Hypertext and Knowledge Management

Francisco J. Ricardo
Harvard Graduate School of Education
201 Longfellow Hall
1 Appian Way
Cambridge, MA 02138, USA
Tel: 1-617-233-1072
E-mail: f_ricardo@post.harvard.edu

ABSTRACT
This paper is a functional survey of knowledge management systems and characteristics from the standpoint of the contribution and relevance of hypertext to this discipline. There is the description of a typical KM architecture as well as some of the current KM and KM-like systems deployed in production at large corporations. This discussion will introduce the perceptions of KM and then emphasize the role of hypertext systems in tackling problems in processing distributed and collaborative knowledge. Although at the moment, hypertext is not seen as an architectural component of KM systems, its potential as an epistemic aid presents opportunities. Finally, I will show the appropriateness of hypertext research to KM development.

KEYWORDS: Hypertext, representation, rhetoric.

This paper is a functional survey of knowledge management systems and characteristics from the standpoint of the contribution and relevance of hypertext to this discipline. Knowledge management (KM) is a relatively recent business practice in which digital content in many forms and formats is brought together under an integrated architecture that permits the underlying semantic data in the corpus to be traversed and used as an aid to strategic understanding and decision making. KM is a growth industry[1] primarily aimed at serving business practices, having originated in the business world as a method for unifying the vast amounts of information generated from meetings, proposals, presentations, analytic papers, audits, engagements, clients and target profiles, training materials, CV's, contracts and agreements, etc. When considered in the context of distributed field offices and practices of a global firm, the challenges of finding, contextualizing, and utilizing this disparate and unstructured data become exceptional.

Situating KM

Because of this, KM is primarily utilized by the large corporation, although the practical problem of navigating and consolidating a large multiform document corpus is likely to be relevant by other enterprises that create and consume distributed knowledge. KM is highly practical and is thus not currently defined as being, nor is associated with the more traditional computer scientific initiatives and research that fall under the abstract umbrella term of "knowledge engineering" or other artificial intelligence work which is much more semantically oriented, often involves natural language processing, and is primarily theoretical or experimental in nature. This is largely due to the fact that business processes and a corporate software tradition promoting business-based

HT'01 8/01 Aarhus, Denmark
® 2001 ACM ISBN 1-59113-420-7/01/0008…$5.00

[1] There is now a Knowledge Management Certification Board (Cf. http://kmcertification.org), a Global Knowledge Economics Council (http://gkec.org/), and even a certification program for Certified Knowledge Manager (CKM) and Certified Knowledge Environment Engineer (CKEE); Cf. http://www.eknowledgecenter.com/certificationco urses/index.htm

decision support strategies currently drive KM. It is not the current central focus of many academic or R&D venues. Being practically derived, KM is more the evolution of and successor to *document management systems* (DMS), whose primary aim has been to administer the creation, indexing, and versioning of documents in distributed environments. Since many DMS architectures predate currently evolving standards like DAV or WebDAV, KM, which is more recent still, is not the focus of any evolving standards outside those built on generic encoding schemes like XML.

I will begin by a brief discussion of why KM is different than other human processes. After describing a typical KM architecture, I will cover of some of the current KM and KM-like systems deployed in production at large corporations, or for sale by commercial software vendors. This discussion will introduce the relationship of KM to predecessor architectures in commercially available DMS. Since, as mentioned, KM is currently developing somewhat outside of research in academic/scientific knowledge engineering, I will illustrate various touchstone points at which a fruitful intersection of both may likely lead to a significant benefit in KM's evolution. Of special significance in this potential bridge is the role of hypertext systems have played in the past in tackling problems in processing distributed and collaborative knowledge. Although at the moment, hypertext is not seen as an architectural component of KM systems, its potential as an epistemic aid presents great opportunities. Finally, I will show the appropriateness of hypertext research to KM development.

