Representing and Retrieving Knowledge Artifacts

Rosina Weber, Sid Gunawardena, and George Abraham

The iSchool at Drexel, College of Information Science & Technology, Drexel University
{rweber,sidath.gunawardena,george.abraham}@ischool.drexel.edu

Abstract. This paper recommends a structure to represent and a method to retrieve knowledge artifacts for repository-based knowledge management systems. We describe the representational structure and explain how it can be adopted. The structure includes a temporal dimension, which encourages knowledge sharing during knowledge creation. The retrieval method we present is designed to benefit from the representational structure and provides guidance to users on how many terms to enter when creating a query to search for knowledge artifacts. The combination of the structure and the retrieval method produces an adequate strategy for knowledge sharing that guides targeted users toward best results.

1 Introduction

Repository-based knowledge management systems (KMS) are a typical knowledge management (KM) initiative employed in today's organizations to promote knowledge sharing between their members. The knowledge to be shared is retained in knowledge artifacts [1], such as lessons-learned, alerts, or best practices. These systems became popularized with the implementation of text databases [2]. Text databases allow contributors to enter knowledge artifacts in a free-text form and can be searched by members seeking knowledge. Unfortunately, such organizations usually lack proper understanding about implementing a strategy for knowledge sharing. Simply making searchable text databases available is not enough to foster knowledge sharing [3].

Consider an example where a scientist has learned that in a commercial airplane, pathogens spread towards the sides and the back–suggesting that passengers in aisle seats in the front will be less likely to become in contact with pathogens spread during a flight. In order to share this knowledge in a KM context, a contributor needs to be guided with answers to questions such as, "Is it the kind of knowledge my organization wants me to share?, "Is it in the right format?", Is it complete or is there anything missing?, "Is it at the right level of specificity?", etc. Providing those answers to KMS contributors has been discussed as a managerial responsibility that organizations have to enforce when implementing a knowledge sharing strategy [4]. Further examples in support of knowledge sharing can be found in, e.g., [5][6][7][8].

Following recommendations in the literature, the contributor of the knowledge in the above example would, for instance, be asked to include the process in which such knowledge becomes relevant, i.e., when booking air tickets. Another structural element to promote sharing would be to include a justification or description of how

T. Yamagchi (Ed.): PAKM 2008, LNAI 5345, pp. 86–97, 2008.

such knowledge was learned, so that other members would be able to determine its validity and decide whether or not reusing it.

This paper advocates the principle that part of this guidance on what and how to submit knowledge artifacts can be embedded in a structure to represent them. The use of the structure is augmented by a reviewing step that provides feedback to contributors helping to educate them about the potential benefits of following proper guidance. We define a structure for artifacts and describe such a structure on a conceptual level that is sufficiently general to be adopted by different target communities in different domains. We also introduce a method that complements the structure to retrieve these knowledge artifacts. We show results from the usage of the representation and a study of the retrieval method.

2 Background and Motivation

This paper extends work that originally focused on lessons-learned, whose concepts are also useful to describe other categories of artifacts, leading to an approach that is general to all. A definition of a lesson-learned was proposed by Secchi, Ciaschi, and Spence [9], which revealed its core concepts. Their definition includes that it is knowledge or understanding that is learned, that it must be validated, and that it has to identify a process or decision where it can be applied [9].

The discussions at the 2000 AAAI Workshop on Intelligent Lessons Learned Systems [10] (see for example, [11] and [12]) focused on the use of intelligent methods to reason with and to retrieve knowledge artifacts. The survey in [3] revealed that the majority of systems used free-text fields to represent knowledge artifacts. It also uncovered the need to include and emphasize the strategy to be learned in fields labeled with names such as *lesson* or *recommendation*. Nevertheless, such a labeled field is usually used as a reminder that a recommendation is expected but there is no reviewing process in place to verify that the artifact meets any necessary requirements.

