Collaborative Tagging and Expertise in the Enterprise

Ajita John
Avaya Labs Research
307 Middletown Lincroft Road
Lincroft, NJ 07738
ajita@avaya.com

Dorée Seligmann
Avaya Labs Research
666 5th Ave.
New York, NY 10103
doree@avaya.com

ABSTRACT

Enterprise communication applications rely on automated reasoning about factors such as expertise for connecting people. Quantifying expertise is necessary for such applications because of the time constraints imposed by communications routing. This paper discusses the potential for collaborative tagging in the enterprise and how it enables the formation of social networks around tags or topics. These social networks are reflective of the interests and expertise of users contributing to the tag. The tagging activity of a user contributes to the expertise of the user and influences the expertise of other users. The paper proposes a ranking mechanism for expertise based on the tagging activity of users for both unstructured and structured tag spaces. The paper also briefly describes a communication platform that will incorporate the ranking mechanism.

Keywords

Tagging, Expert, Recommender, Enterprise Communications, Expert Rank, Expertise

1. INTRODUCTION

Enterprises want to leverage the skills of their entire workforce - from the salesperson on the floor to a retired part time worker - for handling events such as fielding questions in a customer call, making decisions in business workflow processes, helping a fellow employee with a problem, or participating in meetings in routine or emergency situations. People in an enterprise can provide valuable expertise [5] to solve problems. The suitability of a person to participate in an enterprise event depends on a variety of factors such as skills, availability, access to appropriate media, and cohesiveness with other participants. Enterprises are undergoing vast changes with a global and mobile workforce. The problem of selecting suitable people for an event is challenging because people are globally located, with a range of language skills, have easy access to vast amounts of online information leading to a buildup of topic expertise, can communicate through various modes, devices and systems. To address this need, an emerging class of *automated* context-aware communication applications that connect people using user, enterprise, and application context with little or no manual intervention are being developed [10]. This paper focuses on the problem of ascertaining topic expertise for people in an enterprise for such applications.

The task of automatically identifying people with a specified set of skills may rely on profiles that are provided by the people. Skill profiles that have to be updated by users are problematic because users often fail to keep them current. Additionally, in large enterprises, if the granularity in skill levels is too coarse, automated systems have a difficult time selecting the right people and if the granularity is too fine, users have a hard time accurately determining their levels in relation to others. Mining expertise from data has been investigated in a variety of contexts such as authorship by mining public Web documents [11], software source control systems and technical databases [15] and recommendations of Web pages in Usenet news messages [9][22]. Social networks based on references on the WWW have been studied previously to extract communities and authoritative sources [1][12]. Commercial systems use a variety of techniques from static skill profiles to mining of documents and feedback on response to expertise queries [2].

Recommendation systems may, in general, suggest either people or items. However, systems that recommend people based on expertise and systems that recommend items such as books or music based on inferred user preferences and social networks are typically categorized as different classes of systems. A comprehensive survey of both types has been provided in [23].

While authorship in documents and other content are important and valid indicators of expertise, consideration needs to also be given to the fact that users today have easy access to vast amounts of online information. Perusal of information on different topics constitutes, in many cases, a gain in topic expertise. What makes this activity important is that it reflects knowledge gain by users in topic areas whose content may be relatively static or very dynamic such as current news events.

Tagging has received considerable interest recently as a means for improved sharing of content [4][20][26][6] by adding meta-information. Tagging of content enables its organization and facilitates searching and the formation of social networks. Collaborative tagging, where the users do not have to own the content being tagged, offers promise for expertise determination for the following reasons:

- It represents user categorization of shared content that may be presumed to be representative of user interests and expertise - an “I tag, therefore I know” indication by the user.
- Users do not have to be authors of or be referenced in the content being tagged.
- Tags enable the automatic generation of expertise categories.
- It enables the formation of social networks around tags which facilitates identification of expertise communities.
- It provides a way of keeping pace with the user’s changing interests without the user having to update skill profiles.
- The feedback loop leading to asymmetric communication [25] is active in an enterprise environment because users know each other and are aware of reputations.

WWW2006, May 22–26, 2006, Edinburgh, UK.

