

Understanding Email Use: Predicting Action on a Message

Laura A. Dabbish

Robert E. Kraut

Susan Fussell

Sara Kiesler

Human-Computer Interaction Institute

School of Computer Science

Carnegie Mellon University

5000 Forbes Ave., Pittsburgh, PA 15213

{dabbish, robert.kraut, susan.fussell, kiesler}@cs.cmu.edu

ABSTRACT

Email consumes significant time and attention in the workplace. We conducted an organizational survey to understand how and why people attend to incoming email messages. We examined people's ratings of message importance and the actions they took on specific email messages, based on message characteristics and characteristics of receivers and senders. Respondents kept half of their new messages in the inbox and replied to about a third of them. They rated messages as important if they were about work and required action. Importance, in turn, had a modest impact on whether people replied to their incoming messages and whether they saved them. The results indicate that factors other than message importance (e.g., their social nature) also determine how people handle email. Overall, email usage reflects attentional differences due both to personal propensities and to work demands and relationships.

Author Keywords

Electronic mail, email, filtering, messaging, computer-mediated communication, intelligent agents

ACM Classification Keywords

H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces – *computer supported cooperative work, synchronous interaction*

INTRODUCTION

Email is by far the most popular form of computer-mediated communication. Because it has been so widely adopted for communication within organizations, people are receiving an increasing amount of email on the job [8]. To inform the design of technology that may alleviate the communication overload problem, we conducted an empirical study to understand decision rules and strategies people use to reply to, file, or delete email messages.

To help people deal with the deluge of communication,

application designers have been developing more efficient search engines [25], advanced interfaces for navigating contacts [17], and interfaces designed around task management [3]. In addition, researchers have attempted to characterize and develop tools to combat unsolicited commercial email (spam) [4, 20].

Most previous empirical research on email management describes at a general level the functions email serves and the problems associated with email overload. For example, several studies have focused on how people save their email, what purposes it serves for them, and its importance as a tool for coordination in the workplace [6, 7, 12, 14, 21, 22, 24, 26, 27].

In the current study, we build on this previous work by looking carefully at the decision rules people use in dealing with particular messages. Few previous studies have examined, using behavioral data, how people choose to reply to email messages or save or delete them. Analysis of email-related behavior as a function of message and user characteristics is important both for understanding this communication technology and for the development of automated tools to help people manage their email.

In the next section we review the previous literature on email usage. We then present a model of the main functions of email in current organizational contexts, and how these functions relate to important characteristics of email messages. In the remainder of the paper, we describe an email survey in which participants input the characteristics of new email messages in their inbox, and their actions on those messages. We use regression techniques to model message characteristics that influence recipients' action on an email message.

PREVIOUS WORK

Because email is one of the oldest uses of networked computers and one of the most popular, social scientists have long examined how people use it. Sproull and Kiesler [21] provide a summary of much of the early work on the social and organizational aspects of email. Here we will focus on work about email and information management strategies, as well as research dedicated to alleviating the problem of "email overload."

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2005, April 2–7, 2005, Portland, Oregon, USA.

Copyright 2005 ACM 1-58113-998-5/05/0004...\$5.00.

Email as a Task-Management Tool

As early as 1988, Sumner [22] examined how email was being used in organizations, by interviewing and surveying users at an organization with an electronic email system in heavy use. She found that email was displacing previous communication modalities and warned that access to electronic mail systems might lead to information overload.

Mackay [14] observed that people used email in highly diverse ways, and Whittaker and Sidner [27] extended this work. They found that in addition to basic communication, email was “overloaded” in the sense of being used for a wide variety of tasks—communication, reminders, contact management, task management, and information storage.

Ducheneaut and Bellotti [7] performed a study of email usage in three organizations and found, as had previous authors, that email was being used for a wide variety of functions. In particular, they noted that people used emails as reminders for things they had to do and for task management more generally.

Individual Differences in Email Handling

Mackay [14] also noted that people fell into one of two categories in handling their email: *prioritizers* or *archivers*. Prioritizers managed messages as they came in, keeping tight control of their inbox, whereas archivers archived information for later use, making sure they did not miss important messages.

Whittaker and Sidner [27] also examined how people responded to the abundance of electronic mail they received. According to them, people fell into one of three categories, depending upon the strategy they used for handling email: *frequent filers* constantly cleaned their inbox, *spring cleaners* cleaned their inbox once every few months, or *no filers* did not clean up their inbox and used search tools to manage it. Research to identify and compare different strategies for email management has continued.

Extending Whittaker and Sidner [27], Bälter [2] developed a mathematical model using keystroke-level analysis to examine the time necessary to use each organizational strategy. Tyler and Tang in a recent interview study identified several factors that may influence likelihood of response [23]. These empirical studies were qualitative, generally based on 10 to 30 interviews. The current research extends this work, using a larger sample and statistically examining users’ decision rules. Although Venolia et al. [24] used a large-sample quantitative survey technique, their results represented people’s theories about their email habits and not data on specific user behaviors and actions in email, as we present in our results below.

