6-1-2009

From Social Tagging to Social Hierarchies: Sharing Deeper Structural Knowledge in Web 2.0

Harris Wu

Old Dominion University, hwu@odu.edu

Michael D. Gordon

University of Michigan

Follow this and additional works at: <https://aisel.aisnet.org/cais>

Recommended Citation

Wu, Harris and Gordon, Michael D. (2009) "From Social Tagging to Social Hierarchies: Sharing Deeper Structural Knowledge in Web 2.0," *Communications of the Association for Information Systems*: Vol. 24 , Article 45.

DOI: 10.17705/1CAIS.02445

Available at: <https://aisel.aisnet.org/cais/vol24/iss1/45>

This material is brought to you by the AIS Journals at AIS Electronic Library (AISeL). It has been accepted for inclusion in Communications of the Association for Information Systems by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

Communications of the Association for Information Systems

From Social Tagging to Social Hierarchies: Sharing Deeper Structural Knowledge in Web 2.0

Harris Wu

Old Dominion University

hwu@odu.edu

Michael D. Gordon

University of Michigan

Abstract:

Social tagging systems, such as del.icio.us, have helped users find and reuse information by sharing individuals' structural knowledge, i.e., the knowledge of relationships among documents and concepts. Besides being an Internet phenomenon, social tagging can help organizations manage their intranet document repositories. The structural knowledge embedded in tags is flat, shallow, and often ambiguous however. We develop a theoretical model to argue for potential benefits of sharing deeper structural knowledge in an electronic document repository through personal document hierarchies. Based on the theoretical model, we design a "social hierarchies" system. Deployment and exploratory study confirm the benefits of sharing personal hierarchies in a collaborative knowledge work environment and suggest future research directions.

Keywords: social hierarchies, structural knowledge, social tagging, knowledge management, document repository, organizational knowledge creation, design science

Volume 24. Article 45. pp. 785-804. June 2009

I. INTRODUCTION

Electronic document repositories are often the centerpiece of organizational knowledge management systems [Trigg et al., 1999]. Document repositories, also called organization memory systems [Ackerman and Malone 1990], support information storage and retrieval as well as knowledge creation. Each electronic document stored in a repository contains knowledge that can be retrieved, utilized, and combined to create additional knowledge.

Documents should not, however, be viewed as isolated units in a well functioning document repository. Instead, the knowledge of *relationships* among documents, or the structural knowledge of the repository, is critical for individuals to understand and utilize its contents. However, existing document repository systems lack the ability to cultivate and harness individuals' structural knowledge.

This research is intended to make several contributions. First, the paper highlights structural knowledge as an important type of knowledge to be shared in a document repository, and how individuals' structural knowledge can help organize the documents in a repository. Second, the paper develops a theoretical model to show that, while both hierarchies and tags can be used to share structural knowledge and assist organizational knowledge creation, hierarchies may be more effective. Third, we empirically validate such a theoretical model by designing a repository system using a design science approach [Hevner et al., 2004]. Deployment confirms the benefits of sharing structural knowledge and suggests several key challenges. Last, the paper discusses limitations and the next steps of this research.

II. STRUCTURAL KNOWLEDGE

Structure, according to the Merriam-Webster dictionary, is the relationship among elements in a system. Structural knowledge is defined as the knowledge about how elements within a domain are interrelated [Diekhoff and Diekhoff 1982]. Structural knowledge about a document repository is the knowledge of relationships among documents.

Structural knowledge is a building block of cognition. The relationships among information in memory have been viewed as cognitive structures [Preece 1976] or knowledge structures [Champagne et al., 1981], which are essential to recall and comprehension. Schema theory [Rumelhart 1980] claims that knowledge is stored in information packets called schemas, and it is the interrelationships among schemas that give them meaning. "Meaning does not exist until some structure, or organization, is achieved" [Mandler 1983].

Structural knowledge is important to knowledge acquisition, storage, and creation. According to Schema theory, learning is a reorganization of the learner's cognitive structure [Rumelhart 1980]. Structural knowledge has also been called conceptual knowledge, which is the integrated storage of meaningful dimensions in a given domain of knowledge [Tennyson and Cocciarella 1986]. Structural (conceptual) knowledge involves the integration of declarative knowledge (know-that), which leads to the development of procedural knowledge (know-how).

Structural knowledge is gained when a person comprehends and integrates new information. In a document repository, users assimilate structural knowledge through reflection upon the document collection, utilizing contextual information from experiences that go beyond the content of the documents.

While structural knowledge is a theoretical construct, it can be represented in tangible structures. Two representations frequently used in daily life are hierarchy and network [Preece 1976]. A hierarchy groups similar objects together. The complexity of a domain can often be conquered by hierarchical decomposition [Simon 1969], hence, a hierarchy is a good way to represent structural knowledge about a problem domain. According to Quillian's theory of active structural networks [Quillian 1968], concepts are connected in a network with links describing the propositional relationships between the concepts; in hierarchies these links always designate specialization (downward links) or generalization (upward links). While a hierarchy strictly obeys a one-to-many relationship between concepts, a network allows many-to-many relationships that can be much more complex. Hierarchies and networks can complement each other in representing structural knowledge as networked hierarchies.

In an electronic document repository, structural knowledge can be elicited by allowing individuals to categorize documents with hierarchies, construct hyperlink networks, or assign free-formed keywords (tags) to documents. Such an elicitation process itself represents a process of knowledge acquisition and creation [Wellbank 1990]. The

elicitation process also builds structural knowledge as individuals identify and discover new relationships among concepts. For many knowledge workers, categorizing, linking, and tagging are daily routines that add to their knowledge.

Once elicited, one person's structural knowledge can benefit others. Much educational research has been devoted to the conveyance of structural knowledge through instructional materials [Diekhoff and Diekhoff 1982]. An expert's explicit organization of subject matter functions as a scaffold for others to assimilate information they hope to learn. Many studies have also linked conveyance and acquisition of structural knowledge to problem solving performance (e.g., [Larkin et al., 1980]). For example, an expert's categorization of documents in an electronic repository can help novices retrieve and browse documents more effectively and efficiently.

Past research on knowledge codification and reuse has mostly focused on how to manage individuals' knowledge contained in the body of documents [Zack 1999; Markus 2001]. In contrast, though structural knowledge is extremely valuable, how such knowledge can be codified and reused has been largely overlooked. In the next section we look at how structural knowledge can be codified and reused, from a document organization perspective.

III. ORGANIZING DOCUMENTS WITH STRUCTURAL KNOWLEDGE

Organizing documents is critical to information retrieval and reuse. Like the cognitive representations of structural knowledge, documents are typically organized into either hierarchies or networks. These usable structures help people understand the relationships among documents, navigate from one document to another, learn the overall domain, and find information more quickly. The document structures are just as important as the content of documents themselves for information storage/retrieval, learning, and knowledge creation [Van Rijsbergen 1979; Anderson 1995]. Effective organization of documents is critical to the success of an electronic repository and the knowledge community it supports. Most electronic document repositories use hierarchies or hyperlinks to organize documents. Social tagging is emerging as another mechanism to organize Web based repositories.

