

Organizing Corporate Memories¹

Rob van der Spek, Eelco Kruizinga, Gertjan van Heijst

CIBIT Consultants | Educators
Prof. Bronkhorstlaan 10 XII
NL-3720 AA Bilthoven, The Netherlands
{rvdspek,ekruizinga}@cibit.nl

Abstract

This paper presents some preliminary thoughts on how corporate memories should be organized in such a way that they maximally contribute to the competitiveness of an organization. We argue that a corporate memory should support three types of organizational learning, which are described. Then we formulate functional requirements and present an architecture for corporate memories that would satisfy these requirements. The paper ends with pointers to related work and future research issues.

1. Introduction

In the first wave of business computerization, companies mainly focused on automating existing tasks in organizations, usually with the use of special purpose software. In particular, clerical activities such as the registration of stocks, transactions, addresses of business relations and the like were automated. Because automation was often initiated at the level of individual divisions, it was not uncommon that the different databases in an organization contained partially overlapping information. This was an unfortunate situation because duplication and separate maintenance of data may lead to inconsistencies and, as a result of this, uncoordinated and erroneous actions. To make things worse, the databases were often based on idiosyncratic data models, which made it hard to check for such inconsistencies. This state of affairs led to a situation where administration with computers was almost as tedious and error-prone as administration without computers.

To overcome these problems, a lot of effort has since been directed to the integration of the various computer systems that are being used in organizations. For example, a number of data exchange standards have been developed and for some types of organizations there are now integrated software packages on the market which support cooperative work (groupware) and which streamline the flow of information in the organization (workflow management tools). As a result of these developments many companies nowadays have an integrated network of computers through which they share documents and data.

More recently, an increasing number of people have begun to realize that computer networks not only enable large scale data sharing, but that it is also possible to exploit them for *knowledge sharing*, thus enhancing the learning capacity of the organization. With knowledge sharing we mean that insights developed at one place in an organization are made available to other parts of the organization by storing them in an organizational knowledge repository, the *corporate memory*. A number of organizations have begun to implement such corporate memories, usually using either Lotus Notes or Intranet² technology. However, what is still missing at the moment is a methodology for setting up such a corporate memory in such a way that it truly contributes to the effectiveness of the organization.

In this paper we present some initial ideas about how such a corporate memory could be organized. These ideas are based on the assumption that the main function of a corporate memory is that it should enhance the learning capacity of an organization. Therefore, Sec. 3 presents an analysis of organizational learning processes. From this analysis we derive some functional requirements that a corporate memory implementation should satisfy.

Then, in Sec. 4 we review a number of implemented corporate memories and investigate to what extent the requirements are satisfied. Sec. 5 then presents some principles for organizing a

corporate memory in such a way that the requirements can be satisfied. It should be emphasized that this paper only presents a conceptual outline of corporate memories, and it raises more questions than it provides answers. In Sec. 6 these questions are summarized and related to other research efforts in this direction.

2. Corporate Memories

The notion of corporate --- or organizational --- memory has been around for more than a quarter of a century (see Stein (1995) for a review of related concepts), and many definitions have been proposed. Most often, these definitions focus on the persistence of knowledge in an organization, independently of how this persistence is achieved. Therefore the knowledge in the minds of the individual workers is also considered as part of the corporate memory. Because our ultimate goal is to investigate how computer systems can be used to realize corporate memories, we have chosen a narrower interpretation of the concept: "a corporate memory is an explicit, disembodied, persistent representation of the knowledge and information in an organization". Note that definition restricts the form of the corporate memory, but not the contents. Any piece of knowledge or information that contributes to the performance of an organization could (and perhaps should) be stored in the corporate memory. This includes knowledge about products, production processes, customers, marketing strategies, financial results, strategic plans and goals etc.

What is the most suitable organization of a corporate memory, depends on how that corporate memory will be used. As mentioned above, the main function of a corporate memory is to improve the competitiveness of the organization by improving the way in which it manages its knowledge. In this context, the notion of learning organizations as first formulated by Argyris and Schon (1978) is worth mentioning. These authors, and also Prahalad and Hamel (1990) view the knowledge assets and the learning capacity of an organization as the main source of competitive advantage.

In the literature on knowledge management, four basic knowledge processes are distinguished, which we will now briefly describe (Wiig, 1993; Nonaka & Takeuchi, 1995; van der Spek & Spijkervet, 1996).

Developing new knowledge Companies survive by the continuous development new knowledge based on creative ideas, the analysis of failures, daily experiences and work in R&D departments. Corporate memories can support these processes by, for instance, recording failures and successes.

Securing new and existing knowledge Individual knowledge must be made accessible to others in the organization who need that knowledge. This knowledge must be available at the right time and place. By storing knowledge in corporate memories it becomes persistent over time and --- if properly indexed --- it can be found back easily.

Distributing knowledge Knowledge must be actively distributed to those who can make use of it. The turn-around speed of knowledge is increasingly crucial for the competitiveness of companies. To support this process, corporate memories need a facility for deciding who should be informed about a particular new piece of knowledge.

Combining available knowledge A company can only perform at its best if all available knowledge areas are combined in its new products. Products and services are increasingly being developed by multi-disciplinary teams. Corporate memories may facilitate this by making it easier to access knowledge developed in other parts of the organization.

It is argued that good knowledge management involves the continuous streamlining of these processes to improve the learning capacity of the organization. Therefore, a corporate memory must be organized in such a way that it maximally supports each of them.

