DriftCatcher – Understanding and Utilizing Implicit Social Context in Email

Ted Selker, Andrea Lockerd

MIT Media Lab

RUNNING HEAD: DRIFTCATCHER: EMAIL AND SOCIAL CONTEXT

Corresponding Author's Contact Information:

20 Ames St.

E15-322

Cambridge, MA 02139

selker@media.mit.edu

Brief Authors' Biographies:

Ted Selker is a Computer Scientist and constructor with an interest in the use of Artificial Intelligence, Cognitive Science, and Design in creating Human-Computer Interface scenarios; he is an Associate Professor in the Media Arts and Sciences Department of MIT. **Andrea Lockerd** is a Computer Scientist with an interest in Artificial Intelligence and Human-Machine Interfaces; she is a Ph.D. student at the MIT Media Lab and completed her master's work under the direction of Ted Selker.

ABSTRACT

Many social cues that allow people to function naturally with their social networks offline are not obvious in Computer Mediated Communication (CMC). This work uses automatic social network analysis to bring some cues to CMC that foster the user's coherent understanding of the resources of their personal communication network. The DriftCatcher email client analyzes and presents some social qualities of messages, presenting users with the main idea, letting them *catch the drift* of the social undertones of their email. Our user study shows that social context in an email browser significantly increases a person's ability to make decisions about the value and importance of messages, demonstrating the value of these techniques in improving human-human communication. Modelling of social networks, visualization and Artificial Intelligence techniques can all improve people's abilities to communicate effectively with email.

CONTENTS

1. INTRODUCTION

When entering a social setting, you instantly scan the room, see who is there, and make mental notes: 'I haven't seen her lately', 'there's Bob with someone I've not met'. This example characterizes the way people naturally use social network analysis in face-to-face interactions. As people increasingly depend on electronic communication for the maintenance and management of their relationships, computers need to step up to the task of supporting and augmenting this interaction. Computer Mediated Communication (CMC) tools have not included sophisticated social cues, making it harder to maintain ones' social network online. The email client in particular is a tool that has not changed significantly in years. Most of the functionality available in modern email clients (forwarding, folders, sorting by header data) was already present in MSG, written in 1975 and widely considered the first email client (Stewart, 2002).

What do email interactions say about people's relationships and social networks, and can an automated system discover social context cues in people's online communication? In people's daily use of email, an opportunity arises for modeling aspects of their social network. This work promotes the idea of using automatic social network analysis to bring social cues to an email interface that foster the user's coherent understanding of the people and resources of their communication network.

Can a machine, that models social networks, enhance the way email is presented to a person and improve ones ability to recognize and focus on important communications? This paper will show that such goals are achievable and point to work that supports people in working with their social networks effectively. We present an email client that comments graphically on social relevance (closeness of relationship, social nature of the message, and suggested response time) and through this significantly improves a user's ability to attend to messages useful to their tasks. We have thus taken initial steps to show that automatic social network analysis can bring social cues to CMC and support a user's understanding of their communication network.

2. DRIFTCATCHER SCENARIOS

Our overarching goal is to use Artificial Intelligence (AI) of social networks to bring social cues to a user's attention and enhance their email experience. In our system, dynamic social network modeling of the SocNetServer enables the prototype email system, DriftCatcher. Before going into the details we will go through a few examples with fictitious characters to illustrate how users interact with and benefit from the system.

Meet Dr. Lori Adler. Lori uses email regularly to communicate with people from all facets of her life. Like many others, Lori finds that a large part of her day is spent doing "social network maintenance": building, managing, and keeping track of various social and business relationships. Moreover, she does a large portion this maintenance over email. Recently she started using a DriftCatcher email client powered by a SocNetServer,

and has found that it helps her prioritize her email tasks and have a better understanding of her personal social network.

Figure 1 about here

Early for a meeting: Traffic was light this morning, and Lori arrives 15 minutes early for her morning meeting. Having time to check her email quickly, she opens her inbox to find 12 new messages (Figure 1). Using the DriftCatcher CompTime feature (which shows the average time she spends composing messages to the various senders), she is able to prioritize the messages based on how much time they are likely to take her to deal with. She quickly selects and responds to message numbers 6 and 9 in plenty of time for the meeting.

Reciprocating response time: Lori has different response patterns for the various people in her social network. Her friend Felice usually responds within a few days, but her colleague Tyson usually responds within a few hours. She would like to reciprocate these response patterns, and the DriftCatcher client helps her do so with the ResponseTime bar. The time allotted for her reply is based on the response pattern of the sender. Figure 1 shows that Lori has longer to respond to Felice (message 2) than to Tyson (message 10).

Visualizing closeness: As shown in Figure 1, the DriftCatcher client portrays closeness (a function of both symmetry and frequency of communication) in the font size of the sender's name. This lets Lori easily differentiate her online strong and weak ties.

Visualizing context: When Lori is trying to decide which messages are most important, sometimes the subject line is not enough information to determine the social intention of the sender. The DriftCatcher client helps her by color-coding the messages according to their social context. Lori is able to see quickly that most of the messages are informing, but message 1 and message 12 involve planning and message 10 is an inquiry (Figure 1).