A hierarchy allows a document repository to be systematically queried and traversed, as with a hierarchical database. A hierarchy can partition a large document collection into manageable sub collections. One major problem with hierarchies, however, is that a document cannot be easily found through a hierarchy if the document is "miscategorized." This can happen when a prebuilt taxonomy becomes outdated by the new documents coming into the repository, or when the categorizer and the person searching for documents have different perspectives on the topical relationships among documents.

Most repositories utilize a single global hierarchy to organize all documents. A global hierarchy can embed and distribute the structural knowledge of experts. A global hierarchy provides a full, uniform view of the repository as well as a common reference. There are several deficiencies associated with relying on a global hierarchy alone, however.

From a cognitive perspective, a single hierarchy cannot accommodate the conflicting individual viewpoints of knowledge workers or organizations. A common document structure adds cognitive load to individuals by requiring them to map the way they would organize information onto the way the common system does. Further, since individuals cannot personally organize the documents in the system the way they want, users who want to maintain control of the documents they have seen will need to file documents twice: once in the repository and once in their individual document collections. Users either pay the cost of the additional filing effort or face the difficulty of finding information. Users lose either way.

From a knowledge perspective, the knowledge within a document collection is distributed and emergent rather than centralized and static [Nidumolu et al., 2001]. Global hierarchies are typically created, and often maintained, by a central authority. But it is unrealistic to have a central process solely responsible for dynamic local information [Hildebrand 1995].

From a management perspective, the effort to create and maintain a global hierarchy can often be prohibitive. In an interorganizational environment or on the Web, there may not be any single authority responsible for such efforts. Maintenance of a global hierarchy, whether by a central authority or contributors, can be especially difficult in a growing document collection. Over time, even the best structured global hierarchy will deteriorate if it is not well maintained or becomes outdated as topics of interest evolve. The latter concern about a global hierarchy is inevitable for a knowledge community with emerging interests. For example, the classification scheme of Information Systems literature needs periodic updates [Barki et al., 1993].

Instead of hierarchies, many Web based document repositories, and particularly those that contain user contributed content, are organized in a network structure through hyperlinks. Hyperlinks are typically created by authors at the time of document creation. Some advanced repository systems allow readers to annotate documents with hyperlinks, or generate hyperlinks automatically based on textual similarity or/and user navigation. While convenient for casual navigation, the hyperlink organization has several limitations. There is not a guaranteed path to reach any document within a certain number of steps. Hence, a hyperlink network is not appropriate for systematic searches aimed at locating a great percentage of useful documents. In addition, newer documents will not be referenced through hyperlinks by older documents. Also, authors often are insufficiently motivated to add and maintain hyperlinks. Therefore, a manually maintained hyperlink network tends to under represent the relationships among documents.

Social tagging allows users to assign documents free-form keywords, or so called "tags." One of the most popular and advanced social tagging systems, Delicious (del.icio.us), allows a user to bookmark any URL and enter keywords, or tags, for the bookmark. Taking the whole web as a virtual document repository, Delicious allows one user to retrieve documents that others have tagged with certain keywords. Besides search, tagging systems typically allow users to browse tags as hyperlinks. Clicking on the tag hyperlink triggers a search on this tag. Compared to author defined hyperlinks, social tagging allows readers to "create" hyperlinks and contribute to the metadata of documents, which in turn can be used for browsing and search by others. Users are motivated to tag documents to facilitate their own information retrieval, promote the documents being tagged, or to voice their opinion, among other incentives [Ames and Naaman 2007]. Social tagging is a key phenomenon in the Web 2.0 paradigm [O'Reilly 2005], where collective intelligence is harnessed to benefit the whole community.

From a structural perspective, tags are flat – meaning they do not contain the parent, child, and sibling relationships that the hierarchies do. It is difficult to infer relationships among tags due to the problems with uncontrolled vocabulary, such as synonymy (multiple tags for the same concept), homonymy (same tag used with different meaning), and polysemy (same tag with multiple related meanings) [Wikipedia 2008]. In addition, users most frequently use single-word tags [Heckner et al., 2007]. In short, hierarchies contain much richer structural information than tags do. Indeed, the most notable deficiency of popular tagging systems is their inability to organize tags (categories) hierarchically [Walsh 2006]. This severely limits the knowledge that can be created, embedded and reused through tags.

Social tagging also suffers from the heterogeneity of users and contexts. Many tags are personal and difficult or infeasible to interpret in a global context, such as "cute", "to-read", and "me" [Mathes 2004]. Some social tagging systems allow browsing of "users" as well as documents, by discovering who assigned or used a given tag, and what other tags and documents this user has created or accessed. However, the majority of social tagging systems do not provide user information or other contextual information about tags. Most tagging systems do not have access control mechanisms for tags, and sharing user information would lead to privacy issues.

Social tagging, the emerging Web 2.0 way of organizing documents, and global hierarchies, the traditional gold standard, both have their limitations. One alternative or complementary solution would be social hierarchies. That is, the use of individual users' personal document hierarchies to organize documents and then sharing this structural knowledge. (While the owner of the hierarchy can be an organization or a group of users rather than a single person, for simplicity we use "personal hierarchy" instead of "local," "group," or "organizational" hierarchy in the rest of this paper.) In contrast to social tagging, personal hierarchies can capture structural knowledge of complex relationships among documents and concepts. Personal hierarchies also better preserve the personal context of structural knowledge and document organization. In contrast to a single global hierarchy, a personal document hierarchy gives individuals full control over document organization. Personal hierarchies present local views of the repository that suit specific interests of different workers, organizations or knowledge task forces. Personal hierarchies can be modified more easily and rapidly to reflect up to date information. The additional cost to a user of creating and maintaining a personal hierarchy (versus free-form tagging) may be justified by the effort saved in repeated access to documents within the hierarchy, or potential returns from future foraging [Pirolli and Card 1999] of the repository. The benefit of a personal hierarchy need not be limited to the hierarchy owner, either. Sharing personal hierarchies in the repository can help the whole community in locating and utilizing the documents.

State of the art document repository systems have little, if any, support for eliciting, sharing, and utilizing individuals' personal structural knowledge through personal document hierarchies. In the following sections, we will first develop a model of how building and sharing structural knowledge through personal document hierarchies may assist knowledge work in enterprise environments. Based on this model we derive key design requirements for a novel "social hierarchies" system.

IV. A STRUCTURAL KNOWLEDGE SHARING MODEL

The ultimate goal for most knowledge management systems, including document repository systems, is to assist knowledge creation. Other knowledge work including storage/retrieval, transfer, and application of existing knowledge can all be considered as part of the continuous cycle of knowledge creation.

Figure 1. Nonaka's Spiral Model

Nonaka [Nonaka 1994] developed a conceptual framework for organizational knowledge creation. In his “spiral” model (Figure 1), knowledge is created through a cycle of four intertwining modes of conversion between tacit and explicit knowledge: externalization, internalization, socialization and combination. Externalization refers to explication of tacit knowledge into explicit knowledge, which corresponds to the traditional notion of codification. Internalization refers to conversion of explicit knowledge into tacit knowledge, which corresponds to the traditional notions of learning, understanding or sense-making. Socialization refers to creating tacit knowledge through social interactions and shared experience. Combination refers to creating explicit knowledge from explicit knowledge, through merging, categorizing, sorting, and re-contextualizing.

