Visualizing Discourse Architectures with Automatically Generated Person-Centric Social Networks

Shelly Farnham
Social Computing Group
Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Summary

We use a person-centric network metaphor to model social environments, where each person in a conversation space is represented in terms of his or her relationships to other people based on conversation history. A network approach allows us to automatically provide users with dynamic representations of patterns within their social space and a means for navigating through them. We are currently developing two sample applications: one based on a user's email behavior that creates a dynamic interface representing the user's personal social network (*Personal Map*); and another based on distribution list memberships that provides a map of the connections between researchers at MSR (*MSR Connections*). Our goal is to enhance the user's ability to seek out social information and manage online social relationships.

1. Introduction

Asynchronous conversation via email has become a primary use of computers in both the work environment and in the home. We developed the person-centric social network visualizations described below because we were interested in exploring ways to increase the *social adeptness* of such user interfaces to conversation spaces. Socially adept UIs are structured according to people's affective or cognitive mental models of their conversations. Current systems are largely socially *inept*, using a file system metaphor to structure personal contacts and emails. In such systems contacts and emails are represented as objects to be created, modified, or deleted, and organized hierarchically according to name, title, and date, or by static groups. Our research goal was to better model the UI for email clients or member directories after the users' phenomenological experiences of their important online social interactions.

A basic assumption to our approach is that *people* should provide a primary organizational structure to the light-weight conversations found in email. From a phenomenological perspective, conversation is an invisible tool by which meaning and experiences are shared between people. The coherence of a conversation is determined by the location of that conversation in a social space (i.e., with the interactants). We expected that when navigating through conversational spaces people would care more about whom they were talking to than what they were talking about.

In a first, exploratory study of how people naturally model their social environments, we recruited participants in a mall to draw maps of their social spaces (Vronay & Farnham, 2000). We found that people tended to structure social space in terms of their *relationships*. People tended to draw themselves in the center of their maps, place important others near the self, and draw lines to others to represent significant relationships. Others were organized depending on their relationships to each other, forming groups that appeared very dynamic. See Figure 1 for examples.

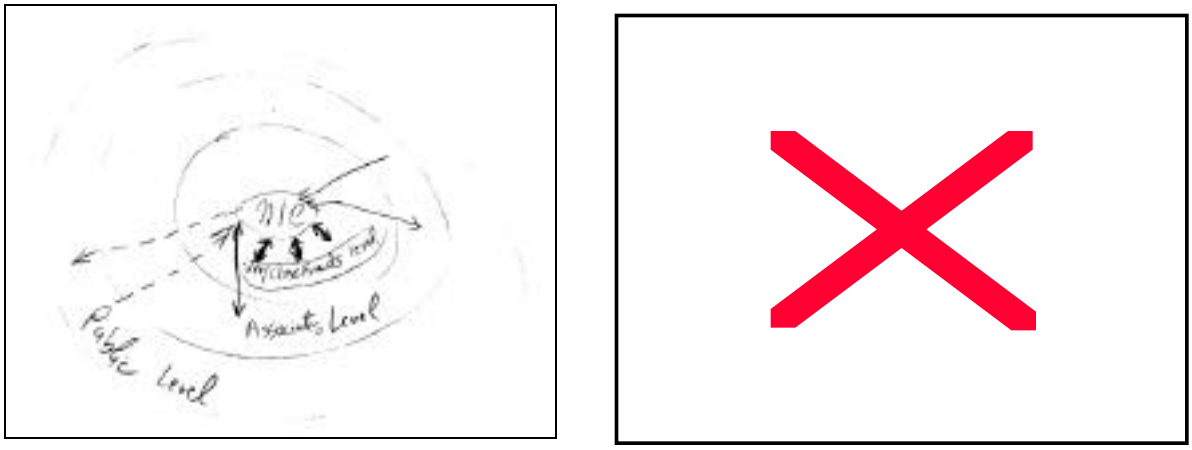

Figure 1: Drawings of "My Social Space" from a User Study

Given the emphasis in these drawings on not only people but the relationships between them, we started drawing from the field of social network analysis to model our users' social spaces. In social network analyses social items are organized depending on their relationships to each other. A network approach allows us to provide users with visualizations of social spaces from which they can infer the importance of people and groups, informal relationships between people, and dynamic grouping.

Another goal of our approach was to expect as little user input as possible. Social networks tend to be fairly dynamic, requiring continual updates to the system. Thus we expected that any user interface that relied heavily on explicitly generated user data was doomed to failure. As a consequence, we have automatically generated network measures from conversational behavior as much as possible.

