

Integrating Task, Role, and User Modeling in Organizational Memories

Ludger van Elst and Andreas Abecker
German Research Center for Artificial Intelligence (DFKI)
Erwin-Schrödinger-Str., Univ. Bldg. 57, D-67663 Kaiserslautern/Germany
{elst, aabecker}@dfki.uni-kl.de

Abstract

In application areas like personal information agents user models typically maintain sophisticated representations of personal interest profiles. These representations can be utilized for personalized information presentation. Information support in Organizational Memories mainly address the same goals, but prevalently derive actual information needs from task and role specific information needs.

In this paper, we present the Frodo architecture for business-process oriented knowledge management. Task-, user-, and role-driven aspects of information needs are modeled separately and then amalgamated into a specific context for information supply. A description of several usage scenarios elaborates on the assembly of a concrete information need for knowledge retrieval. The paper concludes with the description of an application domain that serves as a testbed for the proposed concepts.

Introduction

The more information is gathered and made accessible in organizations through KM initiatives (e.g., in best-practice databases or design rationale documentations), the more important become high-precision information services for the individual user. User profiling provides techniques to individually tailor information search and delivery. In the larger context of business knowledge management, two observations can be made:

1. Task and role models – as a part of enterprise modeling in business process management and reengineering – can contain additional elements useful for sophisticated information services which go beyond what is known at the individual level (Schreiber *et al.* 1999).

2. For the overall KM objective, it is not ideal to mix together all, maybe orthogonal, determinants of an actual information need in one data structure, the personal user profile. Instead it makes sense to factor out different dimensions which influence an actual information need in order to represent, develop, and maintain them separately, as a part of the company's structural knowledge assets.

Copyright © 2001, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

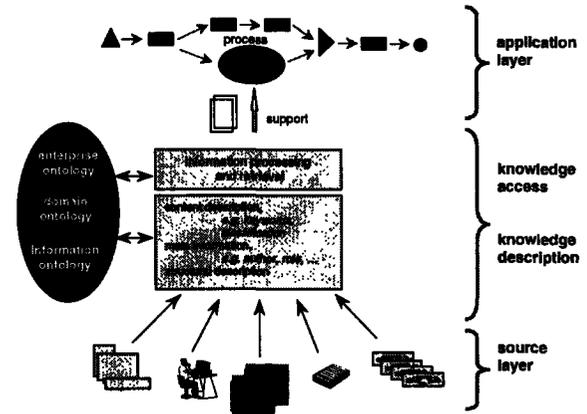


Figure 1: A Reference Architecture for Organizational Memory Information Systems

The paper is organized as follows: First we present a general architectural framework for Organizational Memories. Then we propose to separately model information needs with respect to three determinants, namely task, role, and user. The description of example usage scenarios shows how these information needs may be integrated according to the mode of information transfer. Some remarks regarding future work and a sketch of the envisioned application conclude the paper.

An Architectural Framework for OMs

Organizational Memories are often seen as an adequate means to support knowledge management through information technology. An OM's ultimate goal is to provide knowledge workers with the knowledge they need for their task at hand. Therefore, it stores data, information, and knowledge from different sources of an enterprise. Knowledge dissemination does not only rely on users' queries but also on a proactive knowledge provision based on a highly selective assessment of relevance (cf. (Abecker *et al.* 1997)).

Figure 1 shows the reference architecture that was developed in the KnowMore project (Abecker *et al.*

2000b) for an OM with proactive information services. It consists of a *source layer* that contains information objects, e.g. texts in a document management system, best-practice and lessons-learned archives, yellow pages about expertises and capabilities etc. These information objects with their various formats and degrees of formalization are integrated by *uniform knowledge descriptions*. These knowledge descriptions are formed by linking the information objects to formal structures (ontologies). KnowMore deployed i) an *information ontology*, describing structural properties of information objects (format, author, etc.), ii) an *enterprise ontology* representing the organizational structure (departments, roles, etc.), and iii) a *domain ontology* for modeling the content of information objects.

