

Balancing Formality with Informality: User-Centred Requirements for Knowledge Management Technologies

Simon Buckingham Shum

Knowledge Media Institute
The Open University
Milton Keynes
MK7 6AA
U.K.

S.Buckingham.Shum@open.ac.uk
<http://kmi.open.ac.uk/~simonb/>

Abstract

Numerous disciplines are now trying to analyse and represent the processes and products of organisational memory, in order to clarify what tangible representations future knowledge managers will work with. This short paper begins by reflecting briefly on the nature of systematic representations, as a reminder of the commitments that are made in any classification process. It is argued that there are important political dimensions to such classification, with implications for knowledge modelling. The paper then identifies three processes by which organisational expertise is shared. These processes may represent both a challenge and an opportunity for knowledge modelling approaches. The closing discussion pinpoints *formalisation* as a particularly important process in knowledge management, considers technologies that support *incremental formalisation* as holding particular promise, and proposes the principle that only *stable, sanctioned* knowledge should be formalised, in order to avoid the many problems caused by premature formalisation of organisational knowledge.

Introduction

All forms of management have associated with them particular techniques and tools which allow managers to survey the state of play. These *representations* of the domain of concern afford particular viewpoints and facilitate particular analyses. Such representations include spreadsheets, summary information about organisational performance, and predictive simulations for exploring different scenarios. In order to operationalise the concept of Knowledge Management (KM), one presumably must develop representations of knowledge products and processes to analyse. This short paper explores issues which are unavoidably raised once one begins to explore the nature of representations in support of KM. This is particularly germane to the application of Artificial Intelligence (AI) techniques to KM, since the success of

such approaches rests heavily on finding appropriate *formal representations*, for instance, for knowledge modelling and inference.

The paper begins by reflecting briefly on the nature of systematic representations, as a reminder of the commitments that are made in any classification process. It is argued that particularly in an organisational context, there are political dimensions to such classification processes, with implications for knowledge modelling. The paper then identifies three processes by which organisational expertise is shared:

- via knowledge analysts
- via stories
- via documents and debates

It is proposed that these processes represent both challenges and opportunities for AI approaches.

The politics of formalisation

In selecting any representation we are in the very same act unavoidably making a set of decisions about how and what to see in the world. [...] A knowledge representation *is* a set of ontological commitments. It is *usefully* so because judicious selection provides the opportunity to focus attention on aspects of the world we believe to be relevant. [...] In telling us what and how to see, they allow us to cope with what would otherwise be untenable complexity and detail. Hence the ontological commitment made by a representation can be one of the most important contributions it offers.

(Davis, Shrobe and Szolovits, 1993)

Classification systems provide both a warrant and a tool for forgetting [...] The classification system tells you what to forget and how to forget it. [...] The argument comes down to asking not only what gets coded in but what gets read out of a given scheme.

(Bowker, in press)

The above two quotes, the first from knowledge engineers, and the second from an ethnographer of organisational memory, draw attention to the filtering function that a representation provides, and the problem that through the process of simplifying a problem in order to describe it within a formal scheme, we may also be systematically factoring out certain classes of critical knowledge simply because they are hard to formalise.

Whenever an authoritative body (e.g. corporate management, or a research funding council) declares an interest in certain concepts, it is inevitable that its dependents (e.g. managers, or researchers seeking grants) will seek to align their activities with these concepts in order to maintain a presence. The first point to make therefore is that the introduction of systematic KM (whether or not technology is involved) creates a new *economy of knowledge* and a *knowledge vocabulary*. Any group and their work will remain invisible and thus unresourced unless they can represent themselves within this new economy, using the right language. Bowker presents an illuminating analysis of the impact of 'professionalisation'—systematic classification of skills and courses of action, and management of these via technology—on a profession in which expertise takes the form of hard to codify tacit knowledge and craft skill, in this case nursing:

One of the main problems that [...] nurses have is that they are trying to situate their activity visibly within an informational world which has both factored them out of the equation and maintained that they should be so factored – since what nurses do can be defined precisely as that which is not measurable, finite, packaged, accountable.

(Bowker, in press)

This illustrates clearly the political dimensions to formal classification. The names and labels one uses unavoidably emphasise particular perspectives (see also Suchman (1993) on the politics of computational categories).

