Published in the Proceedings of the 2003 ACM Conference on Human Factors in Computing Systems (CHI'03), Ft. Lauderdale, FL, April 5 – 10, 2003.

Recommending Collaboration with Social Networks: A Comparative Evaluation

David W. McDonald
The Information School
University of Washington, Seattle
Suite 370, Mary Gates Hall
Seattle, WA 98195
dwmc@u.washington.edu

ABSTRACT

Studies of information seeking and workplace collaboration often find that social relationships are a strong factor in determining who collaborates with whom. Social networks provide one means of visualizing existing and potential interaction in organizational settings. Groupware designers are using social networks to make systems more sensitive to social situations and guide users toward effective collaborations. Yet, the implications of embedding social networks in systems have not been systematically studied. This paper details an evaluation of two different social networks used in a system to recommend individuals for possible collaboration. The system matches people looking for expertise with individuals likely to have expertise. The effectiveness of social networks for matching individuals is evaluated and compared. One finding is that social networks embedded into systems do not match individuals' perceptions of their personal social network. This finding and others raise issues for the use of social networks in groupware. Based on the evaluation results, several design considerations are discussed.

Keywords

Recommendation systems, social networks, evaluation, information seeking, expertise locating, CSCW, computer-supported cooperative work.

INTRODUCTION

Dustin is an engineer for a consulting firm. He was recently asked to design a new component for a larger engineering project. Dustin is not very familiar with the project, so he will need help. Perhaps he can find another engineer on the project with whom he can kibitz and thrash out various design ideas. Dustin turns to a system to get a recommendation of another engineer with whom he can collaborate. The system uses a social network to recommend three people that Dustin has communicated with in the past. Dustin stares quizzically at these recommendations wondering why these people were suggested.

The use of social networks in groupware systems is rising in popularity because of their descriptive and analytical power. Designers and implementers are using social networks to visualize small through large group connections and guide users toward collaborative interaction. The hope is to make systems more sensitive to a broad range of social situations. However, the implications of adopting social networks as a system abstraction are poorly understood. There are many factors from both an individual and group perspective that have yet to be evaluated.

This paper describes the evaluation of a system that makes recommendations using two different social networks for the same organization. The evaluation compares recommendations made using a social network to recommendations made without a social network. Evaluations of systems that include a social network have rarely focused on the efficacy of the social network itself. To reiterate, the focus of this contribution is not on the social networks in and of themselves (there are other venues for that specific discussion). Instead the focus is on evaluating two social networks in the context of specific system usage and how users interpret the effectiveness of those social networks.

The paper begins with a review of several systems that include social networks to facilitate new or renewed collaboration. The paper then describes two social networks that were collected and incorporated into a recommendation system. A brief overview of the recommendation system and its social network matching techniques is followed by the evaluation. The results lead to a reflection on current approaches and problems when using social networks in groupware systems.

FROM VISUALIZATION TO COLLABORATION

Social networks often represent groups of people and the connections among them. The strengths of social network analysis have resulted in increasing use for understanding a range of small through large group interaction. In general, social networks originated from a descriptive and analytic discipline, but there is a trend toward embedding social networks into systems with the goal of facilitating new or renewed collaboration.

One common approach is to use social network visualizations as an overview of group participation or group membership. In systems that take this approach the visualization and what it conveys to the user is of primary concern. The Netscan project [24, 25] provides one of the more elaborate visualizations of Usenet news group participation. Netscan mines Usenet news groups across many dimensions. Netscan uses the mined data to present several different social network visualizations of participation in Usenet. Conversation Map [22] is a content browser that also mines Usenet news groups. Conversation Map not only provides a content visualization by analyzing message content, but also displays a social network of participants.

Other systems attempt to use social networks as a mechanism for recommending specific people for collaboration. In this approach a visualization is often a means of finding a specific person. In ReferralWeb [10], co-authoring and co-citation relationships are mined to create a social network and the resulting visualization is used to find a possible expert. The social network in ReferralWeb can also answer queries about how far (in person-person edge links) one researcher is from another and who is between.

Ogata et al. [20] also use social networks to facilitate finding a person with whom to collaborate. Their system focuses on mediating personal connections (PeCo) by mining and analyzing email exchanges among individuals. The mining and construction of the social network relies on a type of speech act model of communication. A user can find another by making different queries, but the user is never shown a visualization of the social network.

Ogata et al. evaluated PeCo through a field experiment and subsequent survey. In the field experiment graduate students were made available (through email) to help undergraduates enrolled in a course on C programming. Briefly, the survey results found that the undergraduates valued PeCo for helping them find help but did not feel that the system accurately reflected the social closeness of their friends in the programming course.