These knowledge descriptions are utilized for *knowledge access* by information processing and retrieval services that present relevant information to the user. These services may extensively use the knowledge description level for concept-based retrieval with domain-specific strategies (e.g. to find employees with certain competencies (Liao *et al.* 1999)).

The users' information needs are modeled at the *application layer*. In the KnowMore approach a conventional Workflow Management System (WfMS) is proposed as the basis for this layer. The traditional workflow model is enriched in two ways:

1. Task-specific generic queries define the information need. Essentially, the ontologies of the knowledge description layer provide the vocabulary for these queries.
2. Additional context variables model the information flow. So decisions and process details can be transported along the enacted business process and thereby be used as search context in later tasks.

Notice that – in contrast to many *personal information agents approaches* – in this architecture great portions of the information need as well as of the specification *when* the knowledge is useful are located at the application layer and *not* at the knowledge access layer. We expect this to result in more precise retrieval and more appropriate presentation times of information.

In the Frodo project (Abecker *et al.* 2000a) we extend this approach for business-process oriented knowledge management (Abecker *et al.* 1998), (Abecker *et al.* 2000b) towards a comprehensive scenario and platform for *distributed organizational memories*. Thereby we want to attend two drawbacks of centralized approaches:

- They neglect the advantages of the distributed nature of knowledge in enterprises (e.g. with respect to development and use). OMs could benefit from balancing both *local expertise* — which may not be represented in globally shared knowledge — and *overall views* on higher levels.
- Centralized approaches are cumbersome in changing environments. For example are OM systems typically not established at once for a whole enterprise

but introduced separately in several departments and must then be integrated. Furthermore, an OM's environment may change due to reorganizations of an enterprise's structure.

In the remainder of this paper we sketch the modeling of information needs in Frodo which is based on a separation of task-, role-, and user-triggered aspects. Then we describe a couple of usage scenarios that show how these aspects can be integrated in different ways, depending on the mode of information transfer in the OM.

Modeling of Information Needs in Frodo

In this section, we briefly present the main constituents of our approach for modeling information needs in an OM to provide a modular basis for high-performance information services.

Similar to conventional workflow systems, tasks (in business processes), organizational roles, and individual users are the primary elements. Now we propose to attach an information need description (or, a profile) to each of these elements (see for figure 2):

1. A task is described by the information and knowledge required to perform it ("In order to configure these products you need component knowledge as described in the product data sheets and catalogues, you might need advice about procedures and strategies as provided by colleague XY, and it might be useful to know something about experiences with prior configurations as documented in our lessons-learned DB.");

2. an organizational role is equipped with a profile of the information usually consumed ("Our project leaders read the Financial Times newsletter about topics relevant for their projects and are not interested in technical details."); and,

3. each user has, of course, still a personal profile which – among other things – inherits, merges, and adopts parts of the profiles of the roles he currently possesses in the company.

In Frodo, these profiles consist of multiple levels of representation (van Elst & Schmalhofer 1999) where vocabulary from the ontologies at the knowledge description level – mainly the domain ontology – as well as informal descriptions (natural language search terms, example documents) can be combined.

Of course, specific representations of information need in existing information services, like a query in an Internet search (if formulated well), already reflect aspects triggered by the task, the role, and the user, but implicitly and in an intermixed manner. Our proposition is that a clear, declarative structuring of information need determinants provides a comprehensive model of all relevant factors, but gives a better chance for organizational take-up and separate evolution of parts over time (e.g., through adaptive mechanisms working on relevance feedback or collaborative filtering).