In a mechanism that utilizes collaborative web tagging for expertise determination, issues such as ambiguity, synonyms, lack of hierarchy in collaborative web tagging [7][16] need to be addressed. Additionally, in mining data about tagging activity, it is important to address issues such as privacy, reputation, and trust.

1.1 Our Contributions

This paper focuses on the identification and ranking of expertise in the context of automated enterprise communications. We discuss the potential for collaborative tagging various types of content in the enterprise and issues related to enterprise tagging. We discuss how the adoption of collaborative tagging by users in an enterprise can form topic-based social networks of users around tags. The social networks are representative of the various categories and levels of user interests and expertise. We propose a rank *ExpertRank* for a user that attempts to quantify a user's expertise in the context of a tag. The paper proposes two models to calculate ExpertRank. The first and simple model assumes an unstructured tag space where tags have no dependencies and expertise gained in a tag's context is independent of expertise gained in another tag's context. The second and more realistic model assumes clustered tag space where each cluster contains set of tags related to each other. The relationships between the tags in a cluster are represented by links between tags and indicate overlapping expertise areas. The model enables a mechanism similar to the personalized version of the PageRank algorithm [18] to propagate expertise through the linked structure of the tags. While the approach is applicable to an Ebay–like service for expert finding, our focus in this paper is on the enterprise. The paper describes the architecture of a mature context-aware enterprise communications platform *Hermes*. Hermes employs a reasoning mechanism to select people based on various factors such as expertise, availability, social cohesion and user profiles. ExpertRank will be integrated with the reasoning process in Hermes to effectively find people with specified areas of expertise.

The rest of the paper is organized as follows. Section 2 describes a communication scenario to motivate the determination of expertise categories and levels in the enterprise. Section 3 describes various types of content that may be tagged in enterprises and how tagging may be integrated with user activities. Sections 4 and 5 describe social networks around tags and ExpertRank, the ranking system for user expertise, respectively. Section 6 briefly describes Hermes. Conclusions are presented in Section 7.

2. COMMUNICATION SCENARIO USING EXPERTISE FINDER

In this section, we discuss a communication scenario that motivates our work. It revolves around a scenario in which person-to-person communication barriers are reduced and users across a global enterprise are connected to solve a specialized and time-sensitive issue. An example of such an issue is the resolution of a customer software issue in a software technology company by people such as an account representative and a software expert.

A business process detects that the deadline specified in a service level agreement is approaching for an issue raised by a customer and triggers a communication application. This application may send invitations for a conference, to be started in say, 2 hours, to an account representative, a technical support person, and the customer. It sends the invitations using a combination of roles, presence and availability information, user rules, and enterprise policies. If a quorum of people accepts the invitation, the application establishes a conference with the three invitees at the scheduled time.

During the conference, the technical support person recognizes the need for an expert on a specific software issue that the customer is raising. There are many potential experts in the global locations of the company and the support person does not know which experts have the necessary skills and are available. The specific type of expertise may be related to a recent issue such as security vulnerability and few experts may have knowledge about it. The technical support person may be uncertain about the specific type of skill to search for and it may be helpful to him to be provided a list of skill categories in a broad topic space, but this list must reflect new and emerging areas and issues. The time constraints of a customer conference dictate that an expert is found programmatically. To automate the expert finder process, the support person initiates a second communication application which selects a possible group of experts based on expertise levels in the desired areas, cohesiveness (i.e. a history of interactions) with participants of the conference, presence, availability, location and environment, and sends out invitations to them. Depending on responses from experts, the application alerts the support person to the availability of an expert. If an expert is available, the expert is bridged into the conference.

The above scenario shows how a system that can intelligently connect people through reasoning about contextual information can result in effective communication to address the needs of the enterprise. The degree of automation of communication tasks in our scenario stands in stark contrast to a traditional communication process where people establish communications manually. Here, the human delays and the potential for errors and suboptimal communications decisions can be costly to the enterprise, especially if the pool of people for selection is large, capabilities of the media used by the people is varied, and the people are geographically distributed. From a user perspective, it protects personal information such as explicit expertise levels from exposure to other users by allowing only a trusted system to access the information.

3. COLLABORATIVE TAGGING IN THE ENTERPRISE

This section discusses the various forms of content that may be tagged in the enterprise, activity that may be integrated with tagging and some issues related to tagging in the enterprise. The issue of integrating tagging with other types of collaborative activity in the enterprise is important for the problem of expertise determination because it will enable better estimations of user expertise.

