Chapter 5

MAINTAINING ONTOLOGIES WITH ORGANISATIONAL MEMORIES

Yannis Kalfoglou*
University of Edinburgh
School of Artificial Intelligence, Division of Informatics
80 South Bridge, Edinburgh EH1 1HN, Scotland
y.kalfoglou@ecs.soton.ac.uk

Abstract In this article I argue for the convergence of core technologies for knowledge management and organisational memories. Most of the work reported in the literature regards knowledge management and organisational memories as intertwined areas. However, the technologies used to implement and support them are not treated in the same fashion. Usually, they are conceived, developed and deployed separately. This prevents us from exploiting their strength. I identify two such technologies: ontologies and experience factories, originating in these different communities. I use an example case to show how their convergence could be of mutual benefit. I then generalise the approach and speculate on the impact of such convergence in the broader context of knowledge management and organisational memories.

Keywords: Ontologies, Experience factories, Experience Bases, Knowledge Maintenance

MOTIVATION

Nowadays, we are witnessing an increased interest in knowledge management (hereafter, KM), and organisational memories (hereafter, OMs). The former refers to: "the formal management of knowledge for facilitating creation, access, and reuse of knowledge"(O'Leary, 1998a). The

*now associated with:
Advanced Knowledge Technologies (AKT), University of Southampton,
Department of Electronics and Computer Science, Highfield, Southampton SO17 1BJ, UK

latter aims to provide the means for storing, retrieving and distributing knowledge from an organisation's repositories. The goal of OMs is to preserve knowledge whereas KM aims to develop and deploy knowledge. Clearly these two areas are intertwined and centred upon the enhancement of an organisation's competitiveness by improving the way it manages its knowledge.

There are various ways of implementing KM systems in an organisation as well as technologies for supporting OMs development. Ontologies and experience factories have been studied as the means to support KM and OMs, respectively. Ontologies were originally investigated in the Artificial Intelligence (hereafter, AI) community as a way to represent consensual knowledge about a domain of interest in reusable and sharable formats. On the other hand, experience factories (hereafter, EFs), studied in the Software Engineering (hereafter, SE) community, have been explored as the means to promote and manage experiences collected throughout the life-cycle of a software project. Recently, both they have begun to applied in KM and OMs as core technologies to support their implementation. Despite the fact that KM and OMs are intertwined the technologies that are employed to implement them are developed and applied separately. For example, most of applications of ontologies do not mention nor use any kind of organisational framework (i.e., in the form of an OM), but merely are deployed in order to achieve knowledge sharing and reuse benefits. In the same line, EFs amid their successful implementations as OMs, do not employ particular types of ontologies, or even when they do, these are often distant from the ontologies reported in the AI literature. This prevent us from exploiting the full potential of these two core technologies.

In this article I argue for their convergence in order to exploit their potential more fully: knowledge sharing and reuse achieved by deploying ontologies could improve the development and deployment of KM, whereas the presence of an OM, in the form of an experience factory, could allow us better to organise, characterise, and store ontologies. O'Leary refers to *knowledge harvesting* as the converting process which "must identify knowledge that it desirable to share, worth converting, and usable by others'. In this article I paraphrase the above definition in the context of ontologies as follows:

Ontology harvesting must identify ontologies that are desirable to share, worth converting, and usable by others.

I argue that this *harvesting* process can be accomplished by employing OMs technologies, such as EFs.

1. ONTOLOGIES & EFs IN SUPPORT OF KM

Many researchers are investigating the role of ontologies in KM. Recently, the reports of O'Leary in (O'Leary, 1998b and O'Leary, 1998a) and of Benjamins and colleagues in (Benjamins et al., 1998), summarise the contributions made and give us an insight for their potential. I briefly recapitulate here the conclusions drawn in these reports: O'Leary identified three factors that lead organisations to use KM. These were classified as *environmental pressures* imposed by the increasingly competitive global market place, *technological advancements* arising from recent developments in Internet technology, and the ability to *create valuable information* by converting individually available knowledge into group or organisationally available knowledge. Whereas the *environmental pressures* and *technological advancements* give an organisation the reason and the means to pursue KM activities, *creating valuable information* is the goal of KM.