Contact Map [17] takes a different approach to supporting collaboration among individuals. Contact Map is a personal communication and contact management application that uses a social network visualization to present contacts. A key difference from prior systems is that Contact Map visualizes an ego–centric social network. Contact Map mines email to gather contact information (e.g. email address, phone number, etc.) of individuals with whom the user has communicated. It generates a metric of the strength of tie between the user and the contact and displays the more heavily weighted ties centrally in the visualization. The user can then color and rearrange contacts to represent work, project, team or other social groupings.

Nardi et al. [17] present a preliminary evaluation that considers the email contact mining metrics and choices that individuals make to include or exclude a given contact from the Contact Map. They found a huge variance in the number of contacts identified by mining and the number that were deemed relevant to each user. As well, they present some qualitative data regarding the choices and rationale of users when deciding to include specific contacts in their ego–centric social network.

This evaluation builds on the prior work by considering, in greater detail, how users perceive and interpret social networks when they are incorporated into a system. Organizations have many different social networks. The two social networks used in this evaluation were collected to represent specific factors that positively influence organizational information sharing and collaboration. Before describing the evaluation it is essential to describe these social networks, how they were collected and what they represent.

ONE ORGANIZATION, TWO SOCIAL NETWORKS

Organizations have complex social structures. It is often difficult to completely understand who knows who, how people cooperate and how they come to know each other. Social networks are one way to analyze groups and the relationships that comprise them.

Studies of collaborative behavior often find that an individual's social network strongly influences information seeking and collaboration behavior [1, 5, 6, 11, 18, 19]. In these studies, social networks describe a broad, complex, range of social and organizational interaction. These results demonstrate the importance of workplace sociability, shared context and physical location, among other properties, as an influence on collaborative activity.

Nardi, Whittaker, and Schwarz [18] is one recent example of this genre. Their qualitative research examines the many factors that influence individuals' choices in initiating collaboration. As well, they provide detailed description of how social ties wax and wane over time, through different projects and varying organizational affiliations. While Nardi et al. did not perform a structural analysis, nor a visualization of the social networks in their study, their results lead to Contact Map [17], a system that does visualize ego–centric social networks.

Others approach information seeking and collaboration research using a more structural approach from the start. Allen [2] and Eveland et al. [7] examined the interaction patterns of information seeking using methods that could visualize the social network of participants. Allen [2] performed several detailed analyses of social networks. Allen mapped intra- and inter- organizational relations for groups of research engineers. These visualizations helped identify key organizational members known as technological gatekeepers. Work by Eveland et al. [7] considered "help networks" as a type of social network. Help networks are arrangements of people who assist each other with adopting new technologies. Eveland et al. visualized the relations that facilitated information seeking during the adoption of new computing infrastructure at a university.

Different collection and analysis techniques can yield different visualizations and interpretations of a given social network. This is not a problem. Indeed organizations and individuals have many differing social networks. This work does not attempt to present a discussion of collection nor analysis techniques; there are established sources for that (c.f. [3, 23, 28, 30]). Rather, it is more important to understand what is represented by a given social network.

The next sections describe two social networks collected at one organization and argue that these networks represent the relationships intended in the collection. One social network was collected using qualitative methods, while the other network was collected using a quantitative approach.

The Organizational Setting

The social networks were collected at a medium sized software development company called Medical Software Company (MSC)[1]. MSC develops, sells and supports medical and dental practice management software. Practice management includes patient demographics, scheduling, patient reminders, insurance reimbursement, and patient billing. MSC's clients include group practices with as few as three doctors or dentists to large managed health care organizations with hundreds of doctors.

MSC has over 100 employees at its headquarters and approximately 60 additional employees who work in wholly owned subsidiaries and remote offices around the US. This study focuses on two key departments, technical development and technical support. These departments are central to MSC's core software business and comprise more than one third of the employees at headquarters.

The studies at MSC cover a three year time span. They began with an in-depth ethnographic study that lasted nine months. The ethnographic study was followed by numerous periods of intensive data collection. MSC has participated in a number of field studies [12, 14-16].

The Work Group Graph: A Social Network of Shared Workplace Context

One important factor of information sharing in an organization is establishing a common context for interaction and communication. The first network represents how participants at MSC view shared contexts. This social network is called a work group graph (WGG) because it prioritizes logical work groups and work context over organizational boundaries.