Knowledge-based systems require the systematic decomposition and classification of expertise; a knowledge-base unavoidably 'holds' an ontological view of the world (ontology with a small 'o'). More recent knowledge-sharing initiatives and other research devoted to formal Ontologies make more explicit the issues faced in knowledge modelling, independent of any particular symbolic implementation as a system. One question that could be productively addressed by this symposium is how developments in ontology building help to manage the inevitable incompleteness and possible inconsistencies in an organisational knowledgebase, due to uncodified, or uncodifiable knowledge. If ontology building is to form part of AI's contribution to KM (as some argue), how can we ensure that areas of uncertainty or incompleteness are made explicit in the ontology, and carried through to the implementation and user interface

of any KM system based on that ontology? If the KM system is to be used by the organisation's managers, then they must be sensitised to the limitations of the tool's ontology, and the potentially misleading 'seductiveness' of manipulating clean computational abstractions. What training is required in order to wield such tools intelligently?

Modelling a changing domain

Our language evolves as the world evolves: new distinctions need to be made, new concepts emerge. If the meanings and inter-relationships of concepts have been codified, whether in a glossary, manual, project lexicon, or executable knowledgebase, they need to be maintained and kept up to date. Musen (1992) describes some of the problems that have arisen in efforts to develop medical ontologies when concepts evolve over time. Knowledge engineers end up using the same concepts in different ways within the same ontology or system.

Maintenance is an acute challenge to any effort to model organisational knowledge, since the modern organisation is the epitome of a changing, and strongly non-deterministic domain. Analysis of knowledge workers in the field (Kidd, 1994) shows that they perform many different roles that resist classification, and rarely conform to their official titles; they build up informal networks of communication that rarely conform to the organisational chart; project teams form and disband in the space of days. Can knowledge modelling techniques handle this kind of flux? Can other AI techniques assist in managing this complexity?

Having set the representational scene, the following sections consider some of the ways in which organisational memory is constructed—how knowledge is shared within and between working communities. After each section, implications and opportunities for AI support are considered.

The importance of knowledge analysts

Organisations devote personnel and other resources to processes that they recognise to be crucial to their operation. Technology infrastructure, accounting and personnel management are divisions found in every large organisation. If knowledge is to become as important a resource to corporations as some analysts predict, then Knowledge Analysts will become as common as Accountants. Their sole task will be to track, audit, and index knowledge, making available the right knowledge in the right media to the right people at the right time. Some companies have already committed staff to precisely this role. In the light of the earlier discussion, maintenance is likely to take substantial effort as the analysts update their knowledge resources to reflect organisational developments.

Conceptually, the role of knowledge analyst becomes interesting when we consider the balance of responsibility given to humans and machines. Davenport (1996) has argued, on the basis of several organisational case studies of successful knowledge sharing, that it will be knowledge analysts who provide the most effective 'intelligent user interface' for staff querying knowledge resources. It is the analysts—with skills very similar to library archivists—who will index, structure, and maintain the webs of information and expertise, so it is they who will be able to provide the most effective guidance to staff who are not skilled in database searching (the majority), and whose queries often have to be reformulated more precisely. It is the 'added value' provided by structuring, updating and prioritising information that often makes the difference between success and failure in organisational memory initiatives (in this sense, organisational knowledge resources are very similar to Web resource sites, where the best ones provide coherent organisation and evaluation of sites, rather than just listing them).

Implications for AI support?

It is unlikely, certainly in the near future, that the demonstrable value of human knowledge analysts will be successfully replicated by knowledge-based agents: the interpretive capabilities required to make sense of new information, structure it for publication, and respond to specific queries is too great. It is human knowledge analysts who will determine the value of new knowledge repositories as they are introduced into organisations, given the high start-up costs of publicising and demonstrating the relevance and potential of systematic KM. A more likely role for software agents is as assistants, drawing analysts' attention to connections between previously unconnected staff or resources through techniques such as data mining, and intelligent information filtering and visualisation techniques.

The centrality of stories

Once war stories have been told, the stories are artifacts to circulate and preserve. Through them, experience becomes reproducible and reusable.

(Orr, 1990, p.156)

War stories preserve and circulate hard-won information within the community. (p.157)

What does it mean to become part of a group, to understand, share in, and contribute to its culture? Sociologists, particularly anthropologists and ethnographers such as Orr, point to the importance of 'stories' in this process. This is not an esoteric finding restricted to preliterate tribes, but is just as relevant to advanced, technological cultures, including knowledge-intensive organisations. 'Stories' are informal accounts of past incidents in the life of the group, or of its members. Through them, members share the latest news, show off

expertise, seek advice and compare notes, in the process maintaining their standing and claim to membership within that group.