However, it is not enough to have a corporate memory that supports each of these processes individually. In real organizations the processes interact in complex ways and the nature of these interactions should be taken into account as well. For example, it will be clear that the ability to combine different types of knowledge depends on the way knowledge is distributed through the organization.

Further, we believe that the basic knowledge processes do not provide the right level of abstraction for formulating the requirements that a corporate memory implementation should satisfy, because they are not directly related to the goals of organizations (Stein & Zwass, 1995). The ultimate goal of organizations is not to maximize knowledge development, storage, distribution and combination per se, but to improve the competitive power by continuously adapting the organization to the external environment (market, social and political climate, customer preferences etc.). The requirements of a corporate memory implementation should therefore be formulated in terms of these adaptation (= learning) processes. The next section presents an analysis of the how learning takes place in organizations and how these learning processes relate to the above mentioned basic knowledge processes.

3. Lessons Learned Processes

On a global level, two forms of learning in organizations can be distinguished: top-down learning and bottom-up learning. With top-down learning, or strategic learning, we mean that (at some management level) a particular knowledge area is recognized as promising, and that deliberate action is undertaken to acquire that knowledge. Bottom-up learning refers to the process where a worker (either on the management level or on the "work floor") learns something which might be useful and that this "lesson learned" is then distributed through the organization. With the term lesson learned we refer to any positive or negative experience or insight that can be used to improve the performance of the organization in the future. This can be a revolutionary idea that completely changes the ways in which business processes are realized, but more often it will be a relatively down to earth idea about how some task could be performed more efficiently with minor changes.

In this paper, we restrict our analysis to bottom-up learning. Because bottom-up learning is centered around the lesson learned concept, we will in the sequel call this type of learning the lessons learned process. It is hypothesized that a full understanding of lessons learned processes requires insight into three forms of learning which occur in parallel: (i) individual learning, (ii) learning through communication and (iii) learning through the development of a knowledge repository. We will first describe these forms of learning individually, and then how they interact.

3.1 Individual Learning

A basic assumption of our model of the lessons learned process is that organizational learning is based on individual learning: workers gain experience with the way they do their jobs and use these experiences to improve the work processes. This form of learning is depicted in Fig. 1.

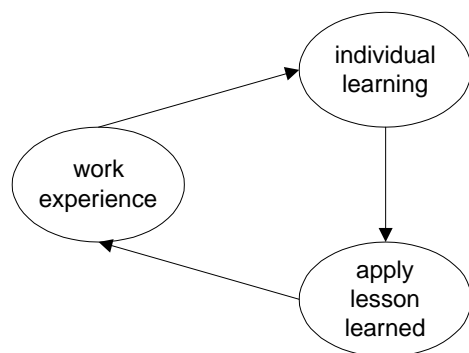


Fig. 1: Individual learning in organizations

In the learning sciences, a large number of different types of individual learning have been distinguished. To name a few: incidental learning, learning by reflection, simulation based learning, case-based learning, learning by exploring, goal directed learning etc. (Reimann & Spada, 1996). Each of these types of learning may occur in the workplace and may lead to useful lessons learned. Therefore, corporate memory implementations should create environments that stimulate these types of learning. For example, case based learning could be supported by maintaining a repository of case descriptions in the corporate memory. Which kinds of individual

learning should actually be supported in an organization depends on a plethora of situational factors.

3.1.1 Organizational requirements

Despite its obvious nature, some organizational requirements have to be satisfied before this type of learning can take place. Firstly, workers need to get feedback about their performance. They need to know the effects of the way they do their job on the processes that they are involved in. Secondly, workers need to have a certain freedom in deciding how they do their job. If this is not the case, workers will not be able to experiment.

3.1.2 Individual learning and the knowledge processes

In this type of learning, new knowledge is *developed* and to some extent integrated in the organizational processes and applied. Knowledge is not *distributed* and it is not *secured* for reuse. For these three processes it is not directly clear how they could be supported by a corporate memory. However, the new knowledge could be the result of *combining* an insight with insights from other parts in the organization. This could be supported by a corporate memory, since this would make it easier to access knowledge from other parts in the organization.

3.2 Learning Through Communication

A second aspect of organizational learning is centered around the concept of communication. This learning process is summarized in Fig. 2. Learning through communication begins with individual learning, but then the individual experiences are shared among co-workers. This may lead to a kind of group learning. Compared with individual learning, learning through communication will be more efficient since lessons learned by one individual can now also be applied by others.

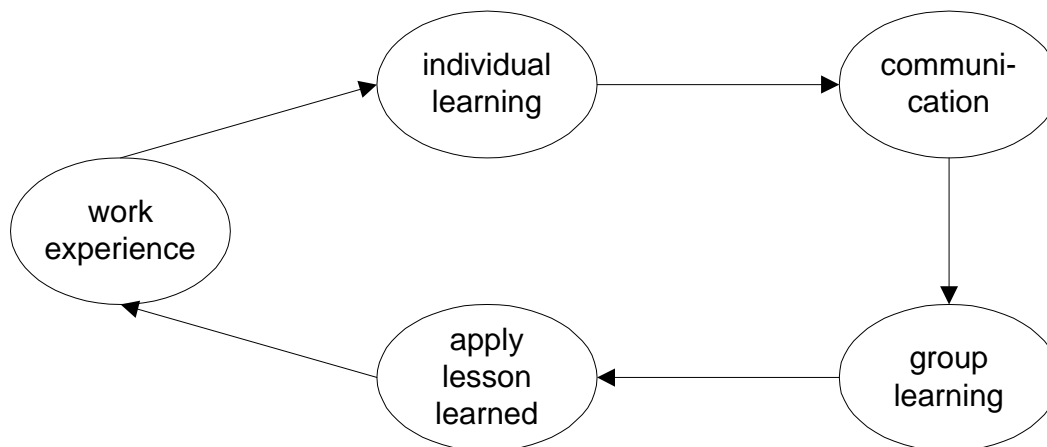


Fig. 2: Organizational learning through communication.

