities to engage in debate and negotiation about their shared work products and their shared knowledge models. These discussions are tightly integrated with the work products, thus enriching the document with greater context.

- enriching mechanisms enabling them to extend knowledge models and interconnect models with specific documents and discussion threads, contributing to a richer document context.

- mechanisms enabling end-users to publish and share their enriched work products with others.

- user interface agents (Lieberman 1997) that point out potential problems or missed opportunities in work products, point out possible links and dependencies to the work of other individuals or groups in the organisation, and inform user activities by providing relevant cases stored in the memory.

To support the activities of <u>knowledge engineers</u>, tools are provided for constructing, editing, and browsing the knowledge models in the archive. The formalism we use for expressing these knowledge models is OCML, the Operational Conceptual Modelling Language (Motta 1998). Our current library of OCML models comprises over twenty five thousand definitions and has been used to model several application domains, such as geology, engineering design, and healthcare, and is used in the Open University knowledge modelling course (Open University 1997). OCML's property of operationality, together with the provision of several mechanisms for integrating OCML models with other software components, facilitates the rapid development of application systems. The CEDAR toolkit builds on our previous research in collaborative knowledge modelling by incorporating model construction and visualization tools developed especially for OCML. These tools enable knowledge models to be constructed through a graphical user interface and model fragments to be shared via the web (Domingue 1998).

The tools provided to support the activities of <u>information providers</u> are based on our prior work with The Digital Document Discourse Environment (D3E) (Sumner and Buckingham Shum 1998). D3E supports the publication of web-based documents with integrated discourse facilities and interactive components and is based on extensive research into how hypertext systems can support critical reflection and the analysis of arguments in writing and software design (Buckingham Shum and Hammond 1994; Buckingham Shum 1996; Kolb 1997; Sumner, Bonnardel et al. 1997). We are adapting the D3E system to make it possible for information providers, such as corporate trainers from BAe's Virtual University, to construct organisational memories with only minimal assistance from knowledge engineers. Using our publishing tools, providers can publish HTML files with special forms of navigation and integrated threaded discussion facilities without requiring them to have detailed knowledge of HTML.

The emerging ENRICH methodology assumes that the application domain satisfies various applicability conditions. These include:

- *Suitable work practices*
  The ENRICH approach is based on asynchronous collaborative problem solving. There must be a need for such work practices. For instance, if collaborators typically resolve problems in face to face discussions and then only document the result, our approach would probably not be accepted. On the other hand, if the collaborators are separated by time and/or space, then a shared discourse environment may provide an important role in the recording of emerging ideas.

- *Availability of the knowledge*
  The knowledge required for creating the archive seeds must be available. Depending on the problem domain, this knowledge is usually a contained in documents which use a variety of formats (e.g. plain text, images, diagrams or segments of computer code).

- *Stability of the knowledge*
  Domain knowledge needs to be stable enough to be formalised as an ontology and associated knowledge bases. If solved examples (cases) are used as part of archive seeds they are annotated in terms of the concepts within the ontologies.

As part of the ENRICH Methodology, we are developing a comprehensive lifecycle guiding the deployment, use, and maintenance of CEDAR archives. One lifecycle we have built on is the seeding, evolutionary growth, and reseeding (SER) model (Fischer, McCall et al. 1994). According to this model, knowledge engineers work with end-users to create an initial memory 'seed.' End-users extend the seed during use, contributing to evolutionary growth. At times after

periods of extensive growth, knowledge engineers return to restructure and re-organise the memory, an activity called 're-seeding.'

Developing archive seeds is an important step of the ENRICH methodology. We often distinguish between two kinds of archive seeds:

- *Domain specific archive seeds*
  These are domain ontologies and knowledge bases. They characterise the context for the problem being solved. The ENRICH methodology supports *domain construction* - a form of group learning in which collaborating users articulate explicit extensions to the domain specific knowledge seeds when reflecting on the problem solving processes.

- *Problem specific archive seeds*
  These are knowledge models of best practices, cases or examples of previous solutions. The knowledge models are described in terms of concepts from domain specific ontologies and their context is provided by the knowledge bases. Relevant examples retrieved from the knowledge model form the building blocks for solving new problems. When appropriate new solutions may be added to the archive. Our methodology is built on the premise that the recording of previous solutions, together with the addition and reuse of elements of the best practice library significantly contributes to organisational learning.

During the case studies we found it productive to start with the problem specific archive seeds and proceed towards domain ontologies and knowledge bases, i.e. first to formulate the problem dependent contents and then to specify a language for describing the domain. The increase in productivity comes from the fact that the problem specific knowledge models drive and focus the knowledge acquisition process.