Technologies to Facilitate Email Handling

Implementation-oriented research has attempted to design and deploy email systems that help people deal with the deluge of email and better support the tasks email serves [3, 15, 17, 18]. Much of this work has focused on intelligently categorizing messages and determining what is important to

the user. For example, Bellotti et al. [3] designed TaskMaster, a system that supports the use of email messages as task reminders. TaskMaster [3] allows users to group their email specifically by its relationship to active tasks. Malone et al. [15] experimented with the ‘Information Lens,’ an intelligent information-sharing system centered around informational aspects of content important to users. Even though particular features of these systems appear valuable, attempting to combine the elements from them may increase the overload problem.

Machine learning and artificial intelligence techniques have also been applied to the problem of email overload. Intelligent agent-based systems have been created to extract content information from email and create meaningful summaries for users [1, 4]. Machine learning techniques have been used to filter out “spam” [20]. Horvitz et al. [9] have developed a system aimed at inferring the criticality of messages and prioritizing email received. The goal of these systems is to facilitate dealing with large numbers of emails in a short amount of time.

None of the intelligent systems discussed here, save spam filters, have been widely adopted. Although much of this implementation research has collected quantitative data on email usage, it typically involves small samples of users and does not provide general models that contribute to our understanding of users’ actions with email.

Message Importance

It is plausible that the perceived importance of email messages influences how quickly people respond to them, whether they delete them or store or file them for later use. However, studies have not examined what aspects of a message influence a user’s evaluation of its importance. Systems such as [9], which prioritize messages based on inferred importance, make assumptions as to what constitutes an important message. Currently there is little evidence about the message characteristics that make a message important. One of the goals of this paper is to assess what makes an email message important and how importance influences the way people act on the message.

Previous researchers as well as the designs of commercial email systems, such as Outlook or Eudora, have assumed that by including a priority field in email messages, users can more effectively manage their email by attending to important and critical messages first. However, email users often ignore the priority field [27]. The problem with this approach is that senders must assign priority to a message manually, a tedious process. In addition, the priority field reflects only the senders’ priorities, giving recipients little reason to attend to it. An intelligent system that could accurately infer the importance of a message to the receiver should help users efficiently attend to critical messages. To understand how a user evaluates the importance of a message we looked to see which message characteristics were associated with users’ assessment of the message’s importance.

A MODEL OF EMAIL ACTIVITIES

Based on our review of the literature, we developed a conceptual model of the main purposes email serves in an organizational context. The literature highlights four distinct uses for email in organizational contexts: task and project management, information exchange, scheduling, and social communication. Each of these uses can be associated with different message content types. The goal here is to identify features of email messages that users are likely to attend to in using email.

Project management, task delegation, and reminders

As Whittaker and Sidner [27] found, users often treat their inbox as an external memory store, with messages in view serving as reminders [3, 7]. Flores et al. [9] also highlight the importance of task management in email. Because email is used for task and project management, the following message content types are likely to determine the importance people attach to messages and the actions they take on them: action requests, status updates, and meeting and deadline reminders. Based on the previous research, we predict that messages that request an action or provide a reminder will be left in the inbox rather than filed or deleted.

Information exchange, storage, and retrieval

Previous research also shows that the ability to archive information is one of the primary reasons users save messages [26, 27]. Email is often used to ask questions, request a document or web link, or reply to requests, as well as to conduct substantive discussion. People are likely to store messages containing important information for later retrieval. This informational use of email suggests that information requests or responses will influence the importance people attach to a message and the actions they take on it. We predict that information requests and responses will be less likely to be deleted by the user, and more likely to be left in the inbox or filed.

Scheduling and planning

As Sumner [22] noted, email is often used for scheduling or planning purposes. This could be to schedule a meeting, event, informal occasion, etc. This purpose of email suggests the following message content types are likely to determine the importance people attach to messages and the actions they take on them: meeting requests, and responses to meeting requests. We predict that a high proportion of email messages will be related to scheduling, and that these messages will be more likely to receive a response. However, some meeting-related messages become obsolete after the meeting occurs, and so these messages are unlikely to be filed.

Social communication

As email becomes increasingly integrated into organizational life, people have come to expect instantaneous delivery of and rapid response to email communication. Though email communication is

asynchronous, in many firms employees read and respond to their email throughout the day, and exchanges can be almost as rapid as one might expect from instant messaging communication. In addition, people now use email to communicate with family and friends as well as work contacts. Thus social content within a message may increase likelihood of response.

By considering the different purposes that email serves in an organization we identified six key message content elements, adapted from the categories presented by [5, 9]. These are: action requests, status updates, reminders, information requests and responses, scheduling requests and responses, and social content. In addition, we made specific predictions about how these types of content should relate to the actions taken on a message.

METHOD

There are several ways to build models of user behavior in any domain. One is the machine learning approach, which builds statistical models from records of users' interaction with their messages, e.g., Horvitz et al. [10]. Another approach is statistical analysis of interviews or surveys. The advantage of the first approach is that models are based on observed behavior, and a great deal of data can be obtained from each respondent, both cross-sectionally and over time. The downside is that machine learning models are often black boxes that provide little insight into the cognitive and work processes that email supports. They might show, for example, that a user saves messages from a certain sender, but not why this sender is important. We chose to use a survey approach to collect data on individuals' actions with their email because the results provide us with general models of human behavior and allow us to see, in human-understandable terms, how work processes mediate between message features (such as the number of addressees on a message) and the actions the recipients take on a message. In addition, machine learning techniques typically use small samples of users (e.g., Lashkari et al.'s research is based on two users [13]) because they are often intrusive. The non-intrusive survey approach allowed us to obtain a larger sample of respondents, increasing the generalizability of our results.