Electronic content in the enterprise is in many forms: external content from the WWW, product documents, enterprise announcements and news articles, news groups, blogs, email, audio broadcasts and messages, two-party and multi-party conversations and text chats. Many enterprises have web portals dedicated for the entire enterprise and for specific organizations where content is accessible by employees.

While some types of content such as email or two-way conversations may be private and may not be available to other users in the enterprise, email forwarding is a way of making content available to a larger audience than the original group it was intended for. Customer support interactions are routinely recorded for quality control. Additionally, multi-party conferences are recorded for future retrieval by participants and non-participants. The key point is that the authors and authorized consumers of an interaction may choose (with the right permissions) to make it available to a larger group, department, and organization as a means of information dissemination.

Stored electronic content that is perused by authorized users can be tagged and the tags and their associated content can be available to all authorized users. While documents and other forms of text can be easily tagged, prior work such as [21] can be used for the tagging of audio. Throughout this paper, we will refer to the pieces of content that get tagged by a user as bookmarks.

It may be worthwhile to integrate implicit tagging activities such as email, document, and voice mail forwarding with an explicit tagging infrastructure. This may be done by giving the user a *tag* option in their email/voice client. Additionally, the ability to tag sections of content in real-time while audio conferences are in session will be useful. In domains such as contact centers and customer support, routine feedback processes can also be integrated with a tagging system.

Tagging holds promise for the enterprise due to its collaborative nature that may contribute to shared knowledge and greater cohesiveness. There have been recent discussions and activities to investigate and provide tagging platforms for the enterprise. Dogear [17] is an enterprise system being designed and developed for collaborative tagging of bookmarks among enterprise users.

Issues such as identity in a tagging environment are interesting to explore in an enterprise domain because users know each other in professional capacities. Anonymous tagging may be an option that some users may wish to opt for. In the context of this paper, it is conceivable that users may want to separate their tagged bookmarks into domains of private, public for viewing and expertise determination, public for viewing only.

Should users rate the content in some fashion as they tag it and should the rating affect the expertise level of a user? After all, a press release about a company releasing the latest version of a piece of software may not be as relevant to a research organization as a research paper describing the algorithm employed in the latest version. But, the press release may be more significant to a business development organization. Incorrect tagging of bookmarks may affect the results derived from a social network of tags.

The influence of organizational constraints and a hierarchical community on tagging behavior needs to be studied. People's tagging activity affect professional reputations and depending on how the tags are used (e.g. automated expertise selection), people may tailor their tagging activity so that they are called upon as experts in areas of their choosing.

4. SOCIAL NETWORKS AROUND TAGS

Tagging activity enables the formation of social networks around tags. Figure 1 illustrates a sample social network that can form around tags that may be of interest in a software technology company. For now, the links in Figure 1 may be ignored. They will be addressed later in the section. Three entities are represented in this social network: (1) Tags, (2) Users, and (3) Tagging Activity in terms of the number of tagged bookmarks. A fourth entity *Time* or age of the tags is not represented in the figure. The age of a bookmark is a tag-specific factor as some bookmarks will never age; however, some others such as a news report about the performance of a stock will age fairly quickly.

The size of the font for the tags indicates the level of tagging activity i.e. the number of bookmarks associated with the tag. The dots inside the circles around the tags indicate users. The size of the dots indicates the level of tagging activity by the users. Larger dots which appear in inner circles correspond to users who have tagged more bookmarks relative to other users who are represented as smaller dots in the outer circle.

Tags such as "*SOA*" (Service Oriented Architecture), "*Web Services*", "*J2EE*", "*J2EE App. Server X*" are central to the professional interests of the users in a software technology enterprise. Tags such as "*Stock Market*" and "*Weekend Escapes*" may fall under personal interests for the users. Note that "*Weekend Escapes*" has approximately the same number of bookmarks as "*Web Services*", but there are fewer overall users and more users in the inner circle for "*Web Services*" than for "*Weekend Escapes*".

The social network in Figure 1 shows the existence of topic-based sub-communities around tags. For instance, the users around the tag "*Java*" illustrate the community of users interested in Java in the enterprise. The network may be enhanced to include lurkers or users who access the tagged bookmarks but do not tag.