Since we address the problem of *learning while working* the ENRICH application must fit into the current work practices. In particular, the ENRICH tools must be integrated with the tools which are currently used. Often the user interface will be a decisive factor in the acceptance of this class of applications.

We have extended the SER model to include two kinds of developers – knowledge engineers and information providers. Knowledge engineers and information providers (such as corporate trainers) work with local developers to create an initial archive seed using the CEDAR toolkit. We use a variant of the VITAL Knowledge Based Systems (KBS) methodology (Jonker, Kontio et al. 1991) to create the initial seed. The VITAL project was a four and a half year research and development enterprise which aimed to provide both methodological and software support for developing large, industrial, embedded KBS applications. Knowledge engineers use CEDAR to create initial knowledge base 'seeds' by re-using and refining existing model libraries. The key contribution of the ENRICH approach is to enable small teams of information providers to efficiently create the bulk of the memory seed by publishing web-based documents and integrated discussion spaces that are linked to these knowledge models, without requiring extensive specialised technical knowledge or dependence on the knowledge engineers.

6. THE BAe TEAM WORKBOOK

The Team Workbook is a planning tool designed to stimulate more effective working of both team leaders and their teams. It is part of an overall educational programme called 'BenchmarkBAe' designed to help the Company achieve its goal to become a benchmark company through the pursuit of excellence. This programme is being led by the CEO and the senior management team, who have identified five Values – Customers, People, Performance, Innovation & Technology, and Partners – as keys to achieving this goal (published in 'Our Value Plan' issue 2 1998). The Workbook is used to prepare, deliver and review 'Value Plans' which are Quality Managment plans that are tied to the aims embodied in the five Values. A Value Plan is an interlinked set of

structured documents where the team sets out their objectives, actions, and methods to achieve their goals for each of the five values.

The current Workbook is paper based and is in the process of being deployed across the entire company (44,000 employees). Using the Workbook, local teams identify their internal customers and suppliers, map their key processes, measure their performance, and collectively work to secure higher levels of performance and greater customer satisfaction. The Workbook includes Total Quality Management (TQM) tools and is designed to be suitable for teams with little or no prior TQM experience. A support structure of trained facilitators is also being implemented. To assist the adoption of the workbook, these facilitators will go into the workplace and work directly with teams conducting planning sessions.

The Workbook is based around a three step process: (1) preparing the Value Plan, (2) declaring and delivering the plan, and (3) reviewing and improving the plan and work practices. These iterative steps provide a learning framework enabling teams to secure continuous improvements in their sphere of responsibility. In the paper-based Workbook, Value Plans are declared in a 'planning sheet'. A planning sheet is a grid where each row represents the plan components for a specific Value and the columns represent generic plan attributes, such as objectives or actions to be carried out, common across all Values. Teams fill out the cells within the planning sheet aided by 'guide' pages which articulate the key issues to be considered.

Typically, in order to fill out a cell, the team must first discuss these issues and engage in specific TQM exercises. Each TQM tool is described in one or two 'tool' pages followed by an example of its use. These TQM tools include activities such as SWOT analyses (placing the Strengths, Weaknesses, Opportunities and Threats surrounding an issue into a quadrant) and Fishbone diagrams where the results of a brainstorming session are placed in a type of tree structure. 'Review' pages enable teams to reflect on what aspects of a planning session worked from an individual and team perspective and how they would consider changing their approach in the future. The paper-based Workbook also contains a few examples of completed exercises and Value Plans. If a facilitator is present during a planning session, he or she provides support and advice to the team on general workbook processes and on the deployment of specific TQM tools and techniques.

6.1. The Value-Enriched Workbook

In this project, we are using CEDAR to construct a digital Value-Enriched Team Workbook underpinned by a knowledge base representing the five Values. While the paper-based workbook helps teams to become communities of practice, it does not directly contribute to organisational learning because the results of the planning activities are not captured, re-used or shared. The digital Value-Enriched Team Workbook tries to address this limitation by enabling:

(1) teams to articulate their plans directly within the Workbook and to discuss and review their plans on-line.

(2) distributed teams (e.g., customers and suppliers) to share experiences and best practices by linking their shared plans and values.

(3) teams and Virtual University staff (such as facilitators) to feed into and augment the corporate knowledge base according to needs arising through use of the Value-Enriched Workbook.

The digital Workbook incorporates the same 'pages' as the paper version and has been additionally 'enriched' with several different forms of knowledge:
- *Informal Knowledge.* Each Workbook section is integrated with a threaded discussion area enabling teams to capture informal aspects of their deliberation and decision-making. These discussion areas are linked to the relevant parts of the Value-Plan currently being considered.

