

Personal Map: Automatically Modeling the User's Online Social Network

Shelly Farnham
Position paper for CSCW 2003

ABSTRACT

The *Personal Map* provides automated visual overviews of the user's online social network, which can be used in lieu of user-generated contact lists. We expected that past email behavior would provide a fairly accurate approximation of who users care about, and how users implicitly cluster or organize contacts in their minds. A user test of the *Personal Map* indicates that people found it both compelling and accurate. We further integrated our underlying model into the new email outlook form, and got mixed results from a user study.

INTRODUCTION

Our ultimate goal was to develop a user interface that enhanced people's ability to easily communicate with individuals and groups in their conversation space. We expected that the best means for doing so would be through models of important social contacts and groups. However, we further wanted to explore whether it was possible to generate such models *automatically*. Given the difficulty people already have organizing and maintaining their contact lists, we expected that people would similarly have difficulty maintaining other forms of contact information. We further wanted to explore how our underlying models might be represented to the user graphically, and how they might be integrated into the communication stream.

We incorporate a social network approach into our designs, which emphasizes that people are organized depending on their relationships to each other.

Design Themes

The design of our automated contact management system, the *Personal Map*, was guided by the following themes.

People Should be Placed in the Context of their Social Networks

People should be shown in the context of their social networks. This is in accordance with the social network perspective, where people are viewed as interdependent and connected by relational ties [13]. Relational ties may vary in strength by context, and form meaningful patterns or regularities that can be assessed automatically.

Users Care Most about how Contacts are Connected to Them

People tend to have an egocentric perspective, and care most about how others relate to themselves. When organizing contacts, others should be placed on the map primarily depending on their importance to the user.

People Should be Organized into Groups

People naturally organize their social network into groups. However while there are consistencies in types of groups across people (e.g., friends, family), there are also a lot of differences.

The Social Network is Dynamic

Patterns or regularities in a network structure are dynamic, needing to be updated automatically depending on ongoing changes in relational ties.

THE PERSONAL MAP APPLICATION

The *Personal Map* application was developed primarily as a contact management tool where people can explore their social network. However, there is little contact management to perform, because contacts are organized automatically. The *Personal Map* infers important people and groups by analyzing the users' email behavior, and then provides the user with a social map of her email social network. See Figure 1. The goal of the *Personal Map* is to automatically model the user's social space, to facilitate communicate with his or her important people and groups.



Figure 2: *Personal Map, sector view. User is in the center, and important people are placed around her in 'pie slice' groups.*

In the *Personal Map* the user is placed in the center of the map, and her most important contacts are placed around her organized into groups (sectors) depending on their tendency to appear together in the user's emails. These groups are dynamic, changing in time depending on changes in the user's emailing behavior. Within each sector, the people are arranged along the radial dimension so that the people

more *important* to the focal point, those most similar to him or her, are placed closer to him or her. User's can adjust both the groups, to break them up into sub groups or merge them into bigger groups, and the number of people on the screen. From this view the user can start an email to any individual or group by right-clicking on the appropriate name.

Personal Map Components

The *Personal Map* application has several underlying components that collect and process email data.

Data Collection. The data collection phase downloads the email headers for all mail (sent and received) in preparation for data processing.

Data Filtering. We use a simple rule to filter out the people from whom the user only receives uninteresting (typically spam) email. We filter out anyone who has not been emailed directly by the user. We further filter depending on the date of the emails. Generally in the work context, where people email 20-40 people a day, we have found that a month's worth of email provides a good mix of short term and long-term groups.

Developing the Similarity Measure. In order to map onto the users' mental model of their groups, we measure similarity between people based on the users' tendency to put them together, rather than based on actual communication frequency between the pair. As such, the similarity measure between two people is measured by the extent to which they co-occur in emails.

People who tend to co-occur on smaller emails with fewer people should be more similar than people who tend to co-occur on larger emails with many other people. We use a nonlinear function to weight the contribution of each email to the similarity among all of the participants on that email depending on the number of people in the email.

Once we have calculated our similarity matrix, people can be clustered into groups. We use hierarchical agglomerative clustering, where newer clusters are formed iteratively by merging smaller clusters [1]. We explored other clustering methods, but found this produced the best results on our own email.

Email Integration

As mentioned earlier, the *Personal Map* application was developed as a tool for exploring automated models of social spaces. Because it is a separate application it is not integrated very well into the communication stream. We decided to incorporate our similarity and clustering information into an email client to see if it could facilitate communication. To do so, we developed a new mail message form for Outlook that incorporates lists of related people and clusters from which users might easily select items to populate the *to:* and *cc:* fields of their emails. See Figure 2.