A preliminary version of the structure presented in this paper for knowledge artifacts was introduced by Weber and Aha [13]. The purpose of that structure was to represent lessons-learned, and to reason with and retrieve them using case-based reasoning. In its essence, it is the same structure that we present in this paper. The difference is that now we have a deeper understanding of its generality and can recommend how it can be used by and adapted to communities from multiple domains. Most importantly, we understand how the formatting of the contributions entered using the structure should be controlled to be complete, effective, and to avoid long texts, which are hard to review, read, and interpret. In this sense, it is controlled because constraints are imposed to the format of contributed artifacts. As a result, the concepts we describe now are more general and more complete. Furthermore, we now have experiences with the use of the structure and have devised a method to retrieve knowledge artifacts represented with it.

The result of the excessive use of free-text forms without guidance in repository-based KMS is collections of knowledge artifacts that are seldom reused. These artifacts may lack vital contents and their unstructured nature can result in difficulties in comprehension. There are many reasons indicating that free-text artifacts are responsible for failure in KMS (e.g., [3], [6]). Nonetheless, we advocate that the most important aspect is how KMS ignore organizational responsibilities for knowledge capture, as discussed by Marshall, Prusak and Shpilberg [4], which are crucial for knowledge sharing. The structure proposed in this paper is aimed at enforcing organizational responsibilities such as guiding contributors on what is to be shared to benefit knowledge sharing.

It may seem that proposing constraints to control the format of artifacts could cause a burden on the contributors whose freedom would be limited. However, our experience has shown that the majority of users are relieved by our concern with not retaining artifacts of poor quality. This resonates with one of the appointed causes of failure in KM approaches that claims that users are not motivated to contribute to a system where they see no value [7]. As a result, our concern with quality of artifacts, even if through some form of control, could potentially increase the confidence users have in the usefulness of the final collection.

3 Structure to Represent Knowledge Artifacts

The time taken by each community member to produce knowledge can vary. Recognizing this variation, we propose a representation for knowledge artifacts that incorporates a temporal dimension associated with the creation of knowledge. We recommend that members share their efforts that are in progress, rather than wait until they are completed and new knowledge is learned. This is particularly useful for communities with tasks spanning long periods.

We first explain concepts in the structures for artifacts that capture artifacts describing learned knowledge, i.e., completed tasks. Next we describe the structure for artifacts that capture the knowledge generation, what we label *in progress* artifacts, i.e., in progress tasks. We also describe how to control their formatting. Then, we present some evidence indicating the usability of the structure, illustrating its benefits for KM tasks.

Table 1. Representation structure for learned knowledge artifacts

Description	Labels	Purpose
It must declare what it teaches	Contribution	Reusable elements
It must state how it was learned	Rationale	
It must explain its usefulness in general terms	Applicable task	Indexing elements
It must explain its usefulness in specific terms	Contexts	

3.1 Elements of the Controlled Structure for Learned Artifacts

The proposed structure for knowledge artifacts consists of four core fields (Table 1). These fields were developed by taking into account the main purpose of knowledge artifacts–knowledge sharing. Consequently, the first field we discuss is the one designated to contain the knowledge to be shared: Contribution.

Contribution. This is the strategy or lesson to be shared, what a contributor learned and believes may be useful to other members of the community. It is a strategy to be either reused or it may even be on to avoid. Part of the concept of the Contribution field is the concept of singleness of an artifact, where each artifact's scope is limited to presenting one strategy. Contribution can be used as a reference to the scope of a knowledge artifact and can also help identify its most suitable level of specificity. The contribution is to be singular in nature and hence it should be communicated by a single statement. Exceptions are granted when a contribution entails technical specifications that may require additional sentences. One rule-of-thumb to identify a contribution is to think of mentioning it and acknowledging its authorship. For example, for the contribution, "*White light is heterogeneous and composed of colors that can be considered primary*" is usually followed or preceded by a statement that it was discovered by Newton. Another form of recognizing the unity in a knowledge artifact is to think of its rationale. Generally, a scientific experiment is designed to demonstrate one hypothesis; on occasion, a single scientific experiment will produce a set of contributions with multiple results. If results had to be repeated to break down the contribution, then this contribution should include all this set. Finally, the unity of a contribution may also be guided by its applicability, which should also be singular. Strategies in contributions are meant to be applicable in an activity or process. Therefore, there should be one single process where it should be applicable. For example, *installing speakers* is a process where a contribution can be applied. Although we would not want to state a contribution that may be useful for more than one process like *installing and selling speakers*, because this would not correspond to a single contribution and would require more than one explanation.