- *Semi-formal Knowledge.* The documents in the Workbook are stylized and structured, with parts of one document being dependent on parts of other documents. While the text entered into the documents is informal, the rich relationships between the structured documents are formally represented in the digital Workbook. Making these relationships explicit helps end-users understand the relationships between different TQM activities and different parts of their Value Plan and enables the system (the Value-Enriched Workbook) to use this knowledge to better support their planning activities.

- *Formal Knowledge.* The Workbook is linked to an extensible knowledge model representing the Values and BAe's organisational structure. These knowledge models, together with the structured document relationships are used by the Workbook to provide case-based retrieval mechanisms and active forms of computational support for the planning process.

Figure 1 shows a screen snapshot of part of an 'enriched Value Plan' in the digital Workbook. As shown, the user is focusing on a cell where the customer-action part of a plan is declared. From this cell, the user is able to move to relevant parts of the workbook, such as appropriate guide pages, the related discussion area or to an appropriate segment of the knowledge model. Most of the top-level threads in the discussion area correspond to an imporant 'question to consider' as listed in the related guide page. Forms are provided that allow users to add new comments or even new topics for discussion.
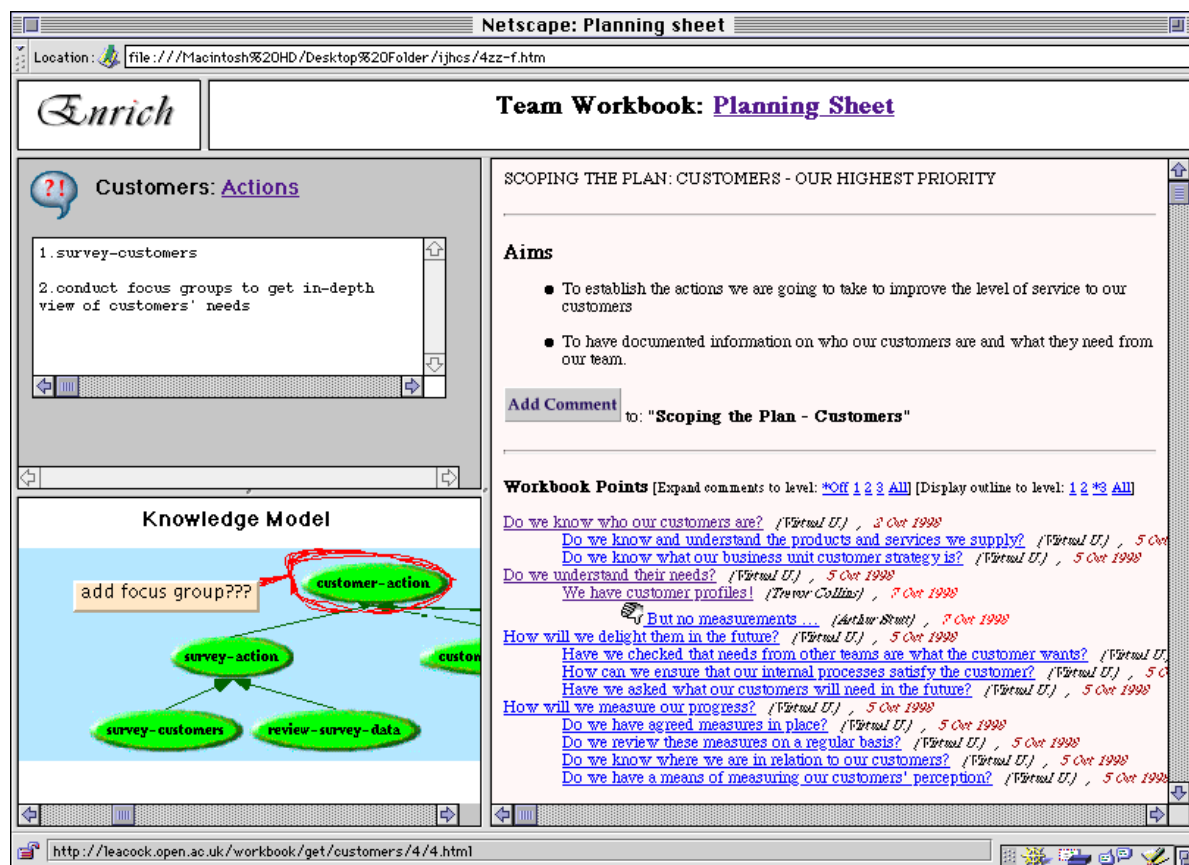

Figure 1: The user is browsing part of the discussion area (shown on the right) and the Customers: Action part of the enriched Value Plan (shown on the left).