A typology of business processes to achieve this goal is described in (O'Leary, 1998b) and in (Kalfoglou, 2000) I argue that ontologies are present in most of these processes by giving extensive references to indicative work from the ontology research literature. A complete listing of them here is out of the scope and size of this article which is to investigate the potential of converging KM and OMs technologies.

While in the area of KM ontologies have being acknowledged by many as a core technology to support KM, there is no such distinguishable technology in OMs. A reason for this may be the variety of activities that OMs are concerned with. For example, Reimer defines OMs as: "the means by which knowledge from the past is brought to bear on present activities, thus resulting in higher or lower levels of organisational effectiveness" (Reimer, 1998). Its purpose, Reimer continues, is to "ensure coherence between organisational units, coherence between cooperating companies, secure knowledge, and provide knowledge". There is no single technology that can accommodate all those activities and recent implementations combine several. For example, in (Abecker et al., 1998), the authors identified potential research areas that contribute to OMs implementation.

However, there is an area of research, stemming from the SE community, which is being promoted as a potential candidate for implementing OMs. This is the area of EFs[1]. These are the means (Basili et al., 1994) to promote reuse of "all-kinds" of artefacts in an organisation. The core of an EF is the experience base which acts as the OM. The key idea is to

[1] In the SE literature, it is also referred to as "experience management" or "experienceware".

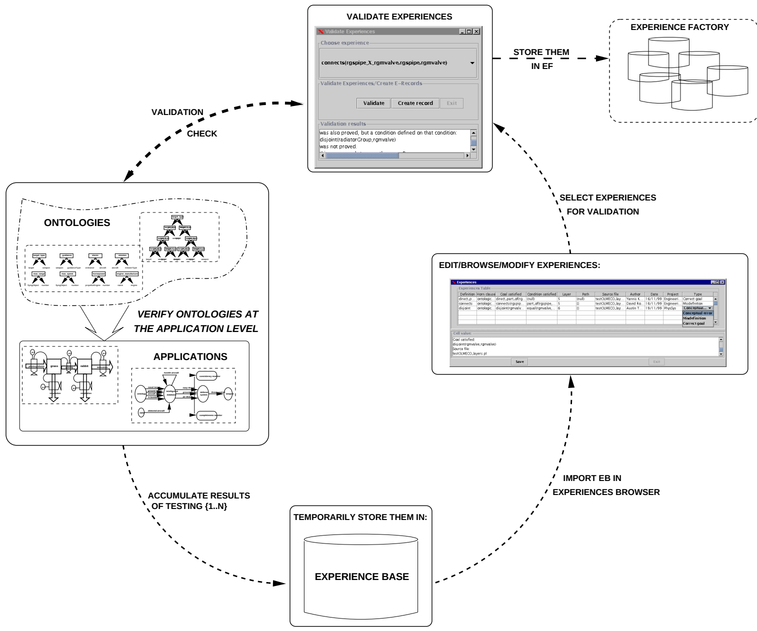

Figure 5.1. Experience-based architecture to support ontology verification (adopted from Kalfoglou and Robertson, 2000).

install an OM to support exchange of all kinds of experiences in the life cycle of a software project. The main focus of an EF is to support 'learning from experience' on a technology-independent organisational level. An EF stores the collected experiences in an experience base. In (Althoff et al., 1998), it was argued that Case-Based Reasoning (hereafter, CBR) plays an important role in the EF paradigm. As CBR provides both the technology and a methodology for 'learning from experience' in the context of case-based knowledge systems, it was natural to use it for implementing continuous learning in an EF style.

2. CONVERGENCE

I argue for the benefits of converging ontologies and EFs via an example case in the area of ontology verification which we originally presented in (Kalfoglou and Robertson, 2000). Most of research in ontologies is focused on development issues and we have witnessed products and tools such as methodologies for building ontologies, online environments for collaborative construction are reported. Even those which are reported do not normally discuss the hidden assumptions and tradeoffs identified during testing. This section explains how OMs technologies, like EFs, may address this issue.

I argue that EFs can be useful in ontology development and deployment as a way of managing the experiences collected from various agents participating in ontology building and usage. This brings, crudely speaking, OMs technology into KM core technology - ontologies - which in turn has important advantages for the organisation that employs KM activities. I elaborate on these in the sequel while here I focus on a simple architecture we built to make the idea more concrete. The architecture, which is centred upon the notion of ontology verification, is given diagrammatically in figure 5.1.