The input into this field should be controlled in length. Though a crisp bound is not defined, the goal is to stick to the contents covered by the guidelines described above. Our experience reveals that 30 or 40 words represent the average, though in a few exceptional cases the text contained close to 300 words. We recommend that humans review these cases and accept exceptions with caution. Long texts are difficult to interpret and may hinder sharing. In case contributors want to include details or background, they can use an additional field that is not a core item in the structure. Nonetheless, excessive background is not necessary because knowledge artifacts are to be communicated to members of the same community, who have some knowledge of the domain.

Rationale. This component of the representation provides an explanation that addresses one of the concerns posed by Szulanski [5] about users being unable to reuse an artifact because they ignore if it is valid. The rationale varies depending on the nature of the contribution. For example, in the Navy Lessons Learned System, one artifact describes a contribution that was learned as advice received from someone else. Advice would be then a type of rationale. Alternative types would be failure or success, meaning that a strategy could have been learned because it was attempted (or not) and the result was success (or failure). A scientific community will obtain results either through quantitative or qualitative methods to support a contribution. A combination of arguments may be used in philosophical contributions, just as in argumentative text.

Control of the Rationale field should follow the same principles adopted to review the Contribution field. The event or source that supports the contribution should correspond to statements in the Contribution field in that the contents of the Rationale field gives the basis for the Contribution field. Descriptions of methodology and experimental design, although relevant, are not meant to be described in Rationale; but results of an event or experiment that substantiate the contribution to be shared.

The two fields Contribution and Rationale belong to the category of reusable elements [13]. They are necessary because they inform the users of the KM system of knowledge they may want to reuse. These fields also provide evidence explaining how the contribution was learned, and help users decide whether to reuse it or not. For effective retrieval of knowledge artifacts, the structure also includes two fields that are indexing elements: Applicable Task and Contexts.

Applicable Task. This is the general activity that one needs to be engaged in for the contribution of a knowledge artifact to become applicable, e.g., installing speakers. While submitting new artifacts, contributors should try to visualize themselves in a position where they could benefit from learning the knowledge they want to convey. Having learned such knowledge warrants them enough understanding to conjecture on what tasks their contribution is applicable.

The Applicable Task is stated in general terms. It is broad enough that it is likely to repeat in the same collection many times. To facilitate its identification, it is limited to only two expressions. The easiest way to control input to the Applicable Task is by identifying verbs and complements that are typical of a community and using drop down lists for user selection. Although the labels of the two dimensions of the Applicable Task will vary depending on the domain, we recommend that the first be a verb and the second be a complement. Of course, not all verbs would be adequate. According to Levin's categorization of verbs [14], we recommend verbs of neutral assessment (e.g., analyze, evaluate) because they have the connotations of exercising actions that are likely to produce a process. These verbs are obviously transitive, so the complement will define the domain the action will impact. The second element of Applicable Task, the complement, is usually a domain-specific term. In the Applicable Task *installing speakers*, *speakers* are associated to the domain because they represent a product sold and installed by a specific organization. It is where the verb's action will be done.

Contexts. The Contexts field is responsible for giving specificity to the Applicable Task. The Applicable Task field presents the general task or process that will match many instances of situations in which community members will be engaged. In contrast, the Contexts field narrows down that generality to help match a description to the usefulness and applicability of the strategy in the Contribution field and exclude contexts that are not applicable. In this sense, Contexts can be filled in with a list of terms that would be potentially used in different situations where the strategy would be applicable. Another way of interpreting Contexts is as a list of state variables. Consider the Applicable Task above *Installing speakers* and think of the contexts of that task and potential variables that could be assigned values. For example, *location* would be a variable and could be assigned values such as home, office, or car; a variable *method* could be assigned values manual or automatic; whereas *type* could be tweeter, back, side, or subwoofer. The ideal list of contexts we are interested in is one that has at least one corresponding value to each variable. The specificity of the artifact would be characterized, for example, by knowledge applicable to manual installation and not applicable to automated installation of speakers; which is useful for home installation and not for in-car installation.