The WGG was constructed through ethnographic methods; interviews, participant observation, and artifact collection. This collection is similar to that used by Nardi et al. [18]. However, the goal here was to develop a network capable of visualizing logical work groups in the organization. The WGG is shown in Figure 1.

The visualization in Figure 1 has much in common with many social network graphs. Each lettered node represents one person. The linear proximity of nodes is less important than the existence of an edge between nodes. The nodes representing individuals are anonymized in a systematic way such that, for the researcher only, a letter can be traced back to a specific person. The WGG visualization represents one way that systematic qualitative data can be visualized from a structural perspective.

Figure 1 shows a hypergraph (a hierarchical graph) containing six main nodes. The main nodes represent common work groups. The six main nodes represent boundaries of participant stated groups. Components within a main node represent logical work similarity within a group.

The subgraph in Node 1 represents members of technical development who work with operating system specific code, utilities, and a PC based prototype system. Node 2 contains two components representing additional members of technical development. The smaller component represents administrative and managerial staff. The larger component represents developers who work directly with the dental and medical system. The small subgraph in Node 3 represents development staff who handle technical writing tasks. Node 4 contains small and large components representing members of dental and medical technical support, respectively. The subgraph in Node 5 represents the members of technical support who perform hardware based field service. Node 6 represents client managing.

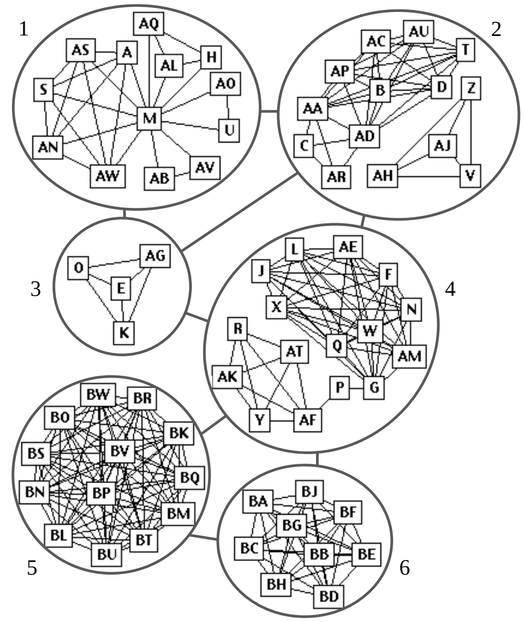

Figure 1. The work group graph (WGG).

The WGG does not represent formal organizational boundaries. In comparison, organizational boundaries at MSC would break this graph into four units; two technical development units (Node 1, separate from the union of Nodes 2 and 3), technical support (Nodes 4 and 5), and client managing (Node 6).

The WGG displays work related context around individuals who do similar tasks. However, it does not represent these contexts equally well. The focus of the ethnographic study (Nodes 1 through 4) reasonably distinguishes individuals with related work context. However, nodes representing other areas (Nodes 5 and 6) were important enough for participants to mention, but not important enough to distinguish relevant tasks. As well, because the WGG focused on work similarity as a means of identifying shared context, it does not effectively identify individuals who span contexts.

The SPS Social Network: A Social Network of Workplace Sociability

Another important factor of information sharing in an organization is the sociability of various individuals. Individuals who socialize frequently are more likely to collaborate in the future. The second social network at MSC was designed to represent and visualize workplace sociability.

1. This is not the real name of the organization. The names of the organization and participants have been changed.

A Successive Pile Sort (SPS) [4, 29] technique was used to collect the second social network. In this technique, the name of every member in the group is written on a card. Participants sort the cards using a high level rubric supplied by the researcher. Each participant is free to interpret the rubric in her own way. The first sort results in a number of "piles" which are, in turn, sorted using the same rubric. The level of the sort at which individuals or groups are broken apart indicates the connection weight between the members. The connection weights are aggregated across all participants to create an edge weighted social network.

The SPS collection again focused on members of technical development and technical support. Participants were challenged to create sorts with the rubric "who hangs out together." This rubric was designed to reveal the social structure rather than work context structure at MSC. Motivating the SPS collection by asking "who hangs out together" was one way to consider the more sociable aspect of interaction at MSC. Each participant required between 45 and 90 minutes to sort 47 cards.

The SPS data were aggregated and a Multi-dimensional Scaling (MDS) was performed. An MDS is useful for identifying related clusters. Many MDS tools generate a visual diagram of these clusters. A byproduct of performing an MDS is a distance matrix. This matrix was fed to a graph layout tool. The MDS diagram and the graph layout displayed equivalent structures. Figure 2 is the graph layout.