Survey

The survey itself was web-based, completed over an Internet browser, and was divided into three sections. The first section collected information about the work context, focusing on the nature of the respondent's job (the number of projects the respondent works on, their number of subordinates, and their feelings of time pressure at work). The second section asked questions about the respondent's general patterns of email use--the number of email messages sent and received, number of messages in the email inbox, and general email habits.

The third section of the questionnaire asked for detailed information about five new non-spam messages in the

respondent's email inbox. For each of the five email messages the respondent indicated the nature of the content (see Table 1), the importance of the message, characteristics of the sender, and the action taken on the message (replied, plan to reply later, do not plan to reply), and what they did with the message (delete, file, or leave in Inbox) (see Fig. 1). We will describe these measures in more detail below.

Measures

Job Complexity

To assess an individual's job complexity, respondents specified the number of their subordinates, the number of their scheduled meetings per day, the number of research grants they were involved with, and the number of courses they taught. The four job-complexity items were combined to form an overall job complexity score ($\alpha = 0.67$).

Message Importance

We hypothesized that users draw on characteristics of the message (sender and content) to determine a message's importance, and that importance in turn influences their action on the message. To assess the importance of a message to a recipient's work we constructed a four-item measure index of message importance for this survey. Respondents rated each of the following questions on a five-point Likert scale:

- *How important is this message for YOUR work?*
- *How important is this message for the SENDER's work?*
- *How urgent are any deadlines associated with this message?*
- *How much work does this message require of you?*

Sender Characteristics

Because we hypothesized that the sender of the message was important in determining likelihood of response, survey respondents described their relationship with each message sender. To evaluate whether the sender was a work contact or not, we had respondents select the sender's role from a list including administrative assistant, co-worker, supervisor, friend, etc. These responses were then coded into a binary variable with 1 for a work-related contact and 0 for non-work-related contacts such as family, friends, etc. Respondents also indicated how frequently they typically communicate with the sender and the number of other recipients on the message.

Message Content

We hypothesized that certain message content, particularly task management and delegation, scheduling, information exchange, and social communication, would change a message's importance. Respondents coded the content type of each message using the categories in Table 1. Content types were not mutually exclusive, e.g., a message could contain both a scheduling request and a piece of information.

What was the content of this message? Select 'yes' for all that apply:	
-	Request for action
-	Request for information (link, contact information, etc.) or a document (file, image, etc.)
-	Status update for an ongoing project or task
-	Request for a meeting or other communication with you, or response to a meeting request
-	Reminder for a meeting, event, or upcoming deadline
-	Social greeting or thank you

Table 1. Message content types.

Message Actions

Finally, respondents described how they acted on the message. For each message respondents indicated whether they deleted it, left it in the inbox, or filed it, and whether they replied immediately or planned to reply later to the message. Through these questions, respondents acted as classifiers, providing detailed categorization of messages in their inbox. We used these data to predict user action on specific messages as a function of message characteristics such as importance, relationship to sender, content, etc.

Participants

The survey was sent via email to a college-wide distribution list at Carnegie Mellon University containing over 1100 email addresses. There was no incentive offered for participation. We do not know the number of inactive email addresses on the list or the number of individuals who actively read it. One hundred and twenty one individuals (11% of the subscribers) completed the survey in its entirety. The survey took an average of 26.7 minutes to complete (standard deviation of 9.7 minutes).

Out of the 124 respondents, 38 were professors or scientists (30.7%), 40 were other staff members, such as research programmers, etc. (32.2%), and 46 were graduate or undergraduate students (37%). Participants ranged in age from 20 to 57, with the average age being 30. A majority of the respondents were male (76%), reflecting the wider population demographic of the college sampled.

RESULTS

Overview

Sender characteristics and message content influenced users' perceptions of message importance. Importance, in turn, influenced how people responded to a message. However, people also responded to messages they did not consider important, suggesting that other sender and message characteristics play a role. In this section, we first present general statistics describing the population with respect to job complexity and email usage, and then discuss our models of importance and message response. We next propose a model of message response incorporating users' perceptions of message importance and message characteristics. Finally we examine influences on where messages end up, and the role of individual differences.

Measure	Mean (S.D.)	Median (N=121)
Messages read per day	30 (17)	25
Messages sent per day	14 (12)	13
Number of Inbox Messages	1336 (2785)	105
Number of Email Folders	22 (12)	25
Times checking email per day	19 (11)	13

Table 2. Summary of general email usage characteristics.

Basic Email Statistics

Table 2 summarizes general statistics about email usage in this sample. Overall, respondents sent 14 messages per day, read 30 of them and kept over 1300 in their inboxes. These numbers are overall means, but email usage varied based on job role. For example, the professor/scientist portion of our sample reported reading significantly more messages per day than the students or other staff ($F [2,118] = 10.12$, $p < 0.0001$; Professor/Scientist $M = 40$ messages, Student $M = 22$ messages, & Staff $M = 27$ messages).