The control of this field is minimal. We recommend keeping a list of terms or expressions rather than full sentences. There are certain situations in which some variables could be treated differently. An example would be to recommend members of a microbial project to enter values for agents as *agent is norovirus*, rather than the word *norovirus*.

Table 2. Representation structure for artifacts in progress

Description	Labels	Purpose
Declare what one is trying to learn	What do you expect to learn?	Reusable elements
State what will be done to learn it	How do you plan to learn it?	
Explain its usefulness in general terms	Applicable task	Indexing elements
…and specific terms	Contexts	

3.2 Elements of the Controlled Structure for Artifacts in Progress

The structure of *in progress* artifacts is laid out in Table 2. The purpose of artifacts that are in progress is to anticipate the timing of contribution, thus increasing the opportunity of sharing.

As presented in Table 2, only the Reusable Elements change by attempting to capture knowledge generation. Rather than asking what someone learned we ask what they expect to learn, like a hypothesis. Similarly, we ask how they plan to learn it instead of asking for results. The way to control them is also analogous to the completed forms.

The decision of the suitability of sharing knowledge artifacts that are in progress should be a community wide decision, based on its potential benefits. For example, in a domain in which this representation is applied, a distributed community of scientists, we learned that sharing in progress knowledge artifacts has many benefits. Interests of members are made available to the community earlier than it would if everyone waited for a completed research, as one may take several months, or even years before it produces an innovation and can be considered completed. Sharing in progress work may help identify collaborations that would otherwise be missed.

Another important benefit to scientific communities is in the question of how members plan to learn a contribution. The benefit is that sharing *in progress* artifacts gives an opportunity to describe the experimental design (or methods) used to reach results that substantiate a contribution; particularly because a substantial portion of the knowledge learned by scientists is embedded in the refining of hypotheses and their respective experimental designs. This may seem to be missing from the completed artifacts structure, but we find that it would make that representation too long to enter and read. This approach that combines this temporal dimension allows us to meet all our objectives in support of knowledge tasks.

3.3 Results from Usage

The most comprehensive application of the proposed structure to represent knowledge artifacts is with a community of scientists, described in PAKM 2006 [15]. This section presents analyses of results from this implementation and attempts to relate it to some of our expected results.

Sufficiency. One metric of quality relates to the structure's ability to represent contents that are in the minds of contributors. Once they understand the concepts of the representation, it is our expectation that they will find the structure sufficient to capture all contents they have in mind. In order to assess how well the proposed structure is meeting this expectation, we computed the number of domain specific expressions in 177 knowledge artifacts (the first 2 years worth of submissions) for the system described in [15]. We then compared that number with the number of domain-specific terms that appeared in the four core fields and the title. The title was used because it is not a core field for the structure but it was requested by users given that it is a habit that may help them organize their ideas. Our assumption is that whenever domain-specific terms show up in a title and not in the core fields, the structure fails to provide sufficient concepts for contributors to communicate knowledge artifacts. For the artifacts contributed, we found that 99.5% of terms in the artifacts can be found within the core fields of the representation; implying that contributors did not find a proper field for only 0.5% of the terms. We will observe this metric in next years to determine consistency.

Temporality. The expectation of adopting knowledge artifacts that vary along the temporal dimension is that it allows knowledge to be shared before it would otherwise be available to the community. More specifically, our assumption is that in progress artifacts would make knowledge available for sharing before completed artifacts. In the first two years, while 22% of the contributed artifacts were completed, 64.4% were in progress. Adopting this temporal dimension provided for roughly three times more knowledge for sharing than it would have been possible without it.

Knowledge Sharing. At the end of the submission of a knowledge artifact, the contributor is asked to search existing knowledge artifacts and indicate, by creating an association between two knowledge artifacts, relationships that only experts can recognize. Our assumption is that, in a community of scientists, a contributor has to understand an existing artifact to be able to recognize an association. When this happens, we are comfortable to claim that knowledge in the existing artifact was shared with this contributor. Consequently, we interpret associations between artifacts contributed by different authors as evidence of knowledge sharing. In the first 2 years, 93 associations were made, 71 between units entered by different authors, for a total of 177 artifacts. This represents that there was knowledge sharing activity that is equivalent to about 40% of the effort in contributing artifacts.