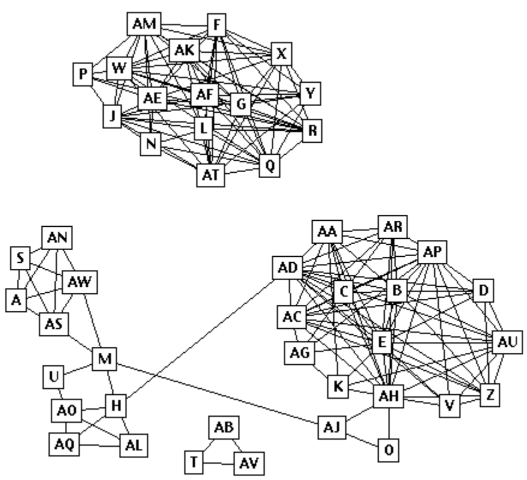

Figure 2. Graph of the SPS social network.

A traditional approach for validating a social network is to take the diagram back to the social group, tell individuals where they are in the graph, and let them 'interpret' the correctness of the graph. At the time of the collection and analysis, the author had been interacting with MSC for over 18 months and had performed a detailed qualitative study. Given the level of familiarity with the participants, the author's understanding of the social groups was used to validate the network. Some social network practitioners may be critical of this approach. However, the whole notion of 'validating' a social network is changing. The trend to collect social networks using automated methods such as mining email or other communication records is pushing the bounds of traditional techniques. In some instances, the size of the networks collected is much too large for traditional validation techniques.

Figure 2 has a high level of face validity for MSC. The clusters are representative of social and work relations. In Figure 2, the cluster at the top represents technical support, who hang out and work together. The clusters at the bottom are technical development. The lower left cluster consists of people who work on operating system code and utilities. The lower right cluster represents developers who work on the dental or medical system. Since data were collected in the workplace, perhaps it is not surprising that the participants notion of "who hangs out together" is guided by who works together. This result is similar to that found when evaluating Contact Map [17]. There is a large overlap between workplace sociability and formal work groups.

EVALUATING TWO SOCIAL NETWORKS

As groupware designers and developers adopt social network techniques, it is important to understand how new systems impact users. Evaluations of systems that include a social network have rarely focused on the efficacy of the social network itself. This evaluation specifically considers the effectiveness of two social networks used in an expertise locating system at MSC. Before describing the evaluation it is necessary to provide an overview of the expertise locating system and how it uses these two social networks.

An Expertise Locating System

Expert finding systems are a type of recommendation system designed to find a person who has specific knowledge of a problem domain. A number of systems fall into this category [8, 10, 13, 15, 26, 27]. Different expertise finding systems approach the problem with different techniques and often have different goals.

The Expertise Recommender (ER) [15] is designed to augment naturalistic information seeking behavior in an organization. ER cannot replace key individuals. Instead, it helps people to locate expertise in unfamiliar parts of an organization and provides alternatives when key individuals are unavailable. ER takes into account many things when making a recommendation. Unique to ER and crucial to this evaluation is that ER attempts to socially tailor recommendations to the individual making a request by using social networks.

Simplifying the actual recommendation process, ER makes recommendations in two steps. First ER finds a set of individuals who are likely to have the necessary expertise. A prior evaluation found that individuals recommended by ER are likely to have expertise [14]. These potential recommendations are then matched to the person requesting expertise using a social network. In this approach, the users are never explicitly shown a visualization of the social network. This social matching approach is similar to that taken by Ogata et al. [20]. The ER matching approach is different from that taken by ReferralWeb [10] and Contact Map [17] which explicitly show a social network visualization.

The MSC implementation of ER includes two socially motivated matching methods; the SPS social network and the work group graph (WGG). The SPS social network reflects workplace sociability. The SPS network is used because the literature on information seeking suggests that people often seek help from individuals they know socially. In the evaluation, the matching option that relies on the SPS network is called "Social Network."

The WGG reflects shared work context. The WGG is used because research literature suggests that shared context is one key factor to effective information seeking and exchange. Sharing a similar context improves the likelihood that two individuals can effectively communicate about a given problem. In the following evaluation, WGG matching is called "Departmental" matching in order to distinguish it from "Social Network" matching and simplify the explanation to the participants.

For the purposes of this evaluation, ER included a way to turn off both socially motivated matching techniques. This "No Matching" simply provided the list of people who ER identified as having potential expertise from the first step of ER's recommendation process. Choosing "No Matching" makes ER generate recommendations similar to many other expert finding systems (cited above). The use of a "No Matching" technique in the evaluation proved useful. Since "No Matching" returns the exact same response for the same query from all participants, it represents a fixed point of comparison.