Implications for AI support?

As an essential process by which culture and knowledge are shared amongst staff, it would be foolish to overlook them in an analysis of KM support. If a veteran member of staff leaves, they take with them their accrued wealth of stories. Is it possible to consider providing technologies which make it very easy to record stories?

Technically, this should not be an obstacle, since we have the digital audio and video recording tools and user interfaces. We can imagine providing easy access to such media, and prompt team members to reflect on key decisions, or relate what they know with respect to a particular problem. Work by Carroll, *et al.* (1994) represents a start in this direction. They captured (as indexed, digital video) designers' recollections of key points in a project, and the rationale behind contentious decisions. Moreover, one could imagine an active software agent prompting designers by email to audio or video record their reflections on key issues, progress reports, project reviews, lessons learned, the genesis of a new idea, or insights from a recent meeting with a client which would benefit colleagues (who would then be alerted). In this context, systems such as Stadium (Eisenstadt, Buckingham Shum and Freeman, 1996) make it tractable for audio-visual briefings/presentations on any scale to be treated as reusable organisational expertise (in Stadium's case, accessible over the Web).

Technically, therefore, capturing and reusing stories is possible. It is, however, important to ask whether stories will be changed by recording them. It is one thing to recount a story to colleagues in the coffee room (and it is in precisely such informal settings that such knowledge sharing occurs), and quite another to record an incident on video which one knows will be added to the company's knowledge repository. Part of the value of stories is that they are elicited and contextualised at a particular moment, for a particular group (often one's peers who share a lot of common knowledge and culture), often to address a particular problem at hand. Recording them so that they can be more robust 'standalone' artifacts, accessible in the future by unknown parties in unknown contexts, will necessitate their modification in certain ways to make them intelligible; the medium moulds the message in this respect. Thus, it will not be as simple as recording the stories that are already told; presentation skills will need to be developed by staff.

As well as the possibility of using agents to prompt for stories on issues of shared concern, there is an opportunity for AI support to track who has told stories about what topics, perhaps inferrable from contextual information such as where and when it was recorded, and to whom.

Document-centred discourse

[S]hared documents within communities are in many ways simply the grounds for a fight, merely the pre-text for agreement. Providing a shared context for constructing meaning, documents are the beginning rather than the end of the process for negotiation.

(Brown and Duguid, 1996)

Documents are central to organisational processes. These are the representations that staff work with—particularly knowledge workers as they consume new information, and translate this into new knowledge, which then shapes, and is disseminated via, new documents. The notion of document includes any digital record: code, mockups, CAD files, etc. It is reasonable on this basis to assume that documents, particularly *digital documents*, have a significant role to play in KM. Given that almost all work involves documents of some sort, the significance accrued by documents in their lifecycle should provide powerful cues for recalling past events and recovering relevant information.

Implications for AI support?

Digital documents can be augmented with facilities to support subsequent annotation and discussion, as they evolve through their lives. Brown and Duguid (1996) have argued that the significance of a digital document may accrue as much from who has read and commented on it, as from the author and the original content. A Web environment that supports this process (in a journal context) is described by Sumner and Buckingham Shum (1996); this provides an environment that tightly integrates a document with discussion about it, providing an 'intellectual trace' of the response the document has provoked, and related resources which readers have linked in.

How could this or a similar environment be augmented with knowledge-based techniques? Since there are invariably standard documents associated with any established organisational process, a knowledge-base should be able to track which documents are currently in circulation, what stage they are likely to be in, dependencies with other documents, and who 'owns' them. It is feasible to imagine agent support for managers wishing to track down and communicate with all 'owners' of documents of a certain class (e.g. related to a particular deadline), where owners might include not only authors, but annotators, reviewers, etc., who have left their mark on it. Often, a document accrues around it a set of 'secondary resources' (assessments; updates; relevant reports or people). In a hypermedia environment, these can be directly linked in. Thus, a search beginning with a single document might lead to many unexpected people and resources within the organisation, that have proven to be significant to the work that the document mediates. Such a web could be monitored and analysed by a software agent. Document-centred discourse is thus a

powerful way to view the discussions and deliberations in an organisation.

