Mapping Knowledge Networks in Organizations: Creating a Knowledge Mapping Instrument

Wayne G. Lutters, Department of Information and Computer Science and Center for Research on Information Technology and Organizations, University of California, Irvine, lutters@ics.uci.edu

Mark S. Ackerman, Department of Information and Computer Science and Center for Research on Information Technology and Organizations, University of California, Irvine, ackerman@ics.uci.edu

James Boster, Department of Anthropology, University of Connecticut, boster@spf1n1.ucc.uconn.edu

David W. McDonald, Department of Information and Computer Science and Center for Research on Information Technology and Organizations, University of California, Irvine, dmcdonal@ics.uci.edu

Abstract

The ability to leverage organizational expertise is a critical success factor in most forms of knowledge work. However, expertise is an exceedingly difficult resource to manage. The design of computer-based support for knowledge management requires extensive, costly and inefficient cycles of knowledge elicitation to generate a reasonable knowledge map. We propose an alternative approximation technique which reduces these costs, while providing functionally equivalent data. In this methodological case study we chronicle the development of the key instrument in this approximation technique.

Keywords: knowledge acquisition; organizational learning; socio-technical approach; IS research methodologies; exploratory study

The Problem

Knowledge work is inherently information intensive and collaborative in nature. As such, it is heavily dependent upon the successful utilization of an organization's accumulated expertise. Clearly the management of this expertise is critical to the success of any knowledge intensive endeavor. However, contrary to a common perception in the popular business press, expertise is a complex, diverse and highly contextualized phenomenon (Suchman and Wynn, 1984). Intrinsically, it is an exceedingly difficult resource to manage.

Within knowledge intensive organizations, one of the most fundamental tasks is expertise location, "how does one locate others with relevant expertise for a problem at hand within an organization?" An information seeker most often finds someone with the required expertise through mutual associates, paper directories, communication technologies, or, more recently, computer-based recommendation systems (McDonald and Ackerman, 1998). These recommendation systems intend to augment the seeker's typical search strategies by including individuals outside of their immediate social environment. For these systems to be of significant assistance, however, they must reasonably reflect an understanding of the greater knowledge network within the organization. While people know at least local portions of the knowledge network intuitively, this knowledge must be built into computer-based recommendation systems.

Generating the requisite understanding to "feed" computer-based recommendation systems is a daunting task. (This is analogous to the well known problem of knowledge elicitation for the development of expert systems.) In general, one needs to inventory the organization's knowledge as well as to map the information flow within the organization. Common approaches to this have involved assessment interviews, skill inventories, and extensive surveys (Hoffman, 1995). Key limitations of these methods are their high cost and their tendency to significantly disrupt daily work. They also tend to collect only fairly flat, one-dimensional assessments of expertise and expertise topics. Most importantly, because of the dynamic nature of expertise networks, these assessments are nearly obsolete the moment they are collected and are very difficult to maintain over time.

Knowledge Mapping Approximations

Since the initial assessment of a knowledge network is unwieldy, the on-going maintenance prohibitively costly, and the results relatively superficial, we have examined a new approach. Our goal is not to completely capture the knowledge network; but instead, to use a reasonable first order approximation. As with any approximation technique in engineering, the motivation is to more easily compute and assess the phenomenon while staying within known error rates.

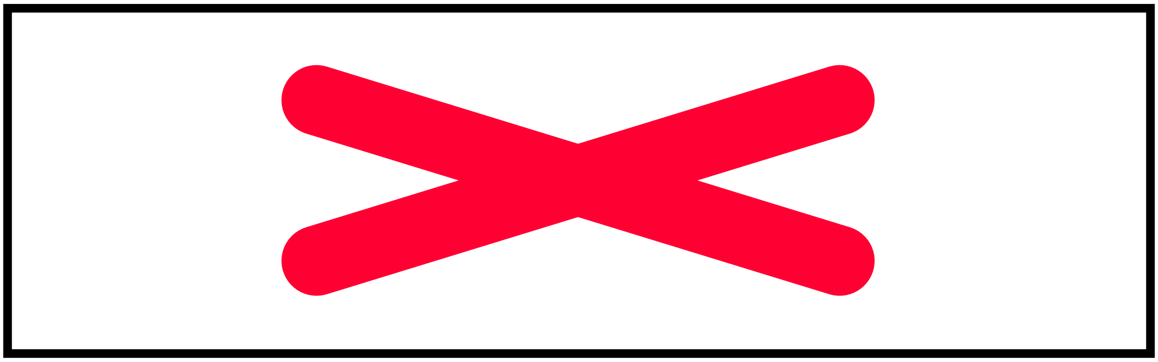

Figure 1. Putting the KMI in perspective, an overview of the knowledge mapping approximation project.

There are myriad possible knowledge mapping approximations; however, the fundamental goals of each should be that its initial data are easily collected (e.g., requiring no more than one hour of each employee's time), it is simple to maintain (e.g., via continuous capture of relevant digital artifacts), and its resultant measures correlate well with the understanding of expertise in the site itself (i.e., high face validity).

This paper describes the pilot study for one such approximation at a medium-sized software development company, Medical Software Corporation (MSC). The study involved three canonical classes of knowledge workers: software developers, technical writers and product support staff. It was performed in conjunction with a larger field study at the site which was examining expertise location behavior and developing an expertise recommendation system. (For an extended description of the site, study, and system, please refer to McDonald and Ackerman, 1998.)

As an exploratory study to find suitable approximations, we worked with three data collection techniques in whose intersection we anticipated finding our approximation. (These are shown in figure 1.) The first involved collecting social network data to augment organizational structure and working relationship data we had already collected. To do this we adapted the successive pile sort method to accommodate larger groups (Boster, 1986; Boster and Johnson, 1989). The second (further discussed below) was an effort to construct a Knowledge Mapping Instrument (KMI) to capture a rough "snapshot" of the current knowledge state of the group. This snapshot would yield an understanding of the expertise topology within the organization, a reasonable approximation of the organization's knowledge network. The third was an evaluation of what people knew about each other. For this we asked each participant to provide the anticipated KMI score for his or her colleagues, providing a rough ranking.

While each of the techniques yielded interesting results, the KMI was the critical instrument. In the following sections we will chronicle its development through the stages of design, elicitation, construction and validation. We will highlight specific lessons learned during the pilot study at MSC.

Knowledge Mapping Instrument

With the KMI, our goal was to have the organizational members tell us what they know best and what they think others around them ought to know. We need to do this for several important reasons. First, this expertise is likely to be highly contextualized to the specific needs of the organization (Orr, 1996). Since the content is provided by the participants themselves, it ensures a high degree of relevance and validity. Second, we wanted to avoid having to acquire domain expertise ourselves, and therefore we relied on organizational members to write the KMI itself. Findings from expert systems engineering and cognitive anthropology suggest how difficult it is to pull the "know-how that cannot be verbalized" from organizational members (Polanyi, 1967). A major component, then, of our research was examining how to obtain this requisite tacit "knowledge of knowledge" by having the organizational members assess expertise.

KMI - Elicitation

The knowledge management, decision support system and expert system literature detail knowledge elicitation techniques (e.g., Davenport and Prusak, 1998; Liou and Nunamaker, 1993). (Hoffman, 1995 provides an excellent review of the literature for these techniques.) All have significant, known limitations, especially in the cost of obtaining the original data for the inventory.

In response to the limitations of these standard techniques, the elicitation procedure for the KMI had to be lightweight, easy to understand, and highly contextualized. Regarding the first criterion, our goal was to take no more than 15 minutes per employee to accomplish this task. In consideration of the second, we wanted the task to be approachable, both so that the participant could feel comfortable contributing immediately and so that we would not waste valuable time with extensive instructions. Lastly, we needed a technique that would ground the KMI appropriately for a given organization, rather than flounder in overly general assessments.

With these three criteria in mind, we developed a knowledge elicitation technique. We needed a way to elicit what knowledge was important to the participants. We did this by telling them that they were helping to create a trivia game, akin to Trivial Pursuit. The end result would be the KMI described below. While appearing to be a trivia contest, and thus reducing the psychological cost to the participants, we would be able to assess, through the game, at least one critical success factor for the organization.

For the pilot study, the game addressed the on-going development, support, and use of MSC's flagship medical practice management system. Drawing from the three technical departments (development, documentation and support) we were able to recruit 35 participants.

We asked each participant in the study for help in generating questions and answers. Asking them to consider their co-workers as future players of the game, we requested questions that would vary from challenging to "only you would know the right answer." Our goal was to have the participants tell us what they think they are best at and what they think others ought to know about their work. (We were also hoping to elicit questions that showed significant differentiation in expertise among the three groups.)

In general, the prompting metaphor of constructing a trivia game provided a meaningful frame of reference for the elicitation process, and it motivated participants. However, we did need to develop prompting aids (described below), and the elicitation task was still difficult for many participants.

Following are some important lessons we learned from this procedure:

- We were able to obtain with reasonably minimal effort nearly 70 questions. This was adequate to create a reliable and valid instrument. Each elicitation interview took approximately 15 minutes. The entire collection effort required eight business days of researcher time.

  Moreover, in asking for questions, we located additional organizational resources at MSC that could generate even more questions. These resources included system documentation, questions within training manuals, and questions generated for user group meetings. We did not use these resources, however, in our pilot study, preferring to test the KMI elicitation process alone.

- Generating trivia question and answer sets is a challenging, abstract cognitive task for anyone. It proved demanding for our participants.

  We asked each participant to produce three trivia question and answer cards. Twenty-one of the 35 participants were able to generate at least one usable trivia card and seven readily supplied two to three times the amount requested.

  Participants did not have equal facility in composing three good multiple choice questions. For example, writing multiple wrong answers (so that they are clearly wrong but not obviously wrong) is quite difficult. This was compounded by time and social pressures (fifteen minutes with the researcher audio recording the process).

- As we began collecting data, we found that we needed to create a formalized elicitation interview and a set of prompting props. We developed sample trivia cards with slots for the question and the multiple choice answers – one correct answer and four challenging incorrect answers. After freeform thought about a particular question we used this form to drive the interaction, ensuring a completed trivia card at the end. The prompting props, such as the trivia cards, were helpful in structuring the data collection and ensuring that participants generated complete sets of answers.

  In retrospect, we had asked people to generate questions in a decontextualized setting (i.e., in a conference room instead of their office). This provided limited environmental cues to prompt question generation and limited local resources to verify the correctness of answers.

- We needed to have a researcher present during the elicitation process. Attempts at having participants generate questions on their own, replying either by e-mail or in person on our next visit, were futile. As co-present researchers, we could maintain motivation, prompt in the case of partial responses, and provide supportive feedback for the iterative design of questions.

- We attempted to get participants to rank the difficulty of their questions. Almost all were unable to do so, noting that they could not rank in the absence of a specific task or referent group. That is, they saw questions as difficult only in relation to specific circumstances – difficult for support but not development, or easy for people who had carried out specific system tasks.

KMI - Construction

Once the elicitation process was complete, we transcribed the question and answer cards, relying on audio recordings for the clarifying nuances. We then collected the cards into an aggregate set sorted by participant. While every card had to be edited from the original, most of these changes were minor, involving a simple rewording or the addition of other incorrect answers (such as "A & B" or "none of the above").

We validated this aggregate set of questions using a three-step procedure. Although our goal is to develop a method that will provide results without needing domain expertise; nonetheless, for our initial pilot study, we had to understand the quality of the questions we received. Therefore, initial validation was performed by the project researcher who spent a total of 18 months observing expertise and knowledge processes at MSC (McDonald and Ackerman, 1998). He reviewed and categorized each card according to the following criteria – knowledge domain (areas of specialization such as users, developers, support, and system administration), perceived difficulty (on a five point scale), and clarity ('clear,' 'ambiguous,' or 'does not make sense').

After repairing the questions to the best of our ability and removing any questions that were too similar, there remained 22 questions (from 13 participants) that were not 'clear.' For each of these we returned to the participants for further clarification or expansion. Of the 22, 14 were revised, six were removed, one was replaced and one remained unchanged. This clarification occurred after a two-month hiatus.

An interesting observation from these return visits is that most participants did not recognize their own questions, suggesting that the material may seem relatively fresh if it is presented to subjects with some time delay.

The questions, after being randomized to evenly spread the topic domains and difficulty levels throughout the instrument, were then handed to two test participants at MSC. These test subjects found the instrument clear, easy to take, challenging in content and, most importantly, enjoyably engaging. Most of the question-answer sets were acceptable as they were or required only minor refinement.

KMI - Validation

At the completion of the construction phase we had selected 58 well-formed question and answer sets for the final version of the KMI. We then administered this to 26 subjects.[1] The data gathered were from a majority of all three departments, as well as key management and technical members. Participation was voluntary, occurring over the lunch hour in small groups of two to eight subjects.

At the conclusion of each session we asked subjects for feedback on the KMI, particularly if they had any problems with any particular questions. In addition, after all data had been collected, we asked technical experts to evaluate our answer key. Through both of these methods we found some additional problems:

- Some questions had minor wording errors (e.g., the name of a program was not FINANBAL, but FINANCBAL). In all cases, it was clear from the question what was meant, and the key was not changed.

- One question asked about an organizational process that involved an employee that no longer filled that role. However, the role had not been given to another person, and the key was not changed.

  In both of these situations, subjects did not seem confused.

- Five questions touched on differences between acceptable organizational practice (i.e., work-arounds and alternative methods of acting) and official practice (Suchman, 1983). In these cases, the correct answer was ambiguous, since subjects could have interpreted either practice as "best." In one of these questions, the acceptable organizational practice was largely historical – the company legend behind a name – but management officially did not recognize that legend. In all of these situations, we allowed multiple correct responses on the answer key.

  In these questions, the occasional alternative answer "none of the above" was the most problematic. Since these questions were not precise, "none of the above" also became a potential answer.

- Two other questions, which were technical, had multiple correct answers. For both of these, there was an additional, obscure way to produce the desired technical result. For these two questions, we allowed multiple correct responses on the answer key.

- One question was removed. Between the time that the KMI was developed and administered, the

[1] While we gained two new participants since the start of the project, we were unable to obtain KMI results from nine of our original 35 due to staffing changes and general unavailability.

referent had changed. The question had involved a data entry procedure based on a form external to the organization. The form had changed four months prior to the administration and there was significant confusion as to which version of the form, old or new, the question referred.

As with the test participants, people found the KMI engaging, challenging and even fun. The average time to completion was just over 30 minutes. The objective scores were midrange (mean = 35.19 out of a possible 57), and there was reasonable variance (s.d. = 9.85). Interestingly, there was significant agreement among participants on the answers, even if they were incorrect.

In evaluating the responses, we found no significant difference in scores among the three organizational groups of technical developers, writers or support. The KMI as developed captured and measured general knowledge well, but did not capture group-specific knowledge. We need to determine whether this was the result of the elicitation technique (e.g., the wording on question capture) or the specifics of this site (e.g., the size of the company).

To examine the robustness of the KMI we ran a sensitivity analysis on scores with and without the problematic questions noted above. This basically examined whether there were any important problems uncovered during the validation. (The rejected question was not included.) We found no statistically significant difference ($p < .01$). This suggests that after due diligence in the construction phase, a small number of ambiguous questions (nearly 15% in this instance) can be tolerated. The KMI, as developed in our pilot study, appears to be a very robust instrument.

Conclusions

We have argued the value of knowledge mapping approximations as a profitable alternative to standard knowledge assessment techniques. Through a pilot study, we have shown the development of a critical instrument for one such approximation. The low cost and highly contextualized KMI allowed organizational members to highlight what knowledge needs to be mapped. Indeed, the study suggested that one may be able to approximate the knowledge map for an entire group by interacting with only a handful of key informants.

Our results warrant further research. Directions for future work include developing, administering and evaluating a KMI in a larger organization, as well as redesigning the elicitation technique to improve differentiation of group expertise in the KMI. There should also be further exploration of alternative knowledge mapping approximations.

Acknowledgements

This research has been supported by grants from the CISE/IIS/CSS Division of the U.S. National Science Foundation and the NSF Industry/University Cooperative Research Center (CISE/EEC) to the Center for Research on Information Technology and Organizations (CRITO) at the University of California, Irvine. Industry sponsors include: ATL Products, the Boeing Company, Bristol-Myers Squibb, Canon Information Systems, IBM, Nortel Networks, Rockwell International, Microsoft, Seagate Technology, Sun Microsystems, and Systems Management Specialists (SMS).

References

Boster, James S., and Johnson, Jeffrey C. "Form or Function: A Comparison of Expert and Novice Judgements of Similarity Among Fish," *American Anthropologist* (91:4), December 1989, pp. 866-889.

Boster, James S. "Exchange of Varieties and Information Between Aguaruna Manioc Cultivators," *American Anthropologist* (88:2), 1986, pp. 428-436.

Davenport, Thomas H. and Prusak, Laurence. *Working Knowledge: How Organizations Manage What They Know*, Harvard Business School Press, Boston, MA, 1998.

Hoffman, Robert R., Shadbolt, Nigel R., Burton, A Mike, and Klein, Gary. "Eliciting Knowledge from Experts: A Methodological Analysis," *Organizational Behavior and Human Decision Processes* (62:2), May 1995, pp. 129-158.

Liou, Yihwa Irene, and Nunamaker, Jay F. Jr. "An Investigation into Knowledge Acquisition Using a Group Decision Support System," *Information & Management* (24:3), Mar 1993, pp. 121-132.

McDonald, D. W., and Ackerman, Mark S. Ackerman. "Just Talk to Me: A Field Study of Expertise Location," in *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (CSCW '98), 1998, pp. 315-324.

Orr, Julian E. *Talking About Machines: an Ethnography of a Modern Job*, ILR/Cornell University Press, Ithaca, NY, 1996.

Polanyi, Michael. *The Tacit Dimension*, 1967, London, UK, Routledge and Kegan Paul.

Suchman, Lucy A. "Office Procedure as Practical Action: Models of Work and System Design," *ACM Transactions on Office Information Systems* (1:4), October 1983, pp. 320-328.

Suchman, Lucy A. and Wynn, Eleanor. "Procedures and Problems in the Office," *Office: Technology and People*(2), pp. 133-154.