Figure 2. Sharing Structural Knowledge Helps Organizational Knowledge Creation

Utilizing Nonaka's framework [Nonaka 1994], we develop a model of how sharing structural knowledge through social hierarchies and social tagging may assist in all four modes of knowledge creation (Figure 2).

- *Externalization*: A user's tacit structural knowledge about the relationship of documents may be externalized into a personal document hierarchy, or tags. Compared to tags, a hierarchy contains richer structural knowledge (knowledge of relationships) of documents.
- *Internalization*: Structural knowledge, in both personal hierarchies and tags, help the author as well as other users retrieve documents and learn from them (internalization) more efficiently. The process of organizing documents (into a personal hierarchy) itself is a sense-making process (internalization).
- *Socialization*: A user's personal hierarchy can help other users locate useful information and transfer structural knowledge to others. Sharing personal hierarchies helps users get to know each other, identify experts, and find like-interested peers. Users may also share experiences through building a document hierarchy collaboratively. While tags may also help socialize, most social tagging systems present tag-document associations without presenting disaggregated information on which users used a given tag. Even when the creators of a given tag are displayed, it is difficult to identify experts and like-interested peers just based on the shallow information contained in tags. It is difficult to establish a user's online reputation through his/her tags, because a flat list is not as effective as a hierarchy in representing domain knowledge.
- *Combination*: Documents may be "combined" into categories and hierarchies, which constitute new explicit knowledge. Personal hierarchies can be connected together. For example, a browsable list of domain experts along with their personal hierarchies would provide a way for browsing the collection. Documents may also be grouped into tags, which can further form a tag cloud or a tag hierarchy. However, tag based combinations suffer from tags' intrinsic issues, such as homonymy.

Through these four intertwining modes, both social hierarchies and social tagging can help individuals enlarge, share, and utilize their knowledge in a document repository. Furthermore, social hierarchies are potentially more effective than social tagging in assisting the knowledge creation cycle. We derive several design requirements for a social hierarchies system:

- R1: Convenient ways of organizing documents into personal hierarchies. Unlike social tagging systems where users simply assign free-form keywords to documents, assigning a category to a document involves recalling the category from a hierarchy. The system needs to minimize user efforts of categorizing a document. R1 supports externalization.
- R2: Easily navigable personal hierarchies. Knowledge internalization often involves browsing the document collection, locating documents and repeatedly accessing certain documents through navigable hierarchies. R2 supports internalization.
- R3: Convenient access to shared personal hierarchies. A user's personal hierarchy can express his/her interests, help others and supports socialization. A personal hierarchy must be conveniently accessible to others. R3 supports socialization.
- R4: Access control of document categories. Flexible assignment of read/write access to document categories allows for group collaboration in categorizing documents. Collaboration not only combines individuals' knowledge but also creates shared experiences among users. Hierarchical structures open the possibility of fine-control of personal information, thus addressing privacy issues. Different levels of voluntary participation are critical to the success of an online community [Preece 2000]. R4 supports combination and socialization.
- R5: Reputation and feedback mechanisms. A reputation mechanism is important to socialization in a community [Resnick et al., 2000]. A reputation mechanism encourages users to build high quality personal document hierarchies and share them with others. To build user reputation, the system should allow for users to provide feedback on the quality of both documents and structural information, including categories. Feedback can also provide a medium for social interaction promoting knowledge exchange. R5 supports socialization.
- R6: Connecting personal hierarchies. Personal hierarchies containing explicit knowledge can be combined and augmented. Isolated personal hierarchies alone are not sufficient to support navigation in a document

repository. The system needs to connect personal hierarchies to allow for user navigation across the repository. R6 supports combination.

Table 1. Design Requirements Support Four Modes of Knowledge Conversion

<i>R1: Personal workspace for reflection and organization</i>	Internalization, Externalization
<i>R2: Easily navigable personal hierarchies</i>	Internalization
<i>R3: Convenient access to shared personal hierarchies</i>	Socialization
<i>R4: Access control of document categories</i>	Combination, Socialization
<i>R5: Reputation management and feedback mechanisms</i>	Socialization
<i>R6: Connecting personal hierarchies</i>	Combination

A system fulfilling the above requirements would allow users to externalize their knowledge into personal hierarchies and internalize knowledge from these shared hierarchies. The system would allow users to socialize by browsing others' interests, providing feedback, establishing identity and reputation. The system would combine knowledge by connecting the personal hierarchies and allowing users to collaborate on hierarchies. Table 1 summarizes how the above design requirements support all four modes of knowledge conversion. In the following section we enhance an existing document repository system to meet these design requirements.

V. DESIGN OF A SOCIAL HIERARCHIES SYSTEM

Based on the theoretically motivated requirements in the above section, we designed a "social hierarchies" system. Starting with personal document hierarchies, the system generates a comprehensive hyperlink network that links documents, categories/hierarchies, and user identities. Document hierarchies and the hyperlink network together capture users' structural knowledge, organize the repository, and in turn, aid users' knowledge acquisition and creation.

Design Base

Our design is based on an open source document repository system called Everything, hereafter also referred to as ES. Many popular online communities, such as PerlMonks.org and Everything2.org, both containing millions of documents and users, use ES as their backend. We chose ES as the design base for its rich, community oriented features and extensibility compared to state of the art commercial systems, such as eRoom™ and OpenText™.

ES can be accessed through any Web browser. Any type of documents can be uploaded to and downloaded from ES, and documents in popular formats, such as .html or .pdf, can be directly displayed in the browser. ES supports the creation of a category, which is simply a special type of document. By default, only system administrators can create categories and organize documents into categories. By creating categories that contain subcategories, the system administrators can build a system wide document hierarchy. Categories are displayed as hypertext documents with a list of links to the documents or subcategories within them. Users can traverse a hierarchy of categories by following these hyperlinks, although it is quite cumbersome to traverse a deep hierarchy this way.

ES has sophisticated access control mechanisms for documents. ES supports user identity in an online community. Each user has a default, system-generated HTML homepage with basic user information gathered during registration. A user can modify his or her homepage, for example, by adding a biography or a photograph. ES has feedback mechanisms such that users, documents and categories can all be rated. ES also has a user reputation module that ranks users based on the quantity and quality of their contributions.

Although ES has a comprehensive list of features compared to other state of the art document repository systems, it does not fulfill all of the requirements we laid out in the previous section. Below we describe the additional features that we added to ES to support the cultivation of structural knowledge, and in turn, knowledge creation. Our major extensions include three modules: Backpack, Hierarchies, and Inspection. Figure 3 shows a screenshot of our extended Everything system. Table 2 summarizes how these extensions, along with existing features in ES, fulfill our design requirements.