KM is a semi-paradigmatic discipline that has been contrasted with more traditional processes (such as business practices) in its emphasis on the interlaced organization of materials and resources, for instance, consider (Processedge.com, 1999) below:

	Feudal Culture	Industrial Culture	Knowledge Culture	Creativity Culture
Organization	Territorial	Hierarchies	Networks	Flow
Focus	Land	Profit	Customer	Innovation
Culture	Domination Control	Control Responsibility	Responsibility Contribution	Contribution Creativity
Key Measure	Quantity	Efficiency	Effectiveness	Quality of Life

Table 1. Comparative Cultures

Here, the KM domain is situated within the notion of a *knowledge culture* that revolves around the *network* as its organizing principle along with *effectiveness* and *contribution* as important values. Emphasis on a socially constructed experience is one reason that KM has developed outside the research venue. Recent work has begun to recognize the practical dimension of creating and fostering knowledge as a social activity. This involves studying the contexts necessary to enable socially constructed knowledge. Wenger (1998) for instance, has posited that computationally enhanced learning, meaning, and identity are overlapping, existing in a relative dynamic of participation and artifact use, and that an important way to approach sustainable organizational learning revolves around communities of practice. This is a computational refinement of Rheingold's notion of the virtual community and Covey's sense of the ideal community (Cf. Hesselbein & Goldsmith for essays on the community of the future; Chaiklin and Lave,1995 for an expanded the notion of learning puts social practice at the center of social, cognitive, and computational activities).

Thinking has thus emerged calling attention to collaborative and participatory learning, particularly as they relate to learning in the corporate environment (Weisbord 1993; Botkin, 1999; Skyrme, 1999; Applehans, Globe & Laugero, 1998). Many perspectives speak to a need for broader data-sharing methodologies than what is designed for small collaborative teams through groupware or other advanced messaging methods, and there is terminology to reflect this new interest, with terms like *distributed mind* (Fisher & Fisher, 1997), the *collaborative enterprise* (Skyrme, 1999), *knowledge communities* (Botkin, 1999), *virtual enterprising* (Savage, 1996). This extends an older (Drucker, 1967) concept of *knowledge organizations*, *knowledge technologies*, and *knowledge societies*. Englebart (1963) tightened this into a notion of the *knowledge workshop* as the place in which knowledge workers do their work.

Academic/scientific examples of knowledge management work are much less socially connected, and include work carried out in labs like KSL, the Experimental Knowledge Systems Laboratory at Stanford University co-directed by Ed Feigenbaum (http://www-ksl.stanford.edu/). KSL focuses on machine-artificial knowledge and reasoning, modeling of physical systems. Likewise, the similarly-named EKSL, the Experimental Knowledge Systems Laboratory of the Department of Computer Science at the University of Massachusetts (http://eksl-www.cs.umass.edu/eksl.html), focuses on knowledge-as-science initiatives: link analysis, predictive systems, autonomous agent development, and machine-based learning. Business programs focus on document management, messaging and collaboration, and search, and all subsets of R&D related to "knowledge" show that KM is limited to business process, not to more abstract endeavors that are variously labeled as knowledge engineering, situation assessment, theory of meaning, language understanding and recognition, plan generation, concept learning, generalized problem solving, simulation generation, causal modeling, or knowledge representation.

On the other hand, it is not entirely accurate to characterize corporate KM as totally outside the scope of scientific computing research advancements. By definition, much of KM involves the need for efficient and intelligent search, and since the typical business corpus is largely comprised of unstructured data, inferential methods become more important to need for locating distributed knowledge efficaciously. Therefore, search subsystems in KM systems and architectures have incorporated many of the advancements of information retrieval (IR) research. This also includes systems not only directly involving search, but also automatic tagging and generation of metadata based on document scans where the metadata is derived from a pre-created metadata space, such as a domain-specific thesaurus, keyword lists, or hierarchical taxonomy tables. The document is located within the space by the application of IR algorithms and the resulting metadata is stored separate from the document, often in a database (examples of this are discussed later). Closely related to this is the role of data mining, discussed later.