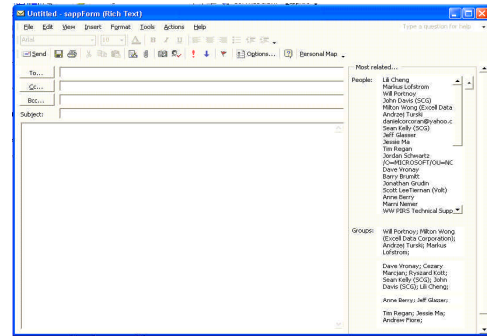


Figure 2: Outlook add-on, with lists of related people and groups that update depending on who is in the TO: and CC: fields.

The people and group lists provide ordered suggestions of the next single person or group of people the user may want to place into the TO: and CC: a new email. To start, the list is sorted by similarity to the user. The list is then updated depending on who is placed in the TO: and CC: fields. We use our similarity measure to display who the user is likely to want to add to an email at the top of the list.

USER STUDY

We completed a user study to get feedback about our underlying model of the user's social network, and test whether our suggestions for next contact in the Outlook add-on related people list (Figure 2) could facilitate email communication.

Methods

Participants

Fifteen company employees (6 male, 9 female) participated in the user study in exchange for a lunch coupon. Participants were on average 33 years of age, had been at the company an average of 4 years, and had job positions that required a fair amount of interaction with others

Procedures

Participants completed the user study in their own offices using their own computers. They were guided through a series of study tasks by the researcher with the help of a packet which outlined the steps of the study. Participants first completed a background questionnaire. Participants then installed the application, including both the *Personal Map* and the Outlook add-on. They were trained in the use of each application, and then completed a structured interview and questionnaire during which they were asked a series of questions about the accuracy and usability of the application. Participants then completed a timed task, where they emailed individuals and groups using both their standard method, and using the Outlook add-on. We then left the application running on the participants' desktops for about a week. Over the course of the week we logged usage patterns in both the Outlook add-on, and the *Personal*

Map application. At the end of the week, participants completed a final questionnaire.

Results

Our results are divided into the following sections: a) pre-questionnaire results, b) the Personal Map accuracy and usability and c) the Outlook new email add-on.

Pre-Questionnaire: Self-reported contacts usage patterns

We first wanted to know whether people tended to work with others in groups, defined as 3-10 people focused on a particular project. People on average reported working with 6 groups ($SD = 4.2$), and 33% indicated that these groups evolved a fair amount over time. People reported emailing about 22 unique people in a typical day. 85% reported spending no time organizing their contacts, and that their contacts were not very organized ($M = 2.3$, where 1 = not at all and 7 = extremely so), and 53% reported not keeping them up to date.

When asked how they normally enter email addresses in the TO: and CC: fields in email, people reported that they generally start typing names, and then use the autocomplete feature ($M = 4.0$, $SD = .85$, on a scale of 1 to 5, where 1 = never and 5 = always). The least common method was to go the contacts folder, and find the address there ($M = 1.5$, $SD = .52$).

In sum, our participants regularly interact with dynamic groups. People tended to not use, organize, or update their contact lists.

Personal Map Usability and Accuracy

After introducing people to the *Personal Map* and allowing them to make adjustments to the user interface, we had them examine the main view and then rate the extent to which it accurately represented their actual email space. Most users rated it as very accurate ($M = 5.7$). See Figure 6.

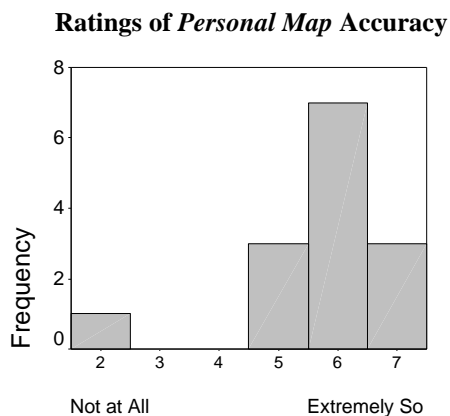


Figure 6: Accuracy of Personal Map

They further reported liking the Personal Map ($M = 5.1$), and did not find it confusing or difficult to use ($M = 2.9$).

Although the *Personal Map* accurately mapped onto how people thought about their groups, the researcher observed

from user comments discrepancies between the users' mental models and the map's underlying model. Users expected the importance and clustering information to be more determined by whom they had *sent* emails to, than whom they had *exchanged* emails with. In cases where people were clustered together through small similarity values, users thought the clusters were erroneous. It would be better to cluster based only on strong or moderately strong similarities.

Participants rated the likelihood of their actually using the *Personal Map* for a number of possible uses. See Table 2. The most highly rated possible use was to look at for fun, and the second most highly rated use was to start an email to a group of people.

To what extent do you think you would use your Personal Map for each of these possible uses?	Mean	SD
Just look at for fun.	4.7	1.65
Start an email to a group of people.	4.0	2.22
Get a sense of who is in my social network.	3.8	1.97
Find information about a recent contact.	3.7	2.30
Figure out whom I have emailed over a certain time period that I specified.	3.3	1.73
Start an email to an individual.	3.1	2.23

Table 1: Extent to which people are likely to use *Personal Map* for a variety of reasons, where 1 = not at all, and 7 = extremely so.