Figure 1 also shows three categories of users in an attempt to show varying interests. Users with a white dot are technology strategists who are in the inner circle of "*Web Services*" and "*SOA*" and in the outer circle of "*weekend escapes*" and "*stocks*". Users with a light gray dot represent the developer community; they appear in the inner circle of "*Java*", "*J2EE*" and "*J2EE App Server X*", and in the outer circle of other tags.

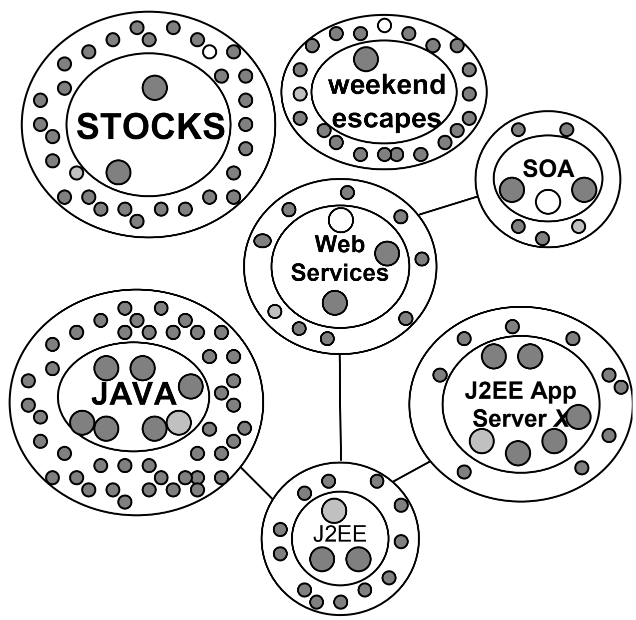

Figure 1: Visualization of Potential Tags, Associated Social Network, and a Relationship Structure expressing dependencies between tags

Numerous studies have shown that social networks in communities are effective channels of expertise dissemination [8][13][24]. The existence of social networks around tags facilitates the sharing of tagged content in the tagging user community. Communities around a tag may be strengthened by each other's tagging behavior. For example, even if user *X* may have marked only 10% of the bookmarks with a tag *T*, she may benefit from another user *Y*'s tagging behavior who has marked 40% of the bookmarks with *T* because user *X*'s tagging activity in the context of *T* may imply that she maintains an active interest in *T* and is likely to read content tagged with *T* by other users.

Some of the issues with tagging are ambiguity and synonyms [7][16] resulting from using different tags for the same bookmark. [3] discusses an algorithm for clustering tags based on strongly related tags. Pairs of tags are strongly related if they have above a threshold of bookmarks marked with both tags. Strongly related tags are relevant for social networks around tags because they represent overlapping areas of interest and expertise. Additionally they partition the tag space into clusters of broad topic domains. Figure 1 shows a potential cluster with tags in the Java-related technologies space.

5. EXPERT RANK

The social networks in tagging communities in an enterprise offer the opportunity to tap into the expertise potential of the enterprise. While a manual search through these social networks is an option for finding a person with a specified set of skills, in scenarios such as the one described in Section 2, automated determination of a small set of potential experts is critical to comply with the time constraints of the application. In order to do this, we propose to quantify the expertise level of a user in the context of a tag. Relevant factors to consider are the number of bookmarks tagged with a particular tag by a user and age of the bookmarks. In this section, we propose a ranking system for experts based on the number of bookmarks tagged by the user.

For each tag, our scheme uses the number of bookmarks that a user contributes as an indication of the user's interests and expertise. The rank of an expert *ExpertRank* is calculated for each tag based on the number of bookmarks that the expert has marked with that tag. Two proposals are presented in this section. The first one is based on a simple model of expertise in the social network where there are no dependencies between the tags. The second one is based on a model where the tag space is partitioned into clusters and each cluster is a group of strongly related tags. Note both models can accommodate the aging of bookmarks by having a decay factor associated with each bookmark.

A search for an expert with a specified set of skills is translated to a search for a user who has contributed bookmarks to a set of tags where the skills are mapped one-to-one to tags in the tag space. The tags may be rated in a priority order. In a search, the rank of a user may be computed as the weighted sum of the ExpertRank of the user for the different tags, where the fractional weights represent the priority order for the tags.

