Designing Multifunctional Knowledge Management Systems

Rosina Weber and Sidath Gunawardena
College of Information Science & Technology, Philadelphia, PA 19104
The iSchool at Drexel University
{rw37, sg349}@drexel.edu

Abstract

Expert locators and repository-based knowledge management systems (KMS) are different architectures proposed to perform different kinds of knowledge management functions. While expert locators can recommend an expert to perform a task, repository-based KMS can share a learned strategy to solve a given problem. In this paper, we describe a framework to develop KMS that can perform multifunctional tasks, usually performed by repository-based and expert locators separately, in one single architecture. The proposed framework follows principles from knowledge engineering and from the management literature on how to prevent failure in KMS. We illustrate the proposed framework with an implementation to support a community of science.

1. Introduction

Among the widely discussed categories of knowledge management systems (KMS) are repository-based and expert locators [8]. Repository-based KMS are typically adopted in support of knowledge sharing and leveraging. The need for such systems is directly proportional to the number of members of an organization. The standard architecture of those systems is a database of knowledge artifacts [28]. Expert locator systems address needs to staff tasks with experts; these systems are typically needed in large organizations that span multiple geographic regions. Like repository-based systems, a well maintained database that stores experts with their skills and competencies is the standard architecture.

Both repository-based and expert locators are typically beneficial in large organizations, but when both are considered, the current view is to build two different systems. The framework we introduce proposes the design of one single architecture to perform the role of both these types of systems.

Although both types of systems under discussion are widely known and used, they still have limitations. For example, repository-based systems may have artifacts that are difficult to interpret [28]; while expert locator systems depend on self-assessment from the experts to keep it to date [6]. The multifunctional framework we describe in this paper allows the implementation of both knowledge sharing and expert location functions, and yet provides distinguishing characteristics that can mitigate the limitations of both kinds of KMS.

The framework we propose in this paper relies on the notion that a KMS to support both knowledge sharing and expert recommendation should collect knowledge artifacts from experts and use them as evidence of expert's competencies.

The distinguishing characteristics exhibited in our framework originate from the multidisciplinary assessment of the knowledge management (KM) literature. Our framework adopts principles suggested by literature published in management science (e.g., [4]), information science and systems (e.g., [1]), and computer science (e.g., [2]) conferences and journals. Some of the synergies between KM and knowledge engineering are described by the leading work of Liebowitz [17].

Although current publications in repository-based and expert locator KMS consider them individually (e.g., [6][20]), our assumption is that a community of knowledge workers who contribute knowledge artifacts is also a community of experts. Furthermore, organizations in need of these systems usually share similar characteristics such as large number of members that are geographically dispersed. Our recommendation is thus to adopt a multidisciplinary approach to attain adequate technological support for KM.

The remainder of the paper is organized as follows. Section 2 presents some background literature and related work. Section 3 describes the elements of the framework. Section 4 describes an implementation followed by discussion. Section 5 concludes and presents future work.

2. Background Literature and Related Work

We start by presenting a brief overview of repository-based and expert locator KMS. Finally, we introduce the field of knowledge engineering, citing types of knowledge the field describes.

2.1. Repository-based KMS

One of the most desired goals in KM is to make organizational knowledge that is individual. Therefore, organizations strive to externalize the hard earned knowledge gained by their employees. One approach large organizations utilize to achieve this goal is the creation of repository-based KMS [4]. Implementations of these are very common in private and public organizations [28].

The limitations of repository-based KMS are usually concerned with aspects of knowledge sharing and distribution. For example, Szulanski [24] lists four characteristics of these systems that he considers obstacles to knowledge sharing. They are the lack of explicit description of how the strategy is adequate for the problem so that users can extrapolate from it; lack of justification of why it is proven to be good (*unproveness*); lack of capacity of users to absorb knowledge, and to retain it.

Another class of limitations relate to the distribution of knowledge retained in these systems. Some of the reasons for problems in distribution given in [29] are: distribution is not in the same context as organizational processes that can benefit from artifacts; users have the burden to take initiative to use standalone tools; users need to have proper training to use the tools.

2.2. Expert Locator KMS

The role of expert locator KMS is to capture knowledge of skills and competencies that are typically missed by the standard human resource department [6]. In this case the implemented system uses many data sources to build a complete profile of each expert. More generally, an expert locator KMS needs to be able to identify, classify, validate, and monitor experts [17].