When participants were asked through an open-ended question which features they expected they might use, four participants said they would use it to email groups. One participant mentioned liking the groups because they might help remind her of individuals she would otherwise forget. A few mentioned it provided a useful tool for viewing recent activities. When asked why they might not use the tool, several mentioned because it was not directly integrated with their email client, it would take too much effort to use it to start a new email.

When asked how they might improve the *Personal Map*, people suggested greater integration with email client, an ability to navigate to specific emails from groups, and an ability to directly modify the people and groupings.

Outlook New Email add-on

We had people install an Outlook new email add-on so that we could assess whether we could facilitate communication by suggesting contacts using our similarity and clustering information.

Prior to viewing the *Personal Map*, we showed participants the related people list in the Outlook add-on and asked them to explain the logic behind the related people list. 80% accurately figured out it was ordered by default by frequency of interaction with the user, 20% did not know how it was ordered. We then asked whether our suggested groups made sense to them. Participants reported that they

were somewhat sensible ($M = 5$, where 1 = not at all and 7 = extremely so). Most participants recognized they were people they tended to email together, but often one or two of the groups did not make sense.

People were fairly neutral in their ratings of the add-on's related people list, on average only somewhat liking it ($M = 4.0$, where 1 = not at all, and 7 = extremely so), being only somewhat likely to use it ($M = 4.1$), and thinking it somewhat confusing or difficult to use ($M = 3.5$). Participants generally indicated that while it was useful to have a list of important, related people, they did not care for the way it updated depending on who had already been selected because it made it difficult to scan for names when the order of the names changed.

We had participants complete a timed task to assess whether using the related people list would aid or hinder communication. We had people list three people they were likely to email every day, three people they were likely to email once every week or so, and then three groups of 3-4 people they were like to email together. We timed from the moment they opened the new email add-on form, to the moment they were ready to hit the send button. Each participant completed the timed email task both using their regular methods to start an email, and using the related people list.

We found no difference in times to start an email across the two conditions for frequently used names ($M_s = 2.7s$ and $2.3s$, respectively, $F(3, 7) = .9$, ns). However participants were significantly slowed down when using the related people list when selecting infrequently used names ($M_s = 8.2s$ and $3.9s$, respectively, $F(3, 7) = 7.29$, $p < .02$). Observation of participants while completing the task indicates selecting infrequent names took more time because of the need to scroll to the bottom of the list.

When entering multiple names, again we found that people were slower when using the related people list ($M_s = 10.8s$ and $17.3s$, respectively, $F(3, 7) = 4.95$, $p < .04$). People found it difficult to scan the list for names after the first one was selected because the list updated. However people were not unilaterally slower when using the contact list to select groups. Overall, people were slower when using the contact list for 73% of the timed tasks, and slower when using their regular methods for 27% of the timed tasks. While most people expected to find the names where last seen in the list, a few people learned to trust that the most related name would appear at the top of the list, would scan for it there, and then would find it.

In open-ended responses to the add-on contact list, people were generally positive about the idea of the related people list. However, they generally held the dynamic nature of

the list in disfavor, and requested that it stay either in an alphabetical order, or ordered by frequency of interaction.

The new email form was instrumented so that we could log user emailing behavior. We found that people on average used the related people list in 7% ($SD = \%10$) of emails sent. Excluding replies (which already have a populated TO field), people used the related list in 12% of emails. We would expect if the contact list helped people email groups by suggested related people, they would be more likely to use it when emailing more than one person. People used the related people list in 5% of emails sent to one person, and 8% of emails sent to more than one person, ($F(1, 10) = 2.28$, $p = .08$, 1-tailed). The groups feature was used only once.

In sum, while people liked the idea of related people list, they did not like that it changed dynamically. Lists that change order are hard for the user to use because the spatial memory of where every item can be found is lost. However, some users learned to trust the list and found it more useful that way. A dynamic interface is not a familiar paradigm and it is possible that users just need time to learn and get used to it.

DISCUSSION AND CONCLUSIONS

The *Personal Map* is an effective tool that allows people to find important people and groups without engaging in any explicit contact management. We were able to automatically generate user models that were very meaningful and accurate to the user. Though some improvements are possible, we have a good infrastructure for generating social networks automatically.

Our first attempt to incorporate automatically generated social networks at the UI level with an Outlook add-on brought mixed results. Apparently, there is a conflict between our goal to adapt the interface dynamically and user expectations that applications have static and predictable interfaces. However some users were able to more easily email groups using the updating list.

For our next steps, we will explore better UIs for integrating dynamic relationships and groups with and user-generated contact lists, and UIs that incorporate other forms of communication such as IM.

REFERENCES

1. Aldenderfer, M. and Blashfield, R. (1974). *Cluster Analysis*. Iowa City: Sage Publications.
2. Wasserman, S., & Faust, K. (1994) *Social Network Analysis: Methods and Applications*. Cambridge, UK: Cambridge University Press.