21 Knowledge Management

21.1 About This Chapter

This chapter aims at educating readers on the potential benefits that CBR can offer to help identify, evaluate, capture, store, and retrieve an organisation's knowledge assets. Understanding of all previous chapters is desirable, but not crucial.

21.2 General Aspects

In a wider sense, KM is about managing certain kinds of knowledge. So, why would a book about case-based reasoning include a chapter on knowledge management (KM)? An answer to this question is the main purpose of this chapter. First, we examine KM problems, the nature of its processes, its goals and its cycle. Then we compare the CBR and the KM cycles. Next, we illustrate how methods from the CBR methodology can be used to implement KM processes. Last, we discuss for what KM tasks CBR should be used.

Knowledge management (KM) concerns methods that aim at organising, coordinating, planning, commanding, and controlling knowledge assets in an organisation. Because organisations can vary from a small team to hundreds of thousands of members, KM appears in multiple scales. Knowledge also has its several facets, making KM an even broader field. Some specific areas such as Library and Information Science (LIS), whose professionals are in charge of managing knowledge in libraries, have a strong and large agenda for KM. Another field with a wide role is Management and Organisational Science, as it comprises specialists in organisations. Other fields playing a significant role in KM are the computing fields such as Computer Science, Information Systems, Software Engineering, and Information Technology, with the role of implementing computational solutions for KM. This last facet is what we present here, specifically, how to use CBR to perform multiple KM processes.

21.3 Knowledge Management

We now want to introduce KM in a bit more detail. Because of the dependency on the concept of knowledge in KM, we will discuss the term *knowledge*. Next, we will associate knowledge with decision making. This is a way to model the context of knowledge, decisions, and problems. Next, we present a few KM problem situations, the nature of KM, its processes and its cycle.

21.3.1 Knowledge and Knowledge Management

Knowledge management is mostly considered as a part of general management in organisations. From this point of view, knowledge is considered as an abstract collection of assets. As with other managed assets, it has to be made clear how the knowledge is obtained, formulated, stored and used for different purposes.

Knowledge has no unique and precise definition; it is used in different perspectives. The literature suggests six different perspectives to conceptualize knowledge. In this chapter, we are mostly interested in and will be using the perspective of, knowledge as that which enables the use of information to make a decision. In knowledge management, decisions are made to deliver organisational processes.

21.3.2 Knowledge and Decision Making

Consider the model of decision making and problem solving given in Fig. 21.1. Decision making comprises three steps. *Intelligence* refers to gathering information about the problem. *Design* is about identifying what approaches could be used to reach a decision that will enable the originating problem to be solved. It is in the *Choice* step that one approach is chosen, which entails determining the potential outcome of each approach so that the one with best expected result is selected. Both *implementation* of the approach and *monitoring* are part of problem solving.

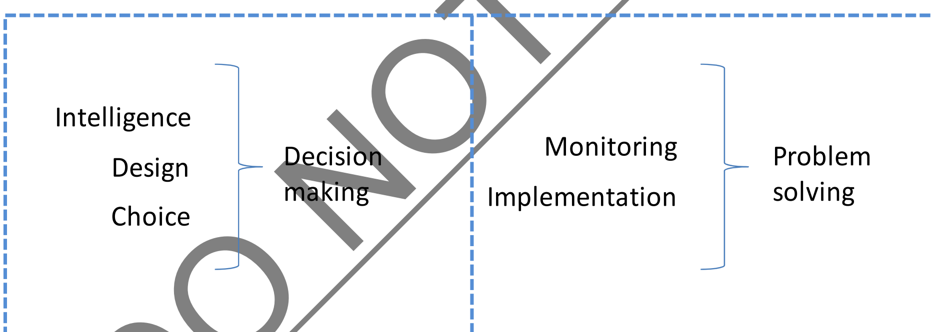

Fig. 21.1 Decision making and problem solving simplified model

The steps in this model of decision making use knowledge in different ways. Knowledge may be used to gather information during intelligence. It is used to execute the design step because it takes knowledge to recognise when an approach has the potential to solve a problem. The choice step in the model also uses knowledge. It entails the prediction of results with the comparison between potential benefits and potential disadvantages.

The problem-solving model allows us to recognise the use of knowledge in knowledge management problems, making it easier to understand the use and reuse of knowledge in the CBR methodology. Next, we discuss some KM problems where KM steps can be implemented to solve knowledge-related problems.

21.3.3 Some Knowledge Management Problems

Consulting companies typically employ qualified personnel to provide knowledge services to their clients. A typical problem is that these companies do not document what they know.

Manufacturing organisations have problems with organizing knowledge about machinery. Experience determines when to stop machinery for maintenance. Members who have this knowledge need to share what they learn so this individual knowledge becomes organisational.

Not specific to any organisation is the need to search for answers to questions or for solutions to problems. Here we refer to the search for answers on topics that are unknown to the searcher. Searches that today are conducted in Web-based search engines not long ago were conducted exclusively by reference librarians. Effectively searching the Web or digital libraries, despite seeming mundane today, requires knowledge of the field and of the resource. Common areas in which laypersons need help are medicine and law.

An important audience of KM solutions is that of scientists. Their work entails production of knowledge and therefore they can immensely benefit from knowledge sharing. Consider identifying open research problems in a field or building a complete literature review on a topic. Given the current view of interdisciplinary scientific challenges and of how collaboration among scientists is viewed as a requirement, sharing of scientific knowledge is a major problem. The other task involving scientists thus becomes knowledge leveraging.