Maybury [19] goes beyond the basic notion of storing skills and competencies and mines all readily available sources to leverage information about experts. The resulting enhanced representation improves the expert location process.

Manually populated skills databases are rare within organizations and even where they exist are of little use as they are not uniformly populated, or the information within them is out of date [19]. In addition they are labor intensive to build, require that experts contribute and maintain skills lists, and the information they contain is on a much more general level than the typical queries made to the systems [33].

Maybury [19] designed an expert finder system that scans all the documents produced by an individual and all documents that refer to that individual. That individual is then associated with the topics located within the documents. The more documents on a topic that are linked to that individual, the higher the expertise rating of that individual on that topic. The types of documents can be given different weights; for example a resumé would be given a higher weight than a newsletter [19]. There are difficulties when attempting to integrate such wide sources of data. One solution is to use correspondence tables that make the semantic connections between fields in differently structured databases [6].

Using self-assessment to build the repository of competencies is convenient, but prone to subjectivity and hard to normalize across the organization. Furthermore, developing the knowledge taxonomies required to augment such a system can prove to be a challenge [6].

The Connection Machine [12] from PricewaterhouseCoopers LLP automates the location of experts associating with it the questions directed to them by the organization's members. By monitoring the question and answer between two members, it keeps the profile of members' interest and expertise constantly updated.

2.3. Knowledge Engineering

Knowledge engineering is the field of study initially dedicated to the development of expert systems. Historically, expert systems have been the first approach chosen by artificial intelligence (AI) when in the the pursuit of a method that would allow a system to produce a solution to a problem. The study of expert systems has grown and expanded, motivating new fields like machine learning, whose science has recently opened the path for the study of knowledge discovery in databases, or data mining. The study of knowledge engineering has kept its focus on manipulating representations of knowledge, from production rules to frames, logic-based methods and heuristics acquired by experts. The limitations of the initial systems, such as lack of commonsense knowledge, have given rise to the study of ontologies, also called knowledge bases. Simultaneously motivated as an alternative to expert systems, as an approach to machine learning, and as a model of memory, the case-based reasoning (CBR) methodology is another field that also deals with knowledge engineering issues. The main steps in knowledge engineering are knowledge acquisition, understanding, and representation. These issues are common to all reasoning methodologies studied under AI. Their study has, among other things, revealed a number of different types of knowledge, e.g., procedural, declarative, heuristic.

Procedural knowledge relates to how a problem may be solved. Declarative knowledge relates to what is known about a problem. Heuristic knowledge is knowledge usually discovered through experience that has specific applicability. There are also two more abstract forms of knowledge. Structural knowledge captures an overall problem and how it is related to other entities. Meta-knowledge is knowledge about the knowledge associated with the problem [11].

3. Multifunctional Framework

A multifunctional framework for designing KMS adopts a single architecture and performs KM functions that originally require multiple architectures. The core of the design lies on two tables: one of structured knowledge artifacts and one of experts (Figure 1). These tables are associated by the authorship of knowledge artifacts. This multifunctional framework aims at performing functions associated with knowledge sharing and expert location.

In order to present the framework, we describe the assumptions and definitions we adopt. Then, we detail the principles we follow from literature in different fields, the characteristics imposed by the framework, and finally its capabilities.

3.1. Assumptions and Definitions

This framework collects knowledge artifacts from experts and assumes that these artifacts are evidence of their expertise.

Knowledge artifacts [14] are the units a repository-based KMS retains. Their purpose is to be shared with members of the organization in which the KMS targets. The purpose of knowledge sharing is to inform an organization's members of relevant knowledge sources, allowing them to perform tasks better than they would by relying solely on their own knowledge. If there is anyone in the organization who has learned how to perform this task in an improved way, then this strategy is to be shared, so everyone can reuse it. When this happens, the intellectual resources from the organization are used to improve its overall result, meeting KM goals. Strategies to perform tasks in an improved fashion are thus the core of organizational knowledge to be retained in knowledge artifacts.

3.2. Principles

Our framework follows principles from researchers in knowledge engineering and in KM. Our analysis of literature discussing how to explain and prevent failure in KMS is discussed with detail in [27].

Our first principle is the integration of humans, processes, and technology, as recommended in [1]. We note that the integration of humans require understanding human roles and the inclusion of humans implementing the KM approach. The processes of each target community also have to be understood for effective integration. Technologies need to be in harmony with the other elements, and they have to be considered by the technologies.

