How to tag it right?

Semi-automatic support for email management

Mateusz Dolata
Department of Informatics
University of Zurich
Zurich, Switzerland
dolata@ifi.uzh.ch

Nils Jeners
Department of Computer Science
RWTH Aachen University
Aachen, Germany
nils.jeners@rwth-aachen.de

Wolfgang Prinz
Cooperation Systems
Fraunhofer FIT
Sankt Augustin, Germany
wolfgang.prinz@fit.fraunhofer.de

Abstract— **Smarting-up email processing is a challenging task. Users file or retrieve multiple messages every day, while receiving little support from most popular email clients. Incorporating semi-automatic sorting into existing applications can help users with their daily work through more efficient organization and more effective search. Successful and seamless integration of tagging into existing email solutions requires exact analysis of user practices, needs and considerations, which are addressed and discussed in this contribution.**

Keywords— email processing; semi-automatic tagging; retrieval; sorting; design principles

I. Introduction

Asynchronous communication plays an important role in everyday work practice. The exchanged information is propagated along multiple media channels. Earlier, traditional mail formed the central element of communication infrastructures [1]. As time proceeded, other tools, such as fax, dominated the field. Till now, *electronic mail* (email) remains the backbone of professional communication [2] while novel technologies like instant messaging or social networks gain continuously growing popularity in the private context (cf. [4]).

The structure offered by electronic mailboxes is by default very spare. Messages are divided into the received (*inbox*) and the outgoing (*outbox* or *sent*) ones, which are chronologically ordered. Therefore, electronic mailbox can be seen as a stack (in case, the newest documents are at the top) or queue (in case, the oldest documents are at the top) of messages, alike traditional mail in the office environment.

Many of the email clients available on the market offer additional paradigms for structuring messages in advance or at the time of retrieval. For instance, they allow for sorting messages on demand, according to such characteristics as sender, recipient, subject, etc. Also, full-text search is available. It relies mostly on indexers that comb through the mailbox' content and generate dictionaries invisible to the user. Automatic threading is another functionality that has recently become popular. More traditional way of grouping messages is filing, which relies either on the manual assignment or statistics- and rule-based filters.

The methods described above involve users at different stages of email processing. As indicated by Whittaker et al. [5] some strategies require more *preparatory* effort, while others can be seen as *opportunistic* ones. In the first case, users create and maintain folder or tag-based structures to facilitate future searches for particular messages. It mostly means categorising of the incoming and outgoing correspondence (cf. [6]). On the contrary, opportunistic email users keep all their messages in a single folder and use, e.g., query-based retrieval whenever they look for a message. In such cases, one needs to recall appropriate phrases or words to find the target (cf. [7]). Frequently, different strategies are mixed or merged, resulting in a gradual rather than discrete classification of individual approaches.

Everyone desires easy, effective, and efficient methods for such routine tasks as processing email. An individual combination of available mechanisms may work well, but often results in confusion or inconsequence (flagging important messages vs. maintaining an “important” folder). A growing variety of sorting mechanisms also leads to creeping featurism in email clients. They support static and dynamic filters in parallel to tags, flags, etc. to maintain compatibility with the most popular standards. Still, users cannot easily migrate between clients, because automatic rules or filters are hardly ever transferable, as well as skills regarding particular interfaces. Those issues can only be tackled with deeper understanding of actual user needs and desires.

This contribution refers to a development project conducted in an iterative manner with strong user involvement. Through a series of state-of-the-art analysis, surveys, interviews, needs-driven development, and prototype evaluation we establish a catalogue of design principles for semi-automatic email processing. Based on the review of relevant

COLLABORATECOM 2013, October 20-23, Austin, United States

DOI 10.4108/icst.collaboratecom.2013.254224

literature and productive systems, we elaborate on the drawbacks of existing mechanisms for email sorting and retrieval. Interview and survey provide an analysis of the current strategies for email processing and lead to a number of design goals followed throughout the project. The participatory development process assures compliance with user desires at the usability level, whereas the final evaluation confirms assumptions regarding the design of semi-automatic tagging approach for email processing, and enables their generalization. This procedure guarantees that the presented solution is driven by the implicit and explicit user needs rather than by the technical state-of-the-art. The remainder of this paper follows the particular stages of the project.

II. Related Work

The research community approaches the topic of email processing in a vivid and still on-going discussion. The focus of the particular studies ranges from understanding the role of email for communication and observation of usage strategies till evaluation of practical systems. A lot of work was done to detect and classify phases in email communication and maintenance of virtual correspondence.

A. *Understanding email processing*

Email is not only a communication tool. It is often used to support coordination tasks, and even asynchronous cooperation. According to Bellotti et al. [8], people use their virtual mailboxes as: a calendar, a to-do repository, a data archive, a contact list, and a message collector. Similar observations bring Whittaker and Sidner [9] to the definition of *email overload*. This term has a twofold meaning. In its roots, it describes the diversity of the functionalities attached to one particular communication channel. Furthermore, it may relate to the large number of messages to be processed [10]. As mentioned, growing complexity of email clients and their functionalities may negatively influence ability to organize email processing effectively. Creeping featurism in email clients may therefore be seen as another form of *email overload*. In the current study, we analyse what sorting and retrieval mechanisms users really like and need, and suggest how to reduce the functional diversity in email applications. At the same time, with semi-automatic tagging we offer a smart method to efficiently cope with the stream of incoming messages. In summary, we directly and indirectly tackle the phenomenon of *email overload*.

To analyse user needs, it is necessary to understand what activities users conduct when approaching email. Numerous studies propose relevant models or frameworks supported by theoretical walkthroughs, previous studies or observations. Accompanied by design implications, those contributions offer a good entry point for further discussions.

Venolia et al. [11], driven by an extensive literature review, suggest a model for email workflow consisting of: *flow*, *triage*, *task management*, *archive*, and *retrieve*. They rely on a company-wide user study for their analysis. Among others, they propose labels as a way to support users at archiving messages, where multiple labels shall be applicable to a single conversational thread. Also, they mention the possibility of supporting users with automatically generated labels. Those generic recommendations are not further exemplified or tested. Suggestions regarding particular interactions with such labels are also quite limited, however the study explicitly stresses the role of supportive and intuitive user interface (UI). Our research, while drawing upon the notion of supported labelling, attaches great importance to the user interaction, which is thoroughly designed, prototyped and tested.

