

That's What Friends are for: Facilitating 'Who Knows What' Across Group Boundaries

N. Sadat Shami¹, Y. Connie Yuan^{1,2}, Dan Cosley^{1,2}, Ling Xia², Geri Gay^{1,2}

¹Information Science Program, ²Department of Communication
Cornell University
Ithaca, NY 14850, USA

{ns293, yy239, drc44, lx33, gkg1}@cornell.edu

ABSTRACT

We describe the design and evaluation of K-net, a social matching system to help people learn 'who knows what' in an organization by matching people with skills with those who need them. Transactive memory theory predicts that K-net will improve individuals' awareness of 'who knows what'. This should lead to improved performance through sharing knowledge across group boundaries. We evaluate K-net in terms of these predictions in an experiment with 41 students in seven groups working on software engineering projects. Accurate recommendations improved awareness of 'who knows what' versus 'random' recommendations, but did not improve performance. Our results highlight issues related to the evaluation of systems for sharing knowledge across group boundaries.

Categories and Subject Descriptors

H5.3 [Information interfaces and presentation]: Group and Organization Interfaces - *Computer Supported Cooperative Work*

General Terms

Management, Design, Experimentation.

Keywords

Transactive memory system, recommender systems, social network analysis, experimentation, expertise, friendship.

1. INTRODUCTION

Recent research on knowledge sharing in groups shows that obtaining knowledge across group boundaries can improve performance [5, 6]. In response, researchers have built expertise recommenders, systems that bridge local 'islands' of knowledge by facilitating information flow across group boundaries. These systems are meant to augment and assist the knowledge discovery process that usually happens through normal socialization processes in organizations.

Terveen & McDonald see expertise recommender systems as a

subset of a broader category they call 'social matching' systems [12]. An important concern for social matching systems is that, because they recommend people, they raise social and cognitive issues such as privacy, trust, emotional closeness, reputation, and availability.

Social matching systems can respect these social and cognitive issues by using friendship networks in making recommendations. The literature on information seeking suggests that when faced with non-routine complex tasks, individuals often consult others with whom they have an existing social relationship [1]. People often form friendships with co-workers, relationships they can draw on when seeking help completing work tasks. In fact, when given the choice, people enjoy working with someone they like over a more competent individual [3].

In this paper, we present K-net, a social matching system that uses both expertise and friendship to match people with others who possess needed skills or share their interests. K-net is a Transactive Memory (TM) system, helping people both to learn 'who knows what' in the organization and to locate others who have needed skills. In an experiment conducted in a real usage setting, K-net successfully fostered awareness of 'who knows what'. We discuss expertise recommendation and TM theory below, and then present our design and evaluation of K-net.

2. EXPERTISE RECOMMENDATION

A number of systems help people locate expertise. ReferralWeb analyzes web documents to identify names associated with topics, infers social relationships from co-authorship data, and presents a referral chain from the seeker to the expert [8]. Expertise Recommender mines software source control systems and technical support databases, associating individuals to specific software modules and helping people contact those who know the most about a given module [9]. Systems such as Answer Garden [1], the Designer Assistant [13], and PHOAKS [7] present relevant information sought by a user and the email address of the person who contributed the information.

K-net is like ReferralWeb and Expertise Recommender in that it uses social networks to provide recommendations. Unlike them, however, it incorporates friendship information when making recommendations. It is not enough to find an expert; the system must find an expert who is likely to respond. A friend is accessible, trustworthy, and has a higher probability of response [18]. Additionally, having a friend as a social conduit between a friend of a friend makes a referral chain more likely to succeed. It also allows multiple interactions creating opportunities to try, err, and seek instruction and feedback.

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 the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GROUP '07, November 4–7, 2007, Sanibel Island, Florida, USA.
Copyright 2007 ACM 978-1-59593-845-9/07/0011...\$5.00.

3. THEORETICAL BACKGROUND

Organizations are transactive knowledge systems in which the bulk of knowledge is distributed among members and exists in individuals' heads, with each individual possessing different areas of expertise [16]. To make full use of the intellectual resources of an organization, its members must know 'who knows what' and interact with each other to utilize and combine knowledge. Knowing who to ask for help with a given task can make work processes efficient. TM theory deals with how a set of individual memories, in combination with the communication that takes place, creates a group memory system [16]. Extant research on TM has been within groups because of the interdependence between group members as well as more opportunity to learn about each other's expertise. A contribution of this study is to show how a social matching system can extend transactive memory across group boundaries by fostering awareness of expertise that can be leveraged outside the group.