KMS are instruments of collaboration in the extent in which they are meant to promote sharing. It is through collaboration that humans share and learn

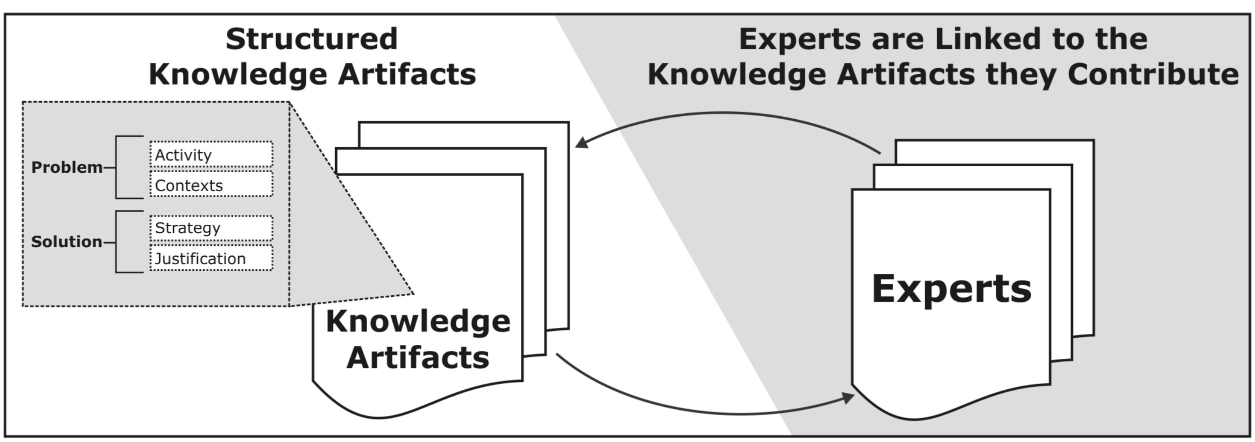

Figure 1. Framework associates knowledge artifacts with experts

knowledge in groups. Transparency is considered cricial to promote collaboration [23]. Therefore, we adopt knowledge artifacts that shall not intervene in a clear and direct communication of contents.

The design of knowledge artifacts include the processes where they are applicable. To guarantee transparency, they are required to be short and focused [28]. This makes it easy for system's users to recognize an artifact's usefulness and applicability.

Another important principle we consider refers to the managerial responsibilities listed by Marshall, Prusak and Shpilberg in [18]. Our framework attempts to guide contributors during collection of knowledge artifacts, taking responsibility for creating adequate infrastructure for its verification, validation, and reuse.

Although our framework focuses on technical aspects of designing KMS, we have taken into account the potential difficulty to attract contributions, which is a cultural aspect believed to be responsible for failure in KMS. Our principle is based on Disterer's [10] argument that contributors may not recognize value for them or others when sharing knowledge. In response to this, we take as a principle to offer something of value to contributors. We will discuss the creation of useful reports as a way to offer a valuable compensation under Section 3.4.

Szulanski [24] listed some impediments to knowledge sharing that are worthy of note (see Section 2.1). Among those, we particularly consider the importance of *unproveness* [24]. For this reason, we utilize a field named *justification* and make it an important requirement in the representation of knowledge artifacts. The role of this field is for the contributor to provide evidence to validate any claims made in the artifact.

The actual contents in the justification field vary in different domains. In a community of science, where members are scientists, the contents of the justification field are results of validation studies that follow the scientific method. Thus, it solves unprovenness by providing scientific validity to the artifact. Using the scientific method to validate knowledge may not be convenient in other communities. Previous work on lessons-learned systems [29] adopted sub-fields to support a learned lesson by explaining how the lesson was learned and providing its origin. One example was that a lesson was learned from someone's advice. Therefore, the strength of the justification will vary depending on the domain environment. We consider enough to require from the contributor to provide a justification. It is though no guarantee that any future member will trust the justification as valid proof of the usefulness of the artifact.

We also consider *absorbency* as an aspect that we can address within our framework. *Absorbency* [24] refers to the capacity a member has to understand the applicability of a knowledge artifact. Thus, our design requires the inclusion of the context in which the knowledge in the artifact should be reused. We do this by including in the representation of knowledge artifacts fields for describing the target task where the strategy applies and additional contexts that can help describe the conditions for reuse of the artifact.