This problem can be tackled with KM. The main steps required are capture, store, and represent knowledge so it becomes available for distribution and reuse. These steps then enable KM processes, commonly referred to as KM tasks, knowledge sharing, leveraging, and organisation.

21.3.4 Knowledge Management: An Organisational Discipline

Knowledge management (KM) inherits the vagueness of the concept of knowledge because it targets the management of knowledge assets. KM is an organisational discipline because it is only needed when more than one individual is involved. Individuals are equipped with internal KM processes, which are apparently seamless. For example, humans do not need an external process to share knowledge with themselves. Most humans are able to remember that touching a very hot surface will burn their skin. This is a form of sharing with oneself, in analogy to sharing between two persons.

Earlier we stated that problems solved by the use of knowledge in KM are organisational in nature. The implication to the model shown in

Fig. 21.1 is that originating problems are organisational and each step of decision making and problem solving requires some form of knowledge. This is an extremely important observation.

Consider that there is an entire field of study dedicated to decision making. Now consider that every organisation — be it for profit or non-profit, private or governmental, — even as informal as a group of friends planning a trip — will reach its goals by executing processes. Many of those processes require decision making. The amount of knowledge involved in reaching goals is thus larger than one can imagine. Now consider that all this knowledge *should* be stored, represented, and ready to be distributed and reused. At least it *should* to the extent that the organisations do not want to make wrong decisions or to reinvent the wheel. Granted that for the friends planning a trip this does not mean much; but for professional organisations, making the wrong decisions or reinventing the wheel means waste.

The overall aggregated ability of an organisation to make decisions and solve problems is reflected in its experience curve. The more experienced an organisation's members are, the better their decisions are likely to be when delivering organisational processes. Recording experiences and providing them when applicable is what is done with CBR.

Through implementing KM steps such as creating, distributing, and reusing knowledge, the target organisation will be more likely to reach its goals and thus achieve its mission. This is because KM steps support other managerial units in better executing their own processes. Consequently, it is of utmost importance that the knowledge used in the KM steps be closely related to the organisation's processes.

As stated before, the knowledge used by an organisation reflects its experience curve. A quick overview of an organisation's activities can reveal whether or not processes are completed. Only a deeper and more subjective view will capture the impact of mature members, who use knowledge learned through experience. The implication is that it is very difficult to demonstrate the value of KM. The outcomes generated by knowledge and experience are not typically included in traditional quantitative methods based on financial statements.

The problem of demonstrating the value of KM has been addressed from the perspective of LIS. This is not surprising, as library services are all KM processes after all. The LIS approach assumes that these hard-to-measure outcomes can be associated with an organisation's mission because they contribute to the organisation's goals. The next section discusses KM goals.

21.3.4.1 Knowledge Management Goals

For any organisational unit, goals are conceived to achieve an organisation's mission. Implementing and maintaining KM goals require change management (see Chap. 11, `Development and Maintenance`). For changing and maintaining a culture that is suitable for KM, the organisation must implement and enforce a series of goals.

The first and more general KM goal is to create an infrastructure for KM. Through the proper infrastructure, the remaining goals can be reached. KM first needs to determine what knowledge is to be managed. It then has to make it transparent to organisation's members who will use KM. They need to be educated on what knowledge needs to be managed.

KM needs to make available to an organisation's members the proper means for knowledge collection, providing them with proper training. The approach shall define how the knowledge is to be represented in order to guarantee better accessibility to it.

For example, it is crucial that an organisation's processes (i.e., of the organisation's units) be included in the captured knowledge to ensure reuse. Knowledge quality is of utmost importance and KM must constantly conduct and maintain processes for validation and verification of knowledge. Note that all these measures require strong leadership support.

Following simple management principles, the entire process shall be monitored so as to guarantee quality and adherence to the chosen approach. Close monitoring of these goals will ensure compliance with the organisation's goals and with the fulfilment of its mission. A better view of KM is given by its cycle.

21.3.5 Knowledge Management Cycle

Like CBR, KM has a cycle too. The cycle varies depending upon the organisation it aims to serve. There is therefore a spectrum of cycles. The level of abstraction of proposed cycles may also vary, and so may the choice of words.

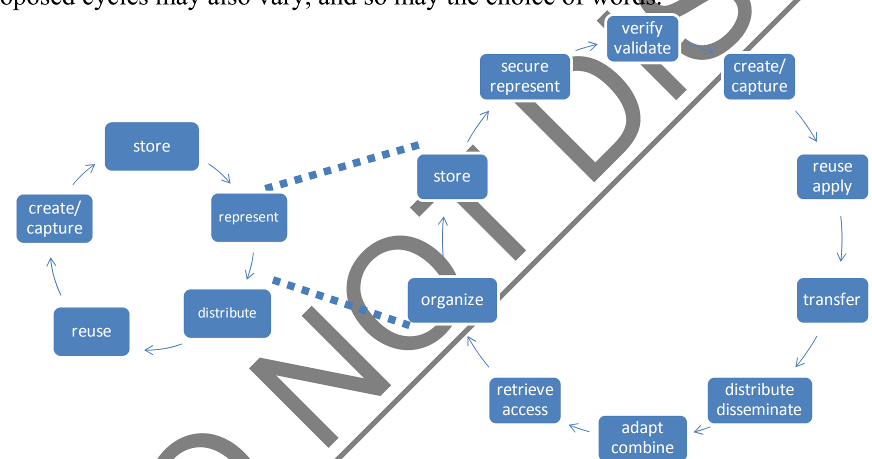

Fig. 21.2. A spectrum of knowledge management cycles

Figure 21.2 shows a minimal cycle on the left. On the right is a more comprehensive cycle. Note that they are, in essence, the same. The cycles may be interpreted as starting in create or capture knowledge, which can be preceded by reuse. They do not end; they continue perpetually, being triggered by organisation's members who, in aggregate, build the organisation's experience curve. One can also observe a relation to the CBR cycle, which we discuss next.