Despite the availability of out-of-the-box solutions, deploying KM is largely a job that requires much integration and customization. This is partly because KM is an enterprise-wide initiative and must consolidate the (legacy) data output of each separate process in an organization, and partly because there is no universal set of requirements for KM service. Some corporations want KM for strategic targeting of clients (KM as a business development enhancement), others for planning (KM as a decision support system enhancement), and still others for data consolidation (KM as intelligent digital repository). Thus, beyond generic and somewhat well-understood functions like collaboration and search, KM's primary value lies in the creation of some semantic or knowledge-discovery service that did not exist before. This is why much of the design and deployment of KM depends on manual development involving portals with more customized extended-functionality modules (e.g., cf. http://www.epistemic.com).

In this connection, the typical KM architecture is a synthesis or integration of at least three storage and/or processing subsystems. The grossest view of a KM system involves (1) a document-based **content repository** normally existing as files in a nest of directories within the hosting file system; (2) a **search subsystem** which indexes the document content into a (normally) proprietary format; and (3) a **metadata store** which normally resides within a database. Optional dimensions of a KM system can include a **security layer**, so as to restrict sensitive content within a hierarchy of role-based authenticated users; a **content management application** to enable users to upload content that is automatically entered into and classified within the repository; a notion of an LDAP or Active Directory based subsystem that enables **personalization** of the front end (and normally also assists the security layer). Finally, there is a development dimension, which comprises all the programmatic elements of the KM system as such. This would include both in-house developed code, in whatever flavor or language, as well as commercial applications that offer scalable feature integration with KM architectures. An example of this is an enterprise-scale document management system like Documentum (http://www.documentum.com).

Functional Perceptions of KM

Being multi-faceted and partly a social end product, KM has four major faces. Some KM systems have been developed confronting KM as a search problem (KM as enhanced search engine). Others have accentuated the need for KM systems to manage large text corpora (KM as a document management system). Other systems go further still and see KM as a problem of drawing useful knowledge from the document corpus (KM as a data mining tool). Finally, some systems have focused on fostering interactive collaboration, creating meeting spaces where knowledge workers can formulate and advance knowledge activities (KM as social nexus).

The KM-as-Search Engine View

A KM-based search based subsystem is normally assembled in one of three configurations. It can be integrated into a KM product (e.g., Search Server in Microsoft Site Server), it can exist as a standalone commercial-strength application that is integrated with the repository programmatically (e.g., Verity Information Server), or it can exist as a combination of both, as in a platform that offers web-hosted integration, as in the case of an application service provider like Eidetica (http://www.eidetica.com). Eidetica's search subsystem, *t-find*, is a component of a larger KM-oriented offering that includes *t-repository*, a textual database and indexing system, as well as *t·mining*, an automated document classifier that applies language recognition patterns to determine relations between authors, subjects, publishers, time frames, and other content categories.

With the strength of search mechanisms in KM systems, a popular view is that a KM is really a multi-format digital library combined with a fine-grained search engine. Some of this perception may be due to (a) the abstract nature of KM as a discipline, whose benefits appear to be strategic rather than concrete, except where the ability to find content is concerned; and to (b) the immaturity of the field, which has not yet seen the emergence of full-blown KM applications that impart a pragmatic identity for KM as its own kind of user experience; and most important (c) lack of familiarity with original KM applications, particularly as they emerge from hypertext-based research. Lacking this, the typical end user sees benefits limited primarily to sophisticated-looking search interfaces, and often KM systems are defined principally like this. For example, the Eidetica's largest deployment of its KM product revolves around its *t-find* module, and its primary interface is a search interface for the library catalogue of CWI, the Dutch national *Centre for Mathematics and Computer Science*:

It is not yet obvious to users or designers of KM systems that a key distinction between KM and search engines lies in how KM systems are (ideally) intended to help *drive* the thinking behind any search for content, whereas search engines alone aim essentially to retrieve the documents most lexically significant to a search query without much consideration of the query's larger use. Systems that propound KM as a search solution are weakened by overlooking the cognitive need of KM for large-scale collaborative thinking.