One of the most distinguishing principles in our framework is our concern with technology. The limitation of using traditional databases has been discussed to reside in that they can only manipulate data and information [21]. Thus, we take as a principle to adopt lessons from over fifty years of studies in knowledge engineering as suggested by Liebowitz in [17]. The main technological principle of our framework is to represent knowledge artifacts after a representation that can be used for reasoning. That is, the knowledge artifact is represented in a way that meets conditions of a knowledge engineering formalism of knowledge representation. More specifically, we recommend a representation that allows a computational system to make decisions and solve problems.

The methodology we use is CBR [3][15]. The reasons to adopt this particular methodology in KMS is discussed in the literature by many authors, mostly indicating the intuitive notion shared with KM of relying on previous episodes to solve new problems [1][5][26][28].

A well-known technological limitation refers to using natural language to interface system with users. Humans use natural language to submit knowledge artifacts when interacting with a computer tool that is limited in its ability to understand the captured artifact. Therefore, if the submitted artifact is not correct, the tool may not be able to determine that and guide the user. The verification of knowledge artifacts requires interpretation, what is neither precise nor easy. For this reason, we recommend that human knowledge facilitators be available to review knowledge artifacts and effectively guide contributors.

Reviews may focus on contents and on compliance to design requirements. Knowledge facilitators are the guardians of design requirements, they have to ensure that each artifact conforms to the guidelines for submission to the repository and to appraise the contributor to any areas that fail to do so. An artifact may travel back and forth several times between reviewer and contributor before it is approved. The specific way to implement reviewing depends on the target community. Artifacts may be kept from appearing in the repository until they are approved when community members are abundant and whose levels of expertise may vary. In communities of science, for example, artifacts can appear in the repository since its submission with a label indicating whether review is pending or complete. Reviewing contents is a more delicate matter and should be kept as an option of the reviewer to prevent social and evaluation related concerns.

In order to use artifacts as evidence of a contributor's expertise and also to meet conditions of interpretability, usefulness, and absorbency, this framework imposes requirements on the representation of knowledge artifacts. We now describe how these principles convert into characteristics of the framework.

3.3. Characteristics

The most distinguishing characteristics of a KM approach that adopts the multifuntional framework we propose are the representation of knowledge artifacts and the utilization of knowledge facilitators. While knowledge facilitators represent a way to integrate humans from the perspective of the KM approach, the fields of knowledge artifacts and its associations embody multiple types of knowledge (see Section 2.3).

To prevent contributors from submitting knowledge artifacts that are not transparent or that are difficult to interpret, our framework requires all submissions of knowledge artifacts to be in a structured form. This structure has four required fields, and other optional fields. As previously stated, the core of a knowledge artifact is a strategy that can be used to solve a problem. Of the required fields, two fields describe the problem, and the two others describe the strategy and its justification (Figure 1); the latter addresses *unproveness* [24]. The first two fields break down the problem into an activity and contextual elements. An activity is an organizational process described at a level of specificity considered by contributors as adequate to be shared. The contextual field provides conditions in which the strategy may be reused.

The representational structure represents an imposition from the organization on how knowledge artifacts are to be shared, thus providing guidance to contributors. This simple set of four fields entails both declarative and procedural knowledge.

The infrastructure for knowledge collection is additionally provided by a reviewing process that assesses the quality of submitted artifacts. Submitted knowledge artifacts can be evaluated on multiple dimensions. The formatting of the fields is the first dimension. For example, to enforce the notion of integration to processes, we recommend contributors inform the applicable activity. Another dimension is length. When we recommend artifacts to be direct and focused to guarantee easy interpretability, it is crucial that we enforce limited length. Formatting and length are evaluated by knowledge facilitators. The reviewing process can also be used for technical evaluation where someone in a supervising position may be required to agree with the technical contents of the artifact.

In addition to the four core fields that embed procedural and declarative knowledge, additional optional fields are included. The most important is when contributors are requested to associate their contributions to existing contributions already retained in the KMS. This imposes an extra step in the submission of knowledge artifacts and in their review. On the other hand, it has the ability to demonstrate how community members envision the integration of their knowledge. Besides, we recommend contributors indicate only associations that are not too superficial, what demands that they understand artifacts they find to be associated with theirs. This can be claimed as evidence that knowledge was absorbed, even if temporarily. The other benefit of these associations is that they capture an additional form of knowledge–structural.

3.4. Capabilities

Our framework provides a number of capabilities to benefit its users and contributors. They include typical knowledge sharing functionalities present in repository-based KMS, the ones typical in expert locator systems, and also reporting capabilities.