The knowledge models support the automatic retrieval of plan components relevant to the plan under construction, and provide consistency checking between plan components. Retrieval is

enabled by a representation of the organisational structure of BAe's business units and of the five Values. Each business unit is described in terms of the activities performed, market segment, customers and suppliers. A number of interface tools exist to support retrieval (see (Domingue and Motta, 1999)). Internally, retrieval requests are transformed into OCML queries. Pre-designated classes of instances within knowledge models contain the URLs of associated documents. Whenever the solution to an OCML query contains one of these instances the relevant web page is automatically displayed.

In Figure 1, the user has elected to view the portion of the knowledge model which describes customer actions. The user can use this hierarchy in three ways:

- as a guide for browsing the workbook guides, or related plans,

- as a knowledge based query interface,

- as a prompt for the current consensus on agreed terminology.

In the formal knowledge model, a corporate value has four main attributes: parameter, objective, action and measure. These attributes are derived from the planning tools within the workbook and are those deemed of importance by the BAe executive management ('Our Value Plan' issue 2 1998). Using our knowledge modelling tools (see section 6.3), a simple algebra describing the relationships between the attributes was defined. This algebra allows the model to be used to specify dependencies and constraints between plan components.

For example, team members need to have a specific set of skills to carry out certain 'actions' related to the Customers Value, such as knowing how to do customer surveys. Thus, filling in one part of a Value Plan (e.g., the customers-actions cell declares 'do customer surveys') can influence how a team should fill in another cell in the Value Plan (e.g., the people-actions cell declares 'learn how to do customer surveys'). In the paper-based Workbook, these dependencies must be explicitly managed by the team during a planning session. This is problematic for several reasons as (1) it is quite a large cognitive load to remember, track, and resolve the numerous dependencies, and (2) successfully managing these dependecies requires a deep understanding of the planning process which teams who are just learning the process will not have. In the digital Workbook, these dependencies between parts of a Value Plan can be established during use time (in a planning session) by enriching the relevant part of the Value Plan with a link to the appropriate part of the underlying Values model. As shown in the following scenario, this form of enrichment enables the system to provide active support during the planning process by pointing out plan inconsistencies and providing relevant knowledge.

Our aim in creating the Value Enriched Workbook is to ensure that the Workbook is not merely a source of knowledge that is referred to from time to time, but rather the means by which teams carry out all plan related actions, discussions and reflection.

6.2. A Scenario using the Enriched Workbook

We'll illustrate the basic functionality of the enriched Workbook, and indeed the core features of any CEDAR archive from the end-user perspective, using a simple scenario. Imagine a planning session in BAe's Regional Aircraft Business Unit where a group in charge of marketing aviation support services is creating their first ever Value Plan, with the assistance of a facilitator from BAe's Virtual University.

The marketing group and the facilitator are gathered together in a meeting room containing a networked personal computer with the monitor projected so that everyone can see the Workbook contents. The marketing group began the session by logging into the Workbook archive. In addition to providing security, logging in enables the system to identify parts of the knowledge base relevant to this particular group's business unit; i.e., the relevant part of the organisational

model, the existing Value Plans from other groups in this Business unit, Value Plans from other similar groups in other units, etc.

After a brief presentation on the planning process by the facilitator, the group first tries to fill in the "Where are we now against our values?" worksheet for the Customers Value. In this activity, the team must decide where they are against a spectrum of customer needs awareness criteria, and where they want to be. These criteria range from 'we have little understanding of our customers' to 'we have documented our customer needs' to 'we have measurable performance criteria with our customers'. As intended by the workbook design, trying to choose the proper criteria sparks a considered debate, as shown in the discussion window in Figure 1. The facilitator acts as a 'scribe' and records the key points of the debate in the discussion area, attributing the added comments to people who made them. After much debate, the group decides that they know about their customers and products and this knowledge is documented, but they have not tried to measure how well they are meeting their customer's needs. They mark boxes in the form to indicate where they are (have documented needs) and where they want to be (have measurable performance criteria).

These activities have taken an hour and the group breaks for coffee. During the break, the facilitator shows the group's team leader how to use the Workbook discussion forum. After the break, the team leader will take over the 'scribe' role, while the facilitator looks on and offers assistance as needed.

After the break, the group begins the Planning exercise and considers what their objectives and actions will be with respect to the Customer Value, given their current level of customer awareness. The facilitator suggests that rather than starting their plan from scratch, the group use the Workbook to retrieve relevant plans they may be able to build on. The team leader chooses to 'retrieve relevant plans' and the Enriched Workbook uses its case-based retrieval mechanism to select similar cases.

The workbook has identified three plans as being very similar – the most similar is a plan from a marketing group in another business unit in the commercial aerospace sector of BAe, the next most similar is a plan from a marketing group in BAe's real estate sector, and the third most similar is a 'prototypical' marketing group plan that the Virtual Univeristy created as part of the Workbook seed. In this case, similarlity is determined by the system's organisational knowledge on products, services, and activities in the different business units and the criteria the group selected to best represent their customer awareness in the "where are we now..." exercise. The retrieved plans are 'enriched' in the sense that the returned Value Plan also includes a digital form of the Planning Sheet with embedded links to parts of the discussion generated by the other teams and parts of the Values model that the other teams chose to link to.