Three processes are vital to TM systems [17]. *Directory updating* is the process by which people learn about 'who knows what'. Accurate directories assist groups in utilizing their resources and improve coordination between workers for task completion [11]. *Knowledge allocation* refers to the process of forwarding requests for specific knowledge to relevant experts. *Knowledge retrieval* describes the process of obtaining knowledge from experts. Brandon & Hollingshead extend TM theory, arguing that TM systems should be geared around a task [2]. Knowing 'who knows what' can be useless if it does not fit the task at hand.

K-net supports all of these processes. It supports directory updating by showing who possesses relevant skills and allowing users to supply expertise information. It supports knowledge allocation and knowledge retrieval through providing names and contact details of individuals with needed skills, and facilitating their responses. Finally, it is task-focused, recommending people with skills a group member deems necessary for completing a task.

4. DESIGN OF K-NET

K-net is a web-based system catering to organizations where groups work on different projects, a common model in modern organizations. Users create profiles describing their relationships, expertise, and needs. K-net uses these profiles to recommend others who share their interests and satisfy their needs. We describe this in more detail below.

4.1 Profiling Users

K-net users provide three types of information: their level of friendship with others in the organization, a list of skills they possess, and following Brandon and Hollingshead, a list of skills they think are necessary for their group to complete their current task. Making these lists of skills serves the *directory updating* function of a TM system and provides the knowledge needed to make recommendations.

Rather than specifying a fixed set of skills, K-net allows input of free-form text. This will sometimes lead to failures in recommendation, because users may use different terms for the same skill, such as "HTML" versus "web design". However, it allows users to request and specify new skills the system does not

yet know about, increasing the system's (and the organization's) flexibility.

Asking users to provide friendship and skill information is a sensitive issue. Such information is not static and will require to be updated. However, when users perceive they will benefit from the system and privacy safeguards are ensured, most people are willing to perform the extra work [12]. In fact, users already provide similar personal information on social networking sites such as facebook, myspace and friendster. Work on discretionary databases also suggests that experts contribute more to public goods when they expect value in return [14], and K-net's directory is like a public good. Furthermore preserving one's reputation should motivate participants to update expertise information so their expertise score is accurate and reputation intact. Future empirical work will involve testing the accuracy of these predictions.

4.2 Computing Matches

To make recommendations for a given user, K-net determines the difference between the skills the user thinks his group needs and the skills the user already has. Group members typically have different conceptualizations of the skills needed by the group [11] and K-net takes this into consideration to provide personalized recommendations. K-net first assigns each other person in the organization an expertise score for a given skill. People can claim they have a given skill, or claim their friends have that skill. K-net also integrates feedback on interaction with a potential expert through a rating system. People can rate their experience with a recommended expert by providing a 'thumbs up' or a 'thumbs down'. The base expertise score for a potential expert is the number of people who claim the expert has the skill plus the number of positive ratings minus the number of negative ratings.

K-net then uses friendship ratings to modify the base expertise score. Users rate others between 0 (don't know this person) and 5 (especially close). The system assigns a score of 1 to people who are rated 0 but who are friends of a friend. A potential expert's friendship score is added to their base expertise score to create a personalized expertise score for the user making the request. It is entirely possible that a close friend with a lower expertise score will be recommended over an expert with a higher expertise score. We make this trade-off because we believe a close friend is more likely to respond to a user rather than an unknown expert. The unknown experts are still displayed, just further down the list, so users can decide who they want to contact. Using friends to find potential experts serves the *knowledge allocation* process of a TM system by providing access to experts while increasing the probability that these experts will respond to a user request.

K-net also identifies individuals looking for the same skills by matching lists of skills needed by the group that were identified by other individuals across the whole system. This 'shared goals' feature helps people looking for the same skills collaborate by comparing notes about how they are finding the expertise they need to complete their task.

4.3 Introducing Users to Each Other

When users request recommendations, K-net provides a list of names for each skill sought by the user in descending order of personalized expertise score. Figure 1 shows an example for a

user who specifies that HTML, socket TCP/IP, database management, and WML are the skills the group needs and that he himself has skills in HTML and Java. K-net found potential experts for database management, but not for socket TCP/IP or WML. Clicking on a name shows the profile of the recommended individual, and the chain of connection between the individual and the user.



Figure 1. K-net interface showing recommendations of people that possess specific skills that a user seeks. Names have been obscured to protect privacy.