21.4 Case-Based Reasoning and Knowledge Management

We have been discussing KM; we now turn our attention back to CBR so we can compare them. CBR is a reasoning methodology that relies on recalling learned and stored experiences and adapting them to solve new problems. KM is an organisational function that aims at embedding knowledge in processes in support of an organisation's mission. It would be also accurate to describe CBR as a methodology that embeds knowledge to make decisions and solve problems. Furthermore, it would also be accurate to describe KM as a function that supports decision making by recalling existing experiences and adapting them to deliver organisational processes. In fact, they are both inherently the same concept, CBR emphasizing the computational aspects and KM the organisational functions. We may note that in KM the experiences need not be, and are usually not, represented as they are in CBR systems.

The affinities between CBR and KM become explicit as we compare both their cycles. Recall Chap. 2, `Basic CBR Elements`, where we introduced the CBR process model through a series of tasks, problem formulation, retrieve, reuse, revise, and retain. These tasks comprise the CBR cycle. Figure 21.3 shows both cycles.

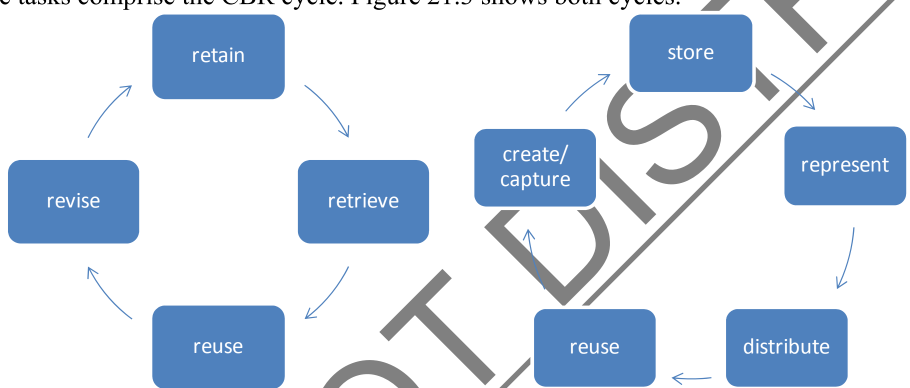

Fig. 21.3. CBR and KM cycles

As previously discussed, the cycles are quite similar. The CBR cycle describes the use of the CBR process while the KM cycle focuses on the processes of a KM organisational unit. The following section explains and illustrates the use of CBR methodology to implement KM cycles.

21.5 CBR Implementing KM Cycles

In this section we illustrate implementations of CBR systems that perform KM steps. Interestingly, as KM relates to knowledge, and all CBR tasks involve knowledge, it could be argued that all CBR implementations perform some form of KM. Nevertheless, if this were true then every intelligent or knowledge-based system could always be seen as a KM system. Therefore, we emphasize here that KM performs knowledge tasks that are processes of an organisation, and their results will likely impact how its goals and mission are achieved. Consequently, CBR should be noted in the KM context when a CBR implementation embeds an organisation's processes in any of its CBR knowledge containers.

21.5.1 Knowledge Infrastructure and Organisation

Code reuse is of great demand in software engineering. For this reason, there are many efforts to utilize CBR techniques for reuse of code in its various forms. The experience factory is one well-structured methodology that goes beyond code reuse. The experiences in the factory are not of the form used in CBR systems as pointed out in Sect. 11.4.3.

21.5.1.1 Experience Factory

When we first mentioned experience factory (EF) in Chap. 11, Development and Maintenance, we presented it as a tool to support development and maintenance of CBR systems. That was a perspective where EF can support CBR. In this chapter we discuss how CBR can be integrated into the EF for managing software engineering experiences to implement a KM cycle. Now CBR is providing support for the EF via KM.

This integration of CBR and EF depends on the fact that EF is an organisational framework for experiences whereas CBR has the techniques to implement reasoning tasks that involve experiences. Using CBR to implement knowledge tasks in the EF demonstrates an important and maybe not so obvious benefit of providing a computational infrastructure for a KM cycle. In this light, the EF is a KM approach.

The resulting integrated framework utilizes a series of experience bases. In this example, CBR is used to perform knowledge creation, analysis, representation, preservation, access, verification, validation, reuse and adaptation, and leveraging. This resulting framework is sometimes referred to as an experience-based information system.

21.5.2 Knowledge Organisation and Retrieval

A well-known profession that specialises in organising knowledge for access is that of reference librarians. A CBR system can potentially realise this entire task. It is, however, not ideal to target a system that replaces humans because this would require the inclusion of several modules such as base ontologies for commonsense reasoning. A rational use of technology is one that complements humans in an integrated environment where the strengths of both computers and humans are maximally utilized. Consequently, the ideal use of a CBR system for knowledge organisation and access would complement the work of reference librarians by providing them assistance. This is a solution for reference librarians in three circumstances: (1) experienced librarians who struggle to keep up with the exponential growth of information resources; (2) novice librarians who are still gaining experience; and (3) one-person reference desks in small libraries, as in schools or law offices.