The KM-as-Document-Management View

Doing precedes searching. Knowledge emerges in human activities, but to be captured, processed, and reused, it must eventually be inscribed and managed in digital form. Many large organizations face the problem of astounding numbers of documents gradually archived and less and less accessible. While KM-search engine hybrids were created to address the problem of *access*, using that information is a problem of *processing*. Certain KM systems go beyond the notion of advanced search engines and incorporate more individual features intended to aid in the exploration and understanding of the repository. Systems in this category, like Microsoft Site Server and supplementary systems like eGain Knowledge Gateway (http://www.egain.com/) comprise processing aids and organizing elements like briefs, interface-level features like personalization.

When the organizational repository is vast, the contribution of this type of KM system is obvious because it organizes all material into a content store similar to an advanced digital library with version and security control. But document management, like search, is a separate problem from that of collaborative intelligence. Such systems, for example, do not normally have the ability to draw inferences from the document repository or can assist in the solution of evolving large-scale problems. A greater degree of sophistication is required at the semantic level.

The KM-as-Data-Mining-Tool View

To address the latter limitation, there has been a rise in specialized web mining tools like Clementine that can scan documents, presentations, even meta-content like Web logs and online registrations, applying forward-looking models to predict what users will do (Cf. http://www.spss.com/clementine/). Some, like the IBM Intelligent Miner, are a collection of modules including text analysis tools, feature extraction tools, clustering tools, summarization tools, categorization tools, and even integrated web crawlers. (Cf. e.g., http://www-4.ibm.com/software/data/iminer/index.html)

Occasionally, the distinction between data mining and KM has been blurred, particularly when the need for active collaboration and socially constructed cognition are minimized (e.g., "knowledge management more and more relies upon the ability to select information relevant to our needs efficiently, and the ability to manage emerging structures of knowledge" Chen & Davis, 1999). Some better-known KM products, such as Fulcrum KnowledgeServer, Knowledge Manager Workstation, and SearchServer (http://www.hummingbird.com/products/dkm/km/fulcrum/), are commercial strength data extraction tools with search capabilities, agent technology and some resource management, extensible via Microsoft COM objects. Another approach, like the Eclipse language in the Inference Engine (Cf. http://www.haley.com/) has been to provide a rule-based toolkit from which a knowledge manager can define declarative constraints in document clustering and retrieval.

This view of KM gives it an extraction and maintenance quality flavor that contrasts with any notion of intelligence or reasoning that might be pursued by more scientific projects discussed later. Moreover, in this approach leaves little room for hypertext as an active element in the construction of knowledge, and reserves hypertext as a method for selecting documents returned from searches or automatic retrievals.

KM as a Social Nexus

Collaboration is in some KM implementations the central feature. These systems are necessarily hybrids of what exists currently because collaboration must include searching, document management, personalization, and other features of a dynamic knowledge-document space. Lotus Notes R5 is an example of such a system, with advanced search, document management, and collaboration. There are multiple features are designed to widen the repository search base (*domain search* lets users search multiple databases, attachments, and files in file systems with a single query; *content mapping* allows users to browse information organized in categories) as well as multiple collaboration features (*TeamRoom* for structured project-oriented discussion; *Sametime* for network-based, real-time communication and collaboration). The Lotus KM flagship product, *K-Station*, extends R5 with even greater emphasis on the collaborative dimension through the implementation of *community places*, online forums for group communication that enable collective access to resources and the overall management and execution of business processes. Users can collaborate on a proposal and then preserve their efforts -- including information about relationships formed, content developed and activities executed -- in a secure space for future use (Cf. http://www.lotus.com/home.nsf/welcome/km).

This approach has the advantage that rules for generating knowledge are internally constructed. It therefore seems better suited for teams than the imposition a large organizational document or practice structure, which can impede momentum and creativity. But the disadvantage of such a team-centric strategy is that what knowledge is created cannot be immediately integrated into a larger epistemic space because it is not in the organizational form that a large-scale system needs. While it may be inscribed into documents, this procedure represents both an secondary transformation and a stagnation of whatever processes were employed in the closed team activity. What's worse, ad hoc collaboration is not aided by sophisticated thinking resources or processing offered to the team; the software is there primarily to aid in communication and management of what is created, not to stimulate the creation of something new. Clearly collaboration must take place in order for KM to be formed and fostered, but this cannot happen outside the providence of algorithmic aids to understanding and tools that support macro-organizational integration of locally derived knowledge.