3. THEORY & RATIONALE

This section details some of the background work that should inspire the 'reinvention of email' and brought about our particular approach. In seeking to create an AI system that helps people communicate, we first looked to the field of Social Network Analysis to better understand both the findings and techniques used to quantify the quality of people's relationships and social interactions. Secondly, we look at prior work and design guidelines from the field of Human-Computer Interaction, focusing on the evolution of email tools.

3.1 Social Network Analysis and Computer Mediated Communication

Social Network Analysis (SNA) is the study of various aspects of the structure and behavior of social networks. A person's social network consists of a set of people (nodes) with whom they have ties, connections between the nodes, and resources that are

exchanged between the nodes. These resources can be information, influence, emotional support, and confidence, to name a few. This work, while not a complete social network analysis, utilizes the theories and findings of social networks as means to improve an online communication interface. Two theories most relevant to the information collected by the SocNetServer include: *social capital* (Lin, 2001), the amount of support (of all forms) which can be called upon from the people in your social network, and *strength of weak ties* (Granovetter, 1982), a group of studies which indicate that the people most important to you in terms of access to information and resources are on the outskirts of your social network.

Over the past decade or two, social network scientists have grown interested in computer networks and to what extent CMC influences social networks. For example, computer networks are especially suited for the maintenance of relationships between people who cannot meet frequently; therefore, de-emphasizing the need for locality in both work and community structure (Wellman, 2001).

3.1.1 Measuring Social Resources in CMC

Interaction Process Analysis (IPA) is an analysis scheme commonly used in studies of small groups (Bales, 1950). It classifies human-human interaction related to group dynamics (in face-to-face interactions). Bales' IPA describes a socioemotional interaction as one that shows solidarity, antagonism, tension, agreement, or disagreement, and a task-oriented interaction involves giving or receiving opinions, information or orientation (Figure 2). In this work, the term social resources will mean any resources exchanged between two people in the social network that has some social significance, covering the whole spectrum of Bales' IPA.

Figure 2 about here

It was not always obvious that the whole range of Bales' IPA can be expressed in email. Some hypothesized that the text-based medium of email would be too constraining to afford the exchange of socioemotional information. A few people addressed this question. In one study, over 2000 email sentences were labeled, by hand, using a slightly modified version of the Bales IPA categories. They showed that CMC does afford the exchange of socioemotional content, and in particular 30% of sentences in their dataset were of a socioemotional nature (Rice & Love, 1987). Another study addressed the existence of social context cues in electronic communication, and discusses how relational cues from face to face communication are translated to text based communication. They found, for example, that when communicating over email a person tends to replace a head-nod indicating agreement with a verbal phrase like 'I definitely agree…' (Walther, 1992).

3.1.2 Applications of SNA in CMC

The work of Bonnie Nardi strongly motivates systems, like SocNetServer and DriftCatcher, which integrate Social Network Analysis with Computer-Mediated Communication. The NetWORKing ethnographic study looked at how people utilize

social networks in the workplace and concluded that success in today's distributed business environment is increasingly dependent on the ability to manage one's social network. They argue that "netWORKing" (the process of building, maintaining, and activating your social network) is an absolute necessity in the modern work environment (Nardi, Whittaker, & Heinrich, 2000).

A few excellent examples of applying SNA in CMC include the following: The Referral Web system (Kautz, Selman, & Shah, 1997), finds a path between two people in a social structure using a closeness metric based on web documents; Yenta (Foner, 1997) is a multi-agent system for matchmaking, based on subject matter of email messages to suggest matches between users; ExpertFinder (Vivacqua & Lieberman, 2000) is an agent system that helps people find an expert to help them in a Java Programming domain; Coordinator (Flores, Graves, Hartfield & Winograd, 1988) is a speech-act application that tries to identify patterns of speech in an organization related to the sort of action that speech tends to induce.

By focusing on the application of automated SNA in an email system we seek to demonstrate the value of AI in supporting human-human communication. The following are some qualities that collectively set our work apart from previous work. We have a personal network approach; rather than take the point of view of a whole organization or community this work understands a social network from the point of view of a single user. While most current applications of social networks and online communication deal with information flow and task-oriented resources, the SocNetServer attempts to recognize all of the social resources exchanged between people in the network in order to better characterize relationships automatically. Additionally, the personalized dynamic social network modeling is automatic and does not require any profile input from the user.

3.2 Approaches to Presenting Email

The previous section addressed the information that might help a CMC interface, but how can this information be made useful to the user? There is a great deal of inspiring work in the field of Human-Computer Interaction (HCI), especially in terms of creative interface techniques for information representation and retrieval. The next sections look at how people handle the information overload of email and various viewpoints from HCI about augmenting this experience.

3.2.1 Organizing Email

What to do with email you have received or read; throw it away or store it? Studies of email habits show that people organize their email in very different ways (Whittaker & Sidner, 1996). To this day, Professor John McCarthy purports to store all the email he has ever received in one flat file that he searches. Some of us keep all of our messages, others like to get rid of them and keep a clean desk. Many organizations refuse to let people store email for longer than a few weeks or months deleting all history or folders of email as they go for a variety of reasons that range from difficulties with managing large

data sets to liability. The technology of email over the years has explored a few simple ways of supporting how we find, store, and view email.