There was an extremely wide spread for number of messages in the inbox, with 50% of individuals having 105 messages or fewer in their inboxes, 25% of individuals having 1050 or more messages in their inboxes, and 2.5% of individuals having 10,000 or more messages in their inboxes. The majority of individuals (75%) had fewer than 1000 messages in their inboxes. This may be due to storage limitations in this environment. However, the small size of the inbox and the high number of email folders suggests that many people are filing their messages into folders.

Email Habits

We gathered data on general email habits, particularly filing behavior. Our results are similar to Whittaker and Sidner's [27], who indicated that their sample of 11 individuals fell into three categories: frequent filers, spring cleaners, and no filers. In our survey, participants were asked about email behaviors on a five-point scale ranging from never (1) to always (5). Table 3 shows the percentage of respondents who report performing each behavior often or always.

Email Habits Item	%	Classification
"I try to keep my inbox size small."	27%	Frequent filers
"I file my messages into folders as soon as I have read them."	32%	Frequent filers
"I leave messages in the inbox after I have read them"	63%	Spring cleaners and no filers

Table 3. Percent respondents reporting often or always for email habit items and correspondence to categories in [27].

Message Level Data

Respondents entered data on each of five email messages in their inbox for a total of 581 messages. Data for each email included sender, content, and action on the message

Message Content Distribution

Figure 1 summarizes the distribution of messages among the various content types. It is important to note that messages could contain more than one type of content. A

high proportion of messages contained requests for action (34%), highlighting the task delegation function of email. Information in the form of an attachment, web link or phone number also characterized a large proportion of messages (36%), pointing to the key role of email for information exchange and storage. Surprisingly, scheduling content was present in only 14% of messages, even though previous research [3, 7, 27] suggested this had become one of the primary functions of email. Only 8% of messages contained social content.

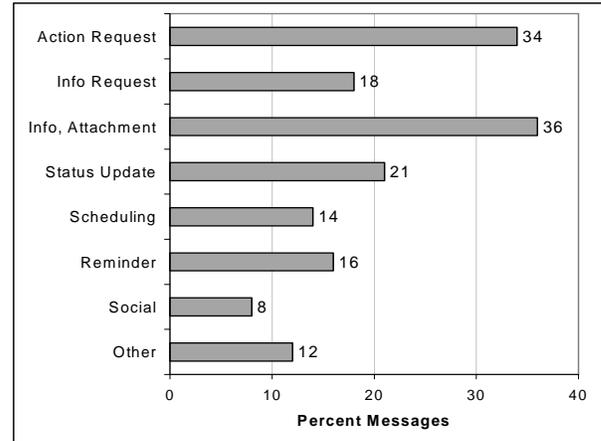


Figure 1. Distribution of message content across types. (Content types are not mutually exclusive.)

Actions on a Message

There were several possible actions people could take on a message. Primarily we were interested in two types of actions: *location actions*-- whether a user ends up storing a message after processing (whether they file, delete, or leave the message in the inbox), and *reply actions*-- whether the user had already replied to, planned to reply to, or did not plan to reply to a message.

There were three possible resting places for a message: users could file a message into a folder, delete it, or leave it in the inbox. Figure 2 shows the breakdown of locations by reply action. Overall respondents deleted 24% of their email, filed 27% into folders, and kept 49% in their inbox. It is notable that for messages where a response was

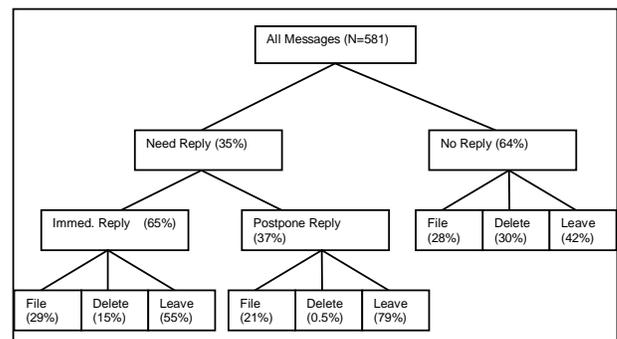


Figure 2. Distribution of messages by reply and location action: categories are mutually exclusive.

postponed, 79% were left in the inbox. This suggests that individuals were using the presence of those messages as a reminder to respond, as previous literature suggests [3, 7, 27]. At the same time, people were twice as likely to delete messages that did not need a reply (30%) as those that did (15.5%). Regardless of reply action, people were more likely to keep messages in their inbox than to file or delete them; a high percentage of messages were left in the inbox (49%). We examine the factors associated with leaving a message in the inbox in a following section.

Out of the 581 messages, respondents felt that a striking 64% did not require any reply, 23% required an immediate reply, and 13% required a reply that they postponed (Figure 2). There may be message characteristics that differentiate messages that respondents felt required a reply from those that did not require a reply. We consider the characteristics that predict message response in detail below.

