

6 References

Boster, James S. 1987. Agreement between Biological Classification Systems is Not Dependent on Cultural Transmission. *American Anthropologist*. 89(4):914-919.

Boster, James S. 1994. The Successive Pile Sort. *Cultural Anthropology Methods* 6(2):7-8.

assessing that knowledge through the evaluations of others (the Social Knowledge Evaluation).

Several important issues remain for future investigation. For use as expertise locators in companies similar to MSC, it may be that it is not worthwhile to develop a full-fledged KMI to directly assess group member's knowledge. The KMI in this pilot study was closely approximated ($r = .85$) by the Social Evaluation. In this organization, people had a good understanding of what other people know - at the least, they had a good understanding of other people's expertise levels for the kinds of knowledge that were valuable to MSC. The fact that the 3 top managers or 7 lower-rank engineers could provide similar evaluations provides another approximation technique that could be valuable in future tests. It may well be that in large organizations, we could collect essentially the same data using far fewer people than we did here.

Second, to help people within similar-sized organizations beyond what they already know, effective routing of questions may require more detail. That is, if people already know a fair amount about one another, what they lack when they are looking for someone with the required expertise is detailed knowledge of what others know - what they need to find answers to specific questions. It may be that the size of this organization may be at the maximum size one might know everyone else's expertise, since, with about 50 people, it is a very large small group.

If this is the case, the KMI approximation may be more suited for larger organizations, as might be expected. While we needed to pilot these approximations within a medium-sized company (in order to have a tractable field study), these methods, particularly the KMI, may offer more substantial help when the organization is large and heterogeneous. We can expect that group member's ability to adequately rate one another will not hold across larger organizations. In larger organizations, one person would not be able to personally evaluate more than a fraction of the other people. In fact, those other people would likely be within the same part of the organization, and therefore know many of the same areas of expertise. In a very large organization (for example, a trans-national company), a person might not even know what groups have expertise within a given area. Still, by systematically aggregating these local evaluations one might be able to develop a robust global expertise locator.

Finally, the technique used in this pilot study cannot distinguish whether participants were rating one another based on the details of their anticipated expertise or their overall sense of expertise. That is, it is possible that raters, especially in a group this small, knew one another's areas of expertise and rated people precisely. Alternatively, group members may have had a sense how good someone technically is without knowing the details of their expertise. From experience on projects, comments in the hallways, and random conversations over lunch and near water coolers, one may come to feel that someone is technically solid without knowing the exact areas of his or her expertise. More work will be required to understand how people were rating one another. In sum, the Knowledge Mapping Approximation Project described in this chapter developed and piloted two promising classes of approximations for finding the data necessary to drive expertise finder systems. These approximations, both require additional refinement, but offer promise for helping solve the data collection problem for sharing and managing expertise.

uniformly those with long experience at the organization, while the over-achievers (those who scored higher on the KMI than expected) are those with relatively short experience at the organization.

Managers show slightly higher agreement in the social evaluation of their employees than do non-managers, perhaps because they have been in the organization for a longer period of time or because their job responsibilities require them to assess the expertise of their employees. Given the strong expectation that managers should be exceptionally skilled at social evaluation, perhaps the more interesting result it is that the difference is relatively small. In fact, aggregating the evaluations of seven engineers provides as robust an estimator of expertise as aggregating the evaluations of the three top managers.

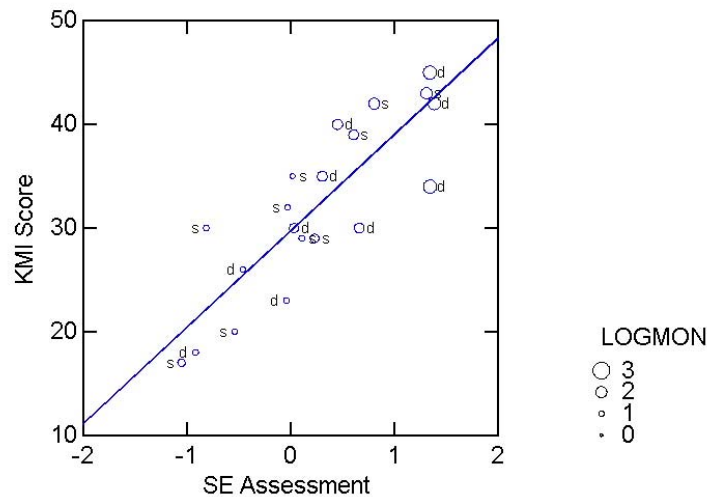


Figure 6: Scores on the Knowledge Mapping Inventory and Social Knowledge Evaluation instruments

5 Conclusions and Future Work

Our prototype approximations, as developed in the pilot study reported here show considerable promise: We were able to obtain a rough estimate of how much each participant knew and also of how others assessed group member’s knowledge. The fact that the instrument was essentially developed by the participants themselves and were rated by each other suggests that our approximation not only captured the critical aspects of the distribution of knowledge in this organization, it also may serve as a prototype for a valuable class of approximation techniques. We were able to uncover a knowledge map, with some limitations at limited expense – this is precisely the beginning of finding adequate approximations. Indeed, we believe that we developed two expertise approximations in this pilot study; one based on directly assessing the knowledge of the group members (the Knowledge Mapping Inventory) and the second based on indirectly

correlated with the duration of employment at the organization ($r = .44, p < .03$), there is no tendency for support people to perform better than development people on the KMI or vice versa, and there is much greater variation in the KMI scores of individuals who have only been at the organization a short time than those who have been there longer. We do not know whether this reduction is due to individuals with long experience at the organization becoming ‘saturated’ (reaching the asymptote of the learning curve) or whether it reflects the selective retention of those who are quick learners and the letting go of those who are slower.




Figure 4 (left): Knowledge agreement among groups (both Development groups aggregated)

Figure 5 (right): KMI score ranked by experience

The social evaluation scores do *not* show a similar pattern – individuals who have been at the organization longer do not agree on the knowledge of others any more than do recent arrivals. There is no relationship between agreement on the social evaluation of others and length of time at the organization ($r = .08, p > .6$). Although the regression line has a positive slope, it is not significant.

It is the case that individuals who have been at the organization longer are presumed to have greater knowledge than the recent arrivals ($r = .5, p < .005$). It is likely that people are using the length of employment of others as a heuristic for guessing other’s expertise as reported in McDonald and Ackerman 1998. Figure 6 shows the relationship between performance on the KMI and perceived knowledge by other members of the group.

Nonetheless, use of this “experience heuristic” alone does not explain how participants evaluate each other’s knowledge, for, in the aggregate, they are much better at it than would be explained by use of the only the experience heuristic. In general, the social evaluation of others very accurately predicts their KMI score; the correlation is .85 ($p < .001$), much higher than the .5 correlation of social evaluation with duration of employment. Nevertheless, use of the expertise heuristic is evident in an examination of the outliers; the individuals who score lower on the KMI than they were expected to are

pile sort (i.e., as an ultrametric). Again, there are three main clusters, corresponding to the three work groups at the organization. (One can think of this as an inductive way of eliciting the organization chart.)

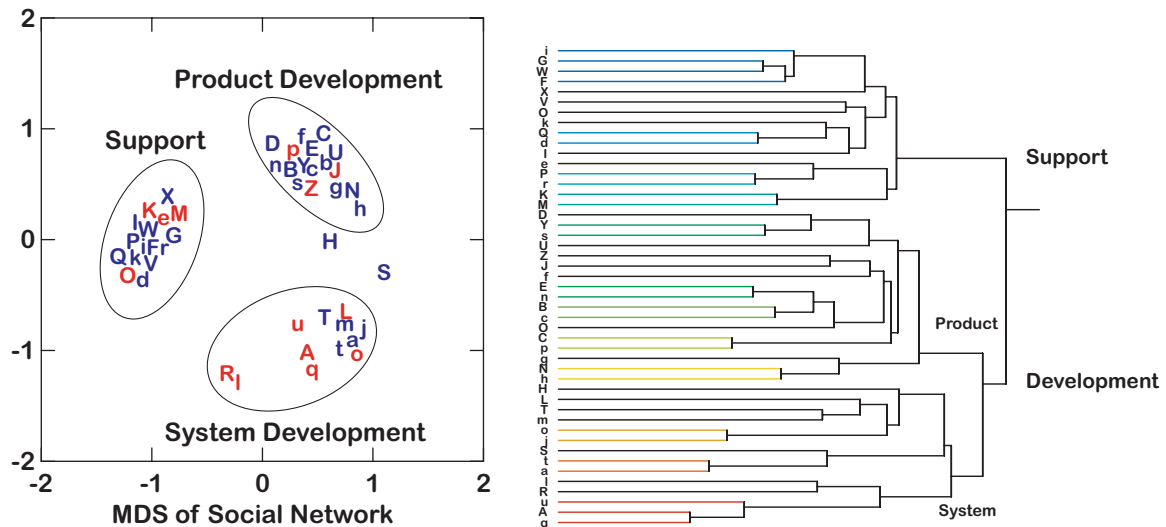


Figure 2 (left): The social network structure in MSC.

Figure 3 (right): The social network structure in MSC (hierarchical cluster diagram)

4.2 KMI

We had originally expected that the pattern of agreement on the KMI would show the division of intellectual labor in the organization – that by and large, the support team would know a great deal about support and little about development, while the development teams would know a great deal about development but little about support issues.

Instead, we found that expertise was spread fairly evenly through all groups, and was determined not so much by group membership as by the number of months that the individual had been in the organization. This may reflect a situation in which success in the organization depends on integrative knowledge – individuals cannot afford to become too specialized and lose sight of what other members of the organization know and the problems they need to solve. Figure 4 shows the relationship between knowledge of development issues on the KMI versus knowledge of support issues.

In general, support knowledge and development knowledge at MSC are highly correlated ($r = .75, p < .001$). However, the outliers make sense in terms of work group membership. Although in general knowledge of the two domains go hand in hand, if there is a marked difference in an individual's knowledge of support and development issues, it reflects his or her work group membership. Figure 5 shows the relationship between overall score on the KMI and the number of months at the organization. (We used a log transformation of the number of months, given that there appears to be a learning curve -- learning is rapid at the beginning of one's employment in the organization and slower latter.) This figure illustrates three findings: performance on the KMI is moderately

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

In both of these situations, participants 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 participants could have interpreted either practice as “best.” In all of these situations, we allowed multiple correct responses on the answer key.
- Two other technical questions also had more than one correct answer. In both of these cases, there existed an obscure way to produce the desired result in addition to the usual way. For these two questions as well, we allowed multiple correct responses on the answer key.
- One question that concerned a data entry procedure based on a form external to the organization was removed. The form had been changed four months before to the administration of the instrument and participants were uncertain as to which version of the form the question referred.

As with the test participants, the rest of the participants found the KMI engaging, challenging, and enjoyable. The average time to completion was just over 30 minutes (longer than our target but not prohibitive). Participants scored a mean of 35.19 out of a possible 57 with substantial variation among participants (s.d. = 9.85). Interestingly, there was significant agreement among participants on the answers even if they were incorrect.

To examine the robustness of the KMI, we compared participants’ scores with and without the problematic questions noted above. (The rejected question was not included.) We found no statistically significant difference in participants’ scores. 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.

4 Findings from the KMA Pilot Study

4.1 Social Network Structure

The social network structure of the group is shown in Figure 2. There are three fairly distinct clusters: the support group, the system development group and the product development group. Two individuals (H and S) are recent arrivals and are not clearly identified with any of the groups. Thus, the overall structure of the group clearly reflects the working groups within the organization: network ties are closer among the members of the support, system and product groups than they are between members of different groups.

This structure can also be represented as a hierarchical cluster diagram, as shown in Figure 3. This is perhaps a more appropriate representation of the data, because a tree diagram more nearly corresponds to the way that the data were collected in the successive

received. Therefore, initial validation was performed by the fourth author, who had 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 sufficiently '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 participants with some time delay.

After randomizing the order of the questions to evenly spread the topic domains and difficulty levels throughout the instrument, we administered the instrument to two test participants at MSC. They found the instrument clear, easy to take, challenging in content, and, most importantly, enjoyable and engaging. Most of the question-answer sets were acceptable as they were or required only minor refinement.

3.5 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 participants.³ 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 participants.

At the conclusion of each session, we asked participants for feedback on the KMI, particularly whether they had any problems with any specific 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.

³ 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.

Following are some important lessons we learned from this procedure:

- We were able to obtain with modest effort nearly 70 questions. Each elicitation interview took approximately 15 minutes. The entire collection effort required eight days of interviews.

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

- While less tasking than standard knowledge elicitation techniques, the elicitation of trivia questions was nonetheless demanding for our participants. 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). We asked each participant to produce three trivia questions together with possible answers. Twenty-one of the 35 participants gave us at least one usable trivia question and seven supplied two to three times the number of questions requested.
- As we began collecting data, we found that we needed to create a formalized interview and a form to prompt participants. We developed sample trivia questions with slots for the question and the multiple-choice answers – one correct answer and four challenging incorrect answers. After brainstorming about a particular question, we used this form to prompt the participant, ensuring a completed trivia question and answers at the end.
- In retrospect, we realize we made a mistake in asking participants to generate the trivia questions in a conference room instead of their offices. This decontextualized setting 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. When we were present, we could maintain motivation, prompt in the case of partial responses, and provide supportive feedback for the iterative improvement 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 Product Development, or easy for people who had carried out specific system tasks.

3.4 KMI - Construction

We validated the aggregate set of questions using a three-step procedure. Although our goal was 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

3.3 KMI - Elicitation

The knowledge management, decision support system and expert system literatures all describe what is often a very difficult process of eliciting knowledge from people (e.g., Davenport and Prusak, 1998; Liou and Nunamaker, 1993). Hoffman (1995**) provides an excellent review of the literature for the various elicitation techniques. All of the techniques have significant, known limitations, especially in the cost of obtaining the original data for an inventory.

With the KMI, we wanted to avoid formally eliciting the knowledge important for members of the group. Our goal, instead, was to have the group members tell us what they know and what they think others around them ought to know. This was important for two reasons: First, what counts as expertise depends on the specific needs of the organization (Orr, 1996). Since the content of the KMI is provided by the participants themselves, we believe we are assured of an instrument with a high degree of relevance and validity. Second, we wanted to avoid acquiring domain expertise ourselves in order to write the instrument because this would have been too costly for a rough approximation. We reduced the cost of pulling out the “know-how that cannot be verbalized” from group members (Polanyi, 1967) by having group members write the KMI themselves.

Thus, an important part of our research was figuring out the best way to prompt group members to assess expertise in order to obtain the requisite tacit “knowledge of knowledge.” We wanted an elicitation procedure for the KMI that was brief (taking no more than 15 minutes of the participant’s time), easy to understand, and specific to the knowledge that was important to group members. With these three criteria in mind, we elicited what knowledge was important to the participants by telling them that they were helping to create a trivia game, similar to Trivial Pursuit. The result was the KMI described below. By framing the instrument for assessing the distribution of knowledge through the group as a trivia contest, we were able to reduce the psychological cost to the participants. For this pilot study, the trivia game assessed participants’ knowledge of 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 recruited 35 participants from the three MSC departments.

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 mildly challenging through very difficult to virtually impossible (questions which “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.

- While there is a constant hum of improvements, bug fixes, and new features, the basic product architecture, underlying feature set, and a substantial code base has been stable for over five years
- Only a handful of technical architects in the company can understand the entire product line and know most of the code base. Other software engineers and support engineers know only bits and pieces of the systems.

The participants in this pilot study worked in three departments: Product Development (the software engineers and administrative personnel responsible for the MSC products), System Development (the software engineers and administrative personnel responsible for supporting the various system platforms including networking), and Technical Support (the support engineers and administrative personnel responsible for solving customers' problems). Each department included entry-level, senior, and management employees.

3 The KMA Data Collection

In the following sections, we will chronicle its development through the stages of design, elicitation, construction and validation of the KMA instruments. These sections will highlight specific lessons learned in our pilot study at MSC about our prototype approximations.

3.1 Social Network Structure

The structure of the social network was elicited with the successive pile sort (Boster, 1987, 1994). In this task, the names of all the members of the social group are written on cards. A participant is asked to sort the members into as many subgroups as they like, based on the participant's judgment of the intensity of interaction of group members. After this initial free pile sort, the participants are asked to successively merge their groups, and then split the groups. The complete order of merges and splits is recorded. The resulting data can be represented as a binary tree expressing the relative social proximity of all members of the group. This method has the advantage of eliciting members' judgments of all the interactions in the group (rather than just those involving ego) and of rank-ordering their intensity. Although the successive pile sort technique is generally used on small groups (that is, groups with 25 members or less), it was applied in this study to a group with 43 members.

3.2 Knowledge Mapping Instrument

To apply the Knowledge Mapping Instrument (KMI), it is necessary first to discover what is important for members of the group to know (*elicitation*), next, to develop an instrument that tests for that knowledge (*construction*), then to assess whether the constructed instrument succeeds at its purpose (*validation*), and finally to document which members of the group know those things (*administration*).

2.1 Pilot Site

This pilot study was conducted at a medium-sized software development company, Medical Software Corporation (MSC)², a company that builds, sells and supports medical and dental practice management software. MSC is a medium-sized company, with just over 100 employees. The pilot study was performed in conjunction with a larger field study at the site that was examining expertise location behavior and developing an expertise recommendation system (McDonald and Ackerman, 1998.)

A short description of MSC's business and products is critical to understanding the types of knowledge that are important to the organization and the employees. MSC has been in the practice management business for over 20 years. Practice management software is primarily concerned with appointment scheduling, treatment planning, patient recalls, insurance billing, patient billing, and payment reconciliation. While these functions are closely related to the clinical management (treatment) of the patient, practice management is often considered the business side of medicine. MSC is the dominant player in the group medical and group dental markets.

MSC sells several software packages, ranging from text-only systems running on high performance Unix servers (largely used by large medical practices) to graphically-oriented Windows systems (largely used by smaller medical and dental practices). The MSC text-oriented system is older and written in a proprietary form of BASIC. Thus, most software engineers must acquire knowledge of fairly arcane details of the MSC systems, the language, and the system platforms. For example, the top level menu of the medical program has over 100 specifically identifiable features, and often there are sub-menus. Moreover, the features are spread through more than 200 medical system specific programs, with numerous programs that are shared by the medical and dental system. Moreover, the standard application is highly customizable and many files have client specific customizations. Customization is often handled with a large number of customization flags that are not always mutually compatible.

The software is in a relatively constant cycle of new feature development and maintenance. New feature development is often the result of requests from existing clients. When enough new features have been added to the software, they are integrated into a complete system and that new system is declared the next version. This new version is then offered to any client for minimal cost. A client is never forced to upgrade to a new version of the software. This development strategy means that many different versions with different customizations are in the field and supported all at the same time.

While there are many idiosyncrasies in the MSC products, MSC is representative of many technical organizations:

²All names and identifiers have been changed in this paper.

approximation technique in engineering, the motivation is to more easily compute and assess the phenomenon while staying within known error rates. For example, the famous bin packing problem in computer science is NP-complete: One cannot determine the optimal placement of a substantial number of three-dimensional items in a pack or two-dimensional items to be cut from a metal sheet in computable time. For real applications, one must use an approximation. The approximation will not provide the optimal solution, but it will provide a satisficing solution within a known tolerance.

Our goal in the Knowledge Mapping Approximation project has been to frame the collection and maintenance of expertise data as a critical research problem. There are probably a myriad of possible ways of mapping the distribution of knowledge in an organization; here, we only wanted to establish the possibility of a rough approximation. However, any means of approximation had to fit within important organizational constraints: the initial data should be easily collected (e.g., requiring no more than one hour of each employee’s time), the database should be simple to maintain (e.g., via continuous capture of relevant digital artifacts), and its resultant measures should correlate well with the understanding of expertise in the site itself (i.e., high face validity).

With these constraints in mind, we designed three data collection techniques that together would generate our rough approximation (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 used the successive pile sort (Boster, 1986; Boster and Johnson, 1989). The second was the construction and administration of a Knowledge Mapping Instrument (KMI) to produce a “snapshot” of the current distribution of knowledge within the group. This snapshot would yield both an understanding of the location of expertise in the organization and a sense of the flow of information through the organization’s knowledge network. The third was a survey of group members’ evaluations of each other’s levels of expertise. For this we asked each participant to guess the likely KMI scores of his or her colleagues.

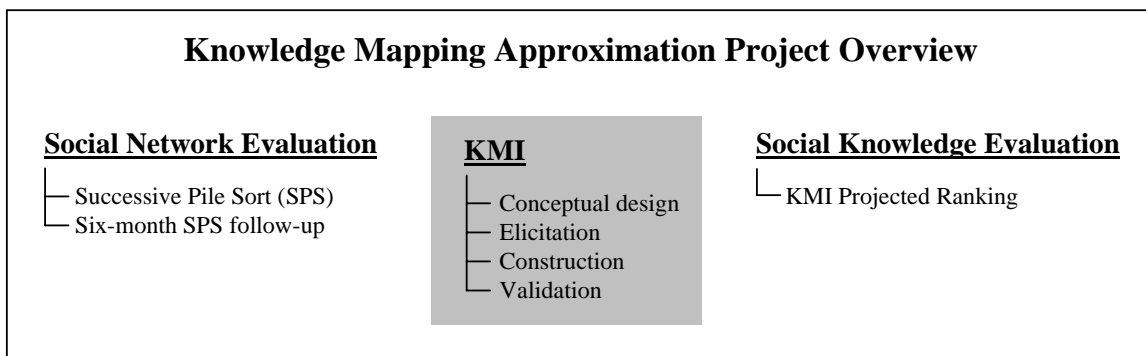


Figure 1. An overview of the knowledge mapping approximation instruments.

seeker's typical search strategies by including individuals outside of the immediate social environment or daily experience.¹

For these expertise finder systems to be of significant assistance, however, they must effectively point at the relevant people for any given problem. Therefore, the expertise finders must reasonably reflect an understanding of the greater knowledge network within the organization. While people are adept at knowing at least local portions of the knowledge network, this knowledge must be built into computer-based recommendation systems.

Trying to provide these data is directly analogous to the well-established problem of knowledge elicitation for the development of expert systems. We call our problem *expertise mapping*, and the required effort to be expertise or knowledge *elicitation*. In expertise mapping, 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 (as measured in time for the organization's members) 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. Maintenance of the data over time becomes a critical issue.

In this chapter, we report on the Knowledge Mapping Approximation (KMA) Project, which concentrates on this problem of providing systems with the type of data needed to adequately determine the people most likely to be able to answer a given question. While it is relatively common to consider system prototypes to find others, less research has been pointed towards finding adequate data for these expertise finders. As Ehrlich (**this volume) points out, generating the requisite data to "feed" computer-based recommendation systems is a daunting task. Because of the difficulty of this problem, we have bracketed off other significant issues, such as how people understand the context surrounding finding someone to help. (This issue is addressed by Fitzgerald in **this volume). We focus here only on the first steps in finding adequate data; this chapter discusses the initial steps in the KMA Project.

2 The KMA Project: Looking for Approximation Techniques

Since the initial assessment of a knowledge network is unwieldy, the on-going maintenance prohibitively costly, and the results relatively superficial, the KMA Project has examined a new approach. Our goal has been to not completely capture the knowledge network; but instead, to find reasonable first order approximations. As with any

¹ We use the term expertise instead of expert because these finders may wish to identify a range of expertise. For example, Paul in the above example might need someone with substantial knowledge of the software systems while someone with more superficial knowledge may help Sarah.

Who's There? The Knowledge Mapping Approximation Project

Mark S. Ackerman¹, James S. Boster², Wayne Lutters³, David W. McDonald⁴

¹Laboratory for Computer Science, Massachusetts Institute of Technology

²Department of Anthropology, University of Connecticut

³Department of Information and Computer Science, University of California, Irvine

⁴Fuji-Xerox Palo Alto Laboratory

word count: 5594

1 Introduction

Within knowledge intensive organizations, one of the most fundamental tasks is expertise location, or how one locates others with relevant expertise for a problem at hand within an organization.

- Sarah, a new employee, wants to find someone who can help her ship a computer to an international tradeshow. She wants to know what shippers are best for this, and what problems she may encounter. She has asked around her group, but everyone left this to the previous assistant. Sarah would like to find another assistant in the company who can help.
- Paul wants to find someone that has used JavaSound before on a new mobile platform. The software does not work properly, and he needs help. Paul is unable to tell whether his problem is because he does not understand how to use the software package, because the package has bugs, or because the package is a new release and conflicts with some of his existing system software.

Both of these people will likely find the answer through one of a small number of ways. If the person cannot find the answer in printed documents, an information seeker may search for someone with the required expertise through mutual associates and gatekeepers [Allen, 1977 #199], paper directories and references, communication technologies [Sproull, 1991 #347] [Finholt, 1993 #486] [Constant, 1994 #550], or, more recently, computer-based systems.

Expertise finders, or expertise recommenders, are a form of recommendation system (e.g. [Resnick, 1994 #399] [Konstan, 1997 #574]). Recommendation systems typically point the user towards books, movies, or other objects of interest, but expertise finders point people to other people. An example is Expertise Recommender (ER) [McDonald, 2000 #645]). It attempts to point the user towards people who know the most about a software module under repair. In general, an expertise finder's intention is to augment the