We are currently developing two sample applications that explore the usefulness of person-centric social networks. The *Personal Map* creates a dynamic interface representing the user's personal social network (*Personal Map*) is based on a user's email behavior. *MSR Connections*, rather than visualizing a personal network, provides a community map for members of Microsoft Research based on connections between researchers at MSR (*MSR Connections*).

2. Person-Centric Social Networks Themes

The goals of our *Personal Map* and *MSR Connections* applications are to help users seek out social information and structure their online interactions based on conversational patterns. In sum, the design of these applications is guided by the following themes:

2.A. *People interact in the context of a social network.*

People interact with each other online in the context of a social network. According to the social network perspective, individual items (persons, etc.) are viewed as interdependent, connected by relational ties (Wasserman & Faust, 1994). The relational ties between units are channels for transfer or "flow" of resources, such as emails, files, web links, and so forth. Relational ties may vary in strength, and vary by context. These relational ties form meaningful patterns or regularities that can be assessed automatically and exploited for the user's benefit. Patterns or regularities in a network structure are dynamic, needing to be updated automatically depending on ongoing changes in relational ties.

2.B. ***People are the primary organizing unit.***

In a social context, people care about other people and their relational ties. Information about people and their relational ties should be actively employed to help people structure, or search/filter/retrieve information and objects that occupy the social space. Furthermore, people and their relationships should be prominent in a user interface.

2.C. The user is a primary organizing unit.

People tend to have an egocentric perspective when conceptualizing their social networks. They most care about how others relate to themselves. When seeking out information, navigating through a social space, or organizing items, people will generally want to start at their own place in a social network, or with people similar to or close to themselves. People in the network should be organized in terms of their relative importance to the self.

3. Sample Application: Personal Map

The *Personal Map* infers the users' personal social networks by analyzing patterns in their email behavior, and then provides the users with a visualization of their personal social networks that they can use as emailing and grouping tools. See Figure 2. Our expectation is that past email behavior will provide a fairly accurate approximation of who users care about, and how users implicitly cluster or organize contacts in their minds. Thus past email frequency and clustering behavior can be used to help future email and clustering behavior. It is also expected that patterns in personal social networks will allow users to search for specific emails or contacts, and organize their contact lists. In accordance with our person-centric social network themes, a) people are placed in the context of their social network, where relationships between people are indicated by group clustering, b) the primary unit within the map is the person, and clicking on a person or a group opens an email to that person or group, and c) the self is placed in the center of the map, and then the map is populated with those people who are most similar to that central person.