Predicting Importance of a Message

Importance of a message may play a significant role in users' action on a message. We used regression techniques to predict a message's perceived importance from characteristics of the recipient, the sender of the message, and the message itself. Because each respondent provided data on several email messages, we used mixed model regression to account for the non-independence of messages

	Base Model	Model 1 Job Char	Model 2 Job Char + Sender Char	Model 3 Job Char + Sender Char + Message Content	Impact on importance score (percent change)
Independent variable	Beta	Beta	Beta	Beta	
Intercept	2.28***	1.66***	1.88***	1.68***	---
Job Complexity		0.26***	0.22***	0.19***	+ 9%
No. of Recipients			-0.23***	-0.22***	- 10%
Comm. Frequency			0.05*	0.05*	+ 4%
Work Relationship			0.56***	0.45***	+ 23%
Action Request				0.38***	+ 20%
Info. Request				0.19*	+ 10%
Status Update				0.35***	+ 18%
Scheduling				0.38***	+ 24%
Reminder				0.25**	+ 13%
Social Message				-0.62***	- 32%
R-squared	0.26	0.27	0.40	0.53	

* $<.05$, ** $<.001$, *** $<.0001$

Table 4. Predicting importance score of a message.

within respondents, including respondent as a random variable. We tested three models of message importance: one with only complexity of respondents' jobs as a predictor (Model 1), one with job complexity and sender characteristics as predictors (Model 2), and one with job complexity, sender characteristics, and message content types as predictors (Model 3). Individual differences accounted for 26% of the variance in the perceived importance of a message, indicating that some people rate their email messages as more important on average than others do.

In Model 1, the complexity of the respondents' jobs significantly increased how much importance they assigned their messages (by 9% in the full model), accounting for approximately a third of the individual difference variance. Possibly, busy people have a self-flattering view that their work is especially important. It is more likely, however, that email messages are more important for people with complex jobs, because they are using the email to coordinate the multiple projects in which they are involved.

Model 2 shows that sender characteristics explain an additional 13% of the variance in message importance scores. Respondents perceived messages with fewer recipients to be more important. In contrast, having a work relationship with the sender and having more communication with the sender in the past increased the importance of a message.

Model 3 adds message content to the model in predicting message importance. The increase in the variance accounted for by Model 3 over Model 2 shows that content plays a major role in determining message importance (Model 3 R-squared = .53 versus .40 for Model 2).

Action request, information request, status update, and reminder message content increased the ratings of a message's importance. In contrast, social content decreased a message's rated importance. This finding indicates that participants rated messages based on their importance to work; social messages were typically non-work-related and rated unimportant.

The strong influence of requests for action and scheduling on a message's perceived importance indicates the impact of others' demands on how people direct their attention on the job. New action requests were a significant proportion of all message content (34% of messages, Figure 1) and increased importance ratings by 20%. These messages may cause people to shift gears and to add new tasks to their current stack. Scheduling was included in just 14% of messages (Figure 1) but increased the importance score of a message substantially, by 24%. Previous literature on managerial behavior has shown the importance of scheduled meetings for managers. Managers spend up to 80% of their day in communication with others, much of this communication in scheduled meetings [16, 19]. Action requests and scheduling content also affect storage; people keep these in their inbox as reminders [3, 7, 27], and they

foresee needing these messages for a longer time period than other types of messages.

Predicting Message Reply Action

We were interested in examining influences on the decisions people make regarding the fate of an email message, particularly whether they will reply to it or not. Recall that respondents indicated that only 36% of the messages they received required a reply.

We created a binary variable for message reply that was equal to 1 if the respondent reported replying to the message immediately or planned to reply to the message, and was equal to 0 if the respondent did not say the message required a reply. We used the same repeated measure mixed-model analysis of variance conducted previously, adding perceived importance as a predictor variable. This analysis treated the response variable as continuous, and thus, for our response variable that ranged from 0 to 1, it predicted the probability of the response reaching a value of 1. The independent variables in the models may be interpreted as increasing or decreasing the probability of reply to a particular message.

We hypothesized that several categories of message characteristics influence the reply decision: perceived importance of the message, sender characteristics, and message content. Using multiple regression, we tested a model including these message characteristics to predict message response (Table 5). Scaled estimates were used for the beta coefficients of the independent variables shown so that the size of each would reflect the relative impact on our response variable, probability of reply. Individual differences accounted for 15% of the variance in the probability of response to a message.

Model 1 in Table 5 indicates that people were more likely to respond to messages they rated as important, but importance ratings accounted for a small amount of the variance in reply action. Holding all other variables at their average levels, a one-standard-deviation increase in the importance of a message increased the probability of reply by only 7%, indicating that whereas importance increased the probability of responding, its influence was small.

Sender characteristics accounted for an additional 15% of the variance in likelihood of replying to a message. As expected, people were less likely to reply to messages with more recipients, over and above the influence of message recipients on importance. Messages sent to a list of individuals were 18% less likely to receive a reply, but messages with only one recipient were 20% more likely to receive a response. Although communication frequency and having a work relationship with a sender increased the importance assigned to a message, these variables did not increase the probability of replying. Indeed, having a work relationship with the sender actually lowered probability of response by 9%. It is possible that these work-oriented messages may have become routine.