In many ways, email user interfaces have only incrementally improved over the years from the original simple layout (a window with a list of emails, the first line or subject line as a place to select what to do, with other windows that show the contents of an email or other lists of emails). Many people still use simple email clients like Rmail on UNIX to this day. Such email clients are basic messaging systems with rudimentary send, receive, and find capability. Early email clients soon began to take this interface further, splitting an Emacs screen into two frames to view an email message in one and the list of messages in a separate window. Indications of state of the message (read or not) were characters lined up on tabs in the text-based email interface. The email client in Smalltalk and other Xerox PARC explorations of the early 1970s are the first truly graphical interfaces. A mailbox icon would change on the desktop to point out that email was available to read. The early Xerox Altos, D machines and Lisp machines also had extensive built-in support for classifying email into folders. This is quite similar to the capability of most modern email systems to automate various aspects of email maintenance.

Approaches to automating the organization of email have been prevalent since its introduction in the 1970s. The main idea for simplifying email has been to categorize things so that people can look though a smaller group of emails to find what they want or need. Early email filters were produced as programs that selected messages based on when the email came, who the email was from, and words on the subject line. This is still true of today's email filter systems, like Elm (Weinstein, 1992) for example, that allow a user to define a rule based on the email header or email contents.

Using an email system's built in features for designing mail filtering seems useful, but it puts a burden on the user. It requires them to build the filters and then remember what those filters are doing in the background so as not to forget about mail that is being filed away. If one is using a filter but forgets about it or forgets exactly what it does, its utility is lost. In general we believe that implicit simplifying systems can immediately help people without taking control away from them; systems that learn and change the social milieu of communication for a person.

A couple of examples of systems heading toward this goal are Maxims (Metral, 1992) and Magi (Payne, 1994) which use corpus based approaches to filtering email. Both of these systems use a collaborative filtering approach to look for similarities between the way a person and people in their peer group look at email. When people are willing to define these groups and associate themselves correctly, these approaches have been somewhat successful. Magi uses a machine learning approach to build up a model of users interest. Experiments found that the subject line did not help classification; however, the From line (a social context field) made an enormous improvement over simply using the body of the email for filtering (Payne, 1994).

Automatically identifying and highlighting social behaviors in email user interfaces frees the user from the burden of having to enter profiles and maintain filtering systems.

Our approach is to identify implicit social behaviors that are basic to online communication that might be more reliable and lasting as a model of human communication desires.

Collecting information implicitly only solves half of the user-burden mentioned previously though. The second problem with filtering is that the user is not able to easily find the filtered information. It is either deleted or put in a folder. An alternative approach to filtering and filing messages away is to mark email with annotations, which act as a visual filter. Lotus notes and many other email systems mark or color the subject line for various purposes; it might not have been read, or is otherwise important. Such annotations have the advantage of expecting the person to be involved in making decisions about it. This paper discusses an approach of using such annotative visualization as a way of presenting information to help a person make actionable decisions.

3.2.2 Improving Electronic Communication Interfaces

When making improvements to a current interface it is important to consider the user's habits with the old interface, and the pros and cons of changing this interface entirely (Raskin, 2000). A new interface can have evolutionary changes compared to the old one, thus taking advantage of the user's familiarity and knowledge of the current interface and hopefully lowering the learning curve. Alternatively, a revolutionary change in an interface could be harder to get used to initially, but reap more benefits in the long term. In our work we chose the incremental approach since the basic email interface is so ubiquitous and we could envision simple changes that introduce our social information to the interface. In this section we'll briefly go through some viewpoints related to enhancing electronic communication interfaces.

The goal of intelligent interfaces has been with us since Eliza and Parry created psychologically charged conversational systems (Weizenbaum, 1966)(Colby, 1975). These are examples of interface agents, interface software that might act as another person to help you. If the agent user interface "knows" something and wants to share it, it can do so in one of two ways. It can do what it knows should be done in your behalf as an *assistant*. Alternatively it could give you the information as an *advisor* to teach and encourage you to do it (Selker, 1994). In such agent interfaces the designer also chooses the approach based on how to least disrupt the user's work flow. All interface design must consider what a person is used to and how the new approach will be accepted. Rather than having the computer manipulate information as an assistant agent, we explore ways to present information in an advisory way to help a person understand and act on social information.

The user interface for presenting such intelligent information has many possible directions as well. A system could communicate with a user as Eliza or Parry did in text. Unfortunately in the Email domain added text will certainly add to the time it takes to read an email and risks being interpreted as coming from the author of the email. Alternative information visualization can be considered.

The field of information visualization demonstrates the possibility of improving a user's performance through a graphical interface. Presented well, the right information can be directly comprehended, making a computer interface more natural and intuitive. Muriel Cooper describes an ideal interface she termed "information landscapes" where a user finds information they need instantly and the experience of navigating is as useful as the information itself (Abrams, 1994). Such a space could make it possible to understand "where" information is in virtual storage. SemNet (Fairchild, Poltrock, & Furnas, 1999) introduced a three-dimensional graphical interface that explores techniques in the presentation of large amounts of data. This graph and node system was evocative but filled all screen real estate with the structure of the relations between the information.

There are a number interface design examples specific to electronic communication, which serve as motivating work: Conversation Map (Sac, 2000) is a Usenet newsgroup browser that does automatic content analysis. Treetables (Newman, 2001) is a tool for visualizing email threads. Babble (Erikson et al., 1999), is a communication tool for small- to medium-sized corporate groups that promotes "social translucence", providing cues about proximity and activity of other participants.