Discussion

Undesirable consequences of formalisation

Software design is the process of moving from vague requirements to executable, computational models. Participatory design approaches to interactive system development emphasise the many stakeholders in a system development project, and the need to involve the system's end-users in order to co-design software and work practices. This is no exception for the construction of KM support systems.

Knowledge-system design, as a particular form of software design, is the construction of computer-manipulable representations of domain knowledge. The process of *formalisation* raises a host of issues, some of which this paper has considered (for other analyses see Bowers, 1991; Subrahmanian, *et al.*, 1993; Shipman and Marshall, 1994). From a participatory design perspective, three of formalisation's most significant features in a KM context are as follows:

1. *Representations can become less flexible*, that is, as layers are added, dependencies on old structures increase, and the whole structure becomes harder to change in response to changes in understanding, or of the domain being modelled. Representations tend also to become less tolerant of incompleteness, inconsistency, or ambiguity. This is of course useful for highlighting weaknesses in an organisation's KM, but it may also be a significant limitation, since the models that different parties hold of a domain may be equally valid, but shaped by competing priorities. It may not be possible to satisfy these with one elegant representation. The cost of formalising too early, even semiformally as hypertext, is that it may be too much effort to revise a representational scheme that turns out to be wrong, so it is left as part of the system. Clearly, the art is in knowing when to formalise.
2. *Representations become less comprehensible* to staff who are not knowledge-engineers. One of the consequences of formalisation is that the contents become increasingly inaccessible to the majority of stakeholders. It is of course common that a profession's language and representations are opaque to outsiders, but extra care needs to be taken in KM-system design, due to the legitimate interests of different stakeholders both during design, and following a system's introduction.
3. *Representations support automated analysis*. Clearly this is the main purpose of formalising, so why should this be a problem? Problems arise when the processes

of decomposition and abstraction, required to create a representation capable of supporting automatic analysis, result in models which strip out important contextual details which are in fact critical to understanding the domain. Models of employees' skills, work processes and interdependencies may not adequately express the true nature of their expertise and coordination of work. If the representation is too incomplete (it will always be incomplete to some degree), then the most powerful manipulations and analyses are meaningless. This of course is not a novel insight, but organisational dynamics are particularly difficult to model.

In previous sections, I have suggested opportunities for knowledge-based agents and systems to support knowledge analysts, the capture of audio/video stories, and the indexing of discussions. In response to the general issue of formalisation as set out above, what responses can be made?

Firstly, there is a growing body of work on *incremental formalisation*, which holds promise in coping with some of these issues. Fischer *et al.* (1995) discuss the implications of multiple, often conflicting, domain models for knowledge-based support, and describe an approach and tools that can support multiple representations and incremental formalisation of domain concepts. The requirement to incrementally formalise has been identified from studies of designers using self-selected tools (Sumner, 1995), and designers testing new design notations for capturing design rationale (i.e. design-specific organisational memory) (Buckingham Shum and Hammond, 1994). However, these issues are relevant to knowledge-intensive work of many types. Shipman and McCall (1994) have described an architecture which has supported the gradual construction of knowledge-bases for a variety of domains, beginning with the diagrams and notes that domain experts (not skilled in knowledge engineering) are familiar with, and gradually adding system-interpretable attributes and more structured-content through understandable user interfaces.

Secondly, the problem of premature commitment to structure can also be ameliorated by intelligent identification of potentially significant patterns within flexible, informal structures. For example, Marshall and Shipman (1995) describe a spatial hypertext system, which allows users the freedom to organise objects representing information resources in nested spaces, and then check for recurring patterns of coinciding node types which may be significant. Spatial hypertext tools could form the basis of a knowledge analyst's personal visualisation of particularly hard to classify or changing knowledge resources.

Stable, sanctioned knowledge

Two principles emerge from this discussion which I contend are worth holding onto as the field of KM develops. These can be summarised as:

KM technologies should formalise only knowledge which is stable, and sanctioned.

"Stability" refers to the rate of change in the domain being modelled, relative to the speed with which these changes can be detected (either by knowledge analysts, or automatically by the KM system), and the underlying knowledge representation then updated. Thus, as organisational structures change, as teams change, as individual's skills change, how will these be reflected in the KM system? This relative notion of stability implies that in principle, as advances in the flexibility of knowledge representation are made, the linkage between the model and the domain being modelled (organisational, group and individual cognitive processes) could become tighter, so that more dynamic classes of knowledge can be managed; the domain will be relatively more stable in relation to what the KM system can cope with.