A CBR system to support reference services is based on its ability to organise information, incorporate expertise, and reuse and adapt previous successful answers.

The input to the system is exactly the same as that received by reference librarians, that is, reference questions (see the example in Table 21.1). The domain-independent nature of this task requires a categorization of case knowledge. This is easily done by grouping cases based on their domains. This generates the need for two-step retrieval, as discussed in Chap. 8, Retrieval. In the first step, a new reference question is used to identify its domain. In the second step, the reference question is used as a new problem. Cases are question-answer pairs.

Table 21.1 shows an example of a typical reference question and its original answer. These actual question-answer pairs are used to create the cases for the case base. Note that the answer includes resources searched by reference librarians. This is the valuable knowledge that can be shared with other reference librarians.

Table 21.1 Question-answer example

Last Update: 2006-04-13 23:12:00.0 Question: Patron needs to know what drugs use equine estrogens. She knows about Premarin. Are there others? Needs drug names and companies that manufacture them. Answer: We were only able to find two other drugs in the Physician's Desk Reference (PDR) in this category. They are named "Prempro" and "Premphase". They are described as (..) manufactures both of these is named Wyeth Pharmaceuticals, Philadelphia, PA. In addition to the PDR, the Seattle Public Library has the following resources which cover this topic. These books are located on the 5th floor of the Central Library. 1. "Complete guide to prescription and non-prescription drugs". by HP Books, c1983- Call #: 615.1 C73865 2. "The Essential guide to prescription drugs". by Harper and Row, c1977- Call #: 615.1 L852E 3. "Prescription drugs" by the editors of Consumer guide. Call#: 615.1 P925 Keywords: estrogen, horses, drugs, nonprescription
This material is from the QuestionPoint Global Knowledge Base.

Question and answers like the one in Table 21.1 are fit to categories of questions, and then simplified for easier matching as a case. Each category shares similar resources for search. Table 21.2 shows examples of two cases originating from actual searches that can be shared among reference librarians.

Table 21.2 Category question and search sources

Question	Answer
Drugs (Other) General (which drugs use equine estrogens)	Physician's Desk Reference (PDR) Complete Guide to prescription and non-prescription drugs (book)
Statistics (death-rate) Treatment (Hepatitis C) Disease (Hepatitis C) General (Hepatitis C, infectious disease)	New York Public Library, Manhattan Branch (http://www.nypl.org/branch/central_units/mm/midman.html; Centre for Disease Control (FAQs); Gale Encyclopedia of Medicine; Medline Plus and HOAH; National Centre for Infectious Diseases; The Hepatitis Info Network, National Institute of Diabetes and Digestive & Kidney Diseases

These elements can be worked as case solutions in multiple ways. Some of the textual elements would need methods as discussed in Chap. 17, Textual CBR. The reuse of previous solution, requires adaptation methods, as discussed in Chap. 9, Adaptation.

With the goal of utilizing the strengths of both humans and computers, such a system would capture expertise embedded in a librarian's answers. Such expertise might otherwise never be explicitly stored. These domain-specific case bases could be potentially shared as cloud resources so librarians anywhere could access and reuse the expertise. This would give more time to reference librarians to use their intelligence to unravel hard cases, that is, searches that are novel and challenging. These unusual and challenging reference questions could also be shared among many human professionals who could use the precious time freed from repeated searches to find new solutions.

This use of CBR in support of reference librarians is also an example of the closeness of knowledge distribution and information retrieval. The association between CBR and information retrieval is later discussed in Appendix B, Relations and Comparisons with Other Techniques.

21.5.3 Knowledge Retrieval and Reuse

Retrieval and distribution are inherently connected. Retrieval functions for CBR have been extensively discussed in Chaps. 8 and 14. The implication for KM is in the distribution because it should be made in a way that motivates other KM processes (e.g., sharing, leveraging, and reuse).

One of the main recommendations in knowledge distribution is to present the knowledge to a potential user when and where it is needed. This implies that KM approaches should be embedded in the environment of the users rather than create new standalone tools for knowledge distribution.

There are a variety of modes for knowledge distribution. The two main categories are passive and active. Passive modes require the user's initiative whereas in active modes knowledge is distributed without the user's request. Another dimension of distribution refers to the number users receiving it; distribution can be broadcast or personalized. Based on the principle mentioned above, the ideal is the active move with personalized distribution delivered in the context of the process for which it is needed, that is, where and when it can be reused.

21.5.3.1 Knowledge Reuse

A properly designed infrastructure for knowledge organisation, distribution and access is the main benefit of supporting knowledge transfer. This support of knowledge transfer is crucial in all fields, so an organisation's members can make the right decisions that will help their organisation achieve its mission and prevent undesired consequences. In some domains the mission of these organisations involves critical aspects like well-being, order, and safety, whereas undesired consequences may include loss of lives. Such organisations are, for instance, dedicated to healthcare and security. Next, we describe an application of CBR for knowledge transfer and reuse for such organisations.

Most organisations today collect lessons learned. One of the early adopters are governmental organisations with members in thousands that use advanced technologies. Examples are space agencies and military organisations. A substantial part of the work in such organisations is in simulated exercises, after which members are asked to describe what they learned and store this information in lessons learned systems.