Whittaker et al. [12] provide the most extensive literature-based study regarding email processing. They aim at describing *personal information management* (PIM) through the activities, users normally conducts when interacting with their mailboxes. They differentiate between four key activities: *allocating attention*, *deciding actions*, *managing tasks*, and *organizing messages*. According to the authors, each of the activities causes specific problems, and is subject to particular improvement. While discussing the future of PIM and email, Whittaker et al. focus on the role of artificial intelligence and predict a growing influence of *natural language processing* (NLP) on email processing. This is in line with the solution presented in this contribution. We acknowledge the role of linguistic analysis for accurate processing of text data, which email indeed is. This seems necessary, as approaches to formalize email correspondence (e.g., in terms of Speech Acts [13], [14]) did not succeed. *Response* and *Forwarding* are the only accepted email acts, that enable non-NLP formalization of email correspondence.

Unlike the above, Matysiak Szóstek [15] assumes two central email activities: *organization* and *retrieval*. She focuses on analysing the dependencies between latent user needs regarding email. *Message annotation* seems to be relevant for organization of virtual correspondence, while *informative overview* and *flexible sorting* play an important role during retrieval. In general, needs linked to retrieval are reported to be more salient than those associated with sorting. This confirms the relation of *email overload* to processing of older messages (cf. [16]). Matysiak Szóstek [15] provides numerous design requirements regarding various email activities, such as: linking between related messages and flexible sorting according to people or projects involved. Those requirements can be addressed by semi-automatic tagging, which enables marking of related messages with a common tag as well as specific, semantic sorting.

B. *Supporting email processing*

Many of the requirements and design solutions resulting from literature review were implemented in prototypical systems over the years. However, they have not yet found broad acceptance in the real world applications. Even though, the diffusion of such solutions as message labelling (e.g., GMail™), semi-automatic classification into categories/tabs (recently made available in GMail™) or automatic detection of appointment times (e.g., Apple Mail) takes place, many of the advances proposed in the academia remained in their original domain. The remainder of this section offers an overview of the most prominent prototypical email clients.

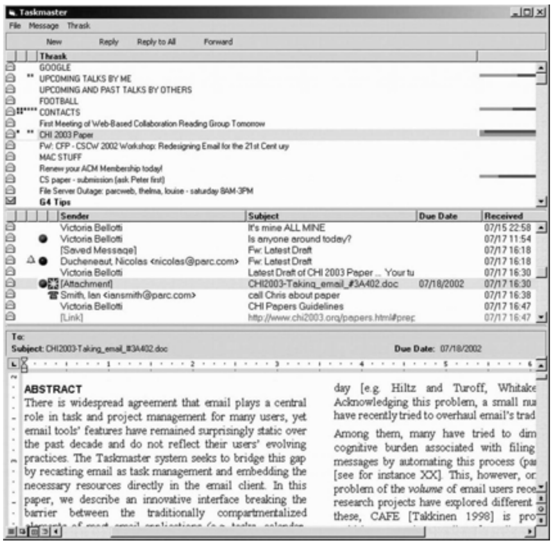

Thrask
GOOGLE
** UPCOMING TALKS BY ME
UPCOMING AND PAST TALKS BY OTHERS
FOOTBALL
!!**** CONTACTS
First Meeting of Web-Based Collaboration Reading Group Tomorrow
* ** CHI 2003 Paper
Fw: CFP - CSCW 2002 Workshop: Redesigning Email for the 21st Cent ury
MAC STUFF
Renew your ACM Membership today!
CS paper - submission (ask Peter first)
File Server Outage: parcweb, thelma, louise - saturday 8AM-3PM
G4 Tips

Sender	Subject	Due Date	Received
Victoria Bellotti	It's mine ALL MINE		07/15 22:58
● Victoria Bellotti	Is anyone around today?		07/17 11:54
[Saved Message]	Fw: Latest Draft		07/17 16:18
△ ● Ducheneaut, Nicolas <nicolas@parc.com>	Fw: Latest Draft		07/17 16:18
Victoria Bellotti	Latest Draft of CHI 2003 Paper ... Your tu		07/17 16:30
● [Attachment]	CHI2003-Taking_email_#3A402.doc	07/18/2002	07/17 16:30
☎ Smith, Ian <iansmith@parc.com>	call Chris about paper		07/17 16:38
Victoria Bellotti	CHI Papers Guidelines		07/17 16:47
[Link]	http://www.chi2003.org/papers.html#prep		07/17 16:47

To:
Subject: CHI2003-Taking_email_#3A402.doc
Due Date: 07/18/2002

ABSTRACT
There is widespread agreement that email plays a central role in task and project management for many users, yet email tools' features have remained surprisingly static over the past decade and do not reflect their users' evolving practices. The Taskmaster system seeks to bridge this gap by recasting email as task management and embedding the necessary resources directly in the email client. In this paper, we describe an innovative interface breaking the barrier between the traditionally compartmentalized

day [e.g. Hiltz and Turoff, Whitake
Acknowledging this problem, a small nu
have recently tried to overhaul email's trad
Among them, many have tried to dim
cognitive burden associated with filing
messages by automating this process (pa
[see for instance XX]. This, however, or
problem of the *volume* of email users rece
research projects have explored different
these, CAFE [Takkinen 1998] is pro

FIGURE I. TASKMASTER'S INTERFACE INCLUDING THRASK PANE AT THE TOP, DOCUMENTS LIST, AND DOCUMENT PREVIEW[17].

Some approaches aim at turning email client into a task manager, e.g., TaskMaster [17]. It aims at unifying the task- and thread-centric view on email processing. Grouping messages works heuristically and uses "reply-to" and "in-reference-to" properties of messages. It is reported to perform well, despite its simplicity. FIGURE I presents the prototypical user interface. Messages can be approached only through *thrasks* (*thread* + *task*), which strongly differs from the known email interaction patterns. In contrary, our aim is to understand usage of current email clients and leverage user experience, by supporting previous usage habits.

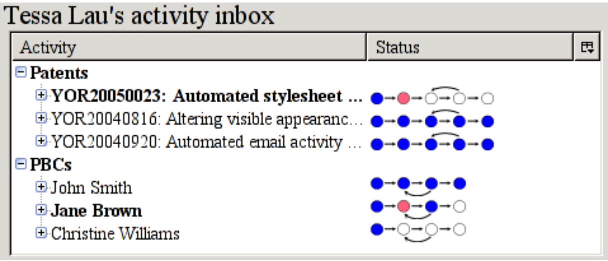
Tessa Lau's activity inbox

Activity	Status
⊟ Patents	
⊞ **YOR20050023: Automated stylesheet ...**	●→●→○→○→○
⊞ YOR20040816: Altering visible appearanc...	●→●→●→●→●
⊞ YOR20040920: Automated email activity ...	●→●→●→●→●
⊟ PBCs	
⊞ John Smith	●→●→●→●
⊞ **Jane Brown**	●→●→●→○
⊞ Christine Williams	●→○→○→○

FIGURE II. TRANSACTION-CENTRIC VIEW ON EMAIL WITH FSA-MODEL OF PARTICULAR TRANSACTIONS AND THE CORRESPONDING ENTITIES[18].