Table 2. Design Requirements Met by Existing System or Extensions

	Base System	Backpack	Hierarchies	Inspection
R1: Personal workspace for reflection and organization		X		
R2: Easily navigable personal hierarchies			X	
R3: Convenient access to shared personal hierarchies			X	
R4: Access control of document categories	X		X	
R5: Reputation management and feedback mechanisms	X			X
R6: Connecting personal hierarchies			X	X

Figure 3. A Screenshot of the Social Hierarchies System

Backpack

Backpack is a user's personal workspace for reflection and document organization. A user will always see the Backpack section in the right section of the browser when using the repository system. When a user comes across an interesting document, he or she can add the document to the Backpack section. Note that a category is just a special type of document and can also be added to the backpack. A user can use Backpack to store documents that he or she is creating or modifying. Typically, these documents contain the contributor's commentary as well as links to other documents within and outside the system. Documents in the Backpack section are conveniently accessible at any time. Using click-and-drag, a user can copy or move documents between the Backpack section and his or her personal hierarchy (refer to below section) or to other users' personal hierarchies if given permission. Backpack fulfills design requirement R1.

Hierarchies

The Hierarchies module allows a user to navigate his or her personal hierarchy, other users' shared personal hierarchies, and any other document hierarchies available in the repository. A user will always see the Hierarchies section in the right frame of the browser when using the repository system. A user can switch between different hierarchies, but at any given time, only one hierarchy is shown in the Hierarchies section. All hierarchies are navigable and managed in a manner similar to file systems by Microsoft Windows Explorer.

The personal hierarchy is a user's own hierarchy. She or he can conveniently create, rename, move, and delete categories in her personal hierarchy. A user can move items from Backpack to her or his personal hierarchy by clicking and dragging. She or he can also copy items from different document hierarchies using copy and paste.

Note that a user can add to her or his hierarchy categories that are created by other users or documents that are already in their personal hierarchies, just as she or he could create shortcuts in a file system.

The others hierarchy in the Hierarchies section contains an alphabetical hierarchy (A-D, E-G, etc.) of users. Expanding a user in the others hierarchy leads to his or her shared personal hierarchy, while clicking on his or her name will lead to her homepage. By default, categories in a personal hierarchy have read permission granted to the public. For privacy, a user can limit category access to selected users only. To facilitate collaboration, a user can grant write permission for his or her categories to selected users, thereby allowing others to add documents to them. Like faceted classification schemes [Broughton 2001], shared local hierarchies provide users multiple ways to locate information. Shared hierarchies also lead users to high quality documents, similarly to other collaborative filtering systems [Resnick and Varian 1997]. By sharing their hierarchies, users establish their identities by expressing what they know and what they are interested in. Users obtain good reputations by sharing high quality personal hierarchies.

The hierarchies section may also contain global hierarchies containing all documents in the repository. In Figure 3, the time hierarchy is a global hierarchy that is divided into months and weeks. Under each week are the documents contributed during that week.

The Hierarchies module helps fulfill design requirements R2, R3, and R4.

Inspection

The alphabetical list of personal hierarchies in the Hierarchies section will not be very useful, however, unless the user knows whose personal hierarchies contain useful documents. This can be addressed in several ways. A user may have a social network of peer knowledge workers whose specialties are familiar to him or her. She or he may obtain a knowledge map, such as an organizational chart or a directory listing, to look up knowledgeable individuals in an organization. Or the document repository itself may assist users in discovering domain experts and user interests.

One way to identify domain experts in a document repository is to look for information “producers” who have authored useful documents on a given topic. ES provides another way to identify potential experts, which is to look for information “consumers” who have categorized useful documents. These categorizers are also “producers” of structural information. The Inspection module constructs a hyperlink network that allows a user to inspect the relationships among users, categories, and documents. These hyperlinks, which we call inspection links, help users identify both domain experts and useful document categories.

Consider the screen shown in Figure 3. We can tell from the Hierarchies section, where the root of the personal hierarchy is pclark, that pclark (Paula Clark) is using the system. The hyperlinks at the bottom of the main section show that a document contributed by btwillia has been personally categorized by John Vadalabene in a category titled “open source databases.” Suppose that, when Paula is browsing through this document she decides to follow the inspection link to John’s “open source databases” category to explore other documents on this topic. Paula finds many other very useful documents in John’s category. Following the inspection link to “John,,” Paula would come to John’s homepage, where the inspection module has appended hyperlinks to John’s personal hierarchy as well as his contributions. By navigating these hyperlinks, Paula realizes that John is an expert in open source. She gets his phone number and calls him with a few questions. As this example demonstrates, inspection helps users identify domain experts or like-minded peers. The inspection feature also encourages users to establish their reputation by contributing high quality documents or document hierarchies. Inspection respects the access control settings on documents and categories so that private information is not revealed.

Inspection combines the knowledge codified in hierarchies. As suggested in Figure 4, inspection links connect document hierarchies via documents. If two categories within different personal hierarchies contain the same document, these categories can be considered related. Following inspection links, a user can navigate among these related categories. When a user arrives at a document within a category in hierarchy A, she may follow an inspection link, as Paula did, to another category in hierarchy B. If there is a global hierarchy containing all documents in the repository, then every document hierarchy is connected to the global hierarchy through inspection links. Note that the Hierarchies section only allows users to explore one hierarchy at a time. In contrast, inspection links allow a user to navigate seamlessly from one hierarchy to another. Such a hyperlink structure across multiple hierarchies has been studied as “multitrees” and has many advantages for information access and reuse [Furnas and Zacks 1994].

Figure 4. Inspection Hyperlinks Connect Documents, Categories and Users

Inspection links also connect documents from different categories, even when there is no explicit connection, such as a hyperlink, between them. The connection is based on transitive association: if documents A and B are in one category (belonging to one user), and B and C in a second category (belonging to a different user), then documents A and C are likely related. Inspection links allow a user to be able to navigate from document A to C (Figure 5). Hence, inspection links allow users to explore documents with underlying associations, which may lead to important serendipitous knowledge discovery [Swanson 1987].

While going from one document to another using inspection links seems cumbersome, it can be surprisingly efficient at times. As shown by research in social psychology and information foraging [Killworth and Bernard 1978; Pirolli and Card 1999], individuals follow cues or so called “information scents” to navigate a social network or hyperlinked environment. Suppose a user is reading document A but wants to find a document on an entirely different topic, B. She or he can look at the categories that contain document A, and browse through a category that appears to be closer to topic B. At that category, she or he can browse through documents or categories that appear to be even closer to topic B. Similarly, the user may browse through the categories and documents of a user associated with one of these documents or categories, if she or he suspects that action will get her or him closer still to topic B. Eventually she or he can reach a document on topic B through inspection alone. In a repository with thousands of documents, one document is usually only a few steps away through such scent-following steps. Under certain assumptions it has been shown that the above scent-following strategy can find any document in the repository within search time $O(\log(n))$, where n is the size of the repository [Kleinberg 2002]. One limitation of inspection links is that for a popular document that has been categorized by many users, the list of inspection links will be long. For scalability, we plan to add a “more” feature to ES so that only a few inspection links are displayed initially, while the rest are accessible by following a “more...” link at the bottom of the webpage.