We can make a distinction between two forms of learning through communication: supply driven learning and demand driven learning. In supply driven learning, an individual worker has found a way to improve the work process and communicates this to his/her fellow workers.

In demand driven learning, a worker has recognized a problem in the current process and asks fellow workers whether they have a solution for this problem.

As shown in Fig. 2, learning through communication requires five steps. We will focus on the communication step because this is the step that distinguishes this type of learning from individual learning.

3.2.1 Communication

There are two types of learning through communication: demand driven learning and supply driven learning. In both cases, the individual has to decide on the recipients of the message and on the medium to use. Two factors should be taken into account when taking this decision: (i) overhead, the amount of unuseful messages that is acceptable for recipients, and (ii) hit rate, the amount of people that get the message compared to the amount of people that should have received the message. There are three options:

Personal casting (Only send the message to people directly involved). This is the most efficient way of communication. Only people who can directly help or can take direct advantage of the new knowledge are informed. This way the communication overhead is kept to a minimum, which is important for keeping the communication channel alive.

Broadcasting (Send the message to everybody in the department or organization). This way, the hit rate is maximized at the cost of a large communication overhead. Usually this is not a good strategy but in cases of direct needs with high associated costs it could be considered. An advantage of sending the message to a large audience is that it creates redundancy in the knowledge assets of the organization, which facilitates knowledge development through combination.

Narrow casting (Send the message to everybody who is interested). This option combines the advantages of the first two options, but it requires that workers decide on beforehand in which kinds of messages they are interested (e.g. by means of a user profile) This in turn requires that there is a predefined set of possible topics or, otherwise, that there are guidelines for creating new topics. An example of a mechanism for creating such topic-specific discussion groups are Usenet newsgroups.

3.2.2 Organizational requirements

Because learning through communication is based on individual learning, it has the same prerequisites, plus two additional ones. Firstly, there must be possibilities to communicate experiences between co-workers on a regular basis. This suggests that a corporate memory should provide a mechanism to discuss work related issues, and to archive these discussions. In the sequel we will call this mechanism a discussion forum. Secondly, the organization should have an atmosphere in which it is rewarding to communicate lessons learned (including failures) to others.

3.2.3 Communication and the knowledge processes

In learning through communication, knowledge is *developed*, *distributed* and possibly *combined* with knowledge from other parts of the organization. However, the knowledge is not *secured*.

3.3 Developing a Knowledge Repository

A third aspect of learning in organizations focuses on storing lessons learned in some information repository so that they can be retrieved and used when needed. This form of learning is summarized in Fig. 3.

The process is similar to learning through communication, but now communication is replaced by collection, storage and retrieval. We will briefly describe each of these steps.

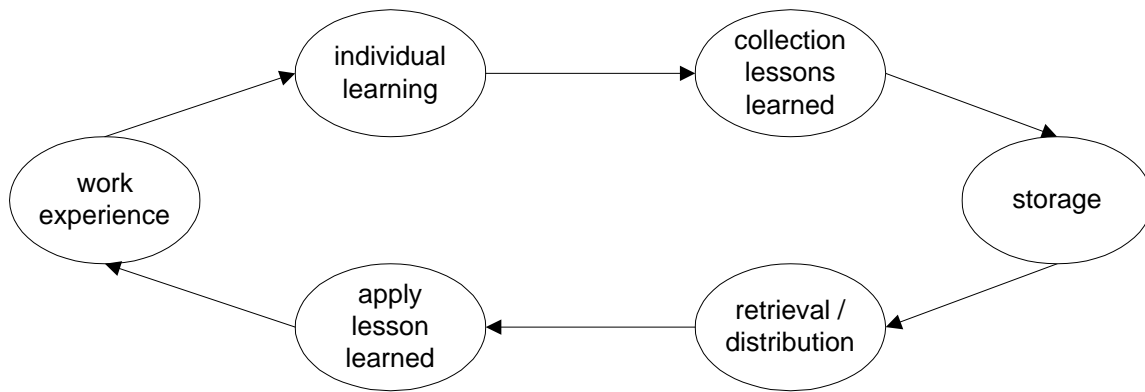


Fig. 3: Organizational learning through maintaining a corporate memory.

3.3.1 collecting knowledge

The knowledge for the corporate memory may be collected actively or passively (from the viewpoint of the corporate memory). Active collection means that some people in the organization are scanning communication processes in order to detect lessons learned. Passive collection means that workers recognize themselves that a piece of knowledge has sufficient value to be stored in the corporate memory. Both forms of knowledge collection require well-defined criteria for deciding if something is a lesson learned.