"Sanctioned" knowledge is knowledge mutually agreed upon by all relevant stakeholders. There is a relevant urban-planning practice to call upon here: after laying a fresh area of grass, wait for the main paths to be trodden down; it is then that one builds proper paths to bear the traffic. In domains where consensus is unclear, formalisation should wait until the daily practises and routines of the organisation—some of which may be too complex to predict in advance—reveal the important, stable patterns that are in most need of support. These might include: regular transformations of knowledge from one medium to another; transfer of knowledge from one party to another; filtering functions; interdependencies between two or more schedules; checklists of action items that always need to be addressed whenever a certain event occurs.

The concept of sanctioning knowledge not only emphasises the multi-perspective nature of representing knowledge, but also the issue of the right to know how one is represented in the KM system. This might take a number of forms, varying in the 'strength' of the 'right to know' policy, and the technical complexity of implementing it:

- the right to know the form and content of one's entry in the knowledgebase (e.g. skills; networks; workflows; responsibilities);
- the right to view, or even update knowledge stored about oneself (accessible user interfaces are required here), or to transform knowledge in one medium to another (e.g. from a video story to a textual summary, or vice-versa);
- the right to know if automatic analysis or inferencing by the KM system forms the basis for management policy (appropriate questions can then be raised if there are concerns about the sufficiency of the representation or reasoning);
- the right to participate in any modelling of one's work domain.

At this early stage, it is hard to predict the implications of a truly established 'knowledge economy' (Stefik, 1986) operating within and between organisations. The wide-ranging issues raised in this section are intended simply to open up the space of legitimate issues that should be of concern within the knowledge systems community. Considerable effort is now being devoted to applying knowledge modelling and systems design techniques for KM, but few papers explicitly recognise the informal and social knowledge processes in the organisations (real or imagined) for which they are designing (though see Vanwelkenhuysen and Mizoguchi, 1994; Euzenat, 1996; van Heijst, van der Spek and Kruizinga, 1996 for promising exceptions). Some might respond that it is too early in this field to see serious inter-disciplinary dialogue. Historically, however, the evidence is that even in much more established domains of interactive system design, the relationship between computing, human and organisational disciplines is complex. We need to learn these lessons, and ensure that from the start, the balance of debate reflects its subject matter: *applied* technology research, intimately tied to the *non-deterministic, tacit, evolving* world of collaborative work.

In conclusion, dialogue between the AI community and other relevant disciplines such as human-computer interaction, collaborative computing, workplace ethnography and organisational learning is essential, in order to begin developing the detailed organisational scenarios of use that are at present conspicuous by their absence. From there, the first design iteration needs to be completed with empirical evidence of the success or failure of knowledge management technologies in action.

Acknowledgements

I am grateful to Geof Bowker, and my colleagues Enrico Motta and Tamara Sumner for fruitful discussions on earlier drafts.