Kushmerick et al. [18] aim at supporting task management too. They model email conversations as *finite-state automata* (FSA) consisting of actions and transitions. It enables tracking the transaction state (cf. FIGURE II). For classification and modelling they apply a mixture of heuristic approaches [19] or use such NLP features as *term frequency-inverse document frequency* (TF-IDF) index [20]. The latter reflects the importance of terms through their distribution in particular messages and in the whole collection. Both methods are reported to be statistically successful and enable fast and automatic classification of messages. Similar methods can be useful for automatic generation of tags, especially when inter-dependency between messages shall be considered. This is, normally, the case of semi-automatic methods – considering relations between messages can significantly higher the reliability of predictions.

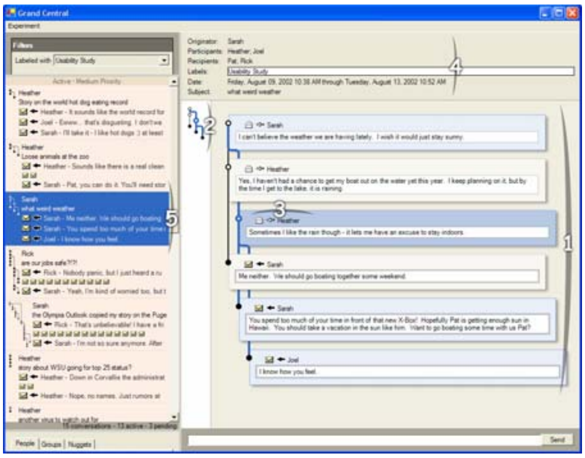

Originator: Sarah
Participants: Heather, Joel
Recipients: Pat, Rick
Labels: Usability Study
Date: Friday, August 09, 2002 10:38 AM through Tuesday, August 13, 2002 10:52 AM
Subject: what weird weather

Heather
Story on the world hot dog eating record
Heather - It sounds like the world record for
Joel - Ewww... that's disgusting. I don't wa
Sarah - I'll take it - I like hot dogs :) at least

Heather
Loose animals at the zoo
Heather - Sounds like there is a real clean
Sarah - Pat, you can do it. You'll need stor

Sarah
what weird weather
Sarah - Me neither. We should go boating
Sarah - You spend too much of your time
Joel - I know how you feel.

Rick
are our jobs safe?!?
Rick - Nobody panic, but I just heard a ru
Sarah - Yeah, I'm kind of worried too, but I

Sarah
the Olympia Outlook copied my story on the Puge
Rick - That's unbelievable! I have a fri
Sarah - I'm not so sure anymore. After

Heather
story about WSU going for top 25 status?
Heather - Down in Corvallis the administrat
Heather - Nope, no names. Just rumors at

Heather
another virus to watch out for

Sarah
I can't believe the weather we are having lately. I wish it would just stay sunny.

Heather
Yes, I haven't had a chance to get my boat out on the water yet this year. I keep planning on it, but by the time I get to the lake, it is raining.

Heather
Sometimes I like the rain though - it lets me have an excuse to stay indoors.

Sarah
Me neither. We should go boating together some weekend.

Sarah
You spend too much of your time in front of that new X-Box! Hopefully Pat is getting enough sun in Hawaii. You should take a vacation in the sun like him. Want to go boating some time with us Pat?

Joel
I know how you feel.

FIGURE III. MODEL OF MESSAGE THREADS INCLUDING SEQUENTIAL AND DISCOURSE DEPENDENCY [21].

Threading is another approach to provide the communication-centred email, as mentioned earlier. Venolia and Neustaedter [21] provide a study on representation of threads. In particular, they focus on trading-off between the sequential model and the tree representation. Whereas trees enable an overview of interdependent messages within a thread, information on their arrival time is missing. The opposite happens in the sequential view. FIGURE III includes the proposed mixed model. Threading is one of the most popular improvements of the last years regarding structuring email messages. It is available in popular web-based clients, desktop and mobile applications, even if it mostly differs from the solution presented above. As an established and popular paradigm for email structuring, it needs to be taken into consideration when designing improvements, such as semi-automatic tagging. It demands decisions on representations of tags in a thread, as well as on the scope of a tag, i.e., whether whole conversations or particular messages shall be tagged.

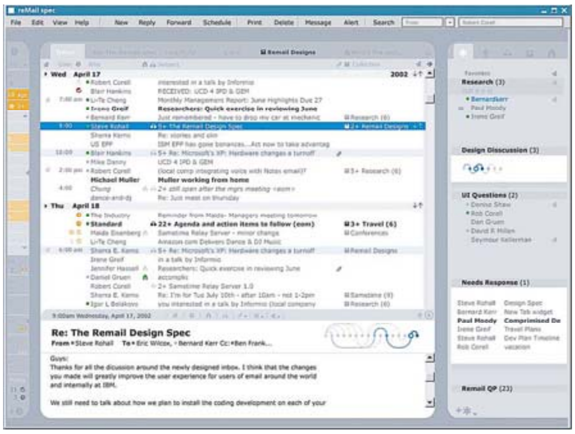

Wed April 17 2002
Robert Corell — interested in a talk by Informix
Blair Hankins — RECEIVED: UCD 4 IPD & GEM
7:00 am Li-Te Cheng — Monthly Management Report: June Highlights Due 27
Irene Greif — **Researchers: Quick exercise in reviewing June**
Bernard Kerr — Just remembered - have to drop my car at mechanic — Research (6)
9:00 Steve Rohall — 5+ The Remail Design Spec — 2+ Remail Designs
Sherra Kerns — Re: stories and clm
US EPP — IBM EPP has gone bananas....Act now to take advantag
10:00 Blair Hankins — 5+ Re: Microsoft's XP: Hardware changes a turnoff
Mike Danny — UCD 4 IPD & GEM
2:00 pm Robert Corell — (local comp integrating voice with Notes email)? — 5+ Research (6)
Michael Muller — **Muller working from home**
4:00 Chung — 2+ still open after the mgrs meeting <eom>
dance-and-dj — Re: Just meet on thursday
Thu April 18
The Industry — Reminder from Maida- Managers meeting tomorrow
Standard — **22+ Agenda and action items to follow (eom)** — **3+ Travel (6)**
Maida Eisenberg — Sametime Relay Server - minor change — Conferences
Li-Te Cheng — Amazon.com Delivers Dance & DJ Music
6:00 am Sherra E. Kerns — 5+ Re: Microsoft's XP: Hardware changes a turnoff — Remail Designs
Irene Greif — in a talk by Informix
Jennifer Hassell — Researchers: Quick exercise in reviewing June
Daniel Gruen — accomplis
Robert Corell — 2+ Sametime Relay Server 1.0
Sherra E. Kerns — Re: I'm for Tue July 10th - after 10am - not 1-2pm — Sametime (9)
Igor L Belakovs — you interested in a talk by Informix (local company — Research (6)
9:00am Wednesday, April 17, 2002

Re: The Remail Design Spec
From Steve Rohall To Eric Wilcox, Bernard Kerr Cc: Ben Frank...