4 Retrieving Knowledge Artifacts

Information retrieval (IR) methods are used to select a subset of records from a collection that are relevant to a given query. Poor IR performance in repository-based KMS is an impediment to knowledge sharing. It may be caused by a) lack of knowledge about the format of the records, b) lack of knowledge of the domain of records, and c) poor query construction. The adoption of the structure presented in Section 3 addresses the lack of knowledge about the format by defining it. It addresses lack of domain knowledge by keeping domain-specific expressions in the fields Applicable Task and Contexts. This allows the construction of a domain-specific taxonomy, which can be used, e.g., to resolve ambiguities in query expansion. Finally, it addresses poor query construction by influencing its length with a recommended cardinality factor: RCF [16].

4.1 Recommended Cardinality Factor: RCF

The distinguishing characteristic of the retrieval method is the RCF. The RCF becomes necessary due to the variable number of terms that may be included in the artifacts. This can be problematic for retrieval quality; therefore, we counteract it by computing the ideal length of a query that leads to a better retrieval performance, which is found from the cardinality of individual artifacts in the repository. In [16] we demonstrated the performance of Equation (1) to compute the RCF from the averaged number of Terms per Artifact (TpA) in a repository,

$$RCF = e^{(0.833 \times \ln(TpA))} \quad . \tag{1}$$

When RCF is used in search, it bounds the comparison between a query and each artifact. The actual comparison between terms may be carried out as described in [16], or via cosine, n-grams, etc. This parameter limits the number of comparison to improve retrieval quality. In other words, if the artifacts have, on average, a large number of terms, submitting a query with few terms will not produce a retrieval of the same quality as a query with a greater number of terms. Next we describe a comparison study with a comparative IR method that does not utilize the RCF.

4.2 Comparison Study: IR vs. RCF

This study utilizes a dataset used in a survey of users of the KMS discussed in [15]. The survey consisted of 6 queries with hypothetical results for the users to score whether they considered them *relevant*, *somewhat relevant* or *not relevant*. Queries and results were designed from actual knowledge artifacts. Each query is created from the Contexts field of an artifact, for example: "modeling, aerosol dispersion, indoor air," and the query result would be the contents of a different artifact. The survey results produced 16 *query result pairs*. For this study, we use only the pairs that were consistently assessed as relevant and somewhat relevant. The users consistently labeled only one query result as *not relevant*, so it was excluded from the analysis.

This study hypothesizes that a search method that adopts the bounding parameter RCF will produce better quality in retrieval than an alternative IR method. We adopt an average performance computed as follows. For each *query result pair*, consisting of a result to a query and a score, we compute the proportion of results that are selected by each method to be in the top *n* results of the retrieved set. We compute *n* from 1 to 10, and then average each of those results, for both the following methods: RCF and IR.

For RCF, each value of the query is compared against each value in the knowledge artifacts, assigning 1 for matching values or 0 otherwise. The results are added until either (a) the number of matches meets or exceeds the RCF or (b) there are no more terms to match in either the query or the knowledge artifact. In (a), a score of 1.0 is assigned; in (b) the score is given by the number of matches divided by the RCF. In other words, the RCF is the number of terms that have to match in order to consider a knowledge artifact to completely match a query.

Table 3. For *n* =1 to 10, the proportion in which a result appeared in the retrieval set

Results rated *relevant* or *somewhat relevant* in survey										
n	1	2	3	4	5	6	7	8	9	10	AVE
IR	0.38	0.62	0.62	0.69	0.77	0.92	1.00	1.00	1.00	1.00	0.80
RCF	0.67	0.80	0.87	0.93	0.93	0.93	0.93	1.00	1.00	1.00	0.91

Or the IR, we used the Indri search engine of the Lemur Toolkit a language modeling IR tool freely available on the web [18]. Indri represents documents as multiple binary feature vectors.

Results and Discussion. Table 3 shows the average and the individual results for each value of *n.* Thus, for instance, when considering the top 5 terms (*n* = 5), the table shows what proportion of the test results were presented by each method (i.e., higher is better). The results are more distinguished at the levels where fewer results considered, i.e., 1 to 5. These are also more strict assessments, making the better values when the RCF is used, more valuable.