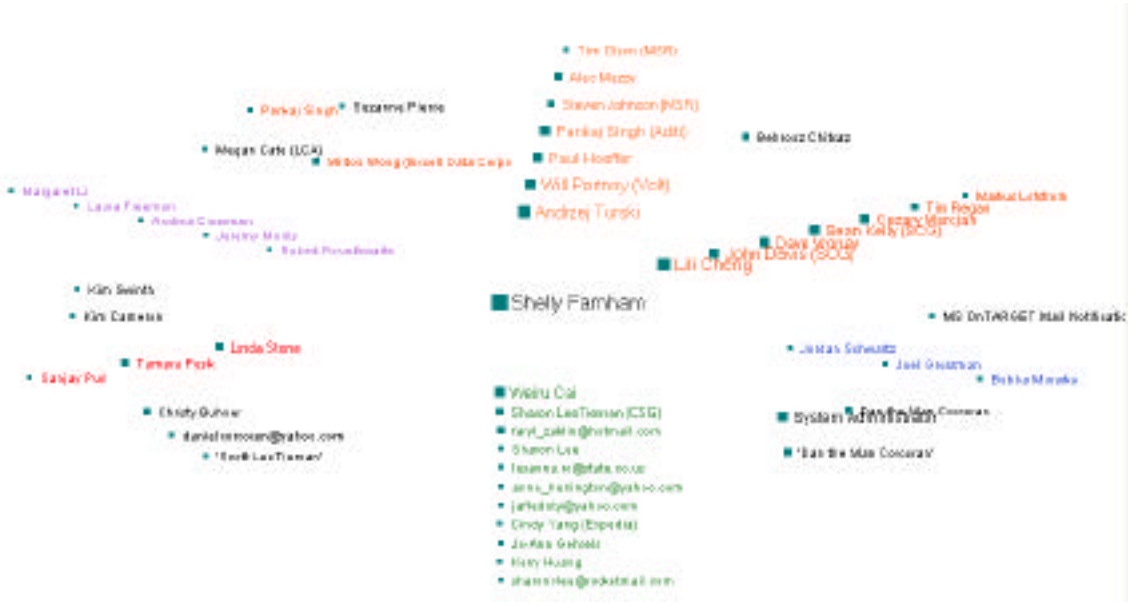

Figure 2: User-centered Personal Map, using Outlook Information

4. Sample Application: *MSR Connections*

The success of Microsoft Research depends on the exchange of knowledge and technology with the product groups. However, any knowledge transfer across the Microsoft research and product groups depends on people's awareness of who's doing what across groups. The challenge of developing an awareness of group activities is made more formidable by the dynamic, informal nature of the organization's social structures. The formal structures found through org charts and internal web pages leave out a large part of the picture. Interpersonal connections, on the other hand, provide more informal, dynamic information. Currently, such knowledge transfer across individuals depends on people's attendance at meetings and conferences or the serendipity of informal encounters. *MSR Connections* provides an online alternative, where users can develop an awareness of the informal social structures of MSR, and exploit interpersonal connections when seeking out information. *MSR Connections* provides users with a web-based, interactive map that allows users to navigate for MSR information using social, relational ties, based for the most part on online conversation patterns. See Figure 3. In accordance with our person-centric social network themes, a) people are placed in the context of their social network, where relationships between people are indicated by lines, b) the primary unit within the map is the person, and navigation through the map leads to information about the person in the center of the map, and c) upon instantiation the map defaults to the person viewing the map, and is populated with those people who are most similar to that central person.

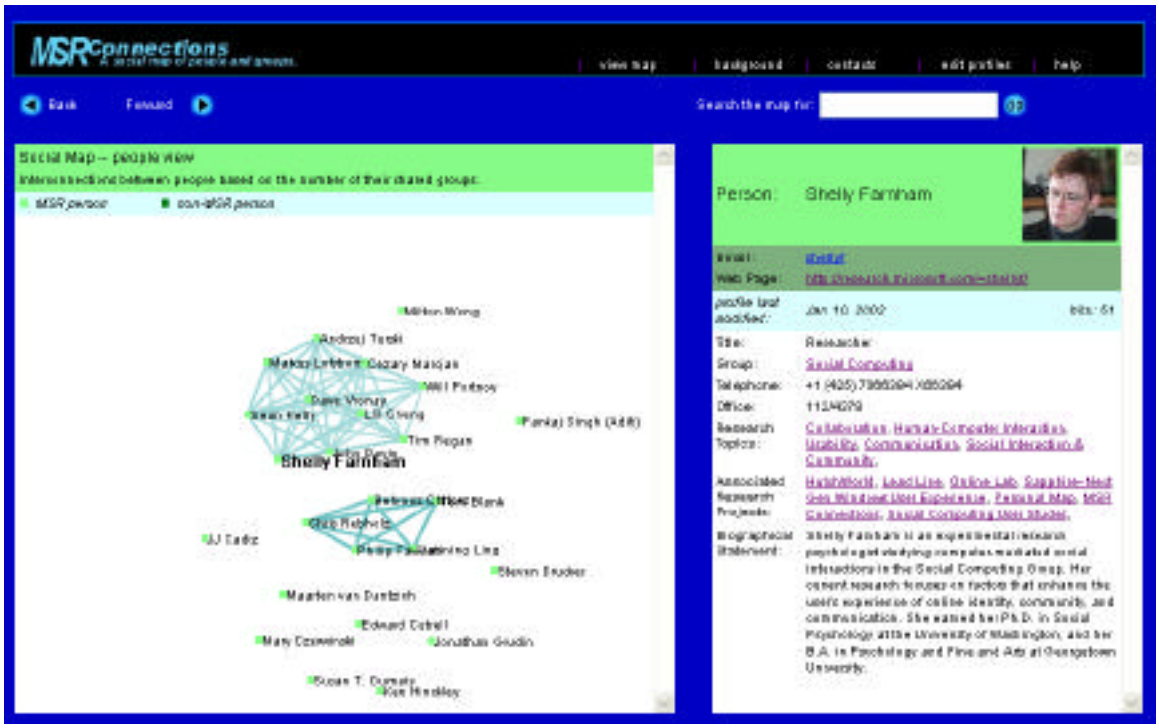

Figure 3. *MSR Connections*

5. Structure of *Personal Map* and *MSR Connections* applications

The *Personal Map* and *MSR Connections* applications have four components: a) data collection, b) processing the data into mathematical models of network and network patterns, c) representing the network and network patterns to the user through graphical or textual visualizations, and d) providing features that allow the user to interact with the network.