In summary, the Inspection module constructs a hyperlink network that connects documents, categories, and users. The network helps users form social networking bonds with each other as well as locate information. Inspection helps fulfill design requirements R2, R3, R5, and R6.

Summary

VI. DEPLOYMENT AND EXPLORATORY STUDY

Volume 24

designed a term paper assignment, demanding complex knowledge creation from shared sources, and studied how our system was able to assist students in this task. While a class environment has obvious limitations, such as a limited number of users and documents and brief duration of use, students were motivated to acquire and create new knowledge and to share knowledge with others.

Most documents in the system were contributed by the 44 students in the class. A typical contribution was a short opinion article, a brief introduction to technology, or a description of an Internet resource. Students were encouraged to add hyperlinks in their contributions to refer to other documents in the Web site. At the end of the semester the repository contained over 1000 contributions in a variety of formats (HTML, Word document, JPEG, etc.). Students were encouraged to provide feedback on the quality of documents and categories by voting for them as “useful” or “not useful.” Providing feedback was voluntary and anonymous.

Students were told at the beginning of the semester that the final term paper would be related to the documents on the class Web site, but the exact assignment was not revealed until three weeks before the end of semester. They knew, however, that all references in their final term paper had to be drawn from the documents within the Web site. Thus, students were motivated to contribute to and make sense of the repository. The paper was a complex case analysis that could potentially utilize most of the class’ contributions. Our research mainly focused on the final three weeks during which students worked intensively on their term paper.

Students navigated the repository using hierarchies, hyperlinks, and search. System hierarchies included a student’s personal hierarchy, the “others” hierarchy (containing an alphabetical list of every student’s personal hierarchy), and a global “time” hierarchy, which classified documents according to the months and weeks when they were contributed (thus mirroring the sequence of topics for the course). Table 3 shows an example of a personal hierarchy created by one of the students. Hyperlinks within the system included both student-created hyperlinks within their contributions and system-created inspection links. The system also contained other navigational hyperlinks, such as “recent contributions,” “previous,” and “next.” All user navigations, or so called “click streams,” were captured by the system with time stamp information.

Table 3. User Pclark’s Personal Hierarchy	
Team Project	
ETL	
Data warehousing	
SQL Server	
DTS Packages	
Scheduling	
Authentication	
Stored procedures	
Visual Basic examples	
Open Source Databases	
MySQL	
PostGres	
Class Information	
Homework	
Syllabus	
Database Security	
Authentication	
Data encryption	
Hashing	
Public key	
IBM	
DB2	
MQ series	

Conjectures

We have derived a number of conjectures based on a theoretical model of how sharing personal structural knowledge may assist the four modes of knowledge creation: internalization, externalization, combination and socialization.

Personal hierarchies should help the internalization process, which corresponds to traditional notions of learning or understanding. Assuming most documents in a repository are relevant to students' learning, personal hierarchies should cover a large portion of the repository. In addition, personal hierarchies should contain a significant proportion of high quality documents that students can easily retrieve for their reference. If personal hierarchies are truly useful to students' learning, they should be extensively used. Our first set of conjectures, therefore, concern the coverage, quality, and usage frequency of personal document hierarchies:

C1. Personal hierarchies cover a large portion of the documents in the repository.

C2. Personal hierarchies cover a large proportion of high quality documents in the repository, where quality is determined based on user feedbacks.

C3. Both the "personal" hierarchy and the "others" hierarchy are used more frequently than the "time" hierarchy. Although all three hierarchies can be navigated in the same way, personal and other hierarchies embed valuable structural knowledge and should be more helpful when users are engaged in knowledge tasks.

The next two conjectures concern the knowledge transfer between experts and novices through personal hierarchies. Experts and novices in the class were identified using a questionnaire at the beginning of the semester, with questions regarding a student's computer and database literacy, academic background, and professional experience. Novices do not have as much structural knowledge as experts do. Therefore novices should likely find others' personal hierarchies to be a source of structural knowledge, and helpful in locating information. In contrast, experts have richer structural knowledge. Experts' hierarchies should be better organized and contain a larger number of useful documents, and therefore more useful for others.

C4. Novices use the "others" hierarchy more often than experts do, in proportion to how often they use their own "personal" hierarchies.

C5. Experts' shared personal hierarchies are more frequently accessed by others than are non-experts' shared personal hierarchies.

The next conjecture concerns the effectiveness of inspection in helping users locate information and experts.

C6. Users are able to discover useful document categories, useful documents, as well as domain experts through inspection.

The next conjecture concerns the overall benefits of hierarchies and inspection for knowledge creation. Since personal hierarchies and inspection links both contain structural knowledge that is important to mastering a subject, each should assist learning.

C7. A student's term paper quality and his or her usage frequency of "personal" hierarchy, the "others" hierarchy, and inspection hyperlinks are positively correlated.

The conjectures above concern the benefits of sharing structural knowledge through personal document hierarchies. Structural knowledge embedded in hierarchies is richer than that in social tagging. The last several conjectures concern the difference between social hierarchies and social tagging. A personal document hierarchy should not contain only first-level categories; otherwise, the hierarchy would not be much different than a flat list of keywords. Empirical studies show that most tags in Bibsonomy, CiteULike, and other social tagging systems are uni-terms (single word tags), whereas controlled vocabularies in ontologies contain mostly two and three word phrases [Heckner et al., 2007; Good and Tennis 2008]. As users have put conscious efforts into constructing the personal hierarchy, we expect that the labels of categories in the hierarchy will contain more two and three word phrases than tags in social tagging systems. But since a hierarchy already indicates the relationships among the categories, the category labels should not be as complex as controlled vocabulary, such as the ACM classification system [Barki et al., 1993]. For the same reason, the compositionality (phenomenon of one keyword containing other keywords) of the set of labels in each personal hierarchy should be lower than non hierarchical controlled vocabulary and tagging systems. Therefore, we have the following final set of conjectures:

C8. Personal document hierarchies contain more than one level of categories, and a majority of documents in personal hierarchies are under second level or deeper levels of categories.

C9. The mean term length (number of terms) of category labels in personal document hierarchies is greater than the mean term length of popular social tagging systems, such as CiteULike and Bibsonomy, but smaller than the mean term length of ACM Thesauri.

C10. The mean percentage of (category) labels that contain other labels in a personal document hierarchy is lower than that of ACM Thesauri and social tagging systems such as CiteULike and Bibsonomy.

Results

We conducted an exploratory study using the system through online feedback, click stream analysis, and user interviews. The click stream data we analyzed was obtained during the final three weeks of the semester, during which time students worked intensively on their term paper.

Over the course of the semester, students categorized 62 percent of documents in the repository. Overall, only 33 percent of the documents in the repository were voted “useful” (received more positive votes than negative votes). Of categorized documents, however, 48 percent were voted as useful. Out of all “useful” documents, 88 percent were categorized. C1 and C2 were supported: Personal hierarchies covered a majority of the repository and included the most high quality documents.