Lessons learned are described and captured in a variety of forms. A complete lessons learned must include a series of contents, namely, the learned strategy, how it was learned, and how it is applicable for reuse. Extensive work has been done on delineating attributes that characterise high-quality repositories. For reuse, it is essential that these contents be included.

Just as there are many ways in which lessons learned can be distributed, there are also different ways they can be reused. In general terms, when processes or operations are to be delivered manually by humans, lessons must be presented to these humans *at the time and in the context* of the processes for which lessons are applicable.

Consider a KM approach that uses CBR integrated into a system for planning operations. The plan is designed one task at a time. The case-based KM system can track each task included in the plan and search its repository for lessons that are applicable to each task.

Table 21.3 New query interpreted by the case-based lessons learned system

Target process:	transport supplies to affected area
Specific contextual indices:	disaster relief

Suppose an operation is needed for disaster relief. The plan being created includes tasks to bring personnel and medical supplies and to rescue the wounded. The step in the plan is a target operation process, "transport supplies to affected area". This is the most important index used by the case-based lessons learned system. This is what we mean by embedding the target process in the knowledge unit. There must be a way that the KM system can recognise these processes to match them against its knowledge repository whenever they are about to be executed by an organisation's member. The case-based KM system creates the new query as in Table 21.3.

Table 21.4 Knowledge unit in form of a case

Target process:	transport supplies to affected area
Specific contextual indices:	disaster relief
Lesson strategy:	include disaster medical supplies, that is, those that direct limited resources to the greatest number of individuals (as opposed to emergency medical supplies that direct maximal resources to a small number of individuals)

Table 21.4 shows the case that was retrieved when the query was submitted. Note that the requirement that knowledge has to be distributed when and where it is useful poses high demands on the similarity threshold. Most or all the indices should match for a case to be distributed. This type of proactive distribution requires caution, as useless interruptions are not tolerated.

This is a simple example of a lesson that changes a facet of a plan based on a specificity that needs to be made available to the user in charge of planning the operation *at the time and in the context* in which it was needed, where it is applicable.

Though it may seem that the lesson content is obvious, when considering the multiple nature of operations that users may be subject to, lessons targeting specificities of each operation context are extremely important. The automated inclusion of lessons learned in simulated operations has been shown to reduce casualties significantly.

21.5.4 Knowledge Sharing

Knowledge sharing may be the most popular KM task. There are two very distinct reasons why knowledge sharing is challenging. The first is due to the nature of people, the second is due to the nature of knowledge.

Knowledge sharing within oneself is the simple remembering of an episode. Once you learn how to better perform a task, you usually remember that. But even the closest human to you will not benefit from the knowledge you have unless you explicitly share it. Knowledge sharing thus is only an issue when more than one member has the potential to reuse the knowledge. Knowledge sharing requires awareness of the possession of the knowledge, complete lack of barriers for sharing, understanding the knowledge needs of others, and opportunity for sharing.

Knowledge can be of a heterogeneous character and will usually have different sources that can be in conflict. Consider an example. Suppose a company can produce a product in two different ways, one that is environmentally friendly and the other that is not but is cheaper to produce. For both environmental friendliness and economic feasibility, the same organisation may have members who will give conflicting recommendations based on the goals of their departments. A final decision has to be made that combines knowledge shared by both.

21.5.4.1 Sharing and Leveraging in Science Collaboratories

Collaboratories are virtual organisations that aggregate individuals working with a scientific purpose. When the users of a KM system are scientists and engineers who produce scientific knowledge, then a KM cycle can support knowledge representation, sharing, transfer, and leveraging.

The main challenge in collaboratories is the vocabulary and the adoption of an agreed format for knowledge representation. Ways to promote knowledge sharing include easy visualisation of colleagues' works, associations between works that have a methodological or topic overlap, and active distribution of knowledge. Above all, interfaces should be simple and contents minimal but sufficient so others can recognise the potential of collaboration.

As previously discussed, as with other KM efforts, demonstrating results is always a challenge. In scientific collaboratories, however, once scientists share their findings, they can also be asked to relate their work to that of others. Those relations when informed by scientists are evidence of knowledge sharing. At times, scientists will explicitly indicate an association between two knowledge units where the latter leverages knowledge of the former.

The contribution of CBR to collaboratories is multifold. Initially, CBR can guide the format of knowledge units that are contributed, linked, and shared. CBR can support a problem-oriented retrieval for search and for opportunities for active knowledge distribution.

The CBR guidance on formatting scientific contributions is to represent them as problem-solution pairs. Usually, KM approaches encourage users to share knowledge once it is learned. For scientific communities, sharing should occur before the process is completed and the novel scientific contribution is learned. This is because sharing motivations and ongoing research questions or hypotheses encourages collaboration. Sharing only after the fact will encourage knowledge leveraging but not collaboration (i.e., at least the opportunity to collaborate on that specific effort has passed). The knowledge format we demonstrate in Table 21.5 has three temporal dimensions for problem-solution cases. Note that the shaded last two rows represent the problem, that is, indexing elements; the two first rows are the solution, the reuse elements.

Table 21.5. Format for scientific contributions

Prior	Ongoing	Completed
State what is known and what needs to be learned	Declare what you are trying to learn, hypotheses	State what you learned
State the support for this	State what will be done to learn it, experimental design	State the support, your results
Where this knowledge is applicable	Explain its usefulness, where this is applicable	Task or process for which this knowledge is applicable
More specific details	More specific details	More specific details

Now we give an example of a prior and a completed unit that reveals another benefit of this format. It is both an example and a statement about this representation. For prior motivation, "it is hard to motivate users to contribute to KM systems". As support, many references can be used. For indexing elements of the prior case, "This is knowledge applicable in designing, developing, and deploying KM systems". The specifics are, "KM system is of the type repository-based". A completed unit would be for what was learned: "by using a structured representation format like the one in Table 21.5, it becomes easier to generate reports about knowledge entered in the KM systems". For completed support we would have, "we learned that providing reports to the users is a motivation for their use of the system".

21.6 For Which KM Tasks Should I Use CBR?

Service organisations like consulting businesses realise their mission through the work of people. Regardless of whether service employees are of high or low level of qualification, there is a lot of knowledge they use that is kept in their minds. This is a situation where building a CBR system to capture and store experiential knowledge can be beneficial.

In Chap. 12, `Advanced CBR Elements`, we discussed contexts and distinguished different levels. We stated that CBR favour a group level. Often, such a level is given when one considers KM in companies.

A well-designed interface for capturing experiences is crucial for success. All goals discussed in Sects. 21.3.4.1 need to be in place. Cases need to be represented as problem-solution pairs, where one index to guide retrieval is the target process for which experiential knowledge is applicable. The similarity will give a high weight for this attribute, while determining its applicability with specific features that will discriminate applicable situations. For example, an experience only applicable in the evening may have been learned; thus we do not want this to be distributed if the process is to be delivered during the day.

Before starting a new project, the employees themselves or their supervisors can search the system for applicable experiential knowledge. This is also an opportunity for knowledge capture. The reuse step after distribution (i.e., retrieval) should allow for adaptation, i.e., for users to enter new experiences they are reminded of when using the system. CBR methodology entails knowledge adaptation as an essential element of its underlying methodology. CBR is based on the notion of reusing and adapting previous experiences. Adaptation methods are discussed in Chap. 9. Adaptation can also lead to knowledge creation. Additionally, a thorough analysis of previous knowledge may reveal gaps that can guide new simulations to learning missing knowledge.

The importance of adapting knowledge is significant. Consider the cycle of knowledge reuse implemented by humans obtaining knowledge from experts to search (i.e., for experts), negotiate (i.e., contact, schedule meetings), retrieve (i.e., elicit knowledge from experts), and adapt. Studies show that the effort allocated for adapting knowledge is greater than the effort needed for the three other individual steps.

It is therefore expected that the improved ability to automate knowledge adaptation will be more beneficial than improvements to search or distribution. This speculation is supported by the current level of sophistication that search methods have reached; it seems that new improvements in search would produce increasingly smaller benefits, while even a small improvement in adaptation may lead to significant benefits.

If human experts, who have learned through experience, are not available, then experiential knowledge might be available in documents. If documents entail problem-solution pairs that contain useful and reusable knowledge, one possibility is to rely on textual CBR methods, discussed in Chap. 17, to create a CBR system to reuse such knowledge.

It may also be the case that previous episodes of problem solving are recorded in structured databases with distinct fields. Consider examining if these fields individually retain problem and solution values. Then, the next step is to conceive a similarity measure and use those records as cases.

Examples in Sect. 21.5 illustrate many ways in which CBR is used to implement KM tasks. These implementations illustrate how CBR can be used as a computational methodology for KM. Furthermore, they have two other benefits. One is to document an organisation's intellectual assets. The other is that, as processes are embedded in the stored knowledge, they can be easily associated with an organisation's mission, making the implementation also an instrument to demonstrate the value of KM.

21.7 Tools

All CBR tools previously mentioned in previous chapters (e.g., CBRWorks, Orenge, jColibri) can be used for knowledge management tasks. Other tools available for knowledge management (e.g., Microsoft® Sharepoint) are widely used in organisations for collaboration and content and document management. These tools, however, are not designed to perform KM tasks themselves; they do not include CBR. The human users of those tools have to perform the reasoning tasks themselves.

21.8 Chapter Summary

Knowledge management (KM) concerns the proper allocation, coordination and planning of an organisation's intellectual assets. Despite being an organisational problem, its solutions span multiple disciplines. The perspective adopted here is that of information technology, using the CBR methodology to perform KM tasks.

The information technology perspective of KM recognises knowledge as what enables decision making and problem solving. From this view, multiple organisational processes in virtually every domain can benefit from KM. KM embeds knowledge for quality decision making while members deliver organisational processes.

One of the main challenges of KM is demonstrating its effectiveness. Because it aims at supporting organisational processes, and these processes can be associated with an organisation's mission, knowledge units can be linked to an organisation's mission. Explicitly representing them in a CBR system will help document KM steps, its efforts, uses, results, and impact.

The relation between KM and CBR is intrinsic. Their cycles reveal that both have manipulation and reuse of knowledge at the core. Consequently, CBR can be used to perform multiple KM steps. This chapter concludes by describing which KM tasks can be implemented with CBR.

21.9 Background Information

A review of the work done through 2005 is discussed in Althoff and Weber (2006).

The six definitions of knowledge are given in Alavi and Leidner (2001). (1) Knowledge can be considered a process that applies expertise. (2) Knowledge is a capability because it can alter the outcome of a process. (3) Knowledge is an object and thus can be manipulated. (4) Knowledge can be a condition when it accesses information. (5) Knowledge is related to understanding and thus it is related to learning. (6) Knowledge can be also described as information that is tailored to a particular individual or situation.