Knowledge sharing capabilities. Whenever a new knowledge artifact is submitted, users are pushed emails from the system triggered in two ways. First, contributors are asked to list names of other members they know that may be potentially interested in learning about the new knowledge artifact (an example of an optional field in knowledge artifacts). Second, all users can customize lists of topics of interest so they can be notified when a knowledge artifact pertaining to such topics is submitted.

Pull methods are also available; users can browse or search for knowledge artifacts using multiple search strategies. The distinguishing characteristic of our framework is that it allows the adoption of an applicability-oriented search given that knowledge artifacts identify activities where they are applicable. This is a benefit of adopting the CBR methodology when designing a representation for knowledge artifacts.

Expert locator capabilities. The expert locator capabilities of the framework are enabled by the knowledge artifacts submitted by contributors. The

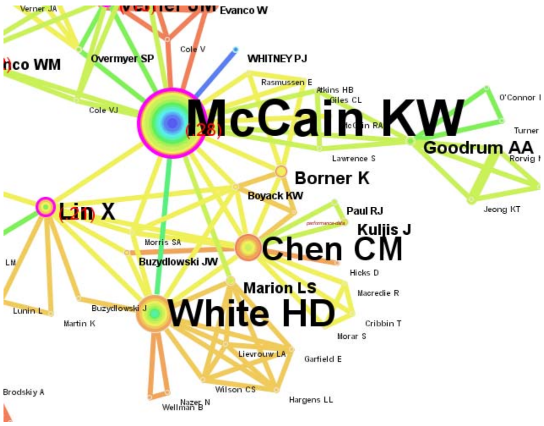

Figure 2. Visualization of experts in a field

quality of expert location functions is driven by the structure adopted for knowledge artifacts. Contributors become associated with the activities and the contextual elements they inform when they submit a knowledge artifact. The association between contributors and knowledge artifacts can be reliably used as an indication of a contributor's expertise in scientific fields. This is because when a researcher submits a scientific contribution, then it is valid to assume this researcher is an expert in the domain of the contribution. In traditional working organizations, such assumption is not as simple. In those instances, we recommend additional constraints to recognize a contributor as an expert because the only valid assumption is that a contributor has an interest in the field of the contribution. Additional constraints may refer to other information about the contributor or the contributed artifacts.

An important benefit of the structure used to represent knowledge artifacts is that it allows the assessment of the similarity between two experts with respect to their areas of expertise. This stems from the use of CBR as underlying methodology. In order to be able to reason with CBR, knowledge artifacts are designed so they can be compared through a measure of similarity. The similarities between experts allow the construction of visualizations to describe experts in maps where experts are distributed along regions of expertise areas. Figure 2 shows a partial visualization of a collaboration map of faculty members from Drexel iSchool based on their publications. This visualization was done with CiteSpace [7]. A line between two authors means that they have published as coauthors. The size of a node represents the citations the author has received over different years. This visualization reveals a type of social network in an organization. Similarly, the same method can be applied to visualize the structure and evaluation of a specialty or a discipline.

Analogously, bigger names can demonstrate experts who submitted more artifacts in a given topic while links between experts can indicate common interests.

Reporting capabilities. As knowledge artifacts are represented with such a problem-solution pair structure (where problem is where the strategy is applicable and the solution is the strategy and its justification) and are associated to other artifacts, the KMS retains not simply artifacts, but knowledge in the sense used in knowledge engineering. Knowledge engineering dictates different types of knowledge, and our framework proposes representations for declarative, procedural, and structural knowledge. The usefulness of such representations permits the manipulation of this knowledge allowing the composition of reports, going beyond knowledge sharing and expert location. Different communities have different needs for reports. Static organizations would benefit from the automated creation of reports for a small portion of their processes. Nevertheless, in inquiring organizations, it is expected that most work delivered by their members is innovative, and thus they have a place in a KMS. Therefore, in the case of inquiring organizations, the fields describing a knowledge artifact can be used as elements for the automated generation of reports. The automated reporting capability is based on the use of templates. Each report is structured with a sequence of templates. Each template is to be selected depending on the contents of the fields of knowledge artifacts. The combination of the templates and the contents form the resulting text in the reports [30].

4. Implementation and Discussion

In this section we illustrate the primary implementation of the framework. The numbers we present reflect the usage of the first four months of the system. In these first four months, only knowledge sharing capabilities were operational. Additional aspects of this implementation such as its design process and particularities of its target community are discussed in [31]. The discussion sub-section proposes potential impacts of the framework and interpretations of its implementation.