This field of information visualization is a likely place to find new ways of helping people understand social connections, but we must find a balance such that the visualizations do not take attention away from the task and use up space. Visual real estate in a Graphical User Interface (GUI) is always at a premium; maximizing its use is to be taken very seriously in design. Simple visualization techniques can help present the value of the information as well. Highlighting information with color, font, and size is a common practice in graphic layout of text on paper and is commonly used for computer interfaces as well.

As mentioned previously, this paper advocates using annotations in an automated system (instead of foldering messages away or using separate visualization imagery) and demonstrates that such a system can focus people's attention on what they want to work with. The Information Lens and Oval (Malone et al., 1987) identified three kinds of filtering for email: cognitive filtering, in which the ideas in the email are paramount; social filtering, in which the relationships between people or organizations are important; and economic filtering, in which the cost benefit analysis of dealing with the email is considered. Our DriftCatcher email client has annotations that cover all three of these domains.

4. IMPLEMENTATION

The system architecture has three main components: the DriftCatcher email client, incoming mail processing, and the social network modeling of the SocNetServer. This system was built as a prototype that demonstrates the effect of automatically analyzing and annotating social cues to improve online communication.

4.1 DriftCatcher: A Social Mail Client

The American idiom ‘catch the drift’ means to understand the main idea of something said. With the DriftCatcher email client, users *catch the drift* of their personal communication network. DriftCatcher is an extension of the Emumail webmail client (www.emumail.com) and is social in two dimensions: data collection and display. It utilizes the SocNetServer to add social context cues, which let users deal with communication in a social context in addition to the temporal context of current mail browsers. It also sends information about a user’s behavior with their social network back to the SocNetServer.

An email client, used regularly, is in a position to collect information dynamically about how the user behaves with the people in their personal network. It sees a number of implicit behaviors that characterize online relationships, such as how long users spend reading and writing messages, or how long they take to reply to a message once it’s been read. Currently, the DriftCatcher client sends information to the SocNetServer about read and compose time for each message (with a timeout if there is no typing for a few minutes), and the SocNetServer incorporates this into its model of the user’s online social network.

DriftCatcher displays the inbox along social dimensions as well as a typical temporal dimension. We considered more radical interface designs (such as a separate social visualization map). We decided that integrating the social information into a standard interface would be more likely to improve the user’s experience. Not only would it preserve people’s familiarity with the current email browser paradigm, but it would not take their attention away from working with the messages themselves. The interface is a typical email browser in which each message line is augmented with social context meta-data (Figure 1), which is detailed shortly.

4.2 Incoming Mail Handling

The email of DriftCatcher users is processed on its way to their mail server with Procmail, a UNIX mail processing utility. The script queries the SocNetServer for *social context* metadata about the message and its sender. It then adds this information to the message header. Once a user puts the script on their mail server, the system starts tracking their personal network and marking email with social context. The information is accessible when they use DriftCatcher to view their email, but the network modeling continues regardless of email client choice.

4.3 SocNetServer

The SocNetServer embodies the personal network analysis behind the DriftCatcher client. While a full model of a person’s social network would include more than email interactions, SocNetServer recognizes elements likely to be relevant to an online communication interface. It models the dynamic personal network of each user and has generic statistical models of social resources exchanged in email. Clients communicate with the SocNetServer using XML-RPC (www.xml-rpc.com).

4.3.1 Modeling Social Resources of Email

We consider whether a general model of online social behavior could automatically improve communication. SocNetServer has static statistical models of social resources people exchange in email (i.e. *informing, inquiring, sharing, planning, intimate, etc.*). The attempt to understand the social resources as well as the dynamic network features is a key element that sets this work apart from related works.

We built Support Vector Machine (SVM) classification models, which are particularly good at learning to classify in a large feature space with sparse data (Witten & Frank, 2000). The process starts with labeled examples of the intended classification. With email, the corpuses available for this are from public mailing lists and newsgroups. While these do have some variance of relationship and social context, we wanted to use a corpus that reflects real life social contact in the domain of personal email. We collected data from 8 subjects emailing each other over a month, yielding 550 email messages, which were then labeled with social resource meta-data. Thirty labels were established by considering of how Bales' IPA best translates from physical to online interaction; eight of these were represented in the corpus: *informing, inquiring, interest, keepInTouch, planning, sharing, intimate,* and *supportive.* Given this labeled training corpus we built SVMs for each category of social resource using the Weka implementation of the SMO algorithm (Witten & Frank, 2000),(Platt, 1998). The final SVM models use the feature set: terminating punctuation, frequency of punctuation, time and date related words, URLs, including the old message in the new, and the frequency of emoticons used. In cross-validation testing with 10% holdout, a standard technique used to judge the success of the models, these models had 50-70% accuracy.

4.3.2 Personal Network Modeling

Some features of a social network are dynamic and unique to a user: network structure, frequency of contact, symmetry of contact, response times, time spent reading and composing messages. In addition to having the social resource models, SocNetServer collects these dynamic elements over time, building a model of each user's personal online social network from their implicit email interactions.

4.4 Social Context in the Interface

The social context information in the message header allows DriftCatcher to augment the presentation of each message in a typical email browser (Figure 1) in the following ways.

The font size of sender's names varies based on social tie strength. Weak ties make the font bigger than strong ties, in accordance with the finding that weak ties are better for finding new information and gaining access to other networks (which are likely to have new resources and social capital) (Granovetter, 1982). Tie strength is based on both frequency and symmetry of contact, and has four resolutions.

With each message, DriftCatcher displays the average time the user has taken to compose messages to this sender (0-30+ minutes), based on messages composed with the

DriftCatcher client. This feature is expected to be useful when, for instance, a user has only a few minutes and wants to find a few messages that can be dealt with quickly.

As shown in Figure 1, the leftmost column is an indication of the time left to respond to this message (0-2 weeks). This indicates how long the user should take to respond to this message if they want to reciprocate the response pattern of the sender. The default for a new contact is two weeks, and this changes as a response behavior is established.

The background color of a message reflects the social resource classification. Green, yellow, pink, orange, and blue represent *Inquire*, *KeepInTouch*, *Interest*, *Planning*, and *Inform/Share* respectively.

5. EXPERIMENTAL RESULTS

We conducted two user studies to explore the following: Did the system learn to classify social resources in a similar way as a group of humans? Do people have a better sense of the value of communication with the DriftCatcher client?

5.1 Social Resource Classification

Six participants labeled a set of 72 messages with the eight categories of social resources; in 82% of the cases the participant's labeling matched each other. This acts as a reminder that we are dealing with a dataset in which ground truth is unclear, in many cases people will disagree about the social resources contained in a particular message.

The SVMs then classified the same set of messages. In 68% of the cases at least one person's labeling matched the machine classification. When looking just at the 82% in which there was participant consensus, the machine matched in 49%. In general, the machine was more generous in giving a message a particular label, yielding many false positives. To some extent, this was expected given our small dataset and the ground truth problem mentioned above. The question posed becomes whether a partial understanding of the social nuances of messages can improve an email interface.

5.2 Effects of the DriftCatcher Client

To test whether people have a better sense of the value of communication when using the DriftCatcher client, volunteers were given this scenario and task:

Today is your 1st day as an administrative assistant. One of the people you support is returning in 5 minutes. Go through the 24 new messages in her inbox, and choose 3 messages she should deal with first. She considers email priorities (in no particular order): People trying to make plans or things that affect her schedule; People asking for a favor or advice; Timely responses, especially to people with whom she has a close relationship

The task involves stepping into the social context of another person's inbox, and making a judgment given too short a time to read all 24 messages. Since participants

spent a short time with the interface, this shows immediate effects; we believe success here is promising for longitudinal effects.

The study was conducted with 30 participants and was counter balanced presenting inboxes with and without social context cues. The results from one-way ANOVA tests with each inbox (Figure 3) support that the DriftCatcher client helped people with the task, not by increasing the number of messages they could attend to, but by increasing the *value* of those attended to. The inboxes were not standardized or tested for equivalence, which we believe explains the non-uniform significance findings; nonetheless, significance was found with two of the three inboxes in each of the following cases. The percentage of messages read from a "close relation" as well as requiring a "quick response" increased significantly with the social client. As did the percentage of "inquiring", "informing", or "planning" messages read. As this is a preliminary result, further work needs to be conducted to confirm or deny these findings.

In addition to these quantitative results we conducted a survey of the participants for qualitative feedback. Almost every participant liked the idea of social context annotation of their email and found our annotated subject line an intuitive addition to the interface. Mentioned most frequently as particularly useful were the reciprocating response time feature and the closeness of relationship indication.

Figure 3 about here

6. CONCLUSIONS

Can a computer learn to understand the value of communication? Computers have been used to classify and separate different email messages for some time, and huge improvements in automating and deepening the value of such classification roles are possible. DriftCatcher shows that even with a small corpus, a first attempt at automatic social resource classification can improve people's prioritization of email. Can a computer use this to help a person relate to other people through technology? Our work found that the addition of these social resource labels and annotations about other personal network features to an email interface had a positive effect on people's actions.

Email is a tool that people use practically every day, making an implicit statement about their relationships with other people, and providing an opportunity for a computer to learn about their social network. Furthermore, over the years people have come to utilize and depend on email more in their daily lives, but the tool has hardly changed to help people deal with the overwhelming amount of information. Many of the social cues that allow people to naturally function with their social network are not inherent or obvious in CMC. In this work, automatic social network analysis is used to improve human-human communication, recognizing social characteristics of human relations in order to achieve a more social online communication interface.

6.1 Contributions

A computer program that recognizes the social context of a message (i.e. *informing, inquiring, sharing, planning, intimate, etc*) is in a better position to determine the value of that communication. The SocNetServer is introduced in this work, and is the social network intelligence of the system. It compiles personal social network information for a user based on their email interactions (who they communicate with, frequency of contact, symmetry of contact, response times, time spent composing/reading messages) It has statistical models of social context of email (SVMs recognize the social resources mentioned above). It also has an XML-RPC interface allowing clients to connect to it and exchange social network information about a user.

DriftCatcher, a prototype email client, serves as an example of an application, built to utilize the AI of the SocNetServer, with the goal of helping users understand and maintain their social network more naturally. DriftCatcher displays email in more than just a temporal context, adding social context cues based on information from the SocNetServer. It completes the loop by sending informing about user behavior back to the SocNetServer to be incorporated in their personal network model.