	Model 1 Import+ Sender Char + Message Content	Model 1 Impact on prob. of reply (percent change)	Model 2 Import+ Sender Char + Message Content	Model 2 Impact on prob. of retention (percent change)
Independent variable	Beta		Beta	
Intercept	0.34***	---	0.63***	---
Importance Rating	0.14*	+ 7 %	0.16**	+ 8%
Many Recipients	-0.21***	- 18 %	-0.07	---
Comm. Frequency	0.008	---	0.02	---
Work Relationship	-0.045*	- 9 %	0.01	---
Action Request	0.019	---	0.03	---
Info. Request	0.11***	+ 22 %	-0.03	---
Status Update	0.009	---	-0.02	---
Scheduling	0.011	---	0.01	---
Reminder	0.004	---	0.02	---
Social Message	0.12**	+ 23 %	0.01	---
Reply Immediately	n/a	n/a	0.03	---
Postpone Reply	n/a	n/a	0.12***	+ 23%
R-squared	0.37		0.55	

*<.05, **<.001, ***<.0001

Table 5. Predicting probability of replying to a message (Model 1) and probability of retaining a message (Model 2).

Message content types accounted for only an additional 4% of the variance in likelihood of response. Information requests were 22% more likely to receive a reply than other messages. Surprisingly, messages with social content were 23% more likely to receive a response, even though respondents considered them relatively unimportant, compared to the work-related messages.

From the analyses presented above, we tested a summary model integrating the relationships among importance, response, and related message factors. This model is presented in Figure 3. Visual inspection of the model reveals that a majority of the message content factors indirectly affect message response through the respondent's perception of message importance. Many of these factors—action requests, status updates, and reminders—relate directly to email's purpose as a task management and delegation tool. Scheduling content also has a major impact on message importance but does not influence probability of response directly. At the same time certain message content such as information requests and social messaging greatly increased the likelihood of response, even though they have small or negative effects on perceived importance.

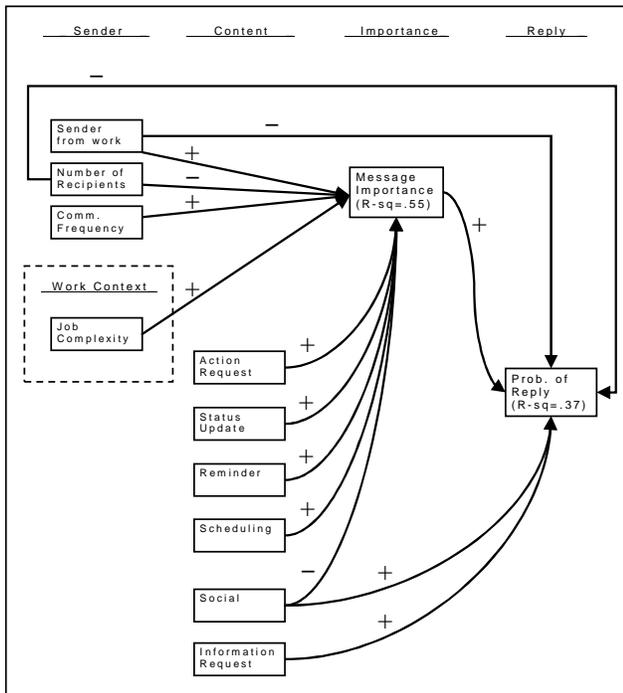


Figure 3. Model showing direction of influence of message characteristics and work context on probability of replying.

Sender characteristics also had direct impact on probability of reply. An increase in the number of recipients directly decreased the probability of response by 18%, whereas its effect as mediated through importance was only a 1% decrease. Interestingly, although a work-related sender increased the importance of a message by 25%, this actually decreased probability of response by 9%.

Predicting Message Retention

As Figure 2 shows, a high proportion of messages in our study were left in the inbox (49%). We were interested in whether respondents were deliberately leaving certain email messages in the inbox, as this behavior has been commonly cited in previous qualitative research and potentially contributes to overload in email use [3, 7, 27]. Several people in our exploratory interviews cited this behavior and in general reported keeping many items in the inbox.

We looked at whether particular message characteristics were associated with leaving a message in the inbox. A binary variable ‘Left In Inbox’ was set to 0 if a user filed or deleted the message, and set to 1 otherwise. We analyzed the data using the same repeated measure mixed-model analysis of variance conducted previously, adding response action as a predictor variable.

Individual differences accounted for a large portion of the variance in the probability of a message being left in the inbox (48%), suggesting stable differences among respondents. For example, some people keep messages in their inbox regardless of message content, as an overarching strategy [27].

We tested a model to see whether message characteristics—the perceived importance of the message, sender characteristics, message content, and reply action—predicted keeping the message in the inbox. Model 2 in Table 5 shows that people were more likely to keep messages they considered important in the inbox, increasing the probability of message retention by 9%. We included dummy variables for whether a user replied to the message immediately or postponed reply to the message. Results are consistent with prior research [3, 7, 27] suggesting that people keep messages in their inbox as a reminder to themselves to act on the message. People were substantially (23%) more likely to retain a message for which they had postponed a reply compared to messages they replied to. As Figure 2 showed, people left 79% of messages in their inbox when they postponed a reply, but only 55% when they replied immediately. This finding implies that people process some comparatively unimportant messages more than once because they have postponed them.

We hypothesized that messages requesting an action and reminder messages would be left in the inbox as a memory aid, but the results did not support these hypotheses. None of the message content types in Model 2 predicted retaining a message in the inbox. These types of message content may only influence message retention in a mediated fashion through their effect on message importance.