4.1. Implementation

Upon implementation the system had 38 users spread in seven different geographic locations. In the first four months since launched, contributors had submitted sixty-four knowledge artifacts (Figure 3).

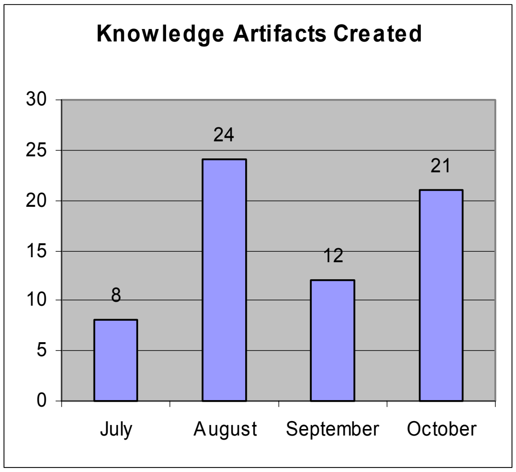

Figure 3. Number of knowledge artifacts during the first 6 months

Of the 64 knowledge artifacts, 26 (40%) were associated with a different artifact. 22% of all associations made linked a knowledge artifact with another contributed by a different author.

Figure 4 depicts system usage, in terms of logins over the first four months. For each month in the period August-October, at least 40% of all users at each of the seven locations logged in to the system at least once. Across that three month period, 87% of all users logged in at least once, and 40% of all users logged in at least 10 times.

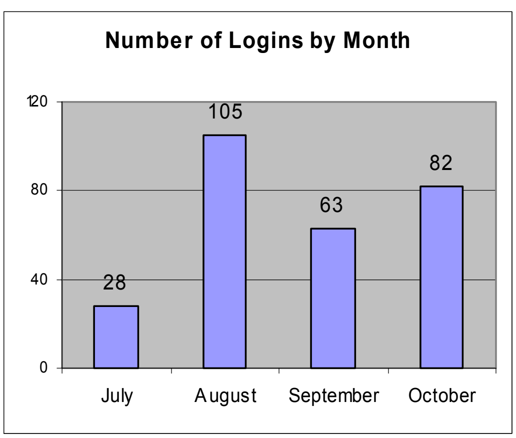

Figure 4. Number of logins

4.2. Discussion

The description of contextual elements with activities (or processes) provides a representation of expertise that can help label contributors based on the artifacts they contribute. As previously stated, this may not be true in any domain; but is a safe assumption among research scientists.

By connecting knowledge artifacts to their contributors, we relieve experts from the burden of describing or assessing their competencies for building expert location capabilities. This is also a way to impose a restriction on expertise because it implies that active experts are more desired than old ones. However, if a target organization opposes this assumption, the use of knowledge artifacts describing previous work can take priority. Besides, a level of quality can be associated to the justification field of a knowledge artifact, creating a metric for prioritizing experts.

Although literature in KMS have discussed the impossibility of successful codification of knowledge (e.g., [9], [22]) the field of Knowledge Engineering has, since the early fifties, codified different knowledge formalisms allowing computer systems to make decisions for problem solving. Regardless of the scientific relevance and significance of the discussion of the depths of human knowledge, for the purposes of organizational knowledge in support of decision-making, the field of Knowledge Engineering provides sufficient expertise and validation to guide KMS to represent knowledge artifacts.

Nevertheless, there are indeed some formalisms studied by knowledge engineering methodologies that may not meet quality requirements such as naturalness. For this reason, we assume that for knowledge sharing purposes, the representation adopted for knowledge artifacts should be natural to facilitate easy interpretation by human users.

One major challenge faced by KMS is the process of capturing knowledge artifacts. Consequently, the representation of knowledge artifacts should be one that facilitates collection as well. This is one of the technical strategies that can be adopted in response to cultural challenges. Another example of addressing culture with technology is the preparation of reports. There are classes of organizations, particularly ones that can be classified as inquiring organizations, whose members can benefit from the automated generation of reports. Among some examples are organizations that produce creative work and scientific research.

The proposed framework relies on the use of a structured representation that separates activities from its contextual elements. Contextual elements may be further broken down into domain dependent or independent. The reason to separate domain dependent from independent elements in the representation originate from the possibility of integrating external sources of information that can enhance the representation. For example, domain independent contexts may be translated or expanded for purposes of automation of report generation or search by integrating their understanding to external ontologies such as Wordnet [32].