The left-hand side of figure 5.1, depicts the task of verification. In particular, we are interested in verification of ontologies at the application level. That is, we verify that ontological constructs are not misused by applications that adhere to an ontology. After applying our verification mechanism we accumulate, temporarily, the results in an experience base. These are code-testing results and we regard them as experiences. The experience base is then imported by an experiences editing tool which allows for further additions and modification of the description of existing experiences. It allows us to customise the experiences to the particular project as it provides a way of expressing information usually not obtainable through code-testing. We then select the experiences we want to validate and send them to a designated tool for verifying their correctness with respect to test results. This tool embodies the verification mechanism we deploy in the first step but here we apply it to verify the correctness of the results themselves. After the selected experiences have been validated we store them in the final experience base to be part of the EF.

This cycle can be repeated as many times as we wish in the same or other ontologies to collect and manage the knowledge accumulated during verification and testing. Ultimately, this will result in an EF of ontology verification and testing that can be deployed in similar projects in order to facilitate ontology use. It is important to mention that the similarity of projects to benefit from this sort of EF dictates the ability to transfer experiences. How easily can we identify similar projects in order to transfer previous experiences? An answer to this question is given by the ontology characterisation framework discussed later in this section.

We applied this approach in an example case where we deployed and verified an existing ontology at the application level. The case is explained thoroughly in (Kalfoglou, 2000) while here I describe it briefly by following the architecture given above. The ontology which we used, the PhysSys, is a formal ontology based upon system dynamics theory and comprises 7 different ontologies related via an inclusion lattice. We deployed two of these, the *mereology* and *topology* ontology in an exemplar application provided by the PhysSys developers themselves. We used the verification mechanism we invented to check the correctness of *mereological* and *topological* definitions used in the exemplar application. We devised and introduced artificial errors in the application in order to demonstrate the usage of the verification mechanism. Surprisingly, we also detected errors occurred in the original ontologies. We accumulated and verified all the discrepancies found, by storing them as experiences in the final experience base mentioned above. The size of such experience base varies in accordance with the project being applied. In our prototypical case we created a small experience base of a dozen of test cases.

Apart from this specific information on testing we also wanted to store additional information with regard to the ontologies used which will allow us to retrieve and adapt the experiences in different settings. We didn't deal closely with retrieval and adaptation techniques as we intend to adopt those made available by the underlying technology, the CBR method . This extra information has already been identified as 'ontology characterisations' and an early attempt to operationalise them is given in (Uschold, 1998). As Uschold points out: "if this kind of description was provided for all applications, we would quickly get an overview of the state-of-the-art of ontology application, and a way to compare applications. It would be relatively easy to see which techniques have been used to apply particular kinds of ontologies in specific contexts". We found the goal of this research similar to our aim, which is to provide the means to organise and collect experiences gained during ontology testing and verification. Therefore, we found it practical to instantiate the proposed framework for characterising ontologies in our verification and testing scenario.

The instantiated framework is described in detail in (Kalfoglou, 2000) whereas here I list the categories used: the *Purpose* of the ontology; *Representation languages and paradigms*, the *Meaning and formality* of terms used; the *subject matter*; the *scale* of the ontology; its *development* status; its *conceptual architecture*; *mechanisms and techniques* used; and its *implementation platform*.

The role of these characterisations is to provide ontology-specific information which can be used for organising and retrieving experiences on testing in similar ontology deployment efforts.

However, the development and use of ontologies in KM hides important caveats. As O'Leary reports, "Each consulting firm we have been