Guys:
Thanks for all the dicussion around the newly designed inbox. I think that the changes you made will greatly improve the user experience for users of email around the world and internally at IBM.

We still need to talk about how we plan to install the coding development on each of your

Steve Rohall	Design Spec
Bernard Kerr	New Tab widget
Paul Moody	**Comprimised De**
Irene Greif	Travel Plans
Steve Rohall	Dev Plan Timeline
Rob Corell	vacation

FIGURE IV. REMAIL USER INTERFACE INCLUDING THREADARCS-VIEW, LABELLING, CALENDAR INTEGRATION, AND OTHER FEATURES [22]

Another system, ReMail, tries to tackle those issues and combines user made annotations with email threading and others structuring approaches [22]. It aims at solving multiple problems in email processing: lack of context, co-opting email, and keeping track of too many things. The prototype includes the ThreadArcs representation of message threads [23], to enable contextualized browsing in the mailbox. The system also enables classification of messages into predefined categories. Furthermore, through incorporation of the calendar, users are given tools to assign calendar markings to messages, such as “To-Do”, which makes it easier to keep track of tasks that depend on email correspondence. The ReMail prototype includes, also, further improvements, such as Message Map, Correspondent Map, Thread Preview [22]. Even though, the prototype enables to test various interesting approaches for email structuring and retrieval, as a whole it extrapolates the tradition of overwhelming the user with additional features. Also, the inter-dependency between the different functionalities may result in uncertainty regarding particular actions.

As discussed above, the different approaches for supporting task management, including collation of related messages, differ strongly from each other. It is notable, that systems trying to induce workflow-based structures on mailboxes remain unpopular, whereas purely heuristic threading of messages is implemented in most email clients. Even though, according to Matysiak Szóstek [15], users are interested in a topically-oriented overview of emails, such functionality is missing in most available email clients. Learning the system how to file or tag messages may solve this issue.

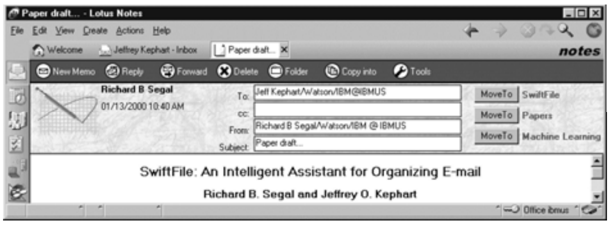

Richard B Segal
01/13/2000 10:40 AM
To: Jeff Kephart/Watson/IBM@IBMUS
cc:
From: Richard B Segal/Watson/IBM @ IBMUS
Subject: Paper draft...

SwiftFile
Papers
Machine Learning

SwiftFile: An Intelligent Assistant for Organizing E-mail

Richard B. Segal and Jeffrey O. Kephart

FIGURE V. USER INTERFACE OF SWIFTFILE PROVIDING THREE FOLDER SUGGESTIONS (TOP, RIGHT) [21].

Academic research produced systems able to learn from users actions and predict their preferences. SwiftFile is using this paradigm while supporting the user at archiving email messages [24], [25]. While using a token-based approach, the system suggests three target folders to the user. FIGURE V depicts the user interface of SwiftFile. The suggestions result from the similarity between each incoming message and each existing folder. Consequently, the system can easily and very fast adapt to a changing message collection, as well as to new users. Instead of filing messages automatically, SwiftFile moves the decision to the user. However, it still does not offer a recovery function and does not enable easy changes to the taken decisions. Also, creation of new folders has to be done manually. The functionality offered by SwiftFile, even if limited, points towards semi-automatic methods and shows how the interaction with users can be designed. Direct representation of the system suggestions leverages the understanding of the system. However, it may also generate additional cognitive load, when the user feels enforced to take any decision.

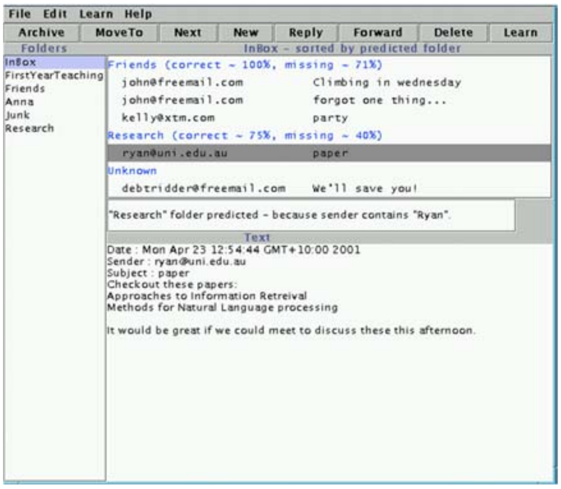

Friends (correct ~ 100%, missing ~ 71%)
john@freemail.com Climbing in wednesday
john@freemail.com forgot one thing...
kelly@xtm.com party
Research (correct ~ 75%, missing ~ 40%)
ryan@uni.edu.au paper
Unknown
debtridder@freemail.com We'll save you!

"Research" folder predicted – because sender contains "Ryan".

Date : Mon Apr 23 12:54:44 GMT+10:00 2001
Sender : ryan@uni.edu.au
Subject : paper
Checkout these papers:
Approaches to Information Retreival
Methods for Natural Language processing

It would be great if we could meet to discuss these this afternoon.

FIGURE VI. INTERFACE OF IEMS INCLUDING THE FOLDER PREDICTION AS WELL AS EXPLANATION OF THE RULE APPLIED FOR PREDICTION [26], [27]

A slightly different approach is taken in the IEMS email client [26], [27]. Here, the user can accept the prediction made by the system or change it (cf. *Archive* and *MoveTo* buttons in FIGURE VI). Additionally, the user can see the rules applied to predict the target folder. IEMS tries to tackle the same issues as SwiftFile. Similarly, it did not become popular and seems to suffer from known problems. IEMS requires additional actions to move a message around or recover from wrong decisions. Both systems, IEMS and SwiftFile, do not fail at classifying messages, but rather at integrating users and their interaction habits into the system [12]. Seamless integration of such semi-automatic tagging or filing may be the key to solving this issue.