An experiment was conducted to measure the extent to which the social context of the DriftCatcher enhances the email experience. The results of this study show that having the social context mail client helped people with an email task that involved stepping into the social context of another person's inbox. In the task instructions (refer to scenario, in section 5.2), people were asked to pay attention to scheduling, inquiries, close relations, and timely responses. The data shows that with the social email client, people read a greater percentage of messages that: were from close relations, needed a quick response, involved planning, or involved an inquiry.

6.2 Future Work

CMC improved by machine intelligence has tremendous implications. This section details a few elements of future work that we find immediately obvious in this endeavor: improvements to the email client, alternate visualization techniques, and explorations of other artificial intelligence techniques for use in an intelligent email system.

6.2.1 DriftCatcher Email Client

Two aspects in particular of the current implementation warrant further exploration: the visual indications of social resource and tie strength. Currently the social resources are indicated with the coloring of messages, and allow only one type to be depicted at a time even though messages rarely fall neatly into only one type of social resource. Therefore, future work should address visualizing multiple resources at once. A second visualization choice in DriftCatcher that could be further optimized is the tie strength indication. Currently weak ties are larger than small ties (a la Granovetter), but in the user study questionnaire many people felt it would make more sense for a weak tie to have the smaller font. Different font or other visual indications of tie strength altogether are worth exploring as well. A third suggestion for a future implementation of the email client is a summary section of the current social context of the inbox. This could be placed in the top right corner of the screen and state, for example:

"You have 20 unread messages. 5 of them are inquiring, 2 are planning, 4 are supportive... Most are from people you don't talk to very often. 3 are ones that you need to get back to today."

6.2.2 Recommendations for the AI

The most significant way in which the modeling of social resources could be improved using the current SVM scheme: more training data. This model building technique would be greatly enhanced with twice as many volunteers, and the resulting data corpus should have on the order of thousands of messages. Additionally, the annotation should happen with multiple people, much like what was done in the evaluation of this work, where the context label comes from a consensus of many human coders.

While this would be the way to get better models using the current technique, this is a tedious process. Other model-building techniques could offer advantages. One promising direction is to look for opportunities to get the examples and the annotation from a user's interaction with the client, allowing the user to explicitly train the system by "showing" it different examples of various email social resources. This could be implemented on top of the current interface by adding a correction module, and might be the most natural form of training: as the system makes classifications, if the user feels strongly enough that it is wrong they will "complain" and thus train the system through correction. In any interface that is changing based on user input it is important for the user to be able to *see* that their input is causing a positive change. In the implementation of the correction interface, there should be a mode where the system shows the user its new classification of some past examples that were misclassified. This would allow the user to see the difference when they are teaching the system explicitly.

Another aspect of letting users train the system is to allow for user defined rules. While we still believe that a pattern recognition approach is the most promising in most aspects of the challenge of social context modeling, there are also cases in which users feel strongly that they know exactly what behavior they want from the system given a particular situation. According to our user study questionnaire, a better filter-maker would be enough to make some users happy. It should be complex enough to let the more advanced user specify regular expressions, but also have the ability to train by example with a more natural interface of the form: "with messages like these"…"do this". This is a delicate business, for we all know that we might act in one way and believe another. Additionally, rule based systems can be hard to maintain as mentioned previously. An automated statistics collection approach on the other hand should easily be able to recognize patterns of behavior. This should be a fascinating research area for evaluating a user's beliefs about other people. Imagine a system that might point out times when our communication might not be consistent with our beliefs. A computer could then literally attempt to save us from, among other things, hypocrisy.

6.3 Discussion

This paper demonstrates a beginning to automated social help for email. The system works in spite of the fact that it is based on a small corpus. A larger "generic" corpus could be made, and would probably improve the social context annotations in breadth and accuracy. Our choice of mappings of annotations to visualization is useful but leaves much to be explored. When would it be useful to show a map of your email related to entities or persons or on a particular topic? What should it look like? When would it be distracting? What more can we do with the presentation of a "subject line"? Should it be text, audio, imagery, animation? How much should we expect/demand of the user in helping the system learn what to do with their email?

Taking a wider lens than email for a moment, technology yields many other ways of aiding human judgment relative to communication. Weblogs, Muds, bulletin boards, RPGs, SMS and instant messaging all have different rhythms of communication use and social rules of acceptable behavior. How will we apply DriftCatcher-like experiences and social cues to these different mediums? As people increasingly depend on electronic communication for the maintenance and management of their relationships, computers will need to evolve into systems that increasingly support and augment this interaction as it changes with time. This work shows a simple and effective example of using AI to enrich the email interface. Working toward aiding people in their most intimate and human activity, social communication, has been an exciting and fruitful experience for us. We hope our experience encourages others to join in the important work that is before us in this field.

NOTES

Background. This paper is based on the Master's thesis of the second author (Lockerd 2002).

Acknowledgments. The authors wish to thank Henry Lieberman, Ronald Rice, Chris Lucas, Matt Mankins, and Win Burleson for their contributions to this work.

Support. This work was sponsored and supported by the MIT Media Laboratory.

Authors' Present Addresses. 20 Ames St., E15-322, Cambridge, MA 02139. Email: selker@media.mit.edu, alockerd@media.mit.edu

HCI Editorial Record. (supplied by Editor)

REFERENCES

Abrams, J. (1994). Muriel Cooper's Visible Wisdom, *I.D.*, Sept./Oct.