Prospects

Some of what is missing in a full view of KM are the larger benefits of the scientific computational work in hypertext research, some systems of whose basic constituents are in rudimentary use by such KM systems. For instance, a notion of *topic creation* is possible in KM-friendly search engines like Verity, which use IR algorithms to cluster documents in hierarchical category spaces. But there is no direct way to evolve topics into other topics, or to use them to generate other subtopics, exactly as would happen in an epistemic exploration. This would be possible only by means of linking documents in an intelligent, flexible manner, as implemented in a system like TextNet. (Trigg &Weiser, 1986), which is based as it is on a semantic network of labeled links attached to nodes that can represent either primitive pieces of text or composite topic hierarchies. Rather, current state of the art in KM systems seems to involve more static or pre-established taxonomies or topic spaces, in which relations among texts can be determined by the appearance of textual markers that are part of the taxonomic or semantic space.

Even so, much of the power of text mining, like IR, is unused partly because it is not implemented in a hypertextual way. Passive processing native to traditional data mining, for instance, executes categorical clustering of documents using compositional criteria like term and feature characteristics. But the inferential power seen in scientific prototypes appears largely unused in commercial systems, such as the ability to *generate* semantic relations and hypothetical inferences by matching operations on term, phrase, or concept occurrences in text corpora. A scientific example is Swanson's work on logical hypothesis extraction from texts, by which lexical inferences enabled logical relationships to be determined across various kinds of medical (Swanson, 1987; Swanson & Smalheiser, 1994) and science (Swanson, 1991; Swanson & Smalheiser, 1997) literatures. In one case a medical correlation between stress and magnesium deficiency and certain kinds of migraine headaches was discovered by inferentially mining the titles of various medical journal articles (Cf. Hearst, 1999). The intertextual links were inferred by the system.

Text mining and IR are latent processes insofar as they work in the background, collecting documents, applying correlative algorithms, and building metadata indexes, catalogs, or collections that form the basis for the inferential processes. This leads to passive knowledge discovery, which would call for hypertextual activity in order to be truly collaborative. The synergy is natural and hints at numerous opportunities for KM-based automated knowledge discovery where text mining can be extended into hypertext mining. The techniques, whether based on hypertext, IR, or data mining, are analogous in their need to *unify* a document space or corpus, whether by criteria that are author-selected (as in the application of the hypertext link), query- or structure-specific (as in IR), or knowledge-centric and inferentially implied across multiple documents (as in TDM).

Some work has been done or proposed to augment knowledge by relating hypertexts by means of their links (e.g.., Ricardo, 1998) and document metadata can be similarly extracted and connected. For instance, some text mining systems can now abstract patterns of text as the first class objects. Information Discovery, Inc.'s Data Mining Suite (http://www.datamining.com/pwhs.htm) uses a p*attern warehouse* as a repository that holds historical patterns rather than historical data.

Discovery tools play an important role in knowledge management, but they do not necessarily contribute to an augmentation of knowledge in the knowledge worker or the enterprise, because the intelligent is abstracted from the user, who only receives the output of an automatic analysis. While this tacit value appears to be the current mode of knowledge management – we see it in the interest in search engines, some resurgence in information retrieval approaches, and the rise of text mining systems – another, more constructivist approach to knowledge management, was established in the hypertext community.

Knowledge Constructivism From Hypertext Research

One cannot have efficacious KM without certain important ingredients. I will mention four derived from hypertext systems.