5. A. Data collection.

The first step is to collect and quantify information about people and their relational ties. Given the volume and complexity of information to be acquired, this should be as automatic a process as possible. The most challenging issue at this stage is operationalizing (making measurable) a "relational tie". Social network analyses also often measure relationships in terms of *similarity* between objects. For example, a count of communications between two people may indicate a strong similarity. Two emails that occur near each other in time might be considered more similar. The *Personal Map* application measures similarity between people through counts of how often their names co-occur in an email. The *MSR Connections* application measures similarity between people through counts of how often they co-occur in groups, research projects, research topics, and distribution lists. Both applications collect data from Outlook Exchange, and *MSR Connections* collects further information from the Microsoft Research internal web site.

5.B. Data processing.

Once the data are collected, they must to be transformed into a *similarity matrix* that represents for each pair of units the degree of their association with each other. Once in a similarity matrix, the data can be further organized or structured using any number of methods. The *Personal Map* uses a fairly simple clustering method--hierarchical cluster analyses (Aldenderfer & Blashfield, 1984)--to automatically organize email contacts into dynamic clusters. The hierarchical method searches the similarity matrix for the two most similar entities and then merges them together until all items that are reasonably similar to other items have been clustered. *MSR Connections* further structures social information by rendering items onto a 2-dimensional plane using an algorithm similar that of the spring model of Kamada and Kawai (1989), which minimizes the error between actual similarity values and distances between items.

5.C. Network Visualization.

People, their relationships, network clusters and network patterns may be represented to the user using a variety of methods (Aldenderfer, & Bashfield, 1984; Card, Mackinlay, & Schneiderman, 1999; Freeman, 2001; Wasserman & Faust, 1994). These methods vary in complexity, ranging from the simple link representations used in web pages, to complex diagrams modeled after neural nets. How networks are represented to the user will depend on the functionality of the application. For example, if the goal is to aid people in sorting their emails, dynamic clustering of emails could be represented to users simply through the order in which emails are listed in their inboxes. Our approach was guided by a desire to show users visualizations of social spaces that allowed them to infer the importance of people and groups, informal relationships between people, and dynamic grouping.

The *Personal Map* application was designed to foster a user's easy navigation to and interaction with important individuals and dynamic groups. As such, it employs a sector view (see Figure 2), which places the user in the center of the map, and places important people around the user based on frequency of interaction. People are clustered into sectors (or pie slices) depending on the users' email patterns, allowing users to interact with the clusters as a whole.

MSR Connections was designed to allow users to explore relationships between researchers and research projects. As such, it uses a graph view of a network of nodes and connections (see Figure 3), which places people and groups on the graph so that the strength of the relational tie between people and between groups is represented by distance on the graph, and lines are drawn for the stronger relational ties. People can easily explore the network for deeper information by clicking on the nodes and the connecting lines.

5.D. Network Interaction.

How people interact with a network, will depend on the goals supported by the application. The four forms of interaction most expected to be employed are those of a) exploratory navigation across the network, b) organization of information using network information (e.g., organize all my emails by dynamic cluster), c) communication with individuals and groups in the network (e.g., email everyone in this dynamic cluster),

and d) searches/filtering/information retrieval using network relational information (e.g., "who am I most similar to in this network?" "Who do I care about the most?" "Who is most important in this network?").

6. Discussion and Next Steps

In limited deployments of the *Personal Map* and *MSR Connections* we have found the users' responses to be fairly positive. In the *Personal Map*, people report finding the dynamic groups to be fairly accurate, and provide a meaningful way to structure the conversation space. People are provided with a tool to adjust the tolerance of the groupings (that is, how similar to a group does a person have to be before that person is classified as a member of the group), which greatly increases the map's accuracy. However, in any given personal map we have found there will be one or two individuals who do not appear to be in the right place. This indicates there is a disjunct between people's phenomenological experience of their social interactions and their actual behavior. For example, a person may not actually email her mother that often, but still feel that her mother should be closer to her on her map than her co-workers. This raises the question should we allow users to directly manipulate the map, and if so how do we incorporate user-generated map adjustments with similarity values derived from user behavior.

With the *Personal Map*, people tend to have an "Ah-hah, that's cool" response to the sector visualization. They immediately grasp the meaning of the clusters with comments such as "hey, that's the Shell team at work", and "hey, those are my friends". However aside from delighting the user with an overview of their personal interactions, it is not clear that the sector visualization adds any value or functionality than would a list view, which uses less space and also allows users to sort by clusters.