After inspecting the three cases, the group chooses to modify the plan from the commercial aerospace unit. The group elects to keep the same objective "to continue to document our customers needs and to establish measurement criteria and processes." However, the group is in disagreement as to whether to adopt the same actions. After some debate (which the team leader records in the discussion area), the group elects to do some customer surveys (as the other group did) but in addition, to conduct a focus group for 'potential' customers who are currently purchasing services from other suppliers.

The team leader now begins to record the group's intentions into the 'Customer – Actions' document (left half of Figure 1). Surveys are a standard technique advocated and hence are represented in the underlying Values model. The team leader selects 'surveys' from the list of techniques provided by the Workbook. Since the relationship between techniques and skills is represented in the underlying knowledge base, the Workbook responds to the selection of 'surveys' by displaying in a separate window the names of people in similar groups that can be contacted for advice about survey techniques. Focus groups are not represented in the underlying knowledge model so the team leader simply types this intention into the text area. The facilitator

asks the team leader to annotate the 'customer-actions' part of the model with a suggestion to include focus groups in the standard techniques. As part of the ongoing Workbook maintenance and release cycle, staff at the Virtual University periodically examine these annotations and decide which suggestions should indeed be formalized in the Workbook's knowledge models.

6.3. Creating the Workbook Seed

To help ensure that the Value-Enriched Workbook and the CEDAR toolkit meets the needs of our users, we are following an iterative and participatory (Greenbaum and Kyung 1991) development strategy where design and implementation activities are driven by use experiences in two sucessive workplace trials. This early deployment will provide important feedback and will drive further development of both the Workbook archive and the CEDAR toolkit.

We are now in the process of revising the seed of the digital Value-Enhanced Workbook for the second user trial. Currently, members of the Knowledge Media Institute (KMi) are acting as 'knowledge engineers' and members of BAe's Virtual University (VU) are serving as the 'information providers' envisaged in the CEDAR Methodology. While we are working together to produce the revised Workbook seed, it is primarily the responsibility of the KMi knowledge engineers to create the underlying knowledge models and the empty Workbook structure and the responsibility of the VU information providers to create the examples of enriched Value Plans and planning techniques that are criticial to the seed's usability.

The existing suite of tools in the CEDAR Toolkit have been used to edit the knowledge models of corporate values and organisational structure and to re-create the threaded discussion forums based on the Workbook's guide pages. The threaded discussion forums are created automatically by a tool which parses basic HTML documents, re-presents them according to a set of style templates defining the look, and creates discussion threads linked to the corresponding document parts in the new style. Knowledge models are edited using the tool shown in Figure 2. Using this tool, knowledge engineers can browse and edit ontologies from Java enabled web browsers. The left panel contains a list of all the classes within the knowledge model. New model elements are created by dragging icons from the central well into the graphical display area on the right. Once the knowledge engineer clicks on the graphical area the textual definition of the corporate-value class, shown in the small central window is updated. The textual definition of the action class is displayed in the small window on the right.
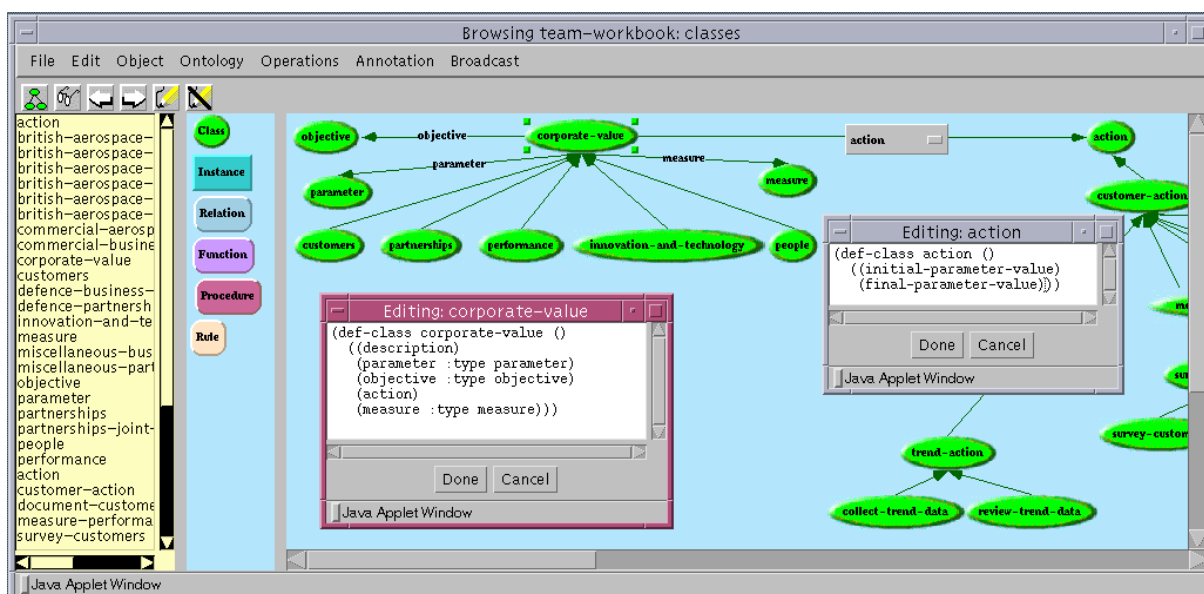