Sender characteristics also did not significantly affect the probability of leaving a message in the inbox, with only minor increase in the variance explained (R-squared= 0.02).

DISCUSSION

Much of the defining work on email usage in organizations was conducted around ten years ago, just as email started to become prevalent in the workplace [11, 13, 21, 22]. Although new communication media such as instant messaging are increasingly used at work, email remains a principal means of organizational communication and information transfer. A main goal in our survey was to understand email behavior in today’s workplace and to define areas of interest for further study. To do so we obtained general information about email usage, and data on how people differentially attend to particular messages. The message level data allowed us to examine how characteristics of individual email messages significantly relate to the action taken on that message.

Implications for HCI

One of our objectives in this data collection was to distill models of people’s email behavior that could inform the direction of future research in Human-Computer Interaction. Our findings relate to the areas of intelligent techniques for email handling and email interface design.

Baseline statistics on email usage support Whittaker and Sidner’s categorizations of email users as no filers, spring cleaners and frequent filers [27]. This result, along with the finding that the respondent’s identity explains most of the

variance in where a message ended up (48% of variance in whether a message was left in the inbox was accounted for by individual differences), suggests that there may be no general *best* solution in terms of placing messages in the proper place. Rather, there are strong individual preferences in filing strategies that are independent of the characteristics of the message or its sender.

Contrary to our expectations, the message level data showed that an overwhelming majority of messages were retained (76% were either filed or left in the inbox, including the 49% that were left in the inbox). The high proportion of messages people retain suggests that technology to aid in the location and viewing of messages is an important area of future research for electronic mail. People may have difficulty finding the messages they need, a problem that was cited in our interviews and in the previous work [2, 3, 7, 24, 27].

In this work we sought to identify features of email messages that influenced attention to the message. One interesting result from our data was that the perceived importance of a message was only one of the influences on the likelihood of replying to it. People responded to information requests, beyond their importance, perhaps because these were easy to attend to. They responded or intended to respond to social messages, even though these messages were unimportant for work, perhaps because they were fun or because of the social obligations they entailed. In contrast, some very important messages may have required a lot of work, so they were retained. It may be that messages considered important are valued for their content and retained to refer back to later. Many types of important messages – e.g., status updates, reminders, or scheduling messages – might not need a reply, but are retained in the inbox for later reference.

Our results suggest ways that design can more efficiently direct user attention in dealing with email. A user interface that makes the importance of a message visible may be more useful to help people find archived messages than to identify messages that require action. The statistical model we presented can be used as a starting point for developing message importance scores. The message features we identified as influencing perceived importance could be used in conjunction with information about sender-receiver relationships as in [9] to prioritize messages for viewing or to allow messages to be sorted by perceived importance.

Interestingly, messages with social content were significantly more likely to receive an immediate response. This result was notable because social content decreased a respondent's rated importance of the message. In the data from our interview respondents, we noticed that messages from friends and family members were often mentioned as things they kept meaning to respond to. For example, one respondent scrolled to the top of his inbox and pointed out four messages from old college friends that were several months old; he saved them because he needed to respond.

This finding suggests that messages with social content may deserve differential treatment in the interface.

The Role of Individual Differences

Differences among individual respondents accounted for 26% of the variance in message importance assessments and almost half of the variance in the probability of leaving a message in the inbox. It may be the case that people have widely different filing and deletion strategies, and previous research may have only documented a subset of these strategies [14, 27]. Factors such as variations in the design of email clients and differences in email account space limitations can influence user action on messages.

Notably, individual differences factored much less into the decision to respond. Only 15% of the variation in decision to respond was explained by differences among respondents. Replying is the most public of the behaviors that we studied. External social factors having to do with communication norms and relationships and organizational structure, may be a much stronger influence on replying than internal factors such as perceived message importance.

LIMITATIONS

There are limitations to the study presented that should be considered when applying the results to practice or design. These limitations include the nature of our sample, contextual factors we did not study, and the representativeness of messages included.

Generalizability

The data reported here come from a sample of 121 people from a single organization, a university. Even though the sample is larger and more diverse than those used in other studies of electronic mail, one would need data from a wider variety of organizations to feel confident about the extent to which the conclusions reported here generalize to other settings.

Controlling for Context

We think a more serious problem is that the data come from respondents' self-reports about the characteristics of their messages and their responses to them. By asking respondents to reflect upon messages in their inboxes, our survey technique approximated the think-aloud protocols common in the HCI community. However, talking about what one has done or will do with a particular message is not the same as actually performing the action. Similarly, asking respondents for their judgments about the content of messages is not the same as capturing the text of messages and coding its contents. Merely asking respondents to reflect on messages in their inbox may have taken them away from their natural use of their email. Finally, our questions may have insufficiently specified the context in which the messages were produced. For example, while we asked respondents to describe their position in the organization and their general workload, we do not have information about the state of the task relevant to particular

messages (e.g., the number of people involved in a project, approaching deadlines, and similar contextual issues).

Despite these issues, our research method has allowed us to gather more detailed information about email messages and actions towards them than has earlier research based on general interviews.