MSR Connections, on the other hand, provides a visualization that might provide information to the user that the user does not already know. It provides an overview of a social space that may allow the user to infer who's important, how people cluster together in their work, and whom in MSR they should be aware of. We are currently preparing an experimental study to test whether the map in *MSR Connections* adds any value to the user. Our concern is that the map may be too complex for the user to understand or use to navigate. We have placed the map next to profile information about the central person in the map (see Figure 3), and users can navigate throughout the social space either by clicking on items in the map, or by clicking on links in the profiles. We may find that people prefer to navigate through the social space through the profiles rather than the map. The site has been instrumented, and will be made public to Microsoft Research in February. In addition to our experimental study, we will have logs of user behavior to help answer whether the visualization provides a meaningful tool for users to infer informal social structures from interaction patterns in Microsoft Research.

References

Aldenderfer, M.S., Blashfield, R. K. (1984). Cluster Analysis. Series: Quantitative Applications in the Social Sciences, 07-044. Sage University Paper.

Backer, E. (1995). Computer-assisted Reasoning in Cluster Analysis. New York: Prentice Hall.

Card, S., Mackinlay, J., & Scheiderman, B. (1999). *Readings in Information Visualization: Using Vision to Think.* San Francisco, Morgan Kaufman Publishers.

Kamada, T., & Kawai, S. (1989). An Algorithm for drawing general unidirected graphs. Information Processing Letters, 31, p. 7-15.

Kelly, Sean Uberoi, Sung, Christopher, & Farnham, Shelly. (2002). Designing for Improved Social Responsibility, User Participation and Content in On-Line Communities. *To appear in Proceedings of CHI 2002, Minneapolis, April 2002.*

Freeman, L. C. Visualizing Social Networks. [http://carnap.ss.uci.edu/vis.html]

Wasserman, S., & Faust, K. (1994) Social Network Analysis: Methods and Applications. Cambridge, UK: Cambridge University Press.

Appendix: Profile of Shelly Farnham

Biographical Statement

I am an experimental social psychologist studying computer-mediated social interactions in the Social Computing Group at Microsoft Research. My current research focuses on factors that enhance the user's experience of online identity, community, and communication. Recent studies explore the use of profiles in online spaces, increasing social presence in shared browsing, increasing social support to cancer patients through an online community site (HutchWorld), and the development of person-centric social networks based on conversational behavior. I received my Ph.D. in Social Psychology at the University of Washington, and my B.A. in Psychology and Fine and Arts at Georgetown University.

Disciplines I am situation in:

My primary discipline is Social Psychology, which emphasizes the use of experimental, quantitative methods to explore how people think about, feel towards, and interact with others. Recently I have been integrating aspects of the fields of sociology (social network analyses) and HCI into my work.

My relevant analytical and/or design work:

For the most part, I have employed quantitative methods to answer research questions about people's online social behavior. In my work I have tried as much as possible to operationalize research questions using implicit user behavior rather than explicit self-report responses. It is a short step to go from operationalizing user behavior by instrumenting software, to feeding the information derived from instrumentation back to the user through visualizations. Relevant work includes:

Farnham, S. D., Cheng, L., Stone, L., Zaner-Godsey, M., Hibbeln, C., Abrams, J., Syrjala, K., Clark, Anne-Marie. HutchWorld: Clinical Study of Computer-mediated Social Support for Cancer Patients and their Caregivers. To appear in Proceedings of CHI 2002, Minneapolis, April 2002.

Farnham, S.D, Zaner, M., Cheng, L. (2001). Designing for Sociability in Shared Browsers. In Proceedings of Interact 2001, Tokyo, July 2001.

Kelly, Sean Uberoi, Sung, Christopher, & Farnham, Shelly. (2002). Designing for Improved Social Responsibility, User Participation and Content in On-Line Communities. To appear in Proceedings of CHI 2002, Minneapolis, April 2002.

Smith, M., Farnham, S., & Drucker S. (2000). The Social Life of Small Graphical Chat Spaces. In Proceedings of CHI 2000, The Hague, Netherlands March 2000.

Other design or analysis that I think is interesting:

Bechtel, W., & Abrahamsen, Adele. (1991). Connectionism and the mind: An introduction to parallel processing in networks. Oxford, Blackwell.

Card, S., Mackinlay, J., & Scheiderman, B. (1999). Readings in Information Visualization: Using Vision to Think. San Francisco, Morgan Kaufman Publishers.

Freeman, L. C. Visualizing Social Networks. [http://carnap.ss.uci.edu/vis.html]

Donath, J., Karahalio, K., Viegas, F. Visualizing Conversation. In Proceedings of HICSS-32, Hawaii, 1999.

Sack, W. Discourse Diagrams: Interface Design for Very Large-Scale Conversations. In Proceedings of HICSS-33, Hawaii, January 2000.