Figure 2. A screen snapshot showing a knowledge engineer editing the Team Workbook knowledge model.

7. DISCUSSION AND FUTURE WORK

As we stated in our introduction, there is a strong need for organisations to move from a traditional 'knowledge transfer' model of on-the-job training to one which integrates working and learning. Specifically, in this project we are trying to rethink current training practices to: (1) move from one-shot training courses to continuous learning processes and (2) move from training focused on the individual to incorporate work group and organisational learning. We will now examine the CEDAR approach and the Team Workbook project with respect to these goals.

Moving from one-shot training to continuous learning processes. The paper-based Workbook programme developed by BAe's Virtual University already goes a long way towards supporting this goal due to its 'teaching by doing' approach and the innovative support of trained facilitators who go into the workplace and help teams with their specifics plans, instead of lecturing on planning techniques. The digital Value-Enriched Workbook tries to further enhance learning by making it easier for novices to engage in planning activities. Specifically, we envision that the case-based retrieval mechanisms will enable 'planning by modification', which should in theory, be easier than starting from scratch for people new to the planning process. Additionally, the linking of the discussion forums with the relevant parts of the Value Plan should support teams to build on their previous decisions and reflect on their previous plans and thus promote a more continuous and incremental learning process.

Moving from training focused on the individual to incorporate work group and organisational learning. In this paper, we have described an integrated learning approach based on three processes: reflection-in-action, domain construction and perspective taking. These three processes are founded on theories of working and learning and our own empirical research, and are embodied in the CEDAR methodology and toolkit. It remains to be empirically evaluated as to whether these processes are indeed realized in specific archives created using CEDAR, such as the Team Workbook. In our future work, we will focus on examining whether these learning processes are supported according to the following working hypotheses:

Hypothesis 1: Individual learning and reflection-in-action will be supported by providing practitioners with active support (i.e., agents) for constructing and reviewing plans using TQM methodologies. We will consider reflection to have occurred if individuals discuss through the workbook their plans or the methodologies. We will consider learning to have occurred if individuals improve their plans in response to agent intervention.

Hypothesis 2: Group learning and domain construction will be supported by linking the plans to the underlying knowledge bases and enabling practitioners to enrich both their own plans and the knowledge base. We will consider domain construction to have occurred if practitioners enrich their own plans with links to the knowledge bases and if practitioners extend or discuss the underlying knowledge bases.

Hypothesis 3: Organisational learning and perspective-taking will be supported by linking groups using the underlying corporate values model. We will consider perspective-taking to have occurred if practitioners engage in debate or negotiation about another team's plan with respect to their own planning objectives.

Consistent with recommended case study approaches, we will collect several types of qualitative data (including memory content analysis, workplace observations, and surveys) during our upcoming user trials to attain convergence in our final data analysis (Yin 1984).

8. SUMMARY

We believe that our approach goes beyond current approaches to supporting learning in the workplace since:

- The *context* in which work products are created and extended is *tightly integrated* with the parent work product,
- Tools and a methodology are provided to facilitate the *continuous growth of the archive*, facilitating domain construction,
- Agent services built on top of the knowledge models *proactively instigate* reflection within the workforce.

The BAe Team Workbook is one of six ongoing trials (3 industrial, 3 academic) involving the ENRICH approach. Preliminary feedback from CEDAR developers and end users in five of the trials were used to create a list of organisational, methodological and software requirements. A new version of the CEDAR toolkit is currently being evaluated in a series of fresh trials which will finish at the end of this year. Future work will involve incorporating the requirements which emerge from these trials into the ENRICH methodology and the CEDAR toolkit.

REFERENCES

Ackerman, M. (1993). Definitional and Contextual Issues in Organizational and Group Memories., Hawaiian International Conference on System Sciences (HICSS 27), Hawaii (Jan), IEEE Computer Press.