Knowing a particular individual has a skill or is seeking a skill that the group needs creates a ‘common ground’ for conversation, which facilitates effective communication [4]. By identifying such individuals, K-net provides an easy platform to start a conversation.

4.4 Interaction

K-net provides the email addresses of potential experts. This facilitates the *knowledge retrieval* process of a TM system. Since K-net is designed to be employed in an organization and email addresses are commonly available in corporate or personnel directories, this seems like a minimal compromise of privacy. Furthermore, individuals are likely to trust others that share their organizational affiliation [9]. However, we do ask users to consent to sharing their email address before they can use the system.

5. EVALUATION

We evaluated K-net in a field experiment, using project groups in a software engineering class. We addressed two questions: does K-net increase awareness of ‘who knows what’, and does it improve task performance?

5.1 Participants

Seven groups composed of 41 individuals participated in our study. Participants were 4th year undergraduates or Master’s students (mostly 5th year undergraduates) enrolled in a software engineering class at a large northeastern university. We hoped the friendship networks these students developed over the duration of their studies would be a reasonable simulation of the kinds of networks that would be present in many organizations.

Participants created groups by themselves. Group size ranged from 4 to 7 students. Participants stayed in the same group throughout the semester to complete a software project. All participants were required to use K-net. On average, participants logged into the system 5.4 times. Users reported having 2.3 skills on average and reported needing 2.5. Skills and needs did not overlap perfectly, so users on average received recommendations for around 1 skill and 9 experts for a given skill.

5.2 Procedure

Groups were randomly assigned to experimental and control conditions. To maximize comparability, we gave both conditions access to a K-net-like system. Providing some with a tool and others no tool within the same class can lead to validity threats like the Hawthorne effect, compensatory rivalry, and demoralization. So we chose to give both conditions access to K-net, but gave ‘false’ (random) recommendations to controls as a placebo. These recommendations were for experts that actually possessed the claimed skills, but may or may not be skills needed by the person requesting them. So they represent true expertise and should not mislead people. There were 3 groups in the experimental condition (N=20) and 4 groups in the control condition (N=21).

5.3 Measures

K-net was given to participants to use towards the later half of the semester. Before it was released, we conducted a survey of ‘who knows what’. In this survey we asked them to list the names of classmates they thought possessed the skills needed for the group to successfully complete its project. We conducted the survey again after the project was completed. Subjects could not use K-net during the second survey. The ‘who knows what’ variable represents the correct number of individuals possessing the needed skill, as determined by a match with self reported or friend reported data on skills reported through K-net. We measured performance as the participant’s final grade in the semester long software engineering project.

6. RESULTS

Our first question is whether K-net improved awareness of who knows what. Since this dependent variable is at the individual level, our data analysis procedures were also at the individual level. Table 1 shows the results of the surveys. We analyzed data using the non-parametric Mann-Whitney test because both ‘who knows what’ and performance data were not normally distributed even after transformations. Before K-net was released, participants in the experimental condition did not significantly differ from controls on how many people they accurately said possessed a skill their group needs. After using K-net however, participants in the experimental condition were significantly more accurate than controls in identifying people with the skills their group needed, $U = 88, p < 0.001, r = -0.7$. Incidental learning of others’ skills through class lecture (there were no labs) did not occur because differences were obtained only post deployment.

Table 1. Who knows what by condition before and after K-net.

	K-net (experimental)		K-net (control)	
	Mean	Median	Mean	Median
Before	2.45	0	1.95	1
After	12.15	8.5	4.38	4

Our second question is whether K-net improved task performance. We did not find any difference in performance between the experimental group ($Mdn = 3.7$) and the control group ($Mdn = 3.7$), $U = 191, ns, r = -0.06$.

7. DISCUSSION

Our findings indicate that participants in the experimental condition had significantly higher awareness of ‘who knows what’ than participants in the control condition after K-net was deployed. If this increased awareness was due to the natural exposure of expertise that occurs through class instruction and interaction, then we would not see a significant difference between experimental and control conditions. K-net was thus successfully able to create awareness of ‘who knows what’ among participants, demonstrating that it supports the *directory updating* process required of TM systems.