FUTURE WORK

In our next study, we plan to examine how contextual factors about a message, such as its relationship to active projects and tasks, affect a person's action on the message. We are interested in whether incoming email messages requesting action derail recipients from their current task. Over time this could result in feelings of overload and stress due to constant interruption from email, as suggested by Jackson et al [11]. We are also interested in experimenting with ways of automatically classifying message characteristics. These could include message importance, work relationships, information centrality, and social network distance or centrality.

ACKNOWLEDGMENTS

We thank Jim Herbsleb, Marina Kobayashi, Sonya Allin, and CSCW 2004 reviewer #1 for their helpful feedback and input. This material is based upon work partially supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. NBCHC030029 and NSF grant IIS-0325351. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA or the Department of Interior-National Business Center (DOI-NBC) or NSF.

REFERENCES

1. Abu-Hakima, S., McFarland, C., & Meech, J.F. An agent-based system for email highlighting. in *Proc. AGENTS 2001*, ACM Press (2001), 224-225.
2. Bälter, O. Keystroke level analysis of email message organization. In *Proc. CHI 2000*, ACM Press (2000), 105-112.
3. Belloti, V., Ducheneaut, N., Howard, M., & Smith, I. Taking email to task: The design and evaluation of a task management centered email tool. In *Proc. CHI 2003*, ACM Press (2003), 345-352.
4. Boone, G. Concept features in Re:Agent, an intelligent email agent. In *Proc. AGENTS 1998*, ACM Press (1998), 141-148.
5. Cohen, W., Carvalho, V., & Mitchell, T. Learning to classify email into "Speech Acts." In *Proc. EMNLP-2004*, (2004), 309-316.
6. Dabbish, L., Venolia, G., & Cadiz, J.J. Marked for deletion: An analysis of email data. *Ext. Abstracts CHI 2003*, ACM Press (2003), 924-925.
7. Ducheneaut, N., & Belloti, V. Email as habitat: An exploration of embedded personal information management. *Interactions* 8,5 (2001), 30-38.
8. Firm bans email at work, <http://www.cnn.com/2003/TECH/internet/09/19/e-mail.ban/index.html>
9. Flores, F., Graves, M., Hartfield, B., & Winograd, T. Computer systems and the design of organizational interaction. *ACM Transactions on OIS* 6, 2 (1988), 153-172.
10. Horvitz, E., Jacobs, A., & Hovel, D. Attention-sensitive alerting. In *Proc. UAI 1999*, Morgan Kaufmann (1999), 305-313.
11. Jackson, T. W., Dawson, R. J., & Wilson, D. The cost of email interruption. *Journal of Systems and Information Technology* 5,1 (2001), 81-92.
12. Kraut, R.E. & Attewell, P. Media use in a global corporation: Electronic mail and organizational knowledge, in *Culture of the Internet*. Lawrence Erlbaum Associates, Mahwah, NJ, USA, (1997), 323-342.
13. Lashkari, Y., Metral, M., & Maes, P. Collaborative interface agents. In *Proc. AAAI 1994*. ACM Press (1994), 444-449.
14. Mackay, W. Diversity in the use of electronic mail: A preliminary inquiry. *ACM Transactions on Office Information Systems* 6,4 (1988), 380-397.
15. Malone, T., Grant, K., Turbank, F., Brobst, S., & Cohen, M. Intelligent information-sharing systems. *Communications of the ACM* 30, 5 (1987), 390-402.
16. Mintzberg, H. *The Nature of Managerial Work*. New Harper & Row, New York, NY, USA, 1973.
17. Nardi, B., Whittaker, S., Isaacs, E., Creech, M., Johnson, J., & Hainsworth, J. Integrating communication and information through ContactMap. *Communications of the ACM* 45 (2002), 89-95.
18. Neuwirth, C.M., Morris, J.H., Regli, S.H., Chandok, R., & Wegner, G.C. Envisioning communication: Task-tailorable representations of communication in asynchronous work. In *Proc. CSCW 1998*, ACM Press (1998), 265-274.
19. Panko, R. Managerial communication patterns. *Journal of Organisational Computing* 2, 1 (1992), 95-122.
20. Sahami, M., Dumais, S., Heckerman, D., & Horvitz, E. (1998). A Bayesian approach to filtering junk email. *AAAI Workshop on Learning for Text Categorization*, July 1998.
21. Sproull, L., & Kiesler, S. *Connections: New ways of working in the networked organization*. MIT Press: Cambridge, MA, USA, 1991.
22. Sumner, M. The impact of electronic mail on managerial and organizational communications. In *Proc. of ACM SIGOIS and IEEECS TCOA Conference on Office Information Systems*. ACM Press (1988), 96-109.
23. Tyler, J., & Tang, J. When can I expect an email response? A study of rhythms in email usage. In *Proc. of ECSCW 2003*, Oulu Univ. Press (2003), 238-258.
24. Venolia, G., Dabbish, L., Cadiz, J.J., & Gupta, A. (2001). Supporting email workflow. MSR Tech Report MSR-TR-2001-88.
25. Vivisimo Clustering Engine, <http://vivisimo.com/>
26. Whittaker, S., & Hirschberg, J. The character, value, and management of personal paper archives. *ACM TOCHI* 8, 2, (2001), 150-170.
27. Whittaker, S., & Sidner, C. Email overload: exploring personal information management of email. In *Proc. of CHI 1996*, ACM Press (1996), 276-283.