3.3.2 storing knowledge

The storage step involves two activities: evaluating the submitted lessons learned (and possibly editing them) and indexing the lessons learned. Both steps may be no-ops but this will make retrieval more difficult. Evaluation might include getting answers to questions such as:

1. Is the lesson learned really new?
2. Is the lesson learned consistent with the information already stored in the corporate memory?
3. Is the lessons learned stand alone or should it be integrated with documents already stored in the corporate memory?
4. Is the lesson learned sufficiently general to be useful?

Depending on the result of the evaluation, the submitted lessons learned, may be accepted with or without editing, or rejected.

3.3.3 retrieving and distributing knowledge

As was the case for collection, the knowledge in the corporate memory may be distributed actively or passively. We will call the former retrieval and the latter distribution. In the case of retrieval, a worker recognizes that (s)he needs a piece of information and consults the corporate knowledge base. In the case of distribution, it is somehow decided that a piece of information stored in the corporate memory should be distributed to a number of workers in the organization (thereby entering the learning through communication cycle!).

3.3.4 Organizational requirements

The obvious organizational requirement for this form of acquiring knowledge is that the organizations maintains some kind of knowledge repository.

Further, workers should be motivated to take time to write down their lessons learned and submit them to the knowledge repository. Especially this latter requirement is often difficult to achieve.

3.3.5 Knowledge repositories and the basic knowledge processes

With knowledge repositories, knowledge is *developed*, *secured*, and *distributed*. The newly developed knowledge is possibly the result of knowledge *combination*.

3.4 Dependencies Between the Learning Processes

It should be emphasized that the three types of learning are complementary and occur in parallel. The three learning cycles may interact in complicated ways. Typically, workers will first test a lesson learned by applying the new knowledge themselves, then they will communicate the results to their fellow workers and finally, if the lesson learned is truly effective, it will find its way into the rule books and manuals that are part of the knowledge repository of the organization. When implementing a corporate memory, also these interactions should be taken into account. Otherwise, the corporate memory will not improve the learning capacity of the organization. Fig. 4 shows how the types of learning are related to the knowledge processes, and how they interact.

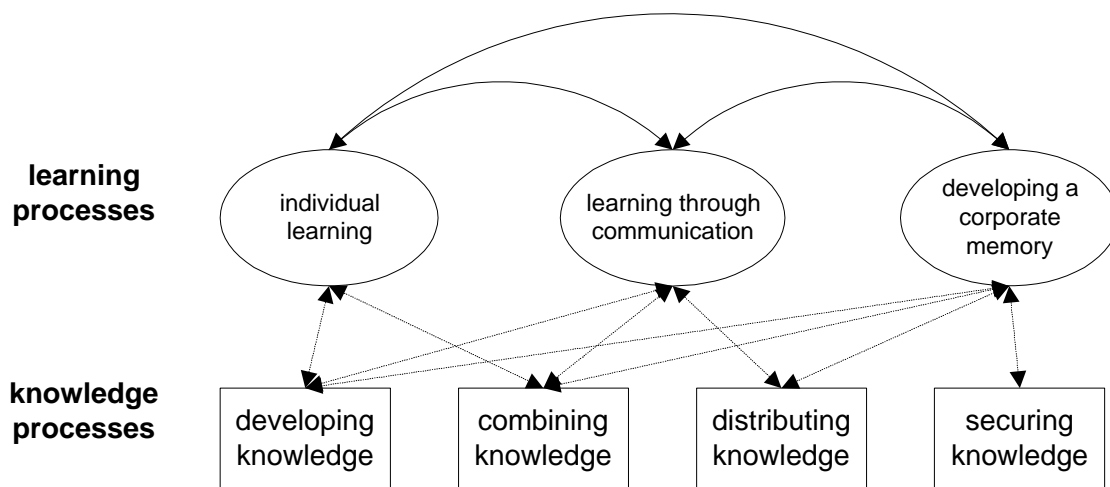


Fig. 4: The types of learning and their interactions and the relation with the knowledge processes

3.4.1 Individual learning and communication

As already shown in Fig. 2, individual learning is a prerequisite for learning through communication. This is particularly true for supply driven learning through communication. Further, individual learning can be the result of combining knowledge acquired through communication with knowledge already available to the individual worker. Thus, the interaction works in both directions.

3.4.2 Individual learning and knowledge repositories

Again, individual learning is a prerequisite for the development of a knowledge repository (see Fig. 3). Individual learning can also be the result of combining available knowledge with knowledge retrieved from the knowledge repository, or even of combining two pieces of retrieved knowledge.

3.4.3 Communication and knowledge repositories

The interaction between communication and the development is also bidirectional. On the one hand, the discussions of communication processes may be archived in the knowledge repository, thus facilitating the learning through communication. On the other hand, the distribution process in the corporate memory cycle is basically a communication process.

3.5 Summary

In this section we have discussed three aspects of bottom-up learning in organization. During this discussion a number of requirements were put forward that a corporate memory should satisfy, which we repeat here.

1. It should be easy for individual workers to access the knowledge in the corporate memory, to facilitate individual learning by combination.
2. It should be easy for workers to decide which of the co-workers could have the knowledge needed for a particular activity.
3. It should be easy for workers to decide which of the co-workers would be interested in a lesson learned.
4. It should be easy (and rewarding) for a worker to submit a lesson learned to the corporate memory.
5. There should be well-defined criteria for deciding if something is a lesson learned, how it should be formulated and where it should be stored.
6. There should be mechanisms for keeping the corporate memory consistent.
7. The corporate memory should have a facility to distribute a newly asserted piece of knowledge to the workers that need that knowledge.