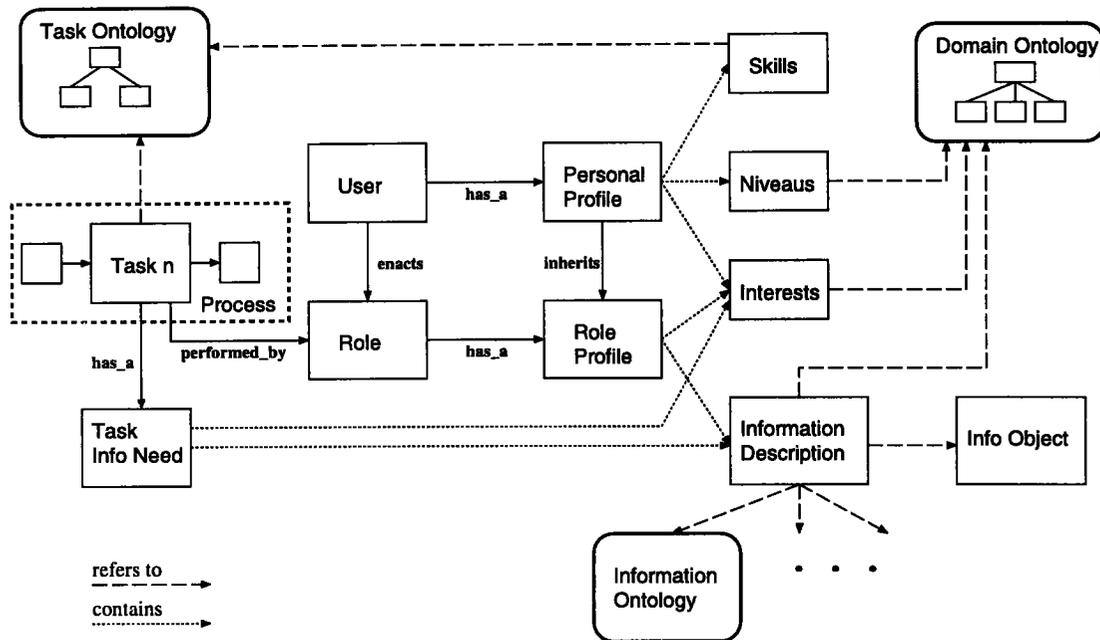


Figure 2: Main Constituents for Modeling Information Needs in OMs

The basic enactment of a business process model then works as follows: The workflow management system goes through the modeled tasks and assigns each task to appropriate actors, according to their role and the skills as specified in the personal profile. In order to actively offer useful information support, the system instantiates the generic task info need, combines it with further restrictions or additional information from the role- and user-profiles and thus achieves a result highly adapted to the given task context, yet tailored for the individual user performing the task.

In the following section, we extend this basic scenario by an overview of different kinds of information services that can be realized on the basis of our architecture. We thereby show that existing information services can be implemented using the model, but also new, innovative services are made possible.

Usage Scenarios: A Comprehensive Overview

Our overview of information services is based on the model of "collection-mediated collaboration" introduced by Winograd and colleagues (Kamiya, Roscheisen, & Winograd 1996) with the Grassroots system. Grassroots unifies personal and organizational information management by integrating several modes of information transfer that are – at the moment – typically provided by independent tools. These modes are characterized by three dimensions, the *regularity* (continuous vs. ad-hoc), the *initiator* (information provider vs. consu-

mer), and the *delivery* (with vs. without notification) of information transfer. These notions allow, e.g., a clear definition of push (information transfer initiated by the source) and pull services (information transfer initiated by the destination).

Ordered according to the predominant information need determinant (task, role, or user), we describe the main "use cases" for information transfer and show the respective roles of the several profile parts defined above.

Task-Driven Information Services

Process-Embedded Ad-Hoc Information Pull: A workflow management system (WfMS) interprets the process logic represented in a business process model. It assigns each task to appropriate actors by matching the possible roles of the task with the roles of the users. When a task is processed the generic task information need is instantiated (e.g. with information from previous activities via flow variables) and modified by the role and user profiles. For example, consider a task "search for literature" in a research project. The task-oriented portion of the information need comes from the project's domain (e.g. "agent technology in KM"). The task itself should be cooperatively solved by the project leader, a researcher, and an assistant. Clearly, each of these roles has different information needs. While the project leader might focus on other finished projects within the project's domain, the researcher would be interested in the latest journal and workshop articles

and the assistant – responsible for the implementation – might need detailed information about agent platforms. The individual employee who performs the task additionally influences the actual information need with his experience and knowledge. Information about a task-related topic where he has much knowledge might be ranked lower; rookies get more basic information and experts more details, etc. Also individual presentation preferences (e.g. language) can be considered.

Process-Embedded Information Prefetch: Process enactment by a WfMS offers knowledge about presumable future tasks. Thus, instead of waiting until a task actually is executed, information about a forthcoming task and the role of its potential processor¹ can be used for information prefetch. This is useful when time consuming, difficult information searches must be performed, e.g., consulting and integrating many sources. After an actual user can be assigned to that task, the result of the information prefetch is at his disposal and may then be individualized (e.g., filtered or re-ranked according to the user's profile).

Task-Oriented Ad-Hoc Information Pull: Often, highly knowledge-intensive processes are not formally modeled because they are too complex or because they are too much ad-hoc. Attentive systems, e.g. personal information agents like Watson (Budzik & Hammond 2000), try to detect the task a user is actually performing, and use this knowledge to retrieve context-oriented information. In contrast to the process-embedded scenarios, only the local work context can be obtained (e.g., the application a user utilizes). Thus, relevant knowledge from preceding tasks is hardly available to better specify the information need. However, the integrated modeling approach allows for an exploitation of the user's roles for more precise assumptions about his current task.

Role-Oriented Information Distribution

Continuous Role-Oriented Information Distribution: Often, regularly information has to be distributed to all users that take a specific role. For example, an up-to-date version of the guideline for preparing project reports should be sent to all project leaders. This is modeled in the role's information need. The role "project leader" can be used as a kind of intensional description of a group of addressees. This description is expanded to generate the extension (e.g., an e-mail list) by resolving the role-user associations. Knowledge about the current task of an recipient then helps to determine the actual presentation strategy. Users that are currently involved in an important other task do not get a high-priority notification while project leaders that are just preparing a report may be interrupted because their actual work is affected. Due to the

¹This means, the actual processor may be unknown at that time.

intension-extension separation the information push is easier to maintain. It is also more user-friendly because of the task-orientated presentation.

Ad-Hoc Notification of User Groups: Additionally, intensional specifications of addressees can be used for an ad-hoc information push with direct notification: E.g. "all secretaries that are preparing an invoice must recognize that the VAT the invoice software generates is no longer valid and must be changed by hand". In this "alarm scenario", the task and the role model are solely used to state who has to get some information.

User-Oriented Information Services

Personal E-Mail: In personal e-mail, the recipient's information need is presumed by the sender. Knowledge about the roles and tasks of a user can be utilized to adequately process the mail. For example, the message can be put into a role- or task-specific mail folder, and the notification mode can be adjusted, depending on the relevance for the task at hand.

Interest-Based Continuous Information Pull: A continuous information pull based on a user's personal interests can be refined by his potential roles. For instance, in a subscribed newsgroup for a specific software system the purchasing agent in a company will mainly be interested in information about new products and prices while a sysadmin is interested rather in installation procedures and troubleshooting.

Summary and Future Directions

In our work, we integrated insights from two research directions that obtain remarkable interest in the AI community, namely user modeling for *personal* information support and organizational memories as parts of comprehensive *knowledge management* endeavours. Though in all scenarios described the actual information need is an *amalgam* from task, role and user, the *discrete modeling* allows for higher reuse of the single models, better maintenance and a more flexible utilization.

From a knowledge management point of view, it is desirable to foster the transitions of information need models from personal user to role and task profiles because the latter two persist even when an employee leaves a company and are therefore better suited to capture corporate knowledge. In our future work we will especially focus on machine learning techniques that utilize traces of the running system to appraise which parts of personal profiles can be leveraged to the role or task, and hence to the organizational level.

The suitability of the Frodo approach is tested in an application scenario in the realm of knowledge management for nuclear power engineering know-how. Here, knowledge is typically distributed over various sites (operators of power plants, federal authorities, etc.),

and there are site spanning processes (e.g. transport of nuclear material) where many roles are involved. Furthermore, in this domain with its long-term scope it is extremely important to have as much knowledge as possible be independent from individuals while for the performance of each process instance a tailoring of information support to the specific knowledge worker is desirable. We think that the presented information infrastructure facilitates the coordination of individual and global concerns with respect to the information needs and thereby can be an enabling factor for comprehensive knowledge management in such a delicate environment.

Acknowledgements. This work was supported by the German Ministry for Education and Research, bmb+f (Grant: 01 IW 901, Project FRODO: A Framework for Distributed Organizational Memories) and by the European Commission (Grant IST-1999-13002, Project DECOR: Delivery of Context-Sensitive Organizational Knowledge). We gratefully acknowledge Knut Hinkelmann, Heiko Maus, and the two anonymous reviewers for valuable input to this paper.

References

- Abecker, A.; Bernardi, A.; Hinkelmann, K.; Kühn, O.; and Sintek, M. 1997. Towards a well-founded technology for organizational memories. In *AAAI Spring Symposium on Artificial Intelligence in Knowledge Management*.
- Abecker, A.; Bernardi, A.; Hinkelmann, K.; Kühn, O.; and Sintek, M. 1998. Toward a technology for organizational memories. *IEEE Intelligent Systems*.
- Abecker, A.; Bernardi, A.; Dengel, A.; van Elst, L.; Malburg, M.; Sintek, M.; Tabor, S.; Weigel, A.; and Wenzel, C. 2000a. FRODO: A Framework for Distributed Organizational Memories. Project Proposal, DFKI GmbH Kaiserslautern. URL <http://www.dfki.uni-kl.de/frodo/>.
- Abecker, A.; Bernardi, A.; Hinkelmann, K.; Kühn, O.; and Sintek, M. 2000b. Context-Aware, Proactive Delivery of Task-Specific Knowledge: The KnowMore Project. *Int. Journal on Information Systems Frontiers, Kluwer* 2(3/4):139–162.
- Budzik, J., and Hammond, K. J. 2000. User interactions with everyday applications as context for just-in-time information access. In *Proceedings of Intelligent User Interfaces 2000*. ACM Press.
- Kamiya, K.; Roscheisen, M.; and Winograd, T. 1996. Grassroots: Providing a uniform framework for communicating, sharing information, and organizing people. In *Proceedings of ACM CHI 96 Conference on Human Factors in Computing Systems*, volume 2 of *SHORT PAPERS: News and Mail*, 239–240.
- Liao, M.; Abecker, A.; Bernardi, A.; Hinkelmann, K.; and Sintek, M. 1999. Ontologies for knowledge retrieval in organizational memories. In *Workshop on Learning Software Organizations (LSO) at SEKE'99*.
- Schreiber, G.; Akkermans, H.; Anjeiwerden, A.; de Hoof, R.; Shadbolt, N.; van de Velde, W.; and Wierlinga, B. 1999. *Knowledge Engineering and Management: The CommonKADS Methodology*. MIT Press.
- van Elst, L., and Schmalhofer, F. 1999. A cooperative comprehension assistant for intranet-based information environments. In Klusch, M.; Shehory, O. M.; and Weiss, G., eds., *Cooperative Information Agents III, Proceedings of the Third International Workshop, CIA '99*, 390–401. Berlin, Heidelberg, New York: Springer.