A two tailed t-test performed on the paired rankings of the query results for both methods showed that there was a significant difference between the IR and RCF methods at $p < 0.1$ (0.072). Moreover, the RCF provides guidance to users. Most users of KMS draw their mental models of search from their interactions with web search engines. Spink et al. [17] show that the average number of terms in web search queries is around 2.5 terms. Knowing the optimal number of terms to include in a search query can help users create better queries and thus give them a better chance at finding relevant results.

5 Related Work

In this section we leverage discussions from design, human-computer interaction, and KM to argue for the persistence and contribution of structure behind knowledge artifacts. Some broad goals of such structures, in these different disciplines, are to promote communication and comprehension. There appears to be considerable overlap between the goals mentioned in other disciplines and what a knowledge management system attempts.

The pattern concept has received much attention from various design communities (e.g., software engineering, HCI, education), and each have adapted and repurposed it for suiting what they do. At some level, these different design disciplines agree on patterns as a representation or structure to share best practices. Alexander et al. [19], who are attributed for introducing the pattern concept in architecture, argue patterns may help designers and non-designers communicate. Another shared, or implicitly agreed upon property of a pattern, is structuring design guidance or best practices as a *... three part rule, which expresses a relation between a certain context, a problem and a solution*. ([20]; p. 247). In addition to these three elements, like our structure, a pattern description also contains a rationale that argues why it is a good solution in the given context.

The EUREKA system at Xerox is often described a win for KM [21]. Bobrow and Whalen point out that their system emerged by learning how Xerox technicians actually share experiences. The experiences were not presented as knowledge per se, but according to Orr's [22] ethnographic study of technicians, as war-stories. Bobrow and Whalen explain that the experiential knowledge, that is embedded in practice, was being exchanged through such rich stories. Each time the story was narrated it was enhanced or contextualized based on the given situation [22]. Understanding the nature and role played by these war-stories showed Xerox a way to structure the experiential knowledge in the EUREKA system. This structure or tips contained a symptom, cause, test and action.

The proposed structure has similarities with structures used in other disciplines. We assert that the elements of the structure described in Section 3 can be mapped, with little effort, into a problem (i.e., Applicable Task and Contexts) and solution (i.e., Contribution and Rationale). Nonetheless, it is novel in how it is presented here; particularly on the use of two fields with different levels of specificity to help discriminate a problem. When we argue that our structure is generalizable to other similar KM efforts, we are referring to our structure in principle. We encourage readers to customize this structure for their specific audience as it is difficult to argue for a one-size-fits-all approach.

6 Conclusions

This paper discusses a structure to represent and a method to retrieve knowledge artifacts in repository-based KMS. Section 3 describes the concepts of the structure for adoption in multiple domains. The recent implementation of this structure indicates that it is sufficient to capture contents that contributors want to share. The approach recommends the incorporation of a temporal dimension to the structure. Its use suggests that it makes knowledge available for sharing sooner than it would without such a dimension. The ease of interpreting a knowledge artifact encourages contributors to make associations between new and existing artifacts. Those associations convey evidence of knowledge sharing.

Section 4 presents the method to retrieve knowledge artifacts that adopts such a structure. The method is characterized by the use of a parameter, RCF, which determines the number of terms that are supposed to match for a knowledge artifact to be considered relevant to a query. A study demonstrates the superiority of this method when compared to an IR method without it.

An essential part of any KM system is the contributors and users of these knowledge artifacts; most importantly whether the users accept it [2]. When approaching KM from the users', or even use, perspective, questions about comprehension and usability of the structure and the knowledge artifacts should receive high priority. We believe that our proposed structure has the potential to answer some of these questions.

Acknowledgements

First and second authors are supported in part by the U.S. EPA-Science to Achieve Results (STAR) Program and the U.S. Department of Homeland Security Programs, Grant # R83236201. The testing of the search is being conducted under IRB protocol #16449.