In ER an *escalation* of a problem relaxes the system constraints on who is considered; both in finding individuals with expertise and in social matching. Escalating a request casts a bigger net, potentially finding less capable people, but maybe people more available and willing to interact.

Evaluating Social Networks for Recommendation

A structured interview protocol served to focus participants' attention on the effectiveness of ER's social network matching. Participants used ER to make expertise requests and were then asked to compare, side-by-side, the results of a request using one matching method to the results without matching. In each instance the participant was specifically asked to judge "Which is better?" and "Why?" Each participant compared four separate requests; two comparing *Social Network* to *No Matching* and two comparing *Departmental* to *No Matching*. Each request was compared, *escalated* and then compared again. Eighteen people participated in the evaluation. The interviews were recorded, transcribed and analyzed. The evaluation was a qualitative investigation, however, the transcripts where coded for 'positive,' 'negative,' and 'non-committal' comments comparing each matching to *No Matching* condition. No significant quantitative differences were found, but the following discussion presents results reflective of the trends.

No Matching

Several participants preferred *No Matching*. Both of ER's socially motivated matching techniques can suppress[2] a person who has expertise because they are too far away in the social network. Some participants did not like the way matching techniques suppressed possible recommendations. In effect, they expressed a type of 'more is better' attitude. One participant was a staunch supporter of the 'more is better' approach. In a discussion about her preference, she commented that the extreme (e.g., just listing everyone in a department or the company) was impractical.

Some people preferred *No Matching* because matching was working too well. At one point, Eyal, a highly respected employee for his technical skills, was asked whether he preferred one set of results to the other. He paused, clearly torn at having to choose between the two. The problem for Eyal was a trade-off.

> Yeah, yeah, because here you have three people where two of them can probably help, the third one not at all [pointing to *No Matching*]. Here I have only one, that is right person, but you don't give me any more choices [pointing to *Departmental*]. Here you might choose the wrong person [pointing back to *No Matching*] … — Eyal

In this particular instance, there was a cost for being correct. When a socially sensitive matching technique performs too well, users can feel like they have somehow lost control. In this case the user wanted the system to be correct, accurate and provide him choices.

Departmental Matching

The *Departmental* matching technique is based on the work group graph (WGG) collected at MSC. Many comments were favorable to the WGG. In the quote below, a junior programmer, is happy with how matching works, but recognizes that a slightly different problem would lead him to choose *No Matching* over *Departmental* matching.

> So the Departmental one is very much accurate for a technical, for me, for a technical question about it. But as far as the setup and how it can work for clients and what different features it may have, as to how certain clients have used it in the past, and how they're using it now, to recommend how a current client should use it, I think I would probably use the no filter one, because that would bring me up people from support and from client managing. — Simon

Simon recognizes that *Departmental* matching is suppressing recommendations of people who have expertise in how customers setup and use the system. For his context, for a technical question, this would be acceptable. However, Simon recognizes that for expertise in a different context, setup and use, he might want people from support and client managing who are in work groups that are, in a sense, further away from him. Simon did not see that the matching technique would adapt. He was anticipating the needs of individuals who would use the system in a different context with different questions and needs.

Still other participants liked how *Departmental* matching worked, but suggested changes to address their concerns.

> … that's a good concept, the Departmental filter. So you could go to the department and then when you

---

2. In ER one or more *escalations* of a recommendation request will eventually expose a person who was previously suppressed. The evaluation protocol included the use of escalation in order to demonstrate this.

> escalate it maybe be able to choose which department you want to escalate it to. Maybe, find somebody in field service or find somebody in development. — Victor

> I guess, ideally, if it's a Departmental, if there was a way to put in support and then another department, training [or] client managing is my two choices and then, then I'd just get other people, so I wouldn't lose Peter because he might be a good contact. — Katie

Both Victor and Katie are support representatives who just want to be able to focus the direction of the matching technique and how it escalates to include more people. They want to control how the system relaxes constraints in the WGG. In this way they could then control which work contexts are considered first. They recognize that they have insight into the problem context and that, with a little control over the way in which the system relaxes constraints, they might improve the effectiveness of matching.

Social Network Matching

Social Network matching relied on the SPS social network collected at MSC. Comments about the SPS social network were mixed. Some participants saw it as working for them and were generally positive. However, there were others who were quite negative. Participant responses to *Social Network* matching were by far the most polarized.

Exemplary of the positive responses was Simon's; a programmer who recognized what *Social Network* matching could provide him.