Figure 6. The Usage of Different Mechanisms for the Critical Task. “Hyperlink” Means Hypertext Links Including Inspection Links

Figure 6 shows the usage of different retrieval mechanisms during the final three weeks in the semester, when students were intensively involved in their final term project. Users could find documents in the repository through their personal hierarchies, other users’ shared personal hierarchies, or the time hierarchy. Users could also find documents through hyperlinks or keyword search. C3 is confirmed: Both the “personal” hierarchy and the “others” hierarchy were used more often than the “time” hierarchy, with the “personal” hierarchy used most often. It is interesting that the “others” hierarchy was used even more often than search, whereas the time hierarchy was not. The information embedded in the hierarchy determines its usefulness, rather than simply the navigational convenience. The more familiar the structural relationship to the user, the more useful it becomes as well, as users strongly prefer using their personally created hierarchies (although the collection of others’ hierarchies certainly contains more structural information and more documents).

By examining the usage data more closely, we found that four students used the “others” hierarchy more often than their own personal hierarchies. By looking at the background questionnaire that students filled out at the beginning of the semester, we learned that three of them were the three students with the least experience with database management. C4 is confirmed. Novices tend to more frequently rely on others’ structural knowledge. However, one of the four students making extensive use of “others” hierarchies had an above average level of database background, and also had high exam scores. Closer examination indicated that this student also had a low participation in terms of number of contributions to and time spent on the Web site. An alternative explanation for her extensive use of others’ hierarchies is that she wanted to reduce her own effort. In an interview after the semester, she admitted that she wanted to be a “free-rider” by relying on her friends’ hierarchies.

The correlation between the usage of a student's personal hierarchy by other students and the expertise of the student owning the hierarchy, measured by either background experience or exam scores, is positive but statistically insignificant ($p > 0.1$ in a two-tailed test) in a linear regression model. C5 cannot be confirmed by our study: Experts' hierarchies may not be the most frequently used. One explanation is that the semester had only three months, making it hard to determine who the experts were. Also, neither exam scores nor students' background were public information, which again, made it difficult for students to identify which of their peers were most knowledgeable. The owners of a few of the most frequently accessed document hierarchies happened to be socially active and popular in the class. This suggests that social networks are important when choosing other users' personal hierarchies. In a larger setting, where expertise is earned over a period of years and experts are well known, we expect experts' structural knowledge to be highly sought.

To explore the benefit of inspection links we analyzed students' click streams during the final three weeks. We reconstructed 524 login sessions by 44 students during that period. A login session contains all click streams from login until logout or a time out by the system due to inactivity. By replaying these sessions we were able to observe several prominent usage patterns within them:

Pattern 1: At an interesting document, the reader follows a link to student A who has either categorized or contributed this document. From A's homepage, the reader proceeds to other documents and categories contributed by A.

Pattern 2: At an interesting document, the reader follows an inspection link to a category that contains the given document. From that category, the reader proceeds to other documents in the same category.

Pattern 3: The reader adds documents discovered through inspection links to his or her backpack and later to his or her personal hierarchy.

Users were able to discover useful documents and document categories through inspection. The documents discovered through the inspection links can be very useful, as users often add these serendipitous discoveries to their backpacks and later to their personal hierarchies. To look at users' behavior from a longitudinal perspective, we linked all login sessions for each user into a single meta session. By replaying these 44 meta sessions, we are able to identify another prominent pattern:

Pattern 4: After user A accessed user B's homepage a few times through inspection links, A started to directly access B's hierarchy through the "others" hierarchy.

It appears that A had discovered B as a subject expert and then started to tap into B's structural knowledge. Interviews confirmed that students did find inspection links a useful tool to identify subject experts. Interviews also informed us that students were able to get to know each others' interests through the document repository, which helped their face to face social interactions. C6 is confirmed.

Figure 7. Students with high paper scores use personal hierarchies more often. The Boxplot shows the minimum, 25th, 50th, and 75th percentiles, and the maximum of access frequency to their personal hierarchy, by students with different term paper scores.

We observed a positive correlation ($p < 0.05$) between students' term paper scores and the access frequency of their own "personal" hierarchy, as shown in Figure 7. Student's term paper scores and the usage of the "others" hierarchy and inspection hyperlinks had similar positive correlations. Therefore C7 is confirmed. Both personal hierarchies and inspection links appear to have helped students' knowledge creation. Students investing in their own personal knowledge structures reaped benefits from efficient access to useful documents and assimilating knowledge by organizing documents. In a complementary fashion, using the "others" hierarchy and inspection links led students to structural information others had created and to documents others had categorized. Similarly to their own hierarchies, using this "external" structural knowledge improved students' performance in the class.

It was surprising to us that users kept building their personal hierarchies even during the three weeks' critical task period, when there was tremendous time pressure to finish the term paper assignment. Their personal hierarchies expanded 20 percent in terms of the number of documents contained, and 30 percent in terms of the number of categories. This seems to suggest two things. First, the time saved from retrieving documents from a personal hierarchy offsets the effort put into building the hierarchy. Second, reflecting on the documents and organizing them is indeed an important step in an individual's knowledge creation process; in this case, writing the course term paper. Interviews with users confirmed both conjectures. In addition, several students mentioned that they used their personal hierarchies as scaffolds to organize references supporting different sections in their paper. In this way, their structural knowledge about the document collection was directly carried over to the term papers – the results of their knowledge creation.

Regarding the comparative studies between personal hierarchies and social tagging systems or controlled vocabularies, C8, C9, and C10 are all confirmed. We captured tagging data from Bibsonomy and CiteULike, two popular social tagging systems containing similar types of documents in our electronic document repository. As of January 2008, Bibsonomy and CiteULike had 58,755 and 238,785 distinct tags respectively. We also obtained the latest ACM classification thesaurus [<http://www.acm.org/about/class/>], a controlled vocabulary targeted for similar types of documents in our repository, which contained 1,207 classes.

The average depth of personal hierarchies is 3.0, meaning that many students have created fairly sophisticated hierarchies with three levels of categories. 84 percent of the documents in personal hierarchies are under second level or deeper categories. This is similar to the ACM classification system, which has three levels of categories, where most items fall below the third level category. Note that the effort of categorizing a document into a first level category is not much different from the effort of tagging. Building a multilevel hierarchy and classifying documents into it, however, requires additional efforts.

The mean term lengths of tags in Bibsonomy and CiteULike are 1.35 and 1.31, respectively. The mean term length of ACM classification thesaurus is 2.44. In comparison, the mean term length of category labels in personal hierarchies is 1.99, which is higher than that of social tagging systems but lower than the control vocabulary's. While a larger number of terms in a label leads to a more precise definition or a narrower scope, too complex labels lead to information overload. The lengths of category labels in personal hierarchies seem to be appropriate and convenient for browsing and search.

The percentages of tags that contain other tags as sub terms in Bibsonomy and CiteULike are 25.8 percent and 23.8 percent, respectively. The percentage of classes that contain other classes as sub terms in the ACM classification system is 40.4 percent. In comparison, only 3.5 percent of the categories in personal hierarchies have their labels contain other category labels in the same personal hierarchy as sub terms. This is clearly due to the hierarchical and personal nature of personal hierarchies. A person can use the location of categories in a hierarchy, rather than redundant words in category labels, to indicate the relationship between categories.