Email clients already exist for a very long period of time. Although they can be called to be the ultimate system in CSCW and groupware research has yielded a number of productive and successful systems, email clients still look the same for the past decades [28], [29]. They have a view of the mailbox structure, the containing mails in a list and a view of a selected email. A need of email redesign exists and is discussed (e.g., [30]). The above review shows that research is mostly attracted by the topic of email management, and recently productive systems appear which slightly change the tradition. The Google GMail™ client, for example, applies the concepts of automatic prioritization and labelling, but the offered features are still away from the suggestions provided in the relevant research. The same holds for Mail Pilot (www.mailpilot.com), another example that allows viewing the inbox as a to-do list to organize a workflow around the incoming emails. It seems that a strong discrepancy between user desires and solutions available on the market and in the academia exist. In the following, we aim to address this issue while providing a study on user behaviour regarding email processing.

III. Study

Our study builds upon the Design Science Research framework for Information Science (IS) as proposed by Hevner et al. [31]. The prototype created in a user-driven fashion forms the central object of our research. This is in line with theories claiming artefacts to be the centre of IS research [32]. On our way to a functional prototype and its evaluation, we incorporate user feedback at numerous stages. The survey conducted in the early phase of the project aims at providing overview of the general tendencies regarding email sorting and retrieval. Interviews offer deeper understanding of user needs and implicit goals, thus extending the results from the questionnaire and literature review. The development process with inherent user feedback enables proper realization of user needs in our study artefact. Final evaluation is then used to ground the observations made throughout the whole process and formulate findings in form of design requirements and principles.

The state of affairs regarding email use is anything but simple. Users apply multiple strategies for processing their emails and organizing their mailboxes. Most popular methods include query-based search, threading, sorting, and manual filing. Some methods demand preparatory effort. Depending on the user, folder-based archives can form complex, nested trees. Also, the decision on what particular branches represent is with the user. The generated folder tree has a long-term character and requires maintenance by the user in order to retain its up-to-date status. On the contrary, query-based search moves the effort of structuring the mailbox towards retrieval situation. The obtained structure has a simple, Boolean character, i.e. it divides all messages into possibly relevant and irrelevant ones. This division is oriented at the goal of the retrieval activity and temporary – as soon as the information need is fulfilled, the structure is unnecessary and is not maintained any more. One can assume, that individuals apply a mixture of the above strategies to balance those different kinds of effort. To better understand the actual situation, the survey and interviews are applied in the early phase of the project. The survey focuses on the popularity of particular approaches, which are then investigated to a further extent in a series of interviews. Their results feed the development process of our prototype. The remainder of this section is structured accordingly.

A. *Investigating popularity of email processing strategies*

The survey took place in spring 2012. At that time, it is available online through a specific web link. The information on the survey along with credentials is propagated through different communication channels like email or social networks (Facebook, Twitter). In the given period of time, 107 users filled out the survey, while answering at least 6 out of 8 questions.

The survey consists of eight main questions. It asks participants on their preferences and characteristics regarding: (1) separating messages across mailboxes, (2) type of messages in the primary mailbox (professional, private, or mixed), (3) preferred structuring strategy, (4) retrieval frequency, (5) favourite retrieval strategy, (6) preferred type of structuring units (folders, tags, etc.), (7) statistical information about mailbox, (8) personal profile. Questions 1, 2, 3, and 5 are single choice questions (with possibility to formulate other answer than predefined ones). Questions 4 and 6 were multiple-choice questions (also with an additional text field when necessary). Questions 7 and 8 ask primarily for discrete information, which can be given in a text field.

The participants declare their background predominantly as Polish (47%), German (25%) and English (8%). The remaining 20% are mixed from different European and Asian nationals. Around 80% of participants are younger than 30, but only 44% of all responses come from students. The remainder is almost equally distributed among researchers, freelancers, professionals and office employees. 72% have at least 3 separate, actively used mailboxes, thus additional maintenance effort. 60% have more than 500 messages in their inbox.

All obtained answers are evaluated and the results are then extensively analysed. Also, some specific subgroups are taken into account during analysis. Five participants extend their answers while using the text field in questions 1 to 6. Their free formulated responses mostly describe a mixture of predefined choices, and are subsumed as “others” in the remainder of this paper.

B. *Investigating individual strategies*

The interviews also took place in spring 2012. Six participants of different age and coming from different professional backgrounds are chosen to participate. Three of the interview sessions were accompanied by observations of user’s interaction with their email client when dealing with standard email tasks. During the interview, memos are taken according to a prepared form including 13 open questions. Interviews are designed around the following areas of interest:

- What do users do when a new message arrives?
- What do they do when looking for previously received message?
- How do they proceed when answering a message or starting a new conversation?

Users assess their techniques and point to their drawbacks. This influences the findings, while providing a good basis for the development of functional prototype. Through questions on alternative courses of action, knowledge about available technologies can be tested. Awareness of features provided by own email client or elsewhere is important to consider the choices people make and the reasoning behind them.

The collected answers are analysed with focus on the requirements regarding a desired email client. In particular, it addressed the obstacles preventing users from fulfilling all their needs. The prototyping phase of the project gives the possibility to further address the drawbacks of known systems and present alternative solutions.

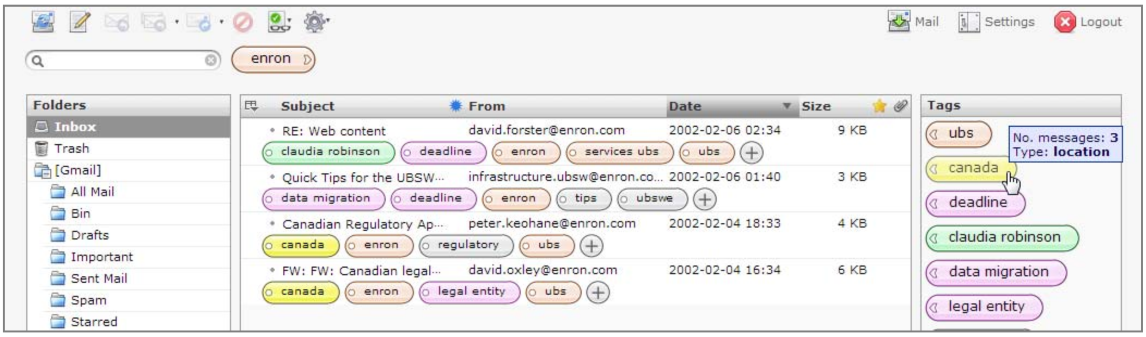

Subject	From	Date	Size
RE: Web content	david.forster@enron.com	2002-02-06 02:34	9 KB
claudia robinson, deadline, enron, services ubs, ubs			
Quick Tips for the UBSW...	infrastructure.ubsw@enron.co...	2002-02-06 01:40	3 KB
data migration, deadline, enron, tips, ubswe			
Canadian Regulatory Ap...	peter.keohane@enron.com	2002-02-04 18:33	4 KB
canada, enron, regulatory, ubs			
FW: FW: Canadian legal...	david.oxley@enron.com	2002-02-04 16:34	6 KB
canada, enron, legal entity, ubs			

FIGURE VII. USER INTERFACE OF THE PROTOTYPE SHOWING THE TOOLBAR, FOLDERS, TAGS, AND MESSAGES WITH GIVEN TAGS.