> … the Social Network one was good in that it gave me somebody I knew I could get a hold of and was close and would be easy for me to talk to. — Simon

The ability to find someone who would be "easy for me to talk to" is one goal of using a social network that represents the sociability of the participants; like the SPS social network at MSC. But the participants saw problems with *Social Network* matching as well.

One set of problems concerned the specific implementation of the social network. Many participants recalled the time spent performing the Successive Pile Sorts. The use of SPS and aggregation techniques may result in an analytically useful social network, however this aggregation was a serious problem for the participants. One example of this was Liz, a junior technical support representative.

> Now is this, when you're talking social, you're just talking physically social? … So it's like a generalization based on data from everybody? … So it's not my personal social network? Just how other people view everybody else in the department. … I probably wouldn't choose the Social Network, just because it wouldn't be appropriate. — Liz

Liz was quite unhappy that the social network was an aggregation. She felt that since the network was a function of other people's opinions it could not result in an appropriate match tailored to her.

Daniel, the V.P. of Development, had a similar reaction. When comparing two recommendations he noticed that the *Social Network* match suppressed a person recommended in *No Matching*. In the interview Daniel asked, "Social network? Whose social network? …" He thought that the match could not be correct, because he felt particularly close to the person who was suppressed. Daniel explained,

> Social network? Whose social network? … I would think that he would show up, I would just think that if you apply a social filter he would be … He'd be closer. He's socially closer. — Daniel

In addition to the problems from an aggregate social network, several participants felt that matching should never be based on social criteria. These participants most often expressed the idea that they just wanted the person who knew the most. For example, Justin, a senior support rep, expressed serious concerns about *Social Network* matching. For Justin,

> 'Hang out' doesn't mean anything. When I need a question answered I'm going to go to the person who knows the answer. — Justin

Liz, the support rep who had a problem with the aggregate social network expressed a similar view. In the quote below, Liz just wants to have the system identify the person who knows the most, not a "friend."

> Certainly the Social Network wouldn't be appropriate. The way I see it when I'm working on an issue I don't want to talk to a friend I want to talk to somebody who's gonna know what they're doing. And if I've never talked to them before, that doesn't matter to me because I'm doing something for a client and that's my revenue. — Liz

Liz was particularly critical of the *Social Network* matching technique and critical, in general, of socially sensitive matching. However, she recognized that there were social situations that are embarrassing because a person may not have a reasonable way to pick one source of help over another. Liz related a story, which she clearly found embarrassing.

> … Like I didn't know until a few months ago that Jake was the VP of our department. I had no idea. And I would call him randomly and just, you know, spit off some stuff that I'd been doing some work for, like, platinum clients and we're supposed to talk to the VP of the department and, and then, I was told to go to Jake and I'm like why? Well, he's the VP of our department. I'm like, oh, really? Oh, my gosh, I've been here for a year and a half and I didn't know that. — Liz

In further discussion, Liz rationalized that there is no good solution to problems like this one. It never occurred to her that a system using socially sensitive matching could assist people in avoiding some embarrassing situations.

Liz and Justin both want to get the expertise as quickly as possible, and seem to be willing to endure embarrassment to get help. For Justin, a senior support representative, the social costs of getting expertise may never be too high. However, for Liz, a junior support rep, the costs could be very high. The comments by Liz and Justin were not focused on the results of any single recommendation, but on the general direction. Their concern was that recommendations might focus on social aspects first and expertise second, possibly recommending "friends" with little expertise.

Lastly, there were a few people who, while still critical of the *Social Network* matching, saw promise in it. These people recognized that socially sensitive techniques presented a range of potential problems; but problems that could be overcome. Some people, like Andreas, recognized that social situations shift over time.

> … people change their social networks from time to time. Departmental is known, probably, so, you should probably make it learn over time and make it easy for people to update their data. — Andreas

The people who saw promise in *Social Network* matching attributed problems to social changes that take place between the time when data are collected and when groupware is deployed. For these participants, a solution requires that the social network adapt over time and that people be allowed to modify their individual social network data.

DISCUSSION

In the context of HCI, social networks, as a groupware design and interaction component, have not been critically scrutinized nor systematically evaluated. This work takes the stance that social networks embedded in systems require user based evaluation similar to any other system or UI component. These results raise a number of issues for the use of social networks in groupware systems.