Finally, when asked what improvement to the system they most wanted for enhancing their learning experiences, nearly half of the students suggested that the system should automatically categorize new documents into personal hierarchies. Several students suggested "merging personal hierarchies into a common hierarchy." In next stage of our research, we will explore these directions.

VII. DISCUSSION AND CONCLUSION

Our research stresses individuals' structural knowledge as an important dimension of knowledge that can be explicated and managed in online document repositories. We developed a model of how sharing personal structural knowledge in an online document repository can assist organizational knowledge creation, based on Nonaka's [Nonaka 1994] framework. Based on this model, we designed a "social hierarchies" system that allows users to categorize documents with personal hierarchies and share their hierarchies with each other. Utilizing these hierarchies, the system automatically generates a hyperlink network among documents, categories and user identities. Deployment of our social hierarchies system confirmed the benefits of sharing personal structural knowledge for organizational knowledge creation. The system helps users learn a domain and locate information. The system also functions as an electronic intermediary for social interaction by helping an individual identify domain experts or peers with similar interests. Once identified, these people can become social intermediaries for retrieving electronic information. The system presents an effective and efficient way of organizing a document repository. Our study suggests that the concept of social hierarchies may be used to improve social tagging in harvesting social knowledge.

Management of structural knowledge can be studied according to Alavi and Leidner's [Alavi and Leidner 1999; Alavi and Leidner 2001] framework for organizational knowledge management, which consists of four sets of socially enacted knowledge processes: creation, storage/retrieval, transfer, and application. In our repository system, users create structural knowledge by reflecting on documents as they organize them. The structural knowledge is stored and retrieved in the form of hierarchies and hyperlink networks. Sharing personal hierarchies facilitates the transfer of structural knowledge among repository users. Users apply structural knowledge in organizing, understanding, and locating documents. The use of personal document hierarchies thus supports all four sets of *structural knowledge processes*.

As we discussed in the beginning of the paper, a global hierarchy is essential but has many limitations. State of the art repository systems have tried to address the deficiency of relying on a global hierarchy alone. For example, Documentum's eRoom allows a document to reside in several categories in a global hierarchy. Some "placeless" document repositories [Huang and Michiels 2000] allow users to assign attributes from multiple classification schemes to the same document, and then documents are located exclusively by an attribute based search. Our study suggests that personal or local hierarchies, which can better capture distributed, up to date knowledge, may complement the global hierarchies. The deployment of our system also suggests that many users would like to have global hierarchies that emerge from local hierarchies. We are experimenting with algorithms to generate global hierarchies from individual personal hierarchies in our next stage of research.

Knowledge management is more than managing documents in a repository. As Markus [Markus 2001] noted, much knowledge reuse involves access to experts, not access to codified expertise, and more and more the identification and selection of experts are mediated through knowledge management systems. Businesses have long strived to develop and maintain knowledge maps of human expertise. Like a common hierarchy in a document repository, a centrally managed knowledge map has limitations, particularly if the domain of the organizational knowledge expands rapidly or shifts over time [Davenport and Prusak 1998]. Hierarchies and inspection in our design lead users to domain experts and dramatically cut the cost of locating information. How to build a dynamic knowledge map based on personal hierarchies is another focus of our ongoing research.

One limitation of our exploratory study is that our system was deployed in an academic setting. We plan to comparatively evaluate the social hierarchies system in a class with two sessions, with one using the social hierarchies system and the other using a social tagging system. We are going to evaluate our design in other knowledge creation environments, especially large enterprise settings. We will try to enhance existing organizational document repositories and social tagging systems using the components or the design concepts of our system.

ACKNOWLEDGEMENTS

This research has been funded in part by the United States National Science Foundation award #0713290, "Collaborative Classification of Large, Growing Collections."

REFERENCES

Editor's Note: The following reference list contains hyperlinks to World Wide Web pages. Readers who have the ability to access the Web directly from their word processor or are reading the paper on the Web can gain direct access to these linked references. Readers are warned, however, that:

1. These links existed as of the date of publication but are not guaranteed to be working thereafter.
2. The contents of Web pages may change over time. Where version information is provided in the References, different versions may not contain the information or the conclusions referenced.
3. The author(s) of the Web pages, not AIS, is (are) responsible for the accuracy of their content.
4. The author(s) of this article, not AIS, is (are) responsible for the accuracy of the URL and version information.

Ackerman, M. and T. Malone (1990). *Answer Garden: a Tool for Growing Organizational Memory* New York, ACM Press.

Alavi, M. and D. E. Leidner (1999). "Knowledge Management Systems: Issues, Challenges, and Benefits," *Communications of the Association for Information Systems* 1(1).

Alavi, M. and D. E. Leidner (2001). "Knowledge Management and Knowledge Management Systems: Conceptual Foundations Research Issues," *MIS Quarterly* 25(1), pp. 107-136.

Ames, M. and M. Naaman (2007). *Why We Tag: Motivations for Annotation in Mobile and Online Media* CHI 2007, San Jose, CA.

Anderson, J. R. (1995). *Cognitive Psychology and its Implications* New York, W.H. Freeman and Company.

Barki, H., S. Rivard, et al. (1993). "A Keyword Classification Scheme for IS Research Literature: An Update," *MIS Quarterly* 17(2), pp. 209-226.

Broughton, V. (2001). "Faceted Classification as a Basis for Knowledge Organization in a Digital Environment," *The New Review of Hypermedia and Multimedia* 7, pp. 67-102.

Champagne, A. B. et al. (1981). "Structural Representations of Students Knowledge Before and After Science Instruction," *Journal of Research in Science Technology* 8, pp. 97-111.

Davenport, T. H. and L. Prusak (1998). *Working Knowledge* Boston, Harvard Business School Press.

Diekhoff, G. M. and K. B. Diekhoff (1982). "Cognitive Maps as a Tool in Communicating Structural Knowledge," *Educational Technology* 22(4), pp. 28-30.

Furnas, G. W. and J. Zacks (1994). *Multitrees: Enriching and Reusing Hierarchical Structure* ACM CHI '94 Conference Proceedings, Boston, MA.

Good, B. and J. Tennis (2008). *Evidence of Term-Structure Differences among Folksonomies and Controlled Indexing Languages* ASIS&T Annual Meeting, Columbus, Ohio.

Heckner, M. et al. (2007). "Tagging, tagging. Analysing User Keywords in Scientific Bibliography Management Systems," *Journal of Digital Information*.

Hevner, A. R. et al. (2004). "Design Science in Information Systems Research," *MIS Quarterly* 28(1), pp. 75-105.

Hildebrand, C. (1995). "Guiding Principles: Information mapping," *CIO Magazine* pp. 2-6.

Huang, J. and J. Michiels (2000). "Exploring Property Based Document Organization in a Collaborative Note Sharing System," *CHI 2000* pp.327-328.

Killworth, P. and H. Bernard (1978). "Reverse Small World Experiment," *Social Networks* 1, pp. 159-192.

Kleinberg, J. (2002). *Small-World Phenomena and the Dynamics of Information* Cambridge, MA, MIT Press.