C. *Prototyping and intermediate testing*

Given the results of the literature review and insights from survey analysis and interviews in form of usage scenarios, a concept for semi-automatic tagging of messages is developed. Particularly, tagging means to add tags to messages: either manually or automatically. Semi-automatic tagging in our prototype is realized by enabling easy and efficient changes to tags, which are generated by incrementally trained tagger.

The system generates tags for a respective message when it arrives. The decisions of the system are understandable and reproducible reflecting the content of the message. Also, the user has the possibility to change the behaviour of the system and adjust it to own needs. Consequently, the system does not only tag incoming messages, but also learns how to tag from the previously labelled messages. The desired functionality along with the insights from the preliminary interviews leads to additional technical requirements. First, the program shall provide tags, even when no tags are available in the mailbox, i.e., no training data exists. Second, it shall adapt to user needs. Third, the system shall be robust and fast.

Under consideration of those requirements, a hybrid solution was chosen to generate tags. Its essence lies in combination of heuristic and *machine learning* (ML) approaches. In particular, the algorithm combines information from linguistically motivated text processing and from a learnable keyword extractor when generating set of tags for a given messages. The heuristics rely on the extraction of nouns and named entities from the text. *Nouns* play an important role in transporting meaning, therefore filling variety of semantic roles in Indo-European languages [33]. The Stanford Part-Of-Speech-Tagger [34] is used to obtain nouns from the text. *Named entities* (NE) are phrases or words that refer to particular, unique entities [35]. As they are mostly names of people, places or organization, they are assumed good candidates for message tags. The Stanford NE Recognizer [36] is employed for extraction. In addition, results of learnable *key phrase extractor* from MAUI indexer [37] are heuristically combined with nouns and named entities and form a candidate set. Each candidate is assigned a weight depending on its frequency and character (noun vs. NE vs. key phrase). The weights change with number of tagged messages in the mailbox, such that the role of the machine learnable key phrase extractor grows with the number of available examples. Further processing, such as removal of stop-words and nearly duplicates, leverages the quality of the candidate set. Finally, the top ranked candidates are assigned as labels to the considered message.

User interface plays an extraordinary role in our approach. Not only the purely technical possibility to change a tag, but also the low burden related to this, stand for adjusting the tagging system to ones needs. It leverages the interaction with tags, makes the user more familiar with them, and finally raises the trust in system decisions. This paper addresses only tagging and not the design of email clients in general. Therefore, efforts were made to test the approach in a traditional, very common email client interface. The prototype presented here builds on top of Roundcube (0.7.2.). FIGURE VII presents the user interface of the prototype.

The most obvious modification is the introduction of a separate frame on the right including all tags used for emails presented in message list. Labels are ordered according to their frequency in the mailbox. In case user wants to use tags for retrieval, a single click suffices to filter messages. FIGURE VII presents the situation where filtering by tag “enron” is applied already. Choosing additional labels can further specify the search. For instance, if the filter was extended by tag “data migration”, only the second message would remain in the view – tags assigned to messages are placed directly below their headers in the message list.

Colours of tags depend on their category (location, topic, time, etc.). Users are, on their own request during the intermediate testing, allowed to choose them freely. For automatically generated tags categories are obtained through the NE Recognizer. It suffices to click the tag only once to reach a menu with tag operations, such as: renaming, deleting or category change. Opposite to email clients like GMail^{TM}, it is not necessary to define labels first before assigning to a message. Opening the “+” dialogue and providing a name suffices. If the name does not yet exist in the mailbox, a new label will be generated and added to the tag list. Otherwise, the message is assigned the already existing tag.

The prototype as presented here was developed in a user-centred process. A stable focus group consisting of four frequent email users was consulted in a cyclic manner throughout the whole development process. The participants are aged from 24 till 38, have different scientific and professional background (two computer scientists, journalist, and political scientist). One of the focus group members is a woman. The focus group meetings are mostly free of strict rules, explicit tasks and time limits. However, all sessions look nearly alike. First, users are informed about the aim of the project if necessary. Second, a short demo of the tested feature is presented. Third, users are given the possibility to try it out by themselves and express their opinion. Driven by the opinions collected in this phase, the prototype is adjusted to best suit user feedback. In parallel, additional features are implemented according to the requirements elicited in all phases of the project.

D. *Final Evaluation*

For evaluating the system, an in-lab experiment with users is conducted. It takes place in the end of 2012 in Germany, at a computer science research institute and involves primarily its employees. The users are asked to solve three basic tasks testing the usability of the system, such as: tagging of two predefined messages, navigational search for a message and summarizing a message given its tags. Between the tasks, short interviews are taken to collect additional opinions. Finally, data regarding acceptance and attractiveness of the system were collected through UTAUT [38] and AttrakDiff2 [39] questionnaires. All 14 participants, aged 24-59, are frequent email users.

The participants are encouraged to think aloud during the testing as well as when filling out the questionnaires. All sessions are voice recorded and in parallel memos are taken. Additionally, a screen capturing software cares for recording the interaction users have with the system. Those recordings allow for further measurements on user performance as well as detailed analysis of particular situations. The test scenario remains constant throughout the whole evaluation phase.

IV. Results

Throughout the study, we make numerous observations regarding email processing in practice. Some of them, originating from the early study phase, influence the design of our prototype and can be confirmed or rejected during the final evaluation. Others become obvious towards the end of the project. All of the findings contribute to the catalogue of design principles, the central output of the project. The remainder of this section reports on the most relevant observations and relates them to the particular phases of the project.

A. *Users who sort manually are less opportunistic during retrieval than others*

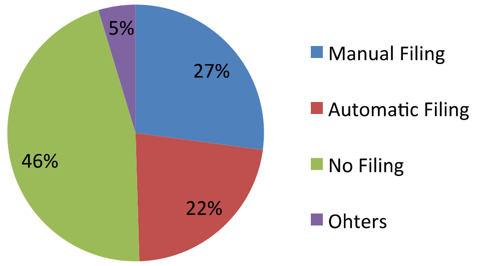

FIGURE VIII. Popularity of email structuring methods.