It seems clear that individuals have mixed feelings about social networks as a tool for finding collaboration. Social networks derive from an analytic and descriptive perspective, whereas their application in groupware is often oriented in a slightly more prescriptive direction. The distinction between the way social scientists *actually use* social networks and the way groupware designers *would like* to use social networks is important and should be highlighted in the growing discussion of social networks in the research literature. Indeed, some of the following results are known by social network practitioners, but have not made their way into the discussion of system design and implementation. Recapitulating some of the results and implications:

Perceived trade-off: Incorporating social networks into a recommendation system, specifically an expertise locating system, results in a conflict between users' perception that a match is purely social and their desire to find a person who has expertise. Effectively, they perceive a trade-off between finding the most knowledgeable person or a person with whom they can easily interact. Actually, the system identifies a potential expert first, then performs social matching; but the users still perceive a trade-off. Groupware systems that incorporate socially sensitive matching techniques, like social networks, will face this dilemma.

Control & transparency of social matching: Users often have insight into their problem. They believe that if given control over how the social network is searched, they can get a better match. This finding is also related to the issue of providing choices when matching. A system that performs too well, recommending a single match, might encounter resistance because of the users' perception that they have less control. Related to control is the way the system explains a match or a recommendation. Users seemed to look to the system to explain a match when the aggregate social network generated recommendations that differed from their view of their ego–centric social network. This argues for some parameterization of socially sensitive matching techniques. Users naturally want the system to augment and assist, not replace their natural behavior.

Ego–centric and/or aggregate social networks: Individuals have a more subtle understanding of their own social network than that attained by aggregation techniques. It would be easy to claim, in hindsight, that it is obvious that individuals would reject aggregate social networks over one that is somehow specific to them. Ogata et al. [20] point to this problem, but it is not a major focus in their evaluation. A majority of systems use aggregate social networks without systematic evaluation. There is clearly value to be derived from both representations, the how, which, and when is an open question from a users' perspective.

Social dynamism: Peoples' social networks are dynamic. The social networks in this evaluation were collected using traditional field methods and were therefore a snapshot of the participants' real social networks. Participants were charitable to the system with regard to this, but it is still a serious problem. As we design groupware that incorporates social networks we need to think carefully about how to handle dynamism in terms of social network membership, link meaning, and strength [9, 21].

Bootstrapping the social network: Although this may not be clearly evident from the evaluation, bootstrapping a social network for a system is amazingly problematic. The system evaluated here relied on traditional field methods to collect the social network data. This approach is time consuming and will not scale to very large groups. However, it is not at all clear that the predominant method for bootstrapping large groups (i.e. mining email or other electronic communications) is any better at producing an accurate social network from an individual's point of view. The preliminary evaluation of Contact Map [17] showed a high variability in the mined result from one individual to the next.

This work examined how users interpret a socially sensitive recommendation system. The evaluation compared two social networks in the same system and found that users have mixed feelings about the effectiveness of the general approach. The application of social networks by groupware designers and implementers will continue. The results and discussion suggest several key issues that designers should consider when building social networks into new systems.

Acknowledgements

This has been funded, in part, by grants from National Science Foundation (IRI-9702904), the UC Irvine/NSF Industry/University Cooperative Research Center at CRITO, a University of California MICRO and a University of California Regents' Dissertation Fellowship. This work benefited from discussions with Mark Ackerman, Jack Muramatsu, Elizabeth Churchill, Jeff Kim, Leysia Palen, Paul

Dourish, Marc Smith, and Danyel Fischer. Special recognition goes to James Boster for consulting and Wayne Lutters for help with data collection. Thanks to the anonymous reviewers who raised important issues, some of which I hope I have adequately addressed. The participants at MSC deserve thanks for their patience and insights.