4. Examples of Corporate Memories

In the previous section it was observed that both the collection and the retrieval of knowledge from the corporate memory can be active and passive. Based on these dimensions, we can therefore distinguish between four types of corporate memories (Table 1).

	Passive collection	active collection
passive distribution	the knowledge attic	the knowledge sponge
active distribution	the knowledge publisher	the knowledge pump

Table 1: types of corporate memories.

We will now briefly describe each of these types of corporate memories, illustrate three of them with a case study, and discuss to what extent they support the three forms of learning described in the previous section. It should be emphasized that we have collected the cases by means of a WWW scan, and our discussion is based on the information published on the web. For none of the cases, we have verified our conclusions with the developers.

4.1 The knowledge attic

This is the simplest form of corporate memory management. The corporate memory is used as an archive which can be consulted when needed. In practice this type of corporate memory will often be the most feasible one. The advantage of this type of corporate memory is that it is not intrusive. It emphasizes the bottom-up nature of organizational learning. However, in order to function well it requires a high discipline of the workers in the company.

4.1.1 The NASA Space Engineering Lessons Learned Programme

The Space Engineering Lessons Learned (SELL) programme³ is an initiative of the System Engineering Group of Goddard Space Flight Center. The mission of this group is to ensure that the flight system undergoing design, development, integration, test and evaluation will meet the requirements. To facilitate this, the group has created a lessons learned database which can be accessed through electronic forms.

collecting the knowledge The initial contents of the lessons learned database were retrieved from a survey conducted in 1989. The survey distinguished between three types of lessons learned: most often repeated mistakes, worst practices, and recommendations. These lessons

learned were further classified according to their domains (general, materials, optics, electrical, mechanical, etc.). New lessons learned can be submitted by filling out an electronic form which includes fields for the name, address, e-mail address and organization of the submitter, the subject of the lesson learned, the category of the subject and a description of the lesson learned.

storing the knowledge From the documentation on the web it is not entirely clear how the lessons learned database is organized. It is evident, however, that submitted lessons learned are reviewed and edited before they are asserted, since there is no direct correspondence between the attributes that must be filled in on submission forms and the attributes that can be specified on retrieval forms.

Retrieving and distributing the knowledge To retrieve lessons learned from the SELL corporate memory, a query form must be filled out. This form includes fields for the development phase, the discipline (general, optics, mechanical, etc.), a project name and a number of keywords. There is however no pre-defined set of keywords to choose from.

4.1.2 The SELL corporate memory and organizational learning

Of the requirements specified in the previous section only requirement 1 and to some extent 4 are satisfied by the SELL architecture. Assuming that the submitted lessons learned are indeed edited, it is likely that requirements 5 and 6 are realized by the editors.

4.2 The knowledge sponge

In knowledge sponge corporate memories, the organization is actively trying to develop a more or less complete corporate memory. Whether the memory is actually used to improve the quality of the organizational processes is left to the individual workers. We have not yet found real life examples of implemented knowledge sponge corporate memories.

4.3 the knowledge publisher

In this type of corporate memory, asserting lessons learned is left to the individual workers. The role of the corporate memory maintainers is to analyse the incoming lessons learned, combine these with knowledge in the corporate memory and forward them to the workers for which the lesson learned might be relevant. This can be done in the form of briefings, newsletters courses etc.

4.3.1 The Department of Energy Lessons Learned Programme

The US Department of Energy (DoE) has developed a standard for lessons learned programmes based on Internet technology⁴. The purpose of this standard is promote consistency and compatibility across the different lessons learned programmes that currently are in place within DoE, thus enabling a DoE-wide sharing of lessons learned. The DoE lessons learned standard includes directives for how the responsibilities for the lessons learned program should be distributed, how lessons learned documents should be structured, how they should be reviewed and how they should be stored and distributed.

collecting the knowledge According to the DoE lessons learned standard, lessons learned may come from any reliable source. There is no unit which has as mission to actively search for lessons learned. Thus, lessons learned collection is a passive process. Lessons learned should be submitted using a standardized (non-electronic) form. Forms have fields for the title of the lesson learned, the originator and the contact person for the lesson learned, the priority and the functional category of the lesson learned, a number of keywords, references, a description of the lesson learned, a discussion of the activities which resulted in the lesson learned, an analysis of the lesson learned and a list of recommended actions.

storing the knowledge Before the lesson learned is stored and disseminated, it is first reviewed by a technical expert and an authorized derivative classifier (someone with the authority to decide

whether information is classified). DoE lessons learned are stored in the DoE Lessons Learned Information System (DOELLIS) which is based on Internet technology.

retrieving and distributing the knowledge In the DoE programme, lessons learned are actively disseminated with an assigned priority description. The lessons learned are distributed to the DoE institutes, where local managers decide who in the institute could benefit from such information. The DoE programme does not include protocols for local lessons learned dissemination. Further, the DoE programme also maintains a mailing list to which individual workers can subscribe.

4.3.2 The DoE lessons learned architecture and organizational learning

Of the requirements summarized in Sec. 3.5, the DoE architecture supports 1, 5 and 7. Requirement 6 (consistency) is implicitly supported by having the lessons learned edited before they are stored.

4.4 The knowledge pump

The knowledge pump is the most complex type of corporate memory. In theory, this model ensures that the knowledge developed in the organization is fully exploited to improve the performance of the organization. This type of architecture more emphasizes the top-down nature of organizational learning than the other architectures.

4.4.1 The center for army lessons learned (CALL)

The Center of Army Lessons Learned⁵, which was founded in 1985, collects and distributes lessons learned in combat missions of the US army. One way in which CALL collects lessons learned is by electronic observation forms, which are then evaluated, analysed and organized.

collecting the knowledge The starting point for the lessons learned architecture of CALL is that all CALL products originate with input in the form of an *observation* by an individual. Therefore it stimulates army members, and in particular soldiers, to report observation regarding tactics, techniques, or procedures that their unit used to work around unfavourable situations or circumstances. To facilitate reporting, CALL maintains a WWW site that includes an electronic observation form.

However, besides this passive form of collection, CALL also has a collection division which performs missions to collect lessons learned with respect to particular subjects. The personnel of this division are trained collectors of information and observations. The collection operations are derived from a Contingency Collection Plan, which is maintained by CALL. Collection operations are guided by a collection plan which states which kinds of observation should be made. The procedures for collection missions are described in a handbook which is also published on the WWW.

storing the knowledge Storage in CALL involves both editing and indexing. Editing is performed by CALL's analysis division. In the case of a collection mission as described above, the analysis team is in continuous interaction with the collection teams. A well defined procedure exists for analysing incoming lessons learned. The collected and edited lessons learned are organized and published in a number of products: newsletters, bulletins and handbooks. To some extent, these products are also published on the WWW.

retrieving and distributing the knowledge The corporate memory consists of a number of publications which are distributed to the soldiers. It is not clear in which ways the soldiers are stimulated to actually read the CALL products. In the case of retrieval, the soldier looking for information can make use of an electronic search form on the WWW.

4.4.2 The CALL corporate memory and organizational learning

The CALL corporate memory supports requirements 1, 4, 5, 6 and 7. However, many of these requirements are not directly supported through the way the CALL corporate memory is

implemented, but through the functioning of CALL as a whole. For example, the decision which employees should be informed of a new lesson learned is taken by CALL workers. Nowhere in the corporate memory is made explicit which employees should be informed about which subjects.

5. Structuring Corporate Memories

Whereas the previous sections concentrated on the requirements that a corporate memory should satisfy, this section will focus on the internal organization. We will first discuss with which types of elements the corporate memory should be populated. Then Sec. 5.2 describes possible indexing schemes that can be used to navigate in the corporate memory. In Sec. 5.3 we will discuss the types of knowledge required for realizing these indexing schemes.

5.1 Elements of the Corporate Memory

As the corporate memory is intended as a mechanism for securing corporate knowledge, it is evident that it should be populated with "knowledge objects". However, knowledge objects can be specified at different levels of details. For example, Wiig (1993) distinguishes between seven levels of detail (see Table 2).

knowledge object	example
<i>knowledge domain</i>	internal medicine
<i>knowledge region</i>	kidney diseases
<i>knowledge segment</i>	diagnosis of kidney diseases
<i>knowledge element</i>	diagnostic strategies (e.g. "first collect all symptoms, then try to explain as many of them as possible with one disease candidate")
<i>knowledge fragment</i>	"If the symptom is excruciating pain, then consider kidney stone"
<i>knowledge atom</i>	"Excruciating pain is a symptom"

Table 2: Knowledge objects at different levels of detail (taken from Wiig, 1993).

Even without agreeing that these seven values on the detail dimension are the appropriate landmarks, it is clear that we have a wide range of options. It is clear that the top most level is too coarse grained. For many organizations this would mean that there would be only one or two objects in the corporate memory. On the other hand, the knowledge atom level is too fine grained. On this level, the objects are propositions. If the objects of the corporate memory are to be propositions, this would require the formalization of all the knowledge in the organization, which is obviously not feasible. Thus the right level of detail is somewhere in the middle.

We propose to use *knowledge items*, as used in the CommonKADS organization model (de Hoog *et al.*, 1994) and also described in (van der Spek & de Hoog, 1995) as the basic objects to populate the corporate memory. Knowledge items are modeled as objects with a number of attributes. These attributes are classified in three groups (see Table 3).

general	name: role description: activity: domain(s):	the role the knowledge is associated with the related organizational task reference to organizational areas/objects/processes
content	generic task type: nature: products / services: functions:	from the CommonKADS library tree heuristic, formal, uncertain ... marketable products of the organization organizational functions involved
availability	time: location: form:	when available where available paper, electronic, mind, collective

Table 3: Knowledge items as used in the CommonKADS organization model (taken from de Hoog et al., 1994)

In terms of Wiig's classification, knowledge items correspond to knowledge segments, as they are specific for a particular activity in the organization. The particular choice of attributes for knowledge items is heavily influenced by their function within the CommonKADS framework: the identification of promising areas for knowledge based systems application in the organization. For the present purpose, organizing a corporate memory, some of these attributes are less important, whilst others should be added. In particular, we think that the attributes role description, generic task type, functions and nature can be dropped whereas an attribute "type" should be added. The attribute activity should be changed in activities, to emphasize that a knowledge item might be related to multiple activities. After having discussed how the knowledge items could be indexed, we will present a more detailed discussion of the possible values of each of the attributes.