Bales, R.F. (1950). *Interaction process analysis*. Reading MA: Addison-Wesley.

Colby, K.M. (1975). *Artificial Paranoia: A Computer Simulation of Paranoid Processes*. New York: Pergamon Press.

Emumail, www.emumail.com.

Erickson, T., Smith, D.N., Kellogg, W.A., Laff, M.R., Richards, J.T. and Bradner, E. (1999). Socially translucent systems: Social proxies, persistent conversation, and the design of Babble. *Proceeding of Conference on Human Factors in Computing Systems(CHI 1999)*. New York, ACM Press.

Fairchild, K., Poltrock, S., & Furnas, G. (1999). SemNet: Three-Dimensional Graphic Representations of Large Knowledge Bases. In Card, S., Mackinlay, J., & Schneiderman, B. *Readings In Information Visualization,* Morgan Kaufmann Publishers.

Flores, F., Graves, M., Hartfield, B. & Winograd, T. (1988). Computer System and the Design of Organizational Interaction. *ACM Transactions on Information Systems*, 6 (2), 153-172.

Foltz, P. W. & Dumasi, S.T. (1992). Personized information delivery: An Analysis of information filtering methods. *Communications of the ACM*, 35 (Dec), 51-60.

Foner, L. (1997). Yenta: Multi-Agent Referral System. *Proceeding of 1st International Conference on Autonomous Agents*.

Granovetter, M. (1982). The Strength of Weak Ties. In Marsden, P. & Lin, N. (Eds.) *Social Structure and Network Analysis* (pp. 105-130). Beverly Hills: Sage.

Kautz, H., Selman, B., Shah, M. (1997). Referral Web: Combining Social Networks and Collaborative Filtering. *Communications of the ACM*, 40 (3), 63-65.

Lin, N. (2001). Building a Network Theory of Social Capital. In Lin, N., et. al. (Eds.) *Social Capital: Theory and Research* (pp. 3-30). New York: Aldine de Gruyter.

Lockerd, A.L. (2002). *Understanding Implicit Social Context in Electronic Communication*. Masters Thesis, Department of Media Arts and Sciences, MIT.

Malone, T.W., Grant, K.R., Lai, K.-Y., Rao, R. and Rosenblitt, D. (1987). Semistructured messages are surprisingly useful for computer-supported coordination. *ACM Transactions on Office Information Systems*. 5, 115-131.

Metral, M. (1992). *A Generic Learning Interface Agent*. S.B. Thesis, Department of Electrical Engineering and Computer Science, MIT.

Nardi, B., Whittaker, S., Heinrich, S. (2000). It's Not What You Know, It's Who You Know. *First Monday*, May.

Newman, P.S. (2001). Treetables and Other Visualizations for Email Threads", (Technical Report in the Human-Document Interaction Group). Palo Alto, CA: Xerox Palo Alto Research Center.

Payne, T.R. (1994). *Learning Email Filtering Rules with Magi, A Mail-Agent Interface*. MSc Thesis, Department of Computing Science, University of Aberdeen, Scotland.

Platt, J.C. (1998). Fast Training of Support Vector Machines using Sequential Minimal Optimization. In Schölkopf, B., Burges, C. & Smola, A. (Eds.). *Advances in Kernel Methods - Support Vector Learning*. MIT Press.

Raskin, J. (2000). *The Humane Interface*. Adison-Wesley.

Rice, R. & Love, G. (1987). Electronic emotion. *Communication Research* 14 (1).

Sac, W. (2000). Conversation Map: a Content-based Usenet Newsgroup Browser. *Proceedings of Conference on Intelligent User Interfaces (IUI 2000)*.

Vivacqua, A. & Lieberman, H. (2000). Agents to Assist in Finding Help. *Proceeding of Conference on Human Factors in Computing Systems(CHI 2000)*. New York, ACM Press.

Selker, T. (1994). COACH: A teaching agent that learns. *Communications of the ACM*, 37.

Stewart, W. (2002). *How Email Was Invented*. www.livinginternet.com.

Walther, J. B. (1992). Interpersonal effects in computer-mediated interaction: A relational perspective. *Communication Research* 19 (1), 52-90.

Weinstein, S. (1992). *The Elm Filter System Guide*, Datacomp Systems Inc.

Weizenbaum, J. (1966). Eliza – A Computer Program for the Study of Natural Language Communication between Man and Machine. *Communication of the ACM*. 9(1), 36-45.

Wellman, B. (2001). Physical Place and CyberPlace. *International Journal of Urban and Regional Research*, 25.

Whittaker, S., Sidner, C. (1996). E-mail overload: Exploring personal information management of e-mail. *Proceeding of Conference on Human Factors in Computing Systems(CHI 1996)*. New York, ACM Press.

Witten, I. & Frank, E. (2000). *Data Mining*. Morgan-Kaufmann.

FOOTNOTES

(Make a copy of all footnotes on a separate page here. This only has to be done for the final submission for production. During the review process, it is okay to just have footnotes at the bottom of pages.)

1. xxx
2. xxx
3. xxx

FIGURE CAPTIONS

Figure 1: (1) CompTime: the time spent composing messages to the various senders indicates that messages 6 and 9 for example are likely to be quick to deal with. (2) Reciprocating response time: the RespTime bar indicates how long you have to respond based on the senders response pattern. (3) Font size of the senders name is an indication of strength of relationship based on frequency and symmetry of contact. (4) Message Context: messages are color-coded according to social context (see the legend). Message 10 is inquiring; messages 1 & 12 are planning.