We first discuss a few privacy issues related to the monitoring of tagging activity followed by a discussion on the calculation of ExpertRank.

5.1 Privacy

Monitoring any kind of user activity is a sensitive issue and any system that monitors and makes inferences based on user activity should be a system that users agree to participate in. Users should have control over what tags can be included in such a system. However, it is to be noted that mining private interaction data such as emails is vastly different from mining an open collaborative activity such as tagging. Collaborative tagging essentially implies that users are comfortable with and possibly want other users to be aware of their tagging activity. Coarse inferences of user interests may be easily done in any open collaborative tagging system unless the system offers anonymous tagging.

Users may be open to participating in a brokered service for expertise where the service makes the matches between the search for an expert and the expert but does not reveal what it knows about the expert. Users may want a brokered service between systems that seek expertise and users because it will enable better matches with experts and users will not be called upon to offer expertise in areas where they may not be the best available experts. Note that a system supporting scenarios such as the one described in Section 2 may utilize a variety of factors such as availability, presence on required media, social cohesion, locality, load-balancing of queries and expertise. Assumptions that an expert recommended by such a system is the best expert for the required topic cannot be made because factors other than expertise have been used to select the person. So, the system offers protection from inferences by users about the specific ranks of other users.

5.2 ExpertRank in an Unstructured Tag Collection

In an unstructured tag collection, where there are no assumed relationships between tags, the ExpertRank for a user *u* for a tag *t* is the (normalized) number of bookmarks contributed by the user *u* to the tag *t* divided by the total number of bookmarks contributed by all the users to the tag *t*:

$$\text{ExpertRank}(u, t) = B_{u,t} / \sum_u B_{u,t}$$

where $B_{u,t}$ is the number of bookmarks contributed by user *u* to the tag *t*.

The expertise of a user in a tag is independent of the expertise of the user in other tags. Note that we are not explicitly capturing the effect of a user being part of the community around a tag in the calculation of her expertise. In other words, her expertise is strictly the number of bookmarks she contributes and not the number of bookmarks that are contributed by others and that she may consume as a result of being in a community. Since we are using a normalized count and users may be presumed to affect each other, the effect of this omission is reduced on any relative ranking among experts. However, we are not capturing individual user behavior as in a user who is part of a community around a tag, does not tag prolifically, but consumes all content related to the tag better than other users. Additionally, we are not considering lurkers or users who do not tag but use the tagging system for consuming content.

5.3 ExpertRank in a Structured Tag Collection

An unstructured tag space is simple and is not a realistic model as many tags may be highly related to each other. In this subsection, we explore tag spaces that have been partitioned into disjoint clusters of tags. Each cluster of tags can be represented by a graph where nodes represent tags and edges between nodes indicate a strong relationship of context between tags. This model implies that expertise gained through tagging activity in a tag's context contributes to expertise in the contexts of all tags that are strongly connected to that tag. An illustrative example is the case of two tags "*Web Services*" and "*SOA*", where expertise gained in "*Web*

Services" contributes to expertise gained in "*SOA*" (See Figure 1) and vice versa. Techniques such as the one described in [3] extract clusters that are undirected graphs (i.e. the relationship between the tags is bi-directional). However, there may exist relationships between tags which are uni-directional such as a small component of a larger system where the context (for expertise determination) is weak from the small component to the larger system but may be strong in the other direction.

There are two aspects that are worth capturing in the calculation of ExpertRank of a user *u* for a tag *T* that belongs to a cluster *C*. First, and most important, the rank should depend on the tagging activity of the user in tag *T*. This aspect may be easily captured by using the normalized number of bookmarks marked by the user with *T* as described in Section 5.2. Second, the tagging activity of the user in the other tags in *C* is also important to consider. While it may be argued that the first aspect will dominate any expertise rank and that users who have tagged a large number of bookmarks with a particular tag will always get selected for expertise in a tag's context, it is to be noted that the use of other factors such as availability, past interactions, presence on certain media may mean that these users may not be potential candidates for expertise. In such cases, the expertise of other users including those who may have tagged only in related tags may have to be considered. To consider the expertise gained from related tags, we use a model similar to that used in PageRank [18] to rank the relative importance of a webpage based on references to it and to rank reputation in [19].

The main idea in PageRank is that the importance of a (web)page is determined by the importance of pages that have a reference to it. The web is represented as a graph where nodes represent pages and a directed edge from node *i* to node *j* represents a reference on the page represented by node *i* to the page represented by node *j*. The pages confer their importance proportionately among their outgoing links. The model is a Markov chain with a primitive transition probability matrix whose principal eigenvector (which is the PageRank vector) can be found by applying the power method to the probability matrix. Many papers have studied the PageRank algorithm and [14] provides a comprehensive survey of issues related to the algorithm.

We consider the tagging activity of a user within the context of a tag model where expertise gained by a user in a tag's context is proportionately divided among outgoing links. We extract tag-specific ranks for each user by using an approach similar to the personalization vector model for PageRank where topic-specific or user-specific rankings for pages may be extracted. The personalization vector in our case is the activity of the user for each tag normalized over the user's activity across all tags.

Assume that there are *n* tags in the cluster under consideration. Let *M* be the n×n normalized matrix that represents the link structure of the related tags in the cluster. We define a damping factor $0 \leq f \leq 1$. ExpertRank of a user *u* for the different tags is the solution to the equation

$$\text{ExpertRank(u)} = (1\text{-f})*\text{A} + \text{f} * \text{ExpertRank(u)} * \text{M}$$

where A is an n×1 vector that represents the normalized tagging activity of user u in the context of all the tags i.e. $A_t = B_{u,t} / \sum_t B_{u,t}$

ExpertRank for a user is the principal eigenvector of the above transformation and can be computed by the power method where an initial guess is made for the ExpertRank and repeated iterations compute successive values for ExpertRank. The damping factor determines how much of the rank is retained in a node and how much is distributed across outgoing links in successive iterations. Different values of the damping factor will yield different values of the ExpertRank vector. The process is guaranteed to converge provided the transition matrix representing the above transformation is irreducible (strongly connected graph) and primitive (matrix P>= 0 is primitive if and only if $P^m > 0$ for some m>0). The latter condition will be satisfied by the clustered tag space. To make the matrix irreducible, nodes that cannot be reached by other nodes in the matrix can be removed.

6. EXPERT FINDER PLATFORM

This section briefly describes the architecture of the communications system Hermes [10] that has been implemented to support scenarios such as the one described in Section 2. More than 200 demonstrations of such scenarios have been done on Hermes at Avaya and at various public forums such as tradeshows. These demonstrations have relied on statically specified expertise profiles for users. Our proposed work is to integrate Hermes with the expert ranking system described in Section 5.

Hermes consists of a communication flow design environment that abstracts away low-level details of communications and allows a designer to easily design application flows and configure contextual parameters such as priorities for tags in the search for an expert. Hermes flows execute on a workflow engine. The workflow engine sequences the execution of tasks in the flows. The communication tasks represent a reasoning layer that accesses contextual information from different sources and use this information to determine who, how, and when to contact and connect people. The reasoning processes use algorithms and rules that operate on the enterprise, application, and user context to make the appropriate selections. The various types of context are accessed from Hermes communication session components as well as LDAP enterprise directories, databases, and other servers that collect and store contextual information such as enterprise rules, policies, user profiles, preferences, and roles. When tasks launch communication activities such as phone calls, conferences, sending of messages, they issue abstract communication requests to a Hermes communication sessions layer which queues up requests, logs them, and issues requests to the servers that offer the communication services.

7. CONCLUSIONS

In this paper we presented a scenario that shows the need for the automated determination of expertise categories and ranks in a modern-day enterprise that has the need for contacting people based on a variety of factors such as expertise, availability, and social cohesion. Collaborative tagging provides an environment where users collaboratively express their categorization of different kinds of content such as documents, blogs, audio, text, and email. The paper notes the need to integrate various implicit tagging activities into an explicit tagging infrastructure to be able to better reflect the expertise of users. By categorizing content through tags, users express their interests and growing expertise in different tags. Tagging activity creates the formation of social networks of users around tags. The paper defines ExpertRank, a quantity that represents the expertise level of a user in the context of a tag. The paper proposes two models that allow the calculation of ExpertRank. We propose to extend a mature and well-tested communication system Hermes to incorporate the computation of ExpertRank for the purpose of selecting experts based on a variety of factors. We plan to conduct user studies that attempt to study user responses and satisfaction to an automated expert system.