Ackerman, M. and E. Mandel, "Memory in the Small: An Application to Provide Task-Based Organizational Memory for a Scientific Community," Technical Report http://www.ics.uci.edu/CORPS/ackerman.html, UC Irvine, 1997.

Argyris, C. and C. Schön (1978). Organizational Learning: A Theory of Action Perspective. Reading, Addison Wesley.

Bannon, L. (1995). "The Politics of Design: Representing Work." Communications of the ACM ,38(9): 66-68.

Bannon, L. and K. Kutti (1996). Shifting Perspectives on Organizational Memory: From Storage to Active Remembering., Hawaiian International Conference on System Sciences (HICCS 29), IEEE Computer Press.

Boland, R. and R. Tenkasi (1995). "Perspective Making and Perspective Taking in Communities of Knowing." Organization Science ,6(4): 350-372.

Brown, J. S., A. Collins and P. Duguid (1989). "Situated Cognition and the Culture of Learning." Educational Researcher ,Jan-Feb: 32-42.

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

Brown, J. S. and P. Duguid (1996). "The Social Life of Documents." First Monday ,1:

Buckingham Shum, S. (1996). Analyzing the Usability of a Design Rationale Notation. Design Rationale: Concepts, Techniques, and Use. Hillsdale, NJ, Lawrence Erlbaum Associates. 185-215.

Buckingham Shum, S. and N. Hammond (1994). "Argumentation-Based Design Rationale: What Use at What Cost?" International Journal of Human-Computer Studies ,40(4): 603-652.

Conklin, J. and M. Begeman (1988). "gIBIS: A Hypertext Tool for Exploratory Policy Discussion." Transactions of Office Information Systems ,6(4): 303-331.

Curtis, B., H. Krasner and N. Iscoe (1988). "A Field Study of the Software Design Process for Large Systems." Communications of the ACM ,31(11): 1268-1287.

Detterman, D. K. (1993). The Case for the Prosecution: Transfer as an Epiphenomenon. Transfer on Trial: Intelligence, Cognition, and Instruction. Norwood, New Jersey, Ablex. 1-24.

Domingue, J., E. Motta and S. Watt (1993). The Emerging Vital Workbench., Knowledge Acquisition for Knowledge-Based Systems 7th European Workshop (EKAW '93), Toulouse and Caylus, France (September), Springer-Verlag.

Domingue, J. (1998). Tadzebao and WebOnto: Discussing, Browsing, and Editing Ontologies on the Web., Proceedings of the 11th Banff Knowledge Acquisition Workshop (KAW´98), Banff, Canada, (April 18-23),

Domingue, J. B. and Motta, E. (1999) A Knowledge-Based News Server Supporting Ontology-Driven Story Enrichment and Knowledge. In the Proceedings of the 11th European Workshop on Knowledge Acquistion, Modelling, and Management. Fensel, D. and Studer, R. Dagstuhl Castle, Germany May 26 - 29, 1999. pp 103-120.

Ehn, P. (1989). Work-Oriented Design of Computer Artifacts. Stockholm, arbetslivscentrum.

Farquhar, A., R. Fikes, et al., "Collaborative Ontology Construction for Information Integration," Knowledge Systems Laboratory, Technical Report KSL-95-63, Department of Computer Science, Stanford University, 1995.

Fischer, G. (1994). "Turning Breakdowns into Opportunities for Creativity." Knowledge-Based Systems Journal ,7(4): 221-232.

Fischer, G., J. Grudin, et al. (1992). "Supporting Indirect, Collaborative Design with Integrated Knowledge-Based Design Environments." HCI ,7 (Special Issue on Computer Supported Cooperative Work)(3):

Fischer, G., A. C. Lemke, et al. (1991). "Making Argumentation Serve Design." Human Computer Interaction ,6(3-4): 393-419.

Fischer, G., R. McCall, et al. (1994). Seeding, Evolutionary Growth and Reseeding: Supporting the Incremental Development of Design Environments., Human Factors in Computing Systems (CHI '94), Boston, MA,

Fischer, G., K. Nakakoji, et al. (1993). Embedding Computer-Based Critics in the Contexts of Design., Human Factors in Computing Systems (Interact '93 and CHI '93), Amsterdam (24-29 April), ACM Press.

Greenbaum, J. and M. Kyung (1991). Design at Work: Cooperative Design of Computer Systems. Hillsdale, NJ, Lawrence Erlbaum Associates.

Grudin, J. (1991). "Systematic Sources of Suboptimal Interface Design in Large Product Development Organizations." HCI ,6: pp. 147-196.

Hidding, G. J. (1997). "Reinventing Methodology: Who Reads it and Why?" Communications of the ACM ,40(11): 102-109.