According to the data collected in the survey, 49% of the participants use any type of filing (see FIGURE VIII), while 27% sort their messages manually. Shape of the resulting structure (plane or nested) does not play a role here. The *automatic filing* subsumes hand-coded and ML-induced filters and rules, whereas the taxonomy (tags or folders) need to be created and specified by the user, as none of the email clients reported in the answers is able to deduce it automatically from the messages. *Manual sorting* means that users do not only create the taxonomy, but also fill it with messages, by moving them from the inbox to particular folders or tagging them appropriately.

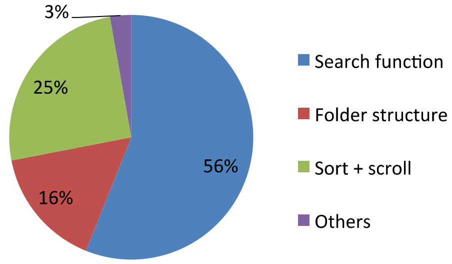

FIGURE IX. Popularity of email retrieval methods

While 49% of the respondents report on using any sorting that can be classified as preparatory methods according to Whittaker et al. [5], only 16% use their folder or tag structure for retrieval. As depicted in FIGURE IX, more than 80% use opportunistic retrieval methods: 56% use a keyword-based search function incorporated in the email client and 25% sort their messages on-demand (e.g., by date or recipient) and scroll through the lists of messages for finding messages older than three weeks.

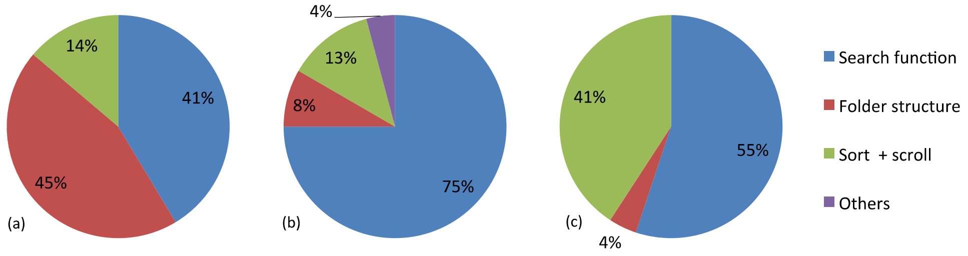

FIGURE X. Popularity of particular retrieval strategies given the preferred structuring approach: (a) manual filing, (b) automatic filing, (c) no filing.

The difference regarding structuring and retrieval preferences is significant (49% vs. 16%). This is contrary to the assumptions made in literature, that structuring is primarily a preparation for retrieval. Further investigation leads to an interesting result depicted in FIGURE X. Among those who apply manual sorting (a), 45% use their taxonomies for retrieval of messages, whereas those who use automatic filters (b) have a much stronger tendency towards opportunistic retrieval methods such as scrolling or keyword search (over 90%). This observation is confirmed by the data collected in the interviews.

Exemplary in this case is the interview with Steven, a 38 years old office employee, who refers primarily to his private mailbox during the interview. He reports that he never uses automatic filtering of incoming messages, as he simply does not trust them. He would feel responsible for checking whether filters work the way he wants, and this would cost him more time than manual filing. For him, it was important that he immediately sees that the message reached the appropriate folder.

In the development phase, focus group members notice the importance of visualization of tags right next to the message. According to the collected opinions, this would enable the users to correct the tag suggestions without large effort, but also develop an understanding of what a tag can mean in different situations. In other words, the meaning of a tag can be seen as emerging from all the messages it is associated with. Similarly, the meaning of a heuristic-based filter, even if carefully prepared, would manifest itself through the relevant messages and not only through its definition. However, in most email clients, user deals with the mailbox rules only at the time of their creation. Attending the results later does not need to happen and does not involve reviewing the underlying rule. Therefore, we postulate the following design requirements to involve the user into semi-automatic sorting: *(1) make the results of automatic processing visible and (2) easy to change*.

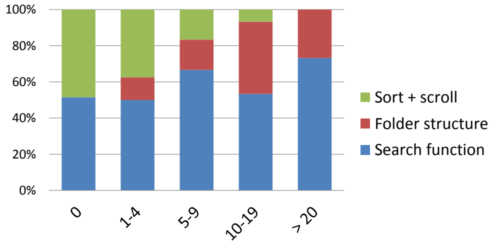

FIGURE XI. POPULARITY OF RETRIEVAL STRATEGIES GIVEN THE NUMBER OF FOLDERS IN THE MAILBOX.

B. Complexity of the structure influences search behaviour

The data obtained through the survey unveils that the complexity of the structure, expressed by the number of tags or folders, correlates with the popularity of preparatory search. As given in FIGURE XI, survey participants with 10 to 20 folders tend to use them for retrieval more often than those with less or more folders (40% vs. 15% on average). This could be seen as a specific "fit theory" between the complexity of the structure and the informative value of its elements.

One of the interviews provides an interesting explanation of this tendency. Paul, a 29 years old software engineer, creates sender-based folders for people who frequently send him highly relevant messages. He exemplifies his strategy while reporting on how he treated a new colleague. At the beginning messages from her are irrelevant, so they are not separately filed. As soon as their relevance increases, Paul creates a filtering rule and the respective folder. This makes the folder list longer than he can display at once, so he removes a folder of another person, who does not work with him at that time. He, himself, observes, that a "fair-minded" solution could be to remove folders of all people with a similar status, but it would have cost him more effort and time. It seems, that his decisions are driven by efficiency of maintenance and visibility.

Furthermore, during the development and evaluation phase, participants often mention the necessity to reduce the complexity, i.e., number of tags, by an appropriate ordering. Our prototype ordered the tags by their frequency, assuming the most important ones to be the most frequent too. However, as reported by the testers, it works as long as the tag list can be overseen at a glance. Larger or more complex structures shall be ordered alphabetically or an easy filtering of tags shall be made possible. This seems to be specifically important if the system creates new tags automatically. In such a case it is required to *(3) prevent the user from becoming overwhelmed by a large number of tags in the mailbox.*

C. Different users have different needs

The survey provides further evidence for specific dependency between mailbox character (professional vs. private) and the structuring approach. As depicted in FIGURE XII, professional mailboxes get sorted more often than private ones.

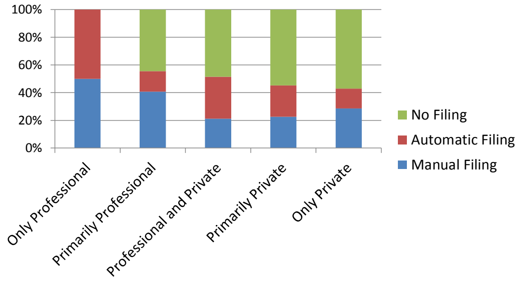

FIGURE XII. POPULARITY OF STRUCTURING STRATEGY GIVEN THE CORRESPONDENCE CHARACTER.