8. REFERENCES

[1] Adamic, L.A., and Adar, E. Friends and neighbors on the web, Social Networks, 25(3):211-230, July 2003.

[2] AskMe. http://www.askme.com/

[3] Begelman, G., Keller, P., and Smadja, F. Automated Tag Clustering: Improving search and exploration in the tag space. Available at http://www.rawsugar.com/lab.

[4] Del.icio.us. http://del.icio.us/

[5] Ehlirch, K. and Cash D. Turning Information into Knowledge: Information Finding as a Collaborative Activity. Digital Libraries '94, 119-125.

[6] Flickr. http://www.flickr.com/

[7] Golder, S., Huberman, B. The Structure of Collaborative Tagging Systems, To appear in Journal of Information Science, 2006.

[8] Granovetter, M. Strength of weak ties. Amer. J. Sociology 78 (1973), 1360–1380.

[9] Hill, W. C. AND Terveen, L. G. Using frequency-of-mention in public conversations for social filtering. In Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'96) (Nov.) Boston MA. 106–112.

[10] John, A., Klemm, R., Mani, A., and Seligmann, D. Hermes: A Platform for Context-Aware Enterprise Communication, 3rd Workshop on Context Modeling and Reasoning (CoMoRea 2006) at PerCoM 2006, March 2006, Pisa, Italy.

[11] Kautz, H., Selman, B., and Shah, M. ReferralWeb: Combining social networks and collaborative filtering. Comm. ACM 30, 3.

[12] Kleinberg, J. M. Authoritative Sources in a Hyperlinked Environment. Journal of the ACM, Vol. 46, No. 5, September 1999, pp. 604 –632.

[13] Kraut, R., Galegher, J. and Edigo C. Intellectual Teamwork: Social and Technological Bases for Cooperative Work. Lawrence Erlbaum, Hillsdale, NJ, 1990.

[14] Langville, A. and Meyer, C. Deeper inside PageRank. Internet Mathematics, Vol 1, No 3: 335-380.

[15] McDonald, D. W. and Ackerman, M. S. Expertise recommender: A flexible recommendation architecture. In Proceedings of the ACM Conference on Computer-Supported Cooperative Work. (CSCW'00). 231–240.

[16] Mathes, A. Folksonomies – Cooperative Classification and Communication Through Shared Metadata. http://www.adammathes.com/academic/computer-mediatedcommunication/folksonomies.html

[17] Millen, D., Feinberg, J., and Kerr, B. Social Bookmarking in the Enterprise. Social Computing, Vol. 3, No. 9. Nov. 2005.

[18] Page, L., Brin, S., Motwani, R., Winograd, T. The PageRank Citation Ranking: Bringing Order to the Web. Tech. Report, Stanford Digital Libraries, 1998.

[19] Pujol, J.M., Sangüesa, R., Delgado, J. Extracting Reputation in Multi Agent Systems by Means of Social Network Topology. AAMAS '02, July 15-19, 2002, Bologna, Italy.

[20] RawSugar. http://www.rawsugar.com/index.faces

[21] Stifleman, L., Aarons, B., and Schmadt, C. The Audio Notebook: Paper and Pen Interaction with Structured Speech. SIGCHI, Vol. 3, No. 1, 2001, Pages 182-189.

[22] Terveen, L. G.,Hill, W., Amento, B.,McDonald, D., and Creter, J. PHOAKS: A system for sharing recommendations. Comm. ACM 40, 3, 59–62.

[23] Terveen, L. and McDonald, D. W. Social Matching: A Framework and Research Agenda. ACM Transactions on Computer-Human Interaction, Vol. 12, No. 3, September 2005, Pages 401–434.

[24] Wasserman, S. and Galaskiewicz, J., Eds. Advances in Social Network Analysis. Sage, Thousand Oaks, Calif., 1994.

[25] Udell, Jon. "Collaborative knowledge gardening." InfoWorld. August 20, 2004. http://www.infoworld.com/article/04/08/20/34OPstrategic_1.html

[26] Yahoo MyWeb 2.0. http://myweb2.search.yahoo.com/myweb