References

[1] Holsapple, C.W., Joshi, K.D.: Organizational Knowledge Resources. Decision Support Systems 31(1), 39–54 (2001)

[2] Davenport, T.H., Prusak, L.: Working knowledge: How Organizations Manage what they Know. Harvard Business School Press, Boston (1998)

[3] Weber, R., Aha, D.W., Becerra-Fernandez, I.: Intelligent Lessons Learned Systems. Expert Systems with Applications 20(1), 17–34 (2001)

[4] Marshall, C., Prusak, L., Shpilberg, D.: Financial Risk and the Need for Superior Knowledge Management. California Management Review 38(3), 77–101 (1996)

[5] Szulanski, G.: Exploring Internal Stickiness: Impediments to the Transfer of Best Practice within Firms. Strategic Management Journal 17, 27–44 (special issue, 1996) (Winter)

[6] Atwood, M.E.: Organizational Memory Systems: Challenges for Information Technology. In: 35th Annual Hawaii International Conference on System Sciences. IEEE Press, Los Alamitos (2002)

[7] Disterer, G.: Individual and Social Barriers to Knowledge Transfer. In: 34th Annual Hawaii International Conference on System Sciences. IEEE Press, Los Alamitos (2001)

[8] Weber, R.O.: Addressing Failure Factors in Knowledge Management. Electronic Journal of Knowledge Management 5(3), 333–346 (2007)

[9] Secchi, P., Ciaschi, R., Spence, D.: A Concept for an ESA Lessons Learned System. In: Secchi, P. (ed.) Alerts and LL: An Effective way to prevent failures and problems (Technical Report WPP-167), ESTEC, Noordwijk, pp. 57–61 (1999)

[10] Aha, D.W., Weber, R. (eds.): Intelligent Lessons Learned Systems: Papers from the AAAI Workshop (Technical Report WS-00-03). AAAI Press, Menlo Park (2000)

[11] Ashley, K.D.: Applying Textual Case-based Reasoning and Information Extraction in Lessons Learned Systems. In: Aha, D.W., Weber, R. (eds.) Intelligent Lessons Learned Systems: Papers from the 2000 Workshop (Technical Report WS-00-03). AAAI Press, Menlo Park (2000)

[12] Everett, J.O., Bobrow, D.G.: Resolving Redundancy: A Recurring Problem in a Lessons Learned System. In: Aha, D.W., Weber, R. (eds.) Intelligent Lessons Learned Systems: Papers from the AAAI Workshop (Technical Report WS-00-03), pp. 12–16. AAAI Press, Menlo Park (2000)

[13] Weber, R.O., Aha, D.W.: Intelligent Delivery of Military Lessons Learned. Decision Support Systems 34(3), 287–304 (2003)

[14] Levin, B.: English Verb Classes and Alternations: A Preliminary Investigation. University of Chicago Press, Chicago (1993)

[15] Weber, R.O., Morelli, M.L., Atwood, M.E., Proctor, J.M.: Designing a Knowledge Management Approach for the CAMRA Community of Science. In: Reimer, U., Karagiannis, D. (eds.) PAKM 2006. LNCS (LNAI), vol. 4333, pp. 315–325. Springer, Heidelberg (2006)

[16] Weber, R.O., Gunawardena, S., MacDonald, C.: Horizontal Case Representation. In: Althoff, K.-D., Bergmann, R., Minor, M., Hanft, A. (eds.) ECCBR 2008. LNCS (LNAI), vol. 5239, pp. 548–561. Springer, Heidelberg (2008)

[17] Spink, A., Wolfram, D., Jansen, M., Saracevic, T.: Searching the Web: The Public and their Queries. Journal of the American Society for Information Science and Technology 52, 226–234 (2001)

[18] Lemur Project. Lemur Language Modeling Toolkit. University of Massachusetts and Carnegie Mellon University, `http://www.lemurproject.org/`

[19] Alexander, C., Ishikawa, S., Silverstein, M., Jacobson, M., Fiksdahl-King, I., Angel, S.: A Pattern Language: Towns, Buildings, Construction. Oxford University Press, New York (1977)
[20] Alexander, C.: The Timeless Way of Building. Oxford University Press, New York (1979)
[21] Bobrow, D.G., Whalen, J.: Community Knowledge Sharing in Practice: The Eureka Story. Journal of the Society for Organizational Learning 4(2), 47–59 (2002)
[22] Orr, J.E.: Talking About Machines: An Ethnography of a Modern Job. ILR Press, Ithaca (1996)