5.2 Indexing the Corporate Memory

Two of the requirements listed in Sec. 3.5 concern the ease of navigating in the corporate memory, which is directly related to the issue of indexing. We can distinguish between three types of navigation, which we will briefly describe.

hierarchical searching Here the knowledge items are organized in a fixed hierarchical structure. The corporate memory is searched by traversing this structure. Following hyper links on WWW documents is an example of this type of searching.

attribute searching Here, the corporate memory is searched by specifying values for attributes. The search mechanism returns the knowledge items that have the specified values on the attributes. Database engines typically use this type of searching.

content searching Here the user enters arbitrary search terms related to the topic of interest. The search mechanism simply returns all the knowledge items in which the terms occur, possibly with a relevance score. The relevance score expresses how relevant the document is for the goals of the user. This is based on simple heuristics like the frequency of the occurrence of the search terms in the document and the distance between the search terms in the document. This technology is applied by the crawlers of most search engines on the world wide web (Alta Vista, Lycos etc.).

Each of the search methods has its own requirements and its own advantages and disadvantages. For example, whereas attribute searching is more flexible than hierarchical searching, it requires a pre-defined set of attributes. Similarly, while content search is the most flexible method, its ability to find the appropriate documents depends heavily on the capacity of the user to formulate appropriate search terms. Further, this method only works for textual knowledge items.

The suitability of each of the methods depends on the types of knowledge items, the amount of structure in the corporate memory and the amount of possibly relevant knowledge items. Therefore, corporate memory implementations should probably be able to support each of the methods.

In the sequel, we will mainly focus on indexing based on attributes. Indexing by means of a hierarchical structure can be viewed as a compiled form of attribute indexing, where the order in which the attributes are specified is predefined. Although the availability of such predefined search paths facilitates navigation, and thus should be supported, they can easily be derived from attribute-based indexes. Indexing based on work frequencies is conceptually a simple well-understood technique (although the technical realization of search engines that implement it is quite complicated).

5.3 Attributes of Knowledge Items

We will now discuss the modified set of attributes that we propose for knowledge items. Because the ability to find knowledge items in the corporate memory depends to a large extent on string matching, it is important that there is a standard set of possible values for every attribute. For

some of the attributes, this set can be standardized over many organizations, but for others every organization must define its own set of standard terms.

5.3.1 Activities

The activities attribute refers to the organizational activities to which the knowledge item is related. Since different organizations have different business processes, the possible values of this attribute must be specified for every organization individually. Thus, the organization must have an explicit model of the activities that are performed as part of the business processes. The names of these activities can then be used as values on this attribute.

5.3.2 Domains

The domains attribute concerns the subject of the knowledge item. In order to use this attribute, organizations must have an inventory of relevant knowledge domains. This inventory is a meta description of the types of knowledge that exist in the organization. It is our experience that organizations usually don't have such an inventory, and that developing one is a difficult issue for which guidelines are needed. Unfortunately, research on knowledge modeling as thus far mainly focused on the development of guidelines for modeling knowledge on the micro level (e.g. concepts and relations). What is needed here are modeling principles and primitives on a more aggregate level. The hierarchy of medical domains which was used to structure an ontology library described in (van Heijst *et al.*, 1995) is an example of such an inventory. However this hierarchy was more the result of a trial and error modeling process than of a systematic application of guidelines.

5.3.3 Form

The form attribute refers to the physical representation of a particular piece of knowledge. De Hoog *et al.* (1994) identify four possible values for this attribute: paper, electronic, mind and collective. We think the latter should be dropped because it does not concern the form of the knowledge but the availability. Further, the value set is specified on a too coarse grained level, considering the currently available possibilities with multi-media technology. For example the value electronic could refer to a electronic text document, an instruction video, a recorded speech etc. Still, the number of possibilities is limited, and it should not be too difficult for an organization to specify the different forms in which knowledge is available in the organization.

5.3.4 Type

This attribute specifies the kind of "document" of the knowledge item. Possible values include concepts such as procedure, guideline, handbook, manual, best (or worst) practice, progress report, failure report, comment, etc. The values for this attribute are assumed to be reusable across a wide range of organizations, although a particular organization may choose to use only a limited subset. A methodology for corporate memory development could therefore provide a standard set of definitions for this dimension, which can then be adapted and refined for specific classes of organizations. For example, the "protocol" knowledge type could be added as a refinement of "procedure" for hospital organizations.

5.3.5 Products/services

Sometimes knowledge items are directly related to the products and services of an organization. By recording these relations, the corporate memory can be used to improve the communication with the customer, which may lead to a better consumer satisfaction. The possible values for this dimension are of course organization-specific, but they are easy to acquire since most organizations will have no difficulties in producing a list of their products and services.

5.3.6 Time and location

These attributes are mainly relevant for knowledge items which have "mind" as value on the form attribute. In cases where the knowledge is only available in a personal form, the corporate memory should support accessing this knowledge by making it easy to find out how this particular person can be contacted.

5.4 Knowledge Profiles of Employees

Three of the requirements listed in Sec. 3.5 have to do with deciding which co-workers would be interested in a particular lesson learned. To facilitate this, the corporate memory should contain knowledge profiles of all the workers in the organization. To facilitate the process of comparing knowledge items with these worker profiles, these should be formulated using the same attributes and attribute values as used for the knowledge items. However, not all attributes are relevant: it is for example unlikely that a worker will only be interested in lessons learned of the type "comment" on the type attribute. Of the attributes listed in the previous section, activity, domain and products/services seem to be the most relevant to be used in knowledge profiles.