The simplified decision making and problem solving model was proposed by Huber (1980). The problem of the relation between different challenges for KM was investigated in the IMCOD project; see Bachmann and Dridi (1994).

The difficulty of demonstrating the value of any KM effort was addressed in Abels et al. (2002 and 2004). They explain that outcomes generated by knowledge and experience are not suitable for traditional quantitative methods based on financial statements. They identify and measure the value of library services by associating them to the organisation's mission because they contribute to the organisation's goals ─ the CLIS method. Examples of those library services are timely support for decision making and for the development of policies.

The knowledge cycle presented in Fig. 21.2 with only four steps was suggested by Weber and Kaplan (2003) as a conceptual cycle. It suggests that these steps are always represented in every proposed knowledge cycle and can entail all variations of the tasks.

The goals we discuss that an organisation must meet for successful KM have been devised by Marshall et al. (1996). They studied the relationship between management success and financial health through the intrinsic relationship between risk management and knowledge management. These authors strongly support the use of technology for implementing KM tasks. They argue that, to be truly effective, KM requires organisational change, making the organisation responsible for directing the change by implementing and enforcing a series of KM goals.

The association between the CBR and KM cycles has been discussed in the literature before (e.g., Watson 2003). In his book, the minimal cycle is described through four steps called acquire, analyse, preserve, and use knowledge.

Aamodt and Plaza (1994) introduce the CBR cycle and names the four R's in the cycle, retrieve, reuse, revise, and retain, which has been discussed in Chap. 2, Basic CBR Elements.

A CBR system to complement the work of reference librarians along the lines we described in Sect. 21.5.2, including its tables, was proposed by Bui (2007), who utilizes examples from QuestionPoint − an online reference service made available to the public courtesy of the New York Public Library. See http://www.questionpoint.org/crs/servlet/org.oclc.home.BuildPage.

A review of knowledge distribution is given in Weber et al. (2001). The review categorizes, describes and exemplifies multiple modes and their uses in KM systems.

Barriers to knowledge sharing have been extensively discussed in multiple publications. See, for example, Weber (2007), Disterer (2001), and Szulanski (1996).

Science collaboratories have been defined in Wulf (1993) and described in Finholt and Olson (1997). Weber et al. (2006, 2008) describe the development of a KM approach to support scientists where knowledge sharing is demonstrated via associations entered by scientists. Weber et al. (2008) describes the use of the format exemplified in Table 21.5.

Jacobsen and Prusak (2006) present the results of a study describing the proportion of effort allocated to different knowledge tasks. Their results are as follows: "Searching for knowledge, 10.2%; scheduling meetings with experts, 6.2%; eliciting knowledge from experts, 37.7%; adapting knowledge gained, 45.9%". They conclude that the future payoffs will be on strategies that facilitate knowledge adaptation.

The experience factory (EF) model (Basili 1995) is an organisational approach for continuously learning from experience. Therefore, CBR is an obvious implementation technology for an EF (Henninger, 1995). It has been integrated with CBR by these authors: Althoff and Wilke (1997), Tautz and Althoff (1997), and Althoff et al. (1998).

Weber et al. (2001) describe and illustrate the potential positive consequences of adopting the CBR methodology. This article integrates ideas collected from the AAAI Intelligent Lessons Learned Systems Workshop (Aha and Weber 2000). For the practical adoption of CBR as the underlying framework for lessons learned, Weber et al. (2001) propose a case representation for lessons learned. The case representation was later used in the monitored distribution (MD) approach for proactive distribution of lessons learned (Aha et al. 2001). A description of lessons learned includes the organisational process that it targets. Therefore, MD can be integrated with organisational systems. MD motivates the reuse of a knowledge artefact by bringing to the attention of the user when and where it is applicable and by including a rationale for its reuse (Weber and Aha 2003). The benefit of the MD approach has been demonstrated in an experiment that simulates military operations planned with and without the reuse of lessons learned, taken from the NLLS (Navy Lessons Learned System) repository.

21.10 Exercises

Exercise 1

Describe a KM task that is seamlessly performed internally by individuals.

Exercise 2

How are the KM and CBR cycles distinguished?

Exercise 3

A user submits a knowledge artefact to a KM system and creates a link that associates this new artefact with a previous one and labels the association "uses". What kind of KM task was performed by the user while creating the new artefact?

( ) sharing ( ) leveraging ( ) creating ( ) associating

Exercise 4

Identify a KM task you are familiar with currently being performed by humans.

Exercise 5

Identify a KM task you are familiar with currently being performed by a computer system.

Exercise 6

What kind of system would you recommend to a KM task you are familiar with currently being performed by humans.

Exercise 7

An industry recently replaced some production line workers with robots. Instead of firing the workers, the management trained them to observe the quality of the produced parts so they could indicate when wear and tear required the robots to be removed for maintenance. The problem was that when the workers realise that parts were coming out with defects, too many had already been produced, causing excessive parts to be rejected. What kind of KM approach would you propose that could potentially decrease the volume of rejected parts?