References

- Bowers, J. (1991). The Politics of Formalism. In M. Lea, Ed. *Contexts of Computer-Mediated Communication*, 232-261. Harvester Wheatsheaf.
- Bowker, G. C. (in press). Lest We Remember: Organizational Forgetting and the Production of Knowledge. *Accounting, Management and Information Technologies* (<http://alexia.lis.uiuc.edu/~bowker/paper.html>)
- Brown, J. S. and Duguid, P. (1996). The Social Life of Documents. *First Monday*:1 (<http://www.firstmonday.dk/issues/issue1/documents/>)
- Buckingham Shum, S. (1997). Representing Hard-to-Formalise, Contextualised, Multidisciplinary, Organisational Knowledge. AAI'97 Spring Symposium on Artificial Intelligence in Knowledge Management (Mar. 24-26, 1997), Stanford University, Palo Alto, CA, AAAI Press. (<http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-40-abstract.html>)
- Buckingham Shum, S. and Hammond, N. (1994). Argumentation-Based Design Rationale: What Use at What Cost? *International Journal of Human-Computer Studies* 40 (4), 603-652.
- Carroll, J. M., Alpert, S. R., Karat, J., Deusen, M. S. V. and Rosson, M. B. (1994). Raison d'Être: Capturing Design History and Rationale in Multimedia Narratives. In Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems, 192-197. New York: ACM Press.
- Davenport, T. (1996). Some Principles of Knowledge Management. Graduate School of Business, University of Texas at Austin, April, 1996. (<http://knowman.bus.utexas.edu/pubs/kmprin.htm>)
- Davis, R., Shrobe, H. and Szolovits, P. (1993). What Is A Knowledge Representation? *AI Magazine* 14 (1), 17-33. (<http://medg.lcs.mit.edu/ftp/psz/k-rep.html>)
- Eisenstadt, M., Buckingham Shum, S. and Freeman, A. (1996). KMi Stadium: Web-based Audio/Visual Interaction as Reusable Organisational Expertise. Workshop on Knowledge Media for Improving Organisational Expertise, 1st International Conference on Practical Aspects of Knowledge Management, Basel, (30-31 Oct.) (<http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-31-abstract.html>)
- Euzenat, J. (1996). Corporate Memory Through Cooperative Creation of Knowledge Bases and Hyperdocuments. Proceedings 10th Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems, Banff, Canada (November, 1996). (<ftp://ksi.cpsc.ucalgary.ca/KAW/KAW96/22euzenat.ps.Z>)
- Fischer, G., Linstaedt, S., Ostwald, J., Stolze, M., Sumner, T. and Zimmerman, B. (1995). From Domain Modeling to Collaborative Domain Construction. DIS'95: 1st ACM Conference on Designing Interactive Systems, Michigan, ACM Press: New York.
- Kidd, A. (1994). The Marks are on the Knowledge Worker. Proc. ACM CHI'94: Human Factors in Computing Systems, Boston, Mass (24-28 April'94), 186-191, ACM Press: New York.
- Marshall, C. C. and Shipman, F. M. (1995). Spatial Hypertext: Designing for Change. *Communications of the ACM* 38 (8), 88-97.
- Musen, M. A. (1992). Dimensions of Knowledge Sharing and Reuse. *Computers and Biomedical Research* 25, 435-467.

- Orr, J. (1990). Talking About Machines. An Ethnography of a Modern Job. Xerox Palo Alto Research Center, *Technical Report SSL-91-07*
- Shipman, F. M. and Marshall, C. C. (1994). *Formality Considered Harmful: Experiences, Emerging Themes, and Directions*. Xerox Palo Alto Research Center, *Technical Report ISTL-CSA-94-08-02*
- Shipman, F. M. and McCall, R. (1994). *Supporting Knowledge-Base Evolution with Incremental Formalization*. *Proc. ACM CHI'94: Human Factors in Computing Systems, Boston, Mass.*, 285-291, ACM Press: New York.
- Stefik, M. (1986). The Next Knowledge Medium. *AI Magazine*: 7 (1), 34-46.
- Subrahmanian, E., Konda, S. L., Levy, S. N., Reich, Y., Westerberg, A. W. and Monarch, I. (1993). Equations Aren't Enough: Informal Modeling in Design. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 7 (4), 257-274.
- Suchman, L. (1993). Do Categories have Politics? The Language/Action Perspective Reconsidered. 3rd European Conference on Computer-Supported Cooperative Work, Milan, Italy (13-17 Sept.), 1-14, Kluwer Academic Publishers.
- Sumner, T. (1995). The High-Tech Toolbelt: A Study of Designers in the Workplace. In I. R. Katz, R. Mack, L. Marks, M. B. Rosson and J. Nielsen, Ed. *Proc. ACM CHI'95: Human Factors in Computing Systems*, 178-185. Denver, Colorado (May 7-11, 1995): ACM Press: New York.
- Sumner, T. and Buckingham Shum, S. (1996). Open Peer Review & Argumentation: Loosening the Paper Chains on Journals. *Ariadne* (Bi-Monthly Magazine of the UK Electronic Libraries Programme) Issue 5 (Sept.), (<http://www.ukoln.ac.uk/ariadne/issue5/jime/>)
- van Heijst, G., van der Spek, R. and Kruizinga, E. (1996). Organizing Corporate Memories. *Proceedings 10th Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems, Banff, Canada* (November, 1996). (<ftp://ksi.cpsc.ucalgary.ca/KAW/KAW96/25vanheijst.ps.Z>)
- Vanwelkenhuysen, J. and Mizoguchi, R. (1994). Maintaining the Workplace Context in a Knowledge Level Analysis. *JKAW'94: Proceedings 3rd Japanese Knowledge Acquisition for Knowledge-Based Systems Workshop, Hatoyama, Japan*, 33-47