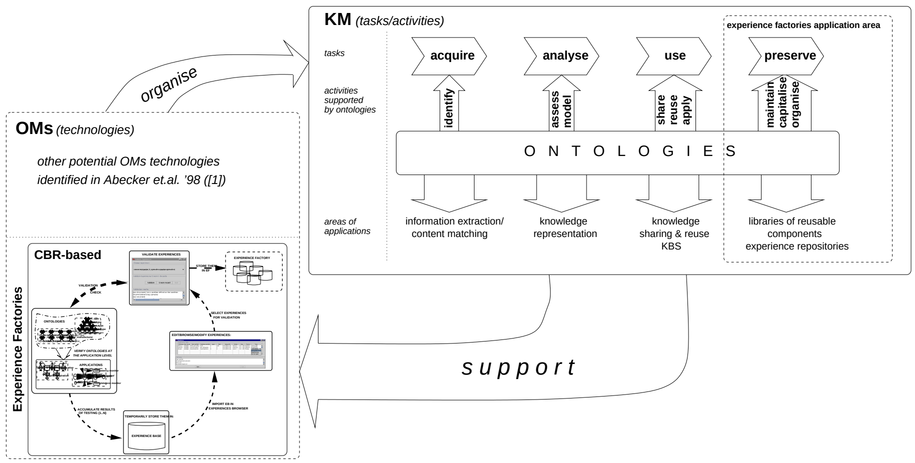

Figure 5.2. OMs in KM: OMs technology used to organise KM tasks/activities which in turn support the implementation of OMs.

examining has built or is building its own ontologies. Because these enterprise ontologies are so costly to develop and maintain and are constantly changing, ontology or taxonomy issues are emerging as some of the most important problems in knowledge management". In the sequel I elaborate on the vision of "bringing OMs technology into KM core technology: ontologies", which could, potentially, alleviate the situation. To explain the idea I illustrate the approach in figure 5.2.

On the left hand side of figure 5.2, I place within a box surrounded by a dash-lined border the OMs technologies. At the bottom of that box I place a potential technology for supporting OMs implementation: EFs. The diagram denotes an example use of EFs in the area of ontology verification as I described earlier. There can be other technologies for implementing OMs. I point the interested reader to (Abecker et al., 1998) for a list of potential technologies.

On the right hand side of figure 5.2, I illustrate the main KM tasks and activities. I identify four main KM tasks: *acquiring*, *analysing*, *using*, and *preserving* knowledge. I argue that these tasks are accomplished by activities which are supported by ontologies. In particular, the knowledge acquisition task, is accomplished by *identifying* activities which are supported by ontologies. This results in the application area of information extraction and/or content-matching. In the same manner, ontologies in the area of knowledge representation are used to *model* and *assess* the environment, which are activities employed in the *analysing* knowledge task. The *using* knowledge task, includes the *apply*, *share*, and *reuse* activities, which are supported by ontologies with such application areas as knowledge sharing and reuse, and KBSs. The last task is *preserving* knowledge. It is accomplished by activities such as *organising*, *maintaining*, and *capitalising* knowledge which are partially aided by ontologies. The resulting application area is that of libraries of reusable knowledge components and experience repositories. The knowledge preservation task and its accompanying activities along with the relevant ontologies are the area of overlap with experience factories as denoted by the dashed box surrounding the task in figure 5.2.

The way in which the two boxes of figure 5.2 are related summarises the linkage I am suggesting in this article. In the introductory section, I argued that OMs and KM are intertwined areas. In this figure I illustrate how the technologies used to implement them can also be intertwined. As can be seen from the curly arrow connecting the OMs technologies with the KM tasks/activities box, technologies such as experience factories, can be employed to *organise* ontologies used to support main KM tasks/activities. The latter, in turn, can *support* OMs implementation by *acquiring*, *analysing*, *using*, and *preserving* knowledge to be processed by an OM.

There are mutual benefits for integrating OMs technology in KM ontologies. On one hand, an OM framework could help to improve ontology development and deployment, facilitate understanding, and ease reuse. A better organised ontology could, in turn, overcome some of the problems identified from the ontology research community: perfect ontology hype, library ontologies, scale-up, interface, and formality. On the other hand, better ontologies could help to meet practical requirements for the implementation of OMs. Among those identified in (Abecker et al., 1999), ontologies could better support the collection and systematic organisation of information from various sources, seamless integrate them into existing work environments and actively present relevant information.

3. CONCLUSIONS

In this article I explored the possibility of employing OMs technology to improve understanding and enhance the usage of ontologies. The latter are a core technology for KM as I discussed in section 5.1. Despite the fact that OMs and KM are intertwined areas and usually are treated as a whole, I argue that the technologies used to implement and support them are not. I discussed the potential convergence of technologies for OMs, in particular experience factories, with ontologies in section 5.2 where I presented an example case in the area of ontology verification and testing. I generalised the approach and speculated on the benefits of such convergence for both OMs and KM.