This observation is in line with results of the interviews. Due to the pre-existing workflows sorting professional correspondence seems to be more natural, than it is in case of private mails. Christina, a 56 years old secretary, maintains the virtual correspondence of a company as if it was traditional post. She directly compares traditional letters and emails: "I've been managing normal post in our company for 20 years. I got used to opening envelopes and giving the letters to respective officers, with or without a comment. I was also responsible for sorting out irrelevant post. I do exactly the same with emails." The dependency between defined workflows and correspondence may leverage sorting in this and other similar cases. Based on the interview with Paul, mentioned earlier, we could further differentiate between two retrieval strategies depending on the workflows implemented at the individual level. In case the retrieval activity is targeted at multiple messages, e.g., when preparing for a meeting, we will talk about *collation*. If the user looks for one specific message, the search has a *navigational* character. Observations reported above lead to the requirement of *(4) supporting diverse structuring and retrieval needs.*

D. Semi-automatic tagging has a high pragmatic quality

Thanks to the final evaluation, specifically the tagging task, we could show that the tag generator in its original mode makes its predictions with high accuracy measures (0.86 recall, 0.73 precision). The opinions regarding the tagger itself are very positive. Users appreciate the easiness of changing a tag, while seeking faster access to the *remove* command. Indeed, there is a strong tendency towards removal, compared to renaming and adding tags (22%, 5%, 7% respectively). Filtering test shows vivid user interest and acceptance, even though performance values for tag-based search do not significantly differ from those for query-based search. The average number of clicks, scrolls and typed signs required for finding the desired message is similar with slight tendency towards the tag-based solution (60 vs. 69 operations). Finally, the last assignment yields to the conclusion that tags facilitate message summarization. 10 out of 14 participants can provide full summary and explain the meaning of tags in the context. Three other participants forget to mention one important characteristic.

The results of the acceptance and attractiveness questionnaires enable further conclusions on semi-automatic tagging. The UTAUT provides very positive values regarding performance and effort expectancy (5.3 and 6.1 out of 7 respectively). In other words, users anticipate the system to fulfil their needs without requiring much work from them. It is in line with the tendency to assist the user at structuring while providing easy-to-use paradigms. The results of the AttrakDiff2 also confirm the high pragmatic value of the proposed solution (1.3 on a scale ranging from -3 to 3). The general attractiveness reaches the same level, while the hedonic quality is graded 0.8, thus suggesting further improvement regarding, e.g., visual elements and speed, as confirmed in the final interviews. Overall, the results of the final evaluation show how important it is to *(5) provide reliable and understandable predictions* and *(6) integrate into processes users are familiar with.*

V. Discussion and design principles

The study presented in this paper aims at exploration of how people can be supported at organizing and maintaining their email correspondence. On the one hand, it provides an analysis of self-declared usage patterns and objectives. On the other hand, the extensive prototyping phase and final evaluation enable observations on how the concept of semi-automatic tagging, created in a user-driven fashion, is adapted by users. The analysis of the findings collected throughout the project enable formulation of abstract requirements listed above. To leverage the development of semi-automatic tagging mechanisms and other paradigms for supported sorting and retrieval, we propose a number of design principles.

Represent the tags directly in the message list view to visualize the results of automatic processing. This principle addresses the first requirement established through the findings analysis. In consequence of making the tags, i.e., results of automatic processing, prominently visible, their role for retrieval will grow, while minimalizing the effort for memorizing the meaning of a tag.

Enable direct operations on the representations of tags to assure that changes can be applied easily. The prominently visible tags shall encourage users to interact with tags, and, by doing this, leverage training of the underlying algorithm for tag elicitation. The functionalities attached to a tag shall be as simple as possible, but as powerful as needed. Removing, renaming and adding seem to form a baseline, and shall be accessible within one click.

Provide efficient filtering of tags to prevent the user from becoming overwhelmed by magnitude of tags. List of tags may easily become large in comparison to a folder structure. Many tags exist only for few or one message; many messages receive more than one tag. Therefore, it is important to maintain the complexity of the structure. Sorting by frequency implemented in the prototype can be seen as a kind of filtering, but works well only for a limited number of tags. Instant filtering or search through tags, as implemented in the bookmarking service Delicious (www.delicious.com) is a desired alternative.

Assure flexibility to support different user needs. Labelling is considered a flexible sorting mechanism due to its many-to-many character. Further features that enhance flexibility include: easy access to tag editing, grouping of tags or search by tags intersection or union. Specifically the latter can be used to support search as navigation and search as collation.

Use keywords to provide reliable and understandable predictions. Appropriate annotation is one of the basic user needs regarding structuring of email. If done properly, applying keywords coming directly from the tagged message or from similar messages enhances the understanding, while keeping the results reliable.

Include tagging into existing email clients to integrate semi-automatic support into known environment. Email clients are sensitive tools used every day. All previous changes, like incorporation of threading, happened slowly, in a reformatory rather than revolutionary style. We think that semi-automatic methods shall enable the user too extend his previous usage patterns, rather than enforcing him to resign from them. This results in increasing pragmatic quality of the product.

VI. Conclusion

Although or precisely because email is the backbone of cooperation processes, it needs to be redesigned. Users are overloading it with functions, which are not supposed to be supported by email. Thereby email communication is overloading the users with an amount of emails, which overstrain the users. Email is the lowest common compromise in cooperation. Therefore it is heavily used.

The necessity of a redesign is commonly accepted, which is expressed in different approaches. The two main approaches are on the one hand an application-centric extension of email and on the other hand the integration of tools on the workspace-level [40]. Application-centric approaches try to support every cooperation scenario with email [12], while the workspace-level approaches support users in their work environment by providing appropriate tools for respective cooperation scenarios and integrating them [41]. Furthermore the usage of social networks and instant messaging tries to substitute email traffic [4].

Email will remain the most used cooperation tool. It will be hardly possible to change cooperation patterns within and especially between organizations. We strongly propose to reuse the underlying infrastructure of existing systems, such as email, and build an integrated view on top of all those systems. Our found requirements are not solely valid for the email system, but rather should integrate seamlessly with other cooperation tools within one common user interface.

Acknowledgement

We thank all participants of the studies for their engagement.

References

[1] L. Sproull and S. Kiesler, "Reducing Social Context Cues: Electronic Mail in Organizational Communication," Manag. Sci., vol. 32, no. 11, pp. 1492–1512, Nov. 1986.

[2] W. Prinz, N. Jeners, R. Ruland, and M. Villa, "Supporting the Change of Cooperation Patterns by Integrated Collaboration Tools," in Leveraging Knowledge for Innovation in Collaborative Networks, L. M. