

Augmenting Expertise Networks through Recommendations

Mark S. Ackerman

Information and Computer Science
University of California, Irvine

I always went straight to [the system] at this point, instead of trying to guess someone who would likely know the answer.... A great deal of time is saved when the questioner does not have to play email-tag to find a person who knows the answer. -- From Answer Garden field data (Ackerman 1993)

Expertise networks are specializations of an organization's social network. They consider not only how people are socially arranged but also what expertise they have and trade. By examining expertise networks, we can investigate how organizational members find other people to answer their questions and obtain the information needed for completing their work activities.

Since knowledge and information flows along these expertise networks, augmentations of the expertise networks could improve organizational knowledge and expertise. Many augmentations are possible. For example, one could facilitate knowledge transfer along the existing social network by providing new communication media or spaces. One could construct systems to bring together new expertise networks, either by finding those interested in a particular topic or by constructing ad-hoc teams with the required knowledge. One could even imagine ad-hoc expertise networks, where agents located people with the required knowledge on an on-going basis.

For the last two years, we have been focusing on suitable augmentations of these expertise networks by either creating new clusters of expertise to solve an immediate problem or by appropriately routing inquiries as needed. Our group is conducting three studies in this area:


1. **Augmenting Expertise Networks** (David McDonald and Mark Ackerman)

This project explores the problem of finding expertise in an organization. The project consists of two steps, the first being a field study of how people seek expertise and the second being the construction of an expertise recommendation system.

The field study examined expertise location in a medium sized software development company (called here MSC) over a six-month period. The study demonstrated that expertise is fundamentally a collaborative activity. Participants employed varying cognitive and social evaluations when searching for expertise to solve problems. The study separated expertise seeking into identifying and selecting behaviors. Identification strategies provided potential candidates for a query, and participants used a variety of extremely situated methods of determining those candidates. Selection strategies helped pick which candidate or candidates were the most appropriate based on workload, psychological, and other considerations. As well, the study showed a number of breakdown strategies including escalation of requests. (Interestingly, escalation did not merely ascend organizational or functional hierarchies; it could also spread across organizational and functional boundaries, drawing in more participants as required.) The details of this field study were presented in a CSCW'98 paper, "Just Talk to Me: A Field Study of Expertise Location."

The results from the field study were then used to inform the design of an expertise recommendation system. The field study showed that a number of identification and selection strategies are heavily situated, specific to MSC and even particular portions of MSC. For example, technical support personnel use the call tracking system to determine others' workload. Development, on the other hand, looks at the "last modified" line in the code to determine who might have the most expertise. A long-standing problem in CSCW system design is how to support situated activities, while providing for generalizability at the same time. This portion of the project provides for a general architecture that allows a range of expertise-seeking strategies. This allows the system to employ heuristics and information sources similar to those used by the field study participants, while also providing for general heuristics that might be found in almost any organization.

The Expertise Recommender (ER) system, then, provides an architecture to incorporate a set of heuristic-based modules that implement identification and selection mechanisms. The architecture of ER enables a designer to describe new heuristics or reuse existing heuristics to solve new expertise identification or selection problems.



Identification is performed in two phases. In the first phase each agent looks to identify everyone who may have knowledge of the specific domain. Parts of this phase may be performed at regular intervals, background to other processes of the system. In the second phase a user makes a request in one or more of the domains supported by the system. An agent applies heuristics specific to the topic and returns a list of most likely candidates. For example, using the "last modified" line in program code and using logging codes in the technical support database are heuristics useful for MSC.

After identification, the list of candidates is passed to ER's selection mechanisms. The selection mechanisms select and filter using techniques similar to the selection techniques discovered in the field study. For example, the selection mechanisms include social filtering modules. As well, selection filters implement the "keep it local" and instantaneous workload selection techniques uncovered in the field study.

The selection mechanisms filter the list of candidates and pass the filtered list to ER's user interaction components. Initially, these will include a communications tracker, an escalation tracker, and a user oriented workload management tool.

This work is Dave McDonald's thesis project.

2. **Approximation Techniques for Mapping Expertise Networks** (Wayne Lutters, David McDonald, James Boster, Mark Ackerman)

A critical bottleneck for constructing suitable expertise recommendation systems is understanding how expertise networks can be measured adequately. A knowledge or expertise management view would require detailing the current expertise within an organization. Mapping people's expertises is, however, hardly an easy task. A precise and detailed mapping of an organization's knowledge would be extremely difficult (in terms of expense and time), if it were even possible. Existing techniques either require a detailed and time-consuming survey process or to elicit the required information for a mapping by examining the natural by-products of working in the digital world, such as email. These latter techniques, however, measure interest in a topic, rather than expertise in that topic.

In this project, we are considering first-order approximations for the mapping. We borrow the language of fluid mechanics and structural engineering in denoting a family of techniques for examining systems that are too complex to model directly. We are attempting to find techniques that require less than one hour of detailed survey from each organizational member, use secondary data collection wherever possible, and have measurable error rates and biases.

We are currently conducting an exploratory study to consider various approximation techniques that might be of value. This preliminary study has two steps:

- A. An elicitation of questions that will be used to measure individuals' expertise about a critical success factor for the organization. Instead of developing the questions ourselves, we wish to elicit questions from the organizational members instead, since our goal is to develop a process to obtain localized instruments at low cost. For our pilot site, the questions are about their main product, a software system. We believe this to be an innovative technique. The end-result of this step is a sixty to seventy question Knowledge Mapping Instrument (KMI).
- B. Measurement of knowledge distribution using the KMI developed in phase A. We are currently collecting data from organizational members at our pilot site.

We are also obtaining secondary data to determine the validity and reliability of our KMI approximation as well as to assess potential factors that might allow us to reduce data collection efforts even more. For example, we have obtained social network data for the relevant organizational members using two cognitive anthropological techniques. We are also obtaining subjective evaluations of scores on the KMI (both one's own and others' scores) in order to assess whether there are key informants who could provide most of this information (rather than requiring a survey of all organizational members).

3. **Architecture Simulations of Expertise Networks** (Anapam Basu, Dan Hirschberg, Mark Ackerman)

Many architectures have been proposed for augmenting expertise networks (using various combinations of email filters and routers, chat areas, newsgroups and discussion databases, expertise agents, social filtering, and so on).

Relatively little work has been done on comparative empirical work, examining the benefits, liabilities, and trade-offs of these varying architectures. In-situ studies of recommender systems are expensive (in terms of system construction and organizational effort), limiting studies to single systems. Furthermore, few organizations are willing to disrupt their normal activities for the introduction of an experimental system, making *any* study difficult to arrange. Accordingly, we have chosen to use simulations to examine comparative architectures, allowing progress to continue at relatively low cost. We believe these simulations to be of great possibility.

This project is just beginning and is looking for sponsors. We have constructed a rudimentary simulator with differentiated agents (both human and machine), expertise (in varying abstract topics), message traffic and cycles, and an input stream of questions. Currently we use a variety of output measures, including time-to-answer, message traffic, and redundancy.

The above three projects together aim to examine how people can seek out expertise more easily and more reliably. A word is perhaps in order about why we chose expertise to examine and how these three projects fit into a general research thrust. Over the last five years, we have been conducting a series of coordinated technical and social investigations of CSCW problems such as information seeking, virtual collectivities, and now expertise. Our goal has been to examine how well we can design for the social world and to determine the technical limits for doing so. We are examining expertise specifically because it is inherently social and situated – while expertise clearly has a cognitive and individual side, it is also socially situated and enacted. The activity of seeking expertise is clearly socially situated and difficult to augment. As such, it serves as an excellent "testbed" for examining the possibilities for and the limitations of CSCW design.

These projects have been funded, in part, by grants from National Science Foundation grant (IRI-9702904), the University of California MICRO program, and the UCI Center for Research on Information Systems and Organizations.