Integration for Organizational Process

Any KM system must tie content creation to a multi-state maintenance of the organization in which it exists at the same time as it helps the knowledge worker interface to that organization's processes. There is an example of this. The first work relating hypertext to knowledge in an organizational context (Engelbart, 1963) was the NLS/AUGMENT system, built by a development team in the Augmentation Research Center (ARC) at Stanford Research Institute (SRI) in the early 1960's (Cf. the Bootstrap Institute, http://www.bootstrap.org/). Influenced by the recognition that knowledge in the corporation is a dynamic necessity, a principle first established by Peter Drucker (Drucker, 1967), the goals of the NLS/AUGMENT system were "to significantly boost individual's, group's and organization's performance (somewhat in that order) by augmenting their ability to work with knowledge" (Stone, 1991).

The effective institutional knowledge focus of NLS/AUGMENT led to its deployment in 47 government agencies and military commands and 21 commercial venues. In addition to support for a wide variety of peripherals in a day largely before standards, the system through its various flavors was comprised of an almost staggering feature set dwarfing any commercial KM system of today. This included tools for *formatting* to create and print official correspondence in the prescribed format; *template-filling*, for forms completion (including calculated totals compounded from data in other fields); project level *guesstimators*, that could project resource expenditure on the basis of what-if scenarios; *online calendars* with workgroup scheduling capability and even automatic notification via email; *intelligent correspondence* including response-due-date reminder capability; file delta *matchers* for document versioning control; *budget* control programs that assisted in filling out budgetary forms; *messaging* subsystems tied to an electronic mail package; *milestone* organization for complex task management; a *digital library* with access to a central repository of military regulations; a *task-assignment* tracking system for further project management. Beyond this, the system eventually developed a notion of *roles* as organizational targets, independent of specific individuals, who would dispatch appropriate responsibilities in agencies; *signatures* as encrypted distributed authority elements; and a system of *disposition-codes* through which the retention value of documents could be reviewed.

Configurable Units of Data, Not Just Complete Documents

Current KM systems work with complete documents, but one needs atomic data units that can be interrelated. Another major hypertext-knowledge project is Knowledge Management System (KMS), developed at Carnegie Mellon University. KMS is, like NLS/AUGMENT a distributed system, allowing the creation of decentralized hypertext networks. The semantic unit in KMS is the *frame*, a holder of text or graphical data that can be linked to other frames in a hierarchical, annotated way to build large-scale referential relationships (Acksyn et al., 1988). KMS was designed to incorporate hyperlinking in a data-rich environment for collaborative work, electronic publishing, and project management and could not achieve this if all content were indivisible below the document level. Documents are collections of thoughts and knowledge, elements of which must be hyper-manipulated for knowledge-creating purposes. Intelligence in configuration of semantic components is necessary. Nonexistent in commercial KM document-centric systems, as mentioned earlier, are semantic network control features like those of a system like *Textnet* (Trigg &Weiser, 1986). This is an example of a system with intelligence in correlating documents and fragments, but whose processing doesn't leave the user without a need to add value by applying associative labor. No true KM system should be totally turnkey.

Propagation of change

Communicating change is important to any KM initiative. Chen and Davies (1999) describe *Jasper*, a knowledge management system (of the scientific kind) as "a virtual reality-enabled multi-user virtual environment with spatialised semantic structures." *Jasper* is built as a content management system where collective recommendations of documents found on the WWW are registered to the system. *Jasper* then invokes agents to notify potentially interested users of each submission on the basis of topic preferences created from user profiles. Notification can take place in various ways: via email, posting to interest groups, a *What's New* section on a web site, or via keyword search on the site.

Balance Between Composition and Structure

A KM system should enable assimilation of emergent knowledge within pre-established processes, but not to the exclusion of knowledge creation itself. An example of the balance between both is SEPIA – Structured Elicitation and Processing of Ideas for Authoring, another early hypertext system involving structured management of ideas (knowledge) as written goals was (Streitz et al, 1989). SEPIA was a knowledge-based authoring and idea-processing engine for creating and managing hypertext documents in five areas: "activity spaces" could interface with a "content space" for collecting information relevant to a content domain; a "planning space" housed the ability to create an agenda for coordinating all sub-activities, plans and goals; an "argumentation space" archived arguments and rationale commentaries; and a "rhetorical space" could assist in structuring the logical form and sequence of the final document. Beyond text mining, the actual intelligent *interrelation* of documents has been largely overlooked by the KM community.