An important issue in KMS is their evaluation. When stating KM goals, it is hard to measure the level in which they are attained given the subjectivity of intellectual work. Besides, traditional evaluation methods suggest asking the users, but this would add to their burden and thus it would not be practical to be implemented on a permanent basis. For this reason, we suggest that associations between knowledge artifacts be adopted as a measure of knowledge sharing. Note that we recommend contributors to make associations that are not superficial, i.e., of the kind that a keyword based function would identify. By tuning their focus to find less superficial associations, we expect to demonstrate that, whenever a contributor finds a valid association and is able to describe it, then the assumption that the knowledge embedded in the artifact they found to associate was indeed shared with this contributor. In terms of the four first months of implementation, we argue that 22% of the activity of contributors was enhanced with knowledge sharing. As presented in Section 4.1, 22% of the knowledge artifacts submitted were associated with an existing artifact that had been submitted by another contributing member.

It is important to note that the implementation we discussed is meant to illustrate rather than validate the proposed design. Nevertheless, the ratio of associations to the number of submitted artifacts can be viewed as promising indications of knowledge sharing.

5. Conclusions and Future Work

This paper describes how to use principles from knowledge engineering in combination with studies in KMS to design one single architecture that can provide capabilities of two different types of KMS. Interestingly, large organizations that are likely to benefit from repository-based KMS are also adequate targets to benefit from expert locators.

Unless this architecture creates unforeseen impediments, it is expected to produce results that are at least comparable to traditional repository-based system and expert locator systems. The impositions of the knowledge artifact we propose conform with the notions of managerial responsibilities in KM [18], whose purpose is to prevent KMS to fail. Thus, the proposed architecture should not create new impediments.

Although a thorough subject evaluation of this framework is still future work, the description in this paper shows how different functionalities can be designed. The discipline dedicated to study the validity of design is design science [13][25]. Thus, we plan to demonstrate the vality of the proposed design in its context.

The current implementation to support a community of science suggests the effort of knowledge facilitators is reasonable because they are required to interact with contributors only sporadically. The reason is that it takes time for scientists to submit knowledge artifacts because producing scientific knowledge is time consuming. Consequently, it is premature to infer on the results of this design and approach to different communities due to the risk that the workload requires training of several knowledge facilitators to meet the demand.

Acknowledgements

The authors would like to acknowledge the helpful comments from the reviewers, the support of members of the CAMRA community, and the contributions of Nick M. Sillik, Kate Ericcson, Kelvin Blasse and Nirali Vora. This work is 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.

Figure 2 has been done as courtesy by Chaomei Chen, Drexel University, using CiteSpace (http://cluster.cis.drexel.edu/~cchen/citespace/).

References

[1] Abecker, A., Decker, S. and Maurer, F. "Organizational Memory and Knowledge Management", Guest editorial, Information Systems Frontiers, 2(3-4): pp.251-252, 2000.

[2] Aha, D.W., Becerra-Fernandez, I., Maurer F., and Muñoz-Avila, H. Eds. "Exploring Synergies of Knowledge Management and Case-Based Reasoning" Papers from the AAAI 1999 Workshop (Tech. Rep. WS-99-10), Menlo Park, CA: AAAI Press, 1999.

[3] Aha, D.W., Marling, C., and Watson, I. "Case-based reasoning commentaries: Introduction," *Knowledge Engineering Review*, 20(3): pp. 201-202, 2005.

[4] Alavi, M. and Leidner, D. E. "Knowledge management and knowledge management systems: Conceptual foundations and research issues," *MIS Quarterly*, 25(1): pp. 107-136, 2001.

[5] Althoff, K.-D., and Weber, R. O. "Knowledge Management in Case-Based Reasoning," *Knowledge Engineering Re*view, 20(3): pp. 305–310, 2005.

[6] Becerra-Fernandez, I.. "The role of artificial intelligence technologies in the implementation of People-Finder

knowledge management systems," *Knowledge-Based Systems*, 13(5): pp. 315-320, 2000.

[7] Chen, C. "CiteSpace II: Detecting and visualizing emerging trends and transient patterns in scientific literature," *Journal of the American Society for Information Science and Technology*, 57(3): pp.359-377, 2006.

[8] Davenport, T.H. and Prusak, L. *Working knowledge: How organizations manage what they know*. Boston, MA: Harvard Business School Press, 1998.

[9] David, P.A. and Foray D. "An introduction to the economy of the knowledge society," *International Social Science Journal*. 54(1): pp. 9-23, 2002.

[10] Disterer, G. "Individual and Social Barriers to Knowledge Transfer," in *Proceedings of the 34th Annual Hawaii International Conference on System Sciences, Los Alamito*s, CA:IEEE Press, 2001.

[11] Durkin, J. *Expert Systems: design and development*. Prentice-Hall, Inc., 1994, New Jersey.

[12] Göker, M.; Thompson, C., Arajärvi, S. and Hua, K. "The PwC Connection Machine: An Adaptive Expertise Provider," In T.R. Roth-Berghofer et al. Eds.: ECCBR 2006, LNAI 4106, pp. 549 – 563, Springer-Verlag Berlin Heidelberg, 2006.

[13] Hevner, S., March, J., Park, and Ram S. "Design Science Research in Information Systems," *Management Information Systems Quarterly*, 28(1): pp. 75-105, 2001.

[14] Holsapple, C. W. and Joshi, K. D. *Decision Support Systems* 31, pp. 39-54, 2001.

[15] Kolodner, J. *Case-Based Reasoning*. Los Altos, CA: Morgan Kaufmann, 1993.

[16] Lam, W. and Chua, A. "The mismanagement of knowledge management," *Aslib Proceedings* 57(5): pp. 424-433, 2005.

[17] Liebowitz, J. K*nowledge Management: Learning from Knowledge Engineering*, CRC Press, 2001.

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

[19] Maybury, M. T. "Knowledge on demand: Knowledge and expert discovery," *Journal of Universal Computer Science* 8(5): pp. 491-505, 2002.

[20] Maybury, M., R. D'Amore, and D. House. "Awareness of organizational expertise," *International Journal of Human-Computer Interaction,* 14(2): pp. 199-217, 2002.

[21] Nissen, M. E. "An Extended Model of Knowledge-Flow Dynamics," *Communications of the Association for Information Systems*, vol. 8, pp. 251-266, 2002.

[22] Spender, J.C. "Organizational knowledge, learning and memory: Three concepts in search of a theory," *Journal of Organizational Change Management*, 9(1): pp. 63-78, 1996.

[23] Stahl, G. *Group cognition: Computer support for building collaborative knowledge*. Cambridge, MA: MIT Press, 2006.

[24] Szulanski, G. "Exploring Internal Stickiness: Impediments to the Transfer of Best Practice within Firms," *Strategic Management Journal,* vol. 17, Special Issue Winter, pp. 27-44, 1996.

[25] Walls, J., Widmeyer, G., and El Sawy, O. "Building and Information System Design Theory for Vigilant EIS," *Information Systems Research*, 3(1), 1992.

[26] Watson, I. "Knowledge management and case-based reasoning: a perfect match?" in *Proceedings of the Fourteenth Annual Conference of the International Florida Artificial Intelligence Research Society*, I. Russel and J. Kolen, Eds. Menlo Park, CA: AAAI Press, pp. 118-122, 2001.

[27] Weber, R. O. "Addressing Failure Factors in Knowledge Management," *Electronic Journal of Knowledge Management*, 5(3): pp. 333-346. Online: http://www.ejkm.com/volume-5/v5-i3/Weber.pdf

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

[29] Weber, R., and Aha, D. W. "Intelligent Delivery of Military Lessons Learned," *Decision Support Systems*, 34(3): pp. 287–304, 2003.

[30] Weber, R. O., Gunawardena, S. and Proctor, J. M. "Generating Reports from Case-Based Knowledge Artifacts," in *Proceedings of the Twentieth International Florida Artificial Intelligence Research Society,* D. Wilson and G. Sutcliffe, Eds. Menlo Park, CA: AAAI Press, 2007.

[31] Weber, R. O., Morelli, M. L., Atwood, M. E. and Proctor, J. M. "Designing a Knowledge Management Approach for the CAMRA Community of Science," in *Proceedings of the Practical Aspects of Knowledge Management,* U. Reimer and D. Karagiannis, Eds.: PAKM 2006, LNAI 4333, pp. 315–325. Springer-Verlag Berlin Heidelberg, 2006.

[32] WordNet: An Electronic Lexical Database, MIT Press, 1998.

[33] Yimam-Seid, D. and A. Kobsa. "Expert-finding systems for organizations: Problem and domain analysis and the DEMOIR approach," J*ournal of Organizational Computing and Electronic Commerce*, 13(1): pp. 1-24, 2003.