21.11 References

Aamodt A, Plaza E (1994) Case-based reasoning: foundational issues, methodological variations, and system approaches. AI Communications 7(1):39–59

Abels EG, Cogdill KW, Zach L (2002) The contributions of library and information services to hospitals and academic health sciences centers: A preliminary taxonomy. Journal of the Medical Library Association 90(3):276–84

Abels EG, Cogdill KW, Zach L (2004) Identifying and communicating the contributions of library and information services in hospitals and academic health sciences centres. Journal of the Medical Library Association 92(1):46–55

Aha DW, Weber RO (eds) (2000) Intelligent lessons learned systems: papers from the AAAI 2000 workshop, technical report WS-00-03, AAAI Press, Menlo Park, CA

Aha DW, Weber RO, Muñoz-Avila H et al (2001) Bridging the lesson distribution gap. In: IJCAI 2001: Seventeenth international joint conference on artificial intelligence, vol 2. Seattle, WA, 2001. Morgan Kaufmann, San Francisco, CA, p 987

Alavi M, Leidner D (2001) Review: knowledge management and knowledge management systems: conceptual foundations and research issues. MIS Quarterly 25(1):107–136.

Althoff K-D, Wilke W (1997) Potential uses of case-based reasoning in experience based construction of software systems and business process support. In: Bergmann R, Wilke W (eds) GWCBR'97: sixth German workshop on case-based reasoning: foundations, systems, and applications, technical report LSA-97-01E, centre for learning systems and applications, University of Kaiserslautern, p 31

Althoff K-D, Birk A, Gresse von Wangenheim C , Tautz C (1998) Case-based reasoning for experimental software engineering. In: Lenz M, Bartsch-Spörl B, Burkhard H-D et al (eds) Case-based reasoning technology: from foundations to applications. Lecture notes in artificial intelligence, vol 1400. Springer, Berlin, p 235

Althoff K-D, Weber RO (2006) Knowledge management in case-based reasoning. Knowledge Engineering Review 20(3): 305–310

Bachmann B, Dridi F (1994) Definition of a communication layer for expert systems. In: Chen JG (ed) 6th international conference on AI and expert system applications, Houston, Texas, 1994.

Basili VR (1995) the experience factory and its relationship to other quality approaches. Advances in Computers 41:65–82

Bui Y (2007) Case-based support for library reference services. In: Weber RO, Richter MM (eds) ICCBR 2007: case-based reasoning research and development. 7th international conference on case-based reasoning, Belfast, Northern Ireland, August 2007. Lecture notes in computer science (lecture notes in artificial intelligence), vol 4626. Springer, Berlin, p 507

Disterer G (2001) Individual and social barriers to knowledge transfer. In: HICSS 2001: 34th annual Hawaii international conference on system sciences. IEEE, Los Alamitos, CA

Finholt T, Olson G (1997) From laboratories to collaboratories: a new organizational form for scientific collaboration. Psychological Science 8(1):28–36

Henninger S (1995) Developing domain knowledge through the reuse of project experiences. In: Samadzadeh MH, Zand MK (eds) SSR'95: ACM SIGSOFT Symposium on Software Reusability, Seattle, WA, April 23–30 1995.

Huber GP (1980) Managerial decision making. Scott, Foresman, Glenview, IL

Jacobson A, Prusak L (2006) The cost of knowledge. Harvard Business Review 84(11):34

Marshall C, Prusak L, Shpilberg D (1996) Financial risk and the need for superior knowledge management. California Management Review 38(3):77–101

Szulanski G (1996) Exploring internal stickiness: impediments to the transfer of best practice within firms. Strategic Management Journal 17(winter):27–44

Tautz C, Althoff K-D (1997) Using case-based reasoning for reusing software knowledge. In: Leake DB, Plaza E (eds) ICCBR 1997: case-based reasoning research and development. Second international conference on case-based reasoning, Providence, RI, July 1997. Lecture notes in computer science (lecture notes in artificial intelligence), vol 1266. Springer, Berlin, p 156

Watson ID (ed) (2003) Applying knowledge management: techniques for building corporate memories. Morgan Kaufmann, San Francisco CA

Weber RO (2007) Addressing failure factors in knowledge management. Electronic Journal of Knowledge Management 5(3):333-346

Weber RO, Aha DW (2003) Intelligent delivery of military lessons learned. Decision Support Systems 34(3):287–304

Weber RO, Aha DW, Becerra-Fernandez I (2001) Intelligent lessons learned systems. International Journal of Expert Systems Research and Applications 20(1):17–34

Weber RO, Gunawardena S, Abraham G (2008) Representing and retrieving knowledge artifacts. In: Yamaguchi T (ed) PAKM 2008: practical aspects of knowledge management. 7th international conference, Yokohama, Japan, November 2008. Lecture notes in computer science (lecture notes in artificial intelligence), vol 5345. Springer, Berlin, p 86

Weber RO, Kaplan RM (2003) Knowledge-based knowledge management. In: Jain R, Abraham A, Faucher C, Zwaag BJ (eds) Innovations in knowledge engineering, international series on advanced intelligence, vol 4. Advanced Knowledge International, Adelaide, p 151–172

Weber RO, Morelli ML, Atwood ME et al (2006) Designing a knowledge management approach for the CAMRA community of science. In: Reimer U, Karagiannis D (eds) PAKM 2006: practical aspects of knowledge management. 6th international conference, Vienna, Austria, November/December 2006. Lecture notes in computer science (lecture notes in artificial intelligence), vol 4333. Springer, Berlin, p 315

Wulf WA (1993) The collaboratory opportunity. Science 261(5123):854–855