The intersection of these two technologies also highlights the role of knowledge engineering. It also reveals a change of focus in KM, from IT-based solutions to Knowledge Technology (KT)-based ones. KT could provide solutions in scoping what knowledge is to be captured and disseminated, dealing with tacit knowledge, and facilitating better understanding. The field of knowledge engineering has been studying and practising solutions to these problems for years.

Applying these solutions to KM could help us to realise the 'ultimate goal' of OMs as envisaged by Kuhn and Abecker: "[...] to provide the necessary knowledge whenever it is needed. To assure this, [OMs] realize an active knowledge dissemination approach which does not rely on users' queries but automatically provides knowledge useful for solving the task at hand. To prevent information overload, this approach has to coupled with a highly selective assessment of relevance" (Kuhn and Abecker, 1997). The CBR-based experience factories coupled with ontologically-based reasoning are a way of achieving this 'selective assessment of relevance'.

In accordance with the 'ontology harvesting' notion I mentioned in the introduction I conclude the article with a motto: *if ontologies are to be a cornerstone for successful KM, we need to manage them by bringing OMs technology into their development, deployment, and maintenance phases, which in turn, will result in a more effective OM.*

Acknowledgments

The research described here was carried out during the author's European Union Marie Curie Fellowship (programme: Training and mobility of Researchers) in Edinburgh University and is continuing in the context of the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University. David Robertson and anonymous reviewers of the ECAI 2000 KMOM workshop provided helpful comments on an early draft.

References

Abecker, A., Bernardi, A., Hinkelmann, K., Kuhn, O., and Sintek, M. (1998). Toward a Technology for Organizational Memories. *IEEE Intelligent Systems*, 13(3):40–48.

Abecker, A., Bernardi, A., and Sintek, M. (1999). Proactive knowledge delivery for enterprise knowledge management. In *Proceedings of the 11th International Conference on Software Engineering and Knowledge Engineering(SEKE'99), Kaiserslauten, Germany*, pages 120–127.

Althoff, K.-D., Bomarius, F., and Tautz, C. (1998). Using Case-Based Reasoning Technology to Build Learning Software Organizations. In *Proceedings of the ECAI'98 Workshop on Building, Maintaining, and Using Organisational Memories (OM'98), Brighton, UK.*

Basili, V., Caldiera, G., and Rombach, D. (1994). Experience Factory. In Marciniak, J., editor, *Encyclopedia of Software Engineering*, volume 1, pages 469–476. John Wiley & Sons.

Benjamins, R., Fensel, D., and Gomez-Perez, A. (1998). Knowledge Management through Ontologies. In *Proceedings of the 2th International Conference on Practical Aspects of Knowledge Management(PAKM'98), Basel, Switzerland.*

Kalfoglou, Y. (2000). On the convergence of core technologies for knowledge management and organisational memories: ontologies and experience factories. In *Proceedings of the ECAI 2000 Workshop on Knowledge Management and Organisational Memories (W11-KMOM'00), Berlin, Germany.*

Kalfoglou, Y. and Robertson, D. (2000). Applying Experienceware to support ontology deployment. In *Proceedings of the 12th International Conference on Software Engineering and Knowledge Engineering, SEKE2000, Chicago, IL, USA*, pages 266–275.

Kuhn, O. and Abecker, A. (1997). Corporate Memories for Knowledge Management in Industrial Practice: Prospects and Challenges. *Journal of Universal Computer Science*, 3(8):929–954.

O'Leary, D. (1998a). Knowledge Management Systems: Converting and Connecting. *IEEE Intelligent Systems*, 13(3):30–33.

O'Leary, D. (1998b). Using AI in Knowledge Management: Knowledge Bases and Ontologies. *IEEE Intelligent Systems*, 13(3):34–39.

Reimer, U. (1998). Building, Maintaining, and using Organisational Memories. Invited talk in the ECAI'98 Workshop on Building, Maintaining, and Using Organisational Memories(OM-98), Brighton, England.

Uschold, M. (1998). Where are the Killer Apps? In Gomez-Perez,A. and Benjamins,R., editor, *Proceedings of Workshop on Applications of Ontologies and Problem Solving Methods, ECAI'98, Brighton, England.*