We expected to see a difference in performance, indicating that K-net effectively supports *knowledge allocation* and *knowledge retrieval*, but did not. This may be because the range of final grades handed out was very narrow, making it hard to detect performance differences. It may be that most groups chose projects that took advantage of skills they had already developed, as suggested by the fact that users on average received recommendations for one skill. Or it may be that K-net is more useful for repeated interactions. In our study, groups only worked on one project. Perhaps the learning of ‘who knows what’ on a first project will pay off in increased performance on subsequent projects. It is hard to perform long-term system evaluations, especially in classroom settings, but such evaluations seem more likely to show the full effects of TM systems on organizational performance.

It is also likely that our measure of performance was too coarse. Turpin and Hersh found no significant difference in performance between participants that used a significantly more accurate IR engine than a normal one [15]. However, users of the less accurate engine had to work harder by issuing more queries. Future evaluations of K-net and other social matching systems might use performance metrics that consider required effort and satisfaction as well as the quality of the completed task.

8. CONCLUSION AND FUTURE DIRECTIONS

Social matching presents a set of interesting research problems, including “what are interesting uses of social matching?” [12]. K-net appears to be useful for supporting at least some processes in developing an organization’s transactive memory. Longer-term evaluations and better performance metrics will help to evaluate whether it can effectively augment an organization’s transactional memory and improve groups’ bottom-line performance. It is also still an open question how to use friendship in recommendation. It is likely that it is easier to get help from friends, but finding the right balance between friendship and expertise when making recommendations promises to be an interesting challenge. We hope K-net may evolve to become both a useful TM tool for organizations and a research platform for social matching.

9. REFERENCES

- [1] Ackerman, M. Augmenting organizational memory: A field study of answer garden. *ACM TOIS* 16, 30, (1994).
- [2] Brandon, D. and Hollingshead, A. Transactive Memory Systems in Organizations: Matching Tasks, Expertise, and People. *Org. Science*, 15, 6, (2004), 633-644.
- [3] Casciaro, T. and Lobo, M.S. Competent jerks, lovable fools and the formation of social networks. *Harvard Business Review*, (2005).
- [4] Clark, H. H. *Using language*. New York: Cambridge University Press, 1996.
- [5] Cross, R. & Cummings, J. Tie and network correlates of individual performance in knowledge-intensive work. *Academy of Management Journal*, 47, 6, (2004), 928-937.
- [6] Hansen, M. Knowledge networks: Explaining effective knowledge sharing in multiunit companies. *Org. Science*, 13, (2002), 232-248.
- [7] Hill, W. and Terveen, L. Using frequency-of-mention in public conversations for social filtering. In *Proc. CSCW 1996*, ACM Press (1996), 106–112.
- [8] Kautz, H., Selman, B., and Shah, M. ReferralWeb: Combining social networks and collaborative filtering. *Communications of the ACM*, 30, 3, (1997).
- [9] McDonald, D. and Ackerman, M. Expertise recommender: A flexible recommendation architecture. In *Proc. CSCW 2000*, ACM Press (2000), 231–240.
- [10] McKnight, D., Cummings, L., and Chervany, N. Initial trust formation in new organizational relationships. *Academy of Management Review*, 23, 3, (1998), 473–490.
- [11] Moreland, R., Argote, L., and Krishnan, R. Socially shared cognition and work: Transactive memory and group performance. In J.L. Nye and A.M. Brower (Eds.), *What’s social about social cognition? Research on socially shared cognition in small groups*, Sage, Thousand Oaks, CA, 1996.
- [12] Terveen, L. and McDonald, D. Social Matching: A Framework and Research Agenda. *ACM Transactions on Computer Human Interaction*, 12, 3, (2005), 401-434.
- [13] Terveen, L., Selfridge, P., and Long, M. Living design memory: Framework, implementation, lessons learned. *Human Computer Interaction*, 10, 1, (1995), 1–38.
- [14] Thorn, B. and Connolly, T. Discretionary Databases. *Communication Research*, 13, 5, (1987).
- [15] Turpin, A. and Hersh, W. Why batch and user evaluations do not give the same results. In *Proc. SIGIR 2001*, ACM Press (2001), 17–24.
- [16] Wegner, D. M. Transactive memory: A contemporary analysis of group mind. In B. Mullen & G. Goethals (Eds.), *Theories of Group Behavior*. Springer-Verlag, New York, 1987.
- [17] Wegner, D. M. A computer network model of human transactive memory. *Social Cognition*, 13, 3, (1995), 319-339.
- [18] Williamson, G., and Clark, M. Impact of desired relationship type on affective reactions to choosing and being required to help. *Personality and Social Psychology Bulletin*, 18, (1992), 10-18.