Kidd, A. (1994). The Marks are on the Knowledge Worker., Human Factors in Computing Systems (CHI '94), Boston, MA (April 24-28), ACM Press.

Kolb, D. (1997). Scholarly Hypertext: Self-Represented Complexity., Proceedings of The Eighth ACM Conference on Hypertext, Southampton,

Landes, D., K. Schneider and F. Houdek (1998). Documentation in Industrial Software Projects., 1st Interdisciplinary Workshop on Building, Maintaining, and Using Organizational Memories (OM-98). In conjunction with the 13th biennial European Conference on Artificial Intelligence (ECAI -98), Brighton, UK (August 23-24), Available at: http://www.aifb.uni-karlsruhe.de/WBS/ECAI98OM.

Lave, J. (1991). Situated Learning in Communities of Practice. Perspectives on Socially Shared Cognition. Washington, DC, American Psychological Association. 63-82.

Lemke, A. C. (1990). Framer-Hypertext: An Active Issue-Based Hypertext System. Proceedings of the Workshop on Intelligent Access to Information Systems (Darmstadt, Germany). Darmstadt, Germany, GMD-IPSI. 34-38.

Lieberman, H. (1997). Autonomous Interface Agents., Human Factors in Computing Systems (CHI '97), Atlanta, Georgia (March 22-27), ACM Press.

Marshall, C., F. Shipman and R. McCall (1994). Putting Digital Libraries to Work: Issues from Experiences with Community Memories., Digital Libraries'94, College Station, Texas (June 19-21),

Motta, E. (1998). Reusable Components for Knowledge Modelling. Knowledge Media Institute, The Open University.

Motta, E. and Z. Zdrahal (1996). Parametric Design Problem Solving., 10th Banff Knowledge Acquisition for Knowledge-Based System Workshop (KAW´96), Banff, Canada, (November 9th-14th),

Nonaka, I. and H. Takeucji (1995). The Knowledge-Creating Company. Oxford, University Press.

Orlikowski, W. (1992). Learning from Notes: Organizational Issues in Groupware Implementation., CSCW '92, Toronto, Canada, ACM Press.

Orr, J. (1995). Ethnography and Organizational Learning: In pursuit of learning at work. Organizational Learning and Technological Change. New York and Berlin, Springer-Verlag. 47-60.

Rein, G., D. McCue and J. Slein (1997). "A Case for Document Management Functions on the Web." Communications of the ACM ,40(9): 81-89.

Rittel, H. and M. M. Webber (1984). Planning Problems are Wicked Problems. Developments in Design Methodology. New York, John Wiley & Sons. 135-144.

Robinson, M. and L. Bannon (1991). Questioning Representations., Second European Conference on Computer-Supported Cooperative Work (ECSCW'91), Amsterdam, The Netherlands (Sept. 25-27),

Ruhleder, K. (1994). "Rich and Lean Representations of Information for Knowledge Work: The Role of Computing Packages in the Work of Classical Scholars." ACM Transactions on Information Systems ,12(2): 208-230.

Sachs, P. (1995). "Transforming Work: Collaboration, Learning, and Design." Communications of the ACM (Special Issue on Representations of Work) ,38(9): 36-44.

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

Suchman, L. (1995). "Making Work Visible." Communications of the ACM ,38(9): 56-64.

Sumner, T. (1995). The High-Tech Toolbelt: A Study of Designers in the Workplace., Human Factors in Computing Systems (CHI '95), Denver, CO (May 7-11), ACM Press.

Sumner, T. and S. Buckingham Shum (1998). From Documents to Discourse: Shifting Conceptions of Scholarly Publishing., Human Factors in Computing Systems (CHI '98), Los Angeles (April 18-23),

Sumner, T., J. Domingue, et al. (1998). Enriching Representations of Work to Support Organisational Learning., 1st Interdisciplinary Workshop on Building, Maintaining, and Using Organizational Memories (OM-98). In conjunction with the 13th biennial European Conference on Artificial Intelligence (ECAI -98), Brighton, UK (August 23-24), Available at: http://www.aifb.uni-karlsruhe.de/WBS/ECAI98OM.

Sumner, T. and M. Stolze (1996). "Integrating Working and Learning: Two Models of Computer Support." AI & Society ,10(Special Issue on Hypermedia and Workplace Learning): 70-78.

Sumner, T. R., N. Bonnardel and B. Kallak Harstad (1997). The Cognitive Ergonomics of Knowledge-Based Design Support Systems., Human Factors in Computing Systems (CHI '97), Atlanta, Georgia (March 22-27), ACM Press.

Yin, R. (1984). Case Study Research: Design and Methods. Newbury Park, CA, SAGE Publications.

Zuboff, S. (1988). In the age of the Smart Machine. New York, Basic Books, Inc.