Conclusion

While KM is a recent phenomenon aiming to solve significant organizational and epistemic problems, hypertext research has already experimented with many KM objectives and needs. Aside from the larger breakthroughs, smaller contributions have been made as well. Versioning of documents has been present in many systems (e.g., Catlin et al., 1989). Collaboration has been long implemented in various systems (even some that are no longer in production). Annotative collaboration involving commenting and questioning others' work, simultaneous multi-user annotation, and contention management was incorporated through InterNote into Intermedia in 1989 (Catlin et al., 1989). More structured information linking, sharing, management in a group environment was implemented in the collaborative environment Virtual Notebook System (VNS) for academic clinical research (Shipman III et al., 1989). Likewise, NoteCards (Halasz, 1988) supported collaborative idea processing where users in a group could create nodes, add links to nodes created by others, and add annotations to other nodes. In its ability to process and manage group idea work, this was an early KM system. These systems and projects are much closer in spirit to organizational KM than the more abstruse scientific KM.

It is difficult to reconcile the computer scientific working definition of *knowledge* is demonstrated behavior that allows an observer to construe the actor as rational, or able to achieve its intended goal (Newell, 1982) with a industry-sensitive goal of socially shared cognition for the purpose of promoting organizational strategy and practice. One is a mentalist definition positioned on naïve realism, the other is entirely constructivist in a Piagetian or Deweyan manner. It is true that, unlike R&D-centric KM, which is primarily concerned with knowledge for abstract or experimental ends, business-based KM must utilize all knowledge for the additional and larger objective of continually maintaining a complex organization. This leads to notions of KM in the business world that are very different from more scientific or abstract ones, and defines KM as primarily a *management* activity that is made possible by the continuous renewal of the organizational knowledge base. Beyond conceptual imperatives, then, business-based KM primarily aims to boost social imperatives and activities like teamwork, collaboration, sharing, and productivity. Insofar as the broad objective of organizational KM is how to produce and optimize skills as a collective entity and this is being expanded into the idea of educational knowledge management systems for networked learning environments (e.g., De Diana & Aroyo, 1998), hypertext research should be the next point of departure as a synthesizer of both KM "cultures." The goals of both camps are similar and reconcilable if seen in the collective sense. Some research initiatives like NLS/AUGMENT went into production environments because they could support collaboration at an organizational level. This social objective is important to organizational KM, which aims to enable a whole organization to "think" (Taylor, 1996) and hypertext research has demonstrated a key role within the society of knowledge.

REFERENCES

Acksyn, R. M., McCracken, D. L., & Yoder, E., A. (July, 1988). KMS: A Distributed Hypermedia System for Managing Knowledge in Organizations. *Communications of the ACM*.

Applehans, W., Globe, A. Laugero, G. (1998). Managing Knowledge: A Practical Web-Based Approach. New York: Addison-Wesley.

Bertel, T. (1996). What is Knowledge Management. The Knowledge Management Forum. http://www.3-cities.com/%7Ebonewman/what_is.htm

Botkin, J. W. (1999). Smart Business : How Knowledge Communities Can Revolutionize Your Company. New York: Free Press.

Catlin, K. S., Garrett, N. L., and Launhardt, J. (1991). A. Hypermedia Templates: An Author's Tool. *Proceedings of Hypertext '91*, ACM Press.

Catlin, T., Bush, P., Yankelovich, N. (Nov. 5-8, 1989). From Ideas and Arguments to Hyperdocuments: Travelling through Activity Spaces. Proceedings of Hypertext '89, pp. 343-364

Chaiklin, S., Lave, J. (1995). Understanding Practice: Perspectives on Activity and Context (Learning in Going : Social, Cognitive and

Computational Perspectives). Cambridge: Cambridge University Press.

Chen, C., Davies, J. (1999). Integrating Spatial, Semantic, and Social Structures for Knowledge Management. HICSS-99. http://www.labs.bt.com/library/archive/hicss_99_1/index.htm

De Diana, I., Aroyo, L. (1998). Knowledge Management for Networked Learning Environments: Applying Intelligent Agents. Online paper: http://projects.edte.utwente.nl/proo/italo.htm.