Figure 2. The breakdown of Bales IPA.

Figure 3. One-way ANOVA tests for each inbox comparing groups presented with social annotations against those without.

FIGURES

(IT'S BEST TO PUT ONE FIGURE PER PAGE)

Figure 1: (1) CompTime: the time spent composing messages to the various senders indicates that messages 6 and 9 for example are likely to be quick to deal with. (2) Reciprocating response time: the RespTime bar indicates how long you have to respond based on the senders response pattern. (3) Font size of the senders name is an indication of strength of relationship based on frequency and symmetry of contact. (4) Message Context: messages are color-coded according to social context (see the legend). Message 10 is inquiring; messages 1 & 12 are planning.

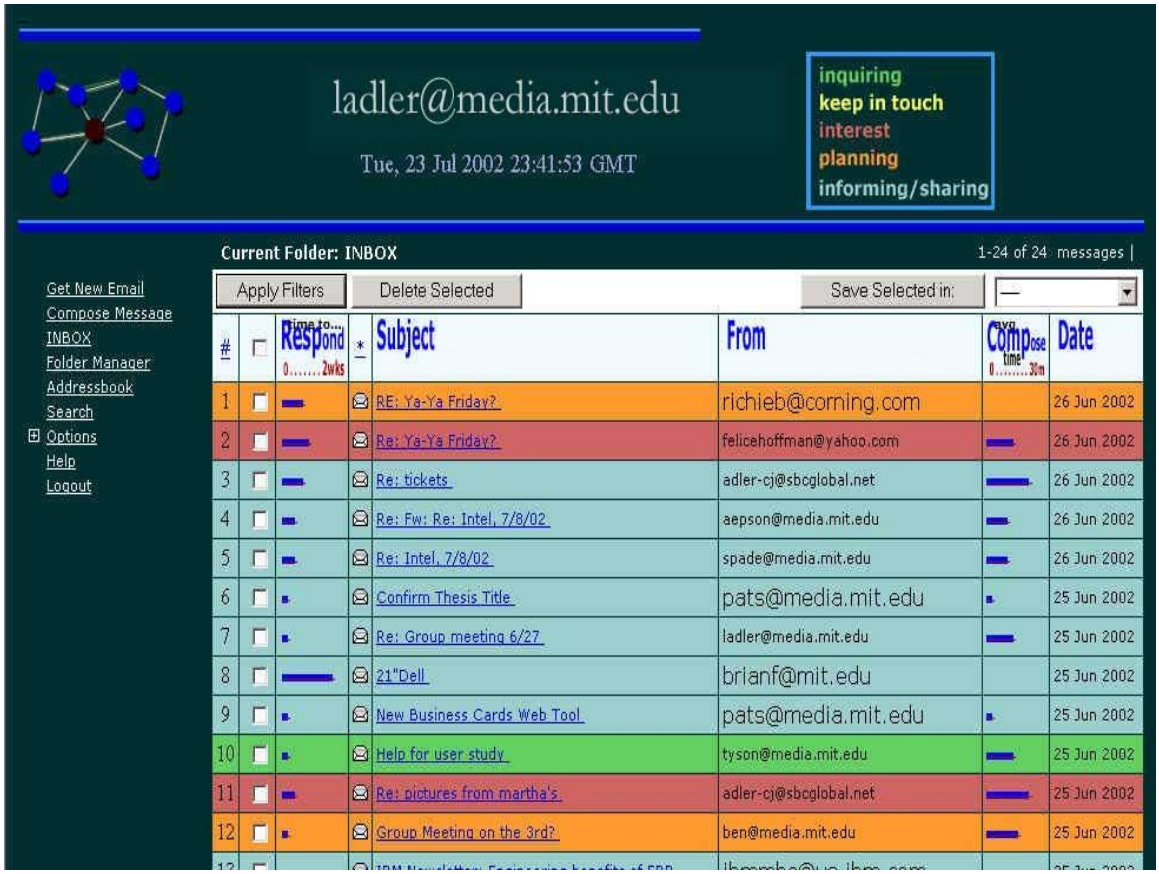

Figure 2. The breakdown of Bales IPA.

SOCIOEMOTIONAL		TASK-ORIENTED	
POSITIVE	NEGATIVE	GIVING	RECEIVING
Solidarity	Antagonism	Suggestion	Suggestion
Agreement	Disagreement	Opinions	Opinions
Releasing Tension	Showing Tension	Orientation	Orientation

Figure 3. One-way ANOVA tests for each inbox comparing groups presented with social annotations against those without.

	Box 1	**Box 2**	**Box 3**
Tot. Read F(2,27)	1.71, p<.2	.32, p<.73	.18, p<.84
%Close F(2,27)	2.71, p<.09	6.1, p<.01	.53, p<.6
%Quick F(2,27)	5.89, p<.01	3.57, p<.04	.52, p<.6
%Inquire F(2,27)	81.38, p<.00	.45, p<.64	7.1, p<.00
%Inform F(2,27)	3.22, p<.06	23.88, p<.00	29.91, p<.00
%Plan F(2,27)	2.24, p<.13	8.71, p<.00	18.51, p<.00