References

1. Ackerman, M.S., Augmenting the Organizational Memory: A Field Study of Answer Garden. in *ACM Conference on Computer-Supported Cooperative Work (CSCW '94)*, (1994), 243-252.
2. Allen, T.J. *Managing the Flow of Technology*. MIT Press, Cambridge, 1977.
3. Berkowitz, S.D. *An Introduction to Structural Analysis: The Network Approach to Social Research*. Butterworths, Toronto, 1982.
4. Boster, J.S. The Successive Pile Sort. *Cultural Anthropology Methods*, *6* (2). (1994) 7-8.
5. Chin, G., Myers, J. and Hoyt, D. Social Networks in the Virtual Science Laboratory. *Communications of the ACM*, *45* (8). (2002) 87-92.
6. Ehrlich, K. and Cash, D. The Invisible World of Intermediaries: A Cautionary Tale. *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, *8* (1/2). (1999) 147-167.
7. Eveland, J.D., Blanchard, A., Brown, W. and Mattocks, J., The Role of "Help Networks" in Facilitating Use of CSCW Tools. in *Proceedings of the ACM 1994 Conference on Computer Supported Cooperative Work*, (Chapel Hill, NC, 1994), 265-274.
8. Foner, L.N., Yenta: A Multi-Agent, Referral-Based Matchmaking System. in *First International Conference on Autonomous Agents (Agent'97)*, (1997).
9. Granovetter, M. The Strength of Weak Ties: A Network Theory Revisited. *Sociological Theory*, *1* (1). (1983) 201-233.
10. Kautz, H.A., Selman, B. and Shah, M. Referral Web: Combining Social Networks and Collaborative Filtering. *Communications of the ACM*, *40* (3). (1997) 63-65.
11. Kraut, R.E., Egido, C. and Galegher, J. Patterns of Contact and Communication in Scientific Research Collaboration. in Galegher, J. and Kraut, R.E. eds. *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*, Lawrence Erlbaum, Hillsdale, NJ, 1990, 149-172.
12. Lutters, W.G., Ackerman, M.S., Boster, J.S. and McDonald, D.W., Mapping Knowledge Networks in Organizations: Creating a Knowledge Mapping Instrument. in *Americas Conference on Information Systems, AMCIS'00*, (Long Beach, CA, 2000), AIS, 2014-2018.
13. Maybury, M., D'Amore, R. and House, D. Awareness of Organizational Expertise, The MITRE Corporation, MacLean, VA, 2000.
14. McDonald, D.W., Evaluating Expertise Recommendations. in *ACM 2001 International Conference on Supporting Group Work (GROUP'01)*, (2001) 214-223.
15. McDonald, D.W. and Ackerman, M.S., Expertise Recommender: A Flexible Recommendation System and Architecture. in *The ACM 2000 Conference on Computer-Supported Cooperative Work (CSCW'00)*, (Philadelphia, PA, 2000), 231-240.
16. McDonald, D.W. and Ackerman, M.S., Just Talk to Me: A Field Study of Expertise Location. in *The 1998 ACM Conference on Computer-Supported Cooperative Work (CSCW'98)*, (Seattle, WA, 1998), 315-324.
17. Nardi, B.A., Whittaker, S., Isaacs, E., Creech, M., Johnson, J. and Hainsworth, J. Integrating Communication and Information Through ContactMap. *Communications of the ACM*, *45* (4). (2002) 89-95.
18. Nardi, B.A., Whittaker, S. and Schwarz, H. NetWORKers and their Activity in Intensional Networks. *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, *11* (1-2). (2002) 205-242.
19. O'Day, V.L. and Jeffries, R., Orienteering in an Information Landscape: How Information Seekers Get From Here to There. in *Interchi '93*, (1993), 438-445.
20. Ogata, H., Yano, Y., Furugori, N. and Jin, Q. Computer Supported Social Networking for Augmenting Cooperation. *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, *10*. (2001) 189-209.
21. Pickering, J.M. and King, J.L., Hardwiring Weak Ties: Individual and Institutional Issues in Computer Mediated Communication. in *The 1992 ACM Conference on Computer-Supported Cooperative Work (CSCW '92)*, (1992), 356-361.
22. Sack, W., Conversation Map: A Content-Based Usenet Newsgroup Browser. in *The 2000 International Conference on Intelligent User Interfaces (IUI 2000)*, (New Orleans, LA, 2000), ACM Press, 233-240.
23. Scott, J. *Social Network Analysis*. SAGE, Thousand Oaks, CA, 1991.
24. Smith, M.A., Cadiz, J. and Burkhalter, B., Conversation Trees and Threaded Chats. in *The ACM 2000 Conference on Computer-Supported Cooperative Work (CSCW'00)*, (Philadelphia, PA, 2000), 97-105.
25. Smith, M.A. and Fiore, A.T., Visualization Componentes for Persistent Conversations. in *The ACM Conference on Human Factors in Computing Systems (CHI'01)*, (2001), 136-143.
26. Streeter, L.A. and Lochbaum, K.E., Who Knows: A System Based on Automatic Representation of Semantic Structure. in *RIAO '88*, (1988), 380-388.
27. Vivacqua, A. and Lieberman, H., Agents to Assist in Finding Help. in *ACM Conference on Human Factors in Computing Systems (CHI 2000)*, (2000), 65-72.
28. Wasserman, S. and Faust, K. *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge, UK, 1994.
29. Weller, S.C. and Romney, A.K. *Systematic Data Collection*. Sage Publications, Newbury Park, 1988.
30. Wellman, B. and Berkowitz, S.D. (eds.). *Social Structures:* A Network Approach. Cambridge University Press, Cambridge, 1988.