- Larkin, J. H. et al. (1980). "Expert and Novice Performance in Solving Physics Problems," *Science* 208, pp.1335-1342.
- Mandler, J. (1983). *Stories: The Function of Structure* Annual Convention of the American Psychological Association, Anaheim, CA.
- Markus, M. L. (2001). "Toward a Theory of Knowledge Reuse: Types of Knowledge Reuse Situations and Factors in Reuse Success," *Journal of Management Information Systems* 18(1), pp. 57-93.
- Mathes, A. (2004). "Folksonomies - Cooperative Classification and Communication Through Shared Metadata," *Computer Mediated Communication (LIS590CMC)*, University of Illinois, Urbana-Champaign, IL.
- Nidumolu, S. R., M. Subramani, et al. (2001). "Situated Learning and the Situated Knowledge Web," *Journal of Management Information Systems* 18(1), pp. 115-150.
- Nonaka, I. (1994). "A Dynamic Theory of Organizational Knowledge Creation," *Organization Science* 5(1), pp. 14-37.
- O'Reilly, T. (2005). *What is Web 2.0*. from <http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>.
- Pirolli, P. and S. K. Card. (1999). "Information Foraging," *Psychological Review* 106, pp.643--675.
- Preece, J. (2000). *Online Communities: Designing Usability, Supporting Sociability* Chichester, UK.
- Preece, P. F. W. (1976). "Mapping Cognitive Structure: A Comparison of Methods," *Journal of Educational Psychology* 68, pp.1-8.
- Quillian, M. R. (1968). "Semantic Memory," M. Minsky (ed.), *Semantic information processing* Cambridge, MA, MIT Press.
- Resnick, P. and H. R. Varian (1997). "CACM Special Issue on Recommender Systems," *Communications of the ACM* 40(3), pp. 56-89.
- Resnick, P. et al. (2000). "Reputation Systems," *Communications of the ACM* 43(12), pp. 45-48.
- Rumelhart, D. E. (1980). *Schemata: The Building Blocks of Cognition* Hillsdale, NJ, Lawrence Erlbaum.
- Simon, H. A. (1969). *The Sciences of the Artificial* Cambridge, MIT Press.
- Swanson, D. (1987). "Two Medical Literatures That Are Logically but not Bibliographically Connected," *JASIS* 38(4), pp. 228-233.
- Tennyson, R. D. and M. J. Cocciarella (1986). "An Empirically Based Instructional Design Theory for Teaching Concepts," *Review of Educational Research* 56, pp.40-71.
- Trigg, R. et al. (1999). *Moving Document Collections Online* Proceedings of the Sixth European Conference on Computer Supported Cooperative Work, Denmark.
- Van Rijsbergen, C. J. (1979). *Information retrieval* London.
- Walsh, T. (2006). "What is Del.icio.us?" Retrieved 2006, from http://www.secretlair.com/index.php/entry/what_is_delicious/.
- Wellbank, M. (1990). "An Overview of Knowledge Acquisition Methods," *Interacting with Computers* 2(1), pp. 83-91.
- Wikipedia. (2008). "Wikipedia: Folksonomy," 2008, from <http://en.wikipedia.org/wiki/Folksonomy>.
- Zack, M. H. (1999). "Managing Codified Knowledge," *Sloan Management Review* 40(4), pp. 45-58.

ABOUT THE AUTHORS

Harris Wu is an assistant professor of information technology and decision sciences at the College of Business and Public Administration, Old Dominion University. His research interests include knowledge management, information retrieval, software engineering, and information economics. His research has appeared in journals such as *Decision Support Systems*, *Communications of the ACM*, *Journal of the American Society for Information Science and Technology*, among others.

Michael Gordon is a professor of business information technology at the University of Michigan. His research interests include information retrieval, especially adaptive methods and methods that support knowledge sharing among groups; information and communication technology in the service of social enterprise (promoting economic development, providing healthcare delivery, and improving educational opportunities for the poor); and using information technology along with social methods to support business education.

Copyright © 2009 by the Association for Information Systems. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and full citation on the first page. Copyright for components of this work owned by others than the Association for Information Systems must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or fee. Request permission to publish from: AIS Administrative Office, P.O. Box 2712 Atlanta, GA, 30301-2712, Attn: Reprints; or via e-mail from ais@aisnet.org.

Communications of the Association for Information Systems

ISSN: 1529-3181

EDITOR-IN-CHIEF

Ilze Zigurs
University of Nebraska at Omaha

AIS SENIOR EDITORIAL BOARD

Guy Fitzgerald Vice President Publications Brunel University	Ilze Zigurs Editor, CAIS University of Nebraska at Omaha	Kalle Lyytinen Editor, JAIS Case Western Reserve University
Edward A. Stohr Editor-at-Large Stevens Institute of Technology	Blake Ives Editor, Electronic Publications University of Houston	Paul Gray Founding Editor, CAIS Claremont Graduate University

CAIS ADVISORY BOARD

Gordon Davis University of Minnesota	Ken Kraemer University of California at Irvine	M. Lynne Markus Bentley College	Richard Mason Southern Methodist University
Jay Nunamaker University of Arizona	Henk Sol University of Groningen	Ralph Sprague University of Hawaii	Hugh J. Watson University of Georgia

CAIS SENIOR EDITORS

Steve Alter University of San Francisco	Jane Fedorowicz Bentley College	Jerry Luftman Stevens Institute of Technology
--	------------------------------------	--

CAIS EDITORIAL BOARD

Michel Avital University of Amsterdam	Dinesh Batra Florida International University	Indranil Bose University of Hong Kong	Ashley Bush Florida State University
Fred Davis University of Arkansas, Fayetteville	Evan Duggan University of the West Indies	Ali Farhoomand University of Hong Kong	Sy Goodman Georgia Institute of Technology
Mary Granger George Washington University	Ake Gronlund University of Umea	Douglas Havelka Miami University	K.D. Joshi Washington State University
Chuck Kacmar University of Alabama	Michel Kalika University of Paris Dauphine	Julie Kendall Rutgers University	Claudia Loebbecke University of Cologne
Paul Benjamin Lowry Brigham Young University	Sal March Vanderbilt University	Don McCubbrey University of Denver	Fred Niederman St. Louis University
Shan Ling Pan National University of Singapore	Jackie Rees Purdue University	Jia-Lang Seng National Chengchi University	Paul Tallon Loyola College, Maryland
Thompson Teo National University of Singapore	Craig Tyran Western Washington University	Chelley Vician Michigan Technological University	Rolf Wigand University of Arkansas, Little Rock
Vance Wilson University of Toledo	Peter Wolcott University of Nebraska at Omaha	Yajiong Xue East Carolina University	

DEPARTMENTS

Global Diffusion of the Internet Editors: Peter Wolcott and Sy Goodman	Information Technology and Systems Editors: Sal March and Dinesh Batra
Papers in French Editor: Michel Kalika	Information Systems and Healthcare Editor: Vance Wilson

ADMINISTRATIVE PERSONNEL

James P. Tinsley AIS Executive Director	Vipin Arora CAIS Managing Editor University of Nebraska at Omaha	Copyediting by Carlisle Publishing Services
--	--	---