Drucker, P. F., (1967). The Effective Executive. New York: Harper and Row.

Engelbart, D. C. (1963). "A Conceptual Framework for the Augmentation of Man's Intellect." In Vistas in Information Handling. P.D. Howerton and D.C. Weeks (eds.). Washington, D.C.: Spartan Books.

Fisher, K. Fisher, M. D. (1997). The Distributed Mind: Achieving High Performance Through the Collective Intelligence of Knowledge Work Teams. NY: Amacom.

Halasz, F. (1998, July). Reflections on NoteCards: Seven Issues for the Next Generation of Hypermedia Systems, *Communications of the ACM*.

Hearst, M. A., (1999, June 20-26). Untangling Text Data Mining. *Paper presented at ACL'99: the 37th Annual Meeting of the Association for Computational Linguistics*, University of Maryland.

Hesselbein F. & Goldsmith, M. (eds.) (1998). The Community of the Future. New York: Jossey-Bass.

Klix, F., Streitz, N.A., Waern, Y. , Wandke, H. (May 1989). *Man-Computer Interaction Research: MacInter II.* Selected Papers of the Man-Computer Interaction Research Network of the International Union of Psycychologists. North-Holland Press.

Newell A. (1982) The knowledge level. *Artificial Intelligence*; 18: 87–127.

Processedge.com (1999). A White Paper on Knowledge Management and Corporate Culture. Online white paper. http://webm32b7.ntx.net/store/paperculture/whitepaperkmculture3.htm

Ricardo, F. J. (1998). Stalking the Paratext: Speculations on Hypertext Links as a Second Order Text. Proceedings of Hypertext '98, 142-151

Savage, C. M. (1996). Fifth Generation Management, Dynamic Teaming, Virtual Enterprising and Knowledge Networking. Butterworth-Heinemann.

Shipman, F.M., Chaney, R.J., Gorry, G.A. (Nov. 5-8, 1989). Distributed Hypertext for Collaborative Research: The Virtual Notebook System. Proceedings of Hypertext '89, 129-135.

Skyrme, D. J. (1999). Knowledge Networking: Creating the Collaborative Enterprise. New York: Butterworth-Heinemann.

Stone, D. L. (1991, May) AUGMENT's Support of Organizations: A Brief History. Online document AUGMENT,132615. Available at http://www.bootstrap.org/augment-132615.htm

Sveiby, K. E. (1996). Managing Knowhow: Add Value by Valuing Creativity.

Swanson, D. R. & Smalheiser. , N. R. (1994). Assessing a gap in the biomedical literature: Magnesium deficiency and neurologic disease. *Neuroscience Research Communications*, 15:1-9.

Swanson, D. R. & Smalheiser. , N. R. (1997). An interactive system for finding complementary literatures: a stimulus to scientific discovery. *Artificial Intelligence*, 91:183-203.

Swanson, D. R. (1987). Two medical literatures that are logically but not bibliographically connected. *JASIS*, 38(4):228-233.

Swanson, D. R. (1991). Complementary structures in disjoint science literatures. In *Proceedings of the 14th Annual International ACM/SIGIR Conference*, pages 280-289.

Taylor, R. (1996). What is Knowledge Management. The Knowledge Management Forum. Online. Available: http://www.3-cities.com/%7Ebonewman/what_is.htm

Trigg, R. & Weiser, M. (January, 1986) TEXTNET: A Network Based Approach to Text Handling, *ACM Transactions on Office Information Systems*.

Weisbord, M. (1993). Discovering Common Ground: How Future Search Conferences Bring People Together to Achieve Breakthrough Innovation, Empowerment, Shared Vision, and Collaborative Action. Berrett-Koehler.

Wenger, E. (1999). Communities of Practice: Learning, Meaning and Identity (Learning in Doing: Social, Cognitive and Computational Perspectives), Cambridge: Cambridge University Press.