As a default, the profiles of the workers should specify that they are interested in knowledge items that are about the activities, domains, and products and services that are directly related to their jobs. However, in order to facilitate learning through combination, it is important that workers have control over their own knowledge profile. This way they can tune the amount and type of information that they receive to their own interests, needs and information processing capacities.

6. Discussion and Future Work

This paper presents some initial ideas about how corporate memories could be organized in such a way that they contribute to the competitiveness of organizations. In our view, the main function of corporate memories is that they contribute to the learning capacity of organizations. Therefore we have presented an analysis of learning processes in organization and identified ways in which a corporate memory could facilitate these processes. Based on this analysis, we then identified a number of requirements that a corporate memory implementation should satisfy and verified to what extent this was the case for a number of actually implemented corporate memories. Then in Sec. 5 we have presented a framework for organizing corporate memories in such a way that all requirements could be satisfied.

The here presented work only forms an initial step to a complete methodology for corporate memory development, and the major contribution is probably that it has identified a number of issues that need to be resolved. The probably most difficult issues is that there is a need for a method for structuring knowledge at the "macro level" (the level of knowledge items). In order to identify relevant knowledge items in a corporate memory, it must be possible to characterize them on a global level. Thus far, the AI community has mainly focused on characterizing knowledge on the micro level. At this level ontologies have been identified as a powerful tool for specifying knowledge pieces without enumerating them. A similar mechanism is needed for the macro level.

Another interesting issue is how the here presented work relates to the work on enterprise modeling (e.g. Fraser, 1994) and enterprise ontologies (Gruninger & Fox, 1994). We expect that the standard vocabulary for describing several aspects of organizations developed in those projects can be used to standardize the allowed values for some of the attributes of knowledge items. For example, the enterprise ontology described in (Uschold *et al.*, 1995) contains definitions of activities and how they relate to other aspects of the enterprise. Such models could be used to refine the attribute structure of knowledge items as proposed in Sec. 5. However, we don't expect that enterprise ontologies can be used to resolve the issue of characterizing knowledge at the macro level. Whereas enterprise ontologies can be used for modeling the organization, what is needed here is a vocabulary and guidelines for modeling the *knowledge* in the organization.

References

- Argyris, C. and C.A. Schon** (1978) *Organizational Learning: a Theory of Action Perspective*, Reading Massachusetts, Addison Wesley.
- Fraser, J.** (1994) "Managing Change through Enterprise Models", in *Proceedings of Expert Systems '94, the 14th Annual Conference of the British Computer Society Specialist Group on Expert Systems*, Cambridge, United Kingdom.
- Gruninger, M. and M. S. Fox** (1994), The Design and Evaluation of Ontologies for Enterprise Engineering, in *Proceedings of the ECAI Workshop on Comparison of Implemented Ontologies*, Amsterdam, the Netherlands.
- van Heijst, G., S. Falasconi, A. Abu-Hanna, A.Th. Schreiber, M. Stefanelli** (1995) A Case Study in Ontology Library Construction, *Artificial Intelligence in Medicine* 7:227-255.
- de Hoog, R., B. Benus, C. Metselaar, M. Vogler, W. Menezes**, (1994) Organisation model: Model definition document, technical report, the University of Amsterdam, Deliverable DM6.2c of ESPRIT Project P5248 KADS-II.
- Nonaka, I. and H. Takeuchi** (1995) *The Knowledge Creating Company; How Japanese Companies Create the Dynamics of Innovation*, Oxford University Press, New York, USA.
- Prahalad C. K. and G. Hamel** (1990) The Core Competence of the Organization, *Harvard Business Review*, may-june, pages 79—91.
- Reimann, P. and H. Spada**, editors (1996) *Learning in Humans and Machines: Towards an Interdisciplinary Learning Science*, Elsevier, London
- van der Spek, R. and R. de Hoog** (1995) A Framework for a Knowledge Management Methodology, In Wiig, K. editor, *Knowledge Management Methods*, pages 379-393, Schema Press, Arlington, Texas, USA.
- van der Spek, R. and A. Spijkervet**, (1996) *Knowledge Management, Dealing Intelligently with Knowledge*, Kenniscentrum CIBIT, Utrecht, The Netherlands.
- Stein, E. W.** (1995) Organizational Memory: Review of Concepts and Recommendations for Management, *International Journal of Information Management*, 15 (2):17-32.
- Stein, E. W. and V. Zwass** (1995) Actualizing Organizational Memory with Information Technology, *Information Systems Research*, 6(2): 85-117.
- Uschold, M., M. King, S. Moralee, Y. Zorgios** (1995) The Enterprise Ontology, Technical Report, AIAI, the University of Edinburgh, United Kingdom.
- Wiig, K.** (1993) *Knowledge Management, The Central Management Focus for Intelligent-Acting Organizations*, volume 2, Schema Press, Arlington, Texas, USA.

¹ Paper presented at the 10th Knowledge Acquisition Workshop KAW '96, Banff, AB, Canada

² Intranet = Internet technology used for in-house purposes

³ <http://www710.gsfc.nasa.gov/704/lesslrnd/lesslrnd.html>

⁴ <http://www.tis.eh.doe.gov/others/II/II.html>

⁵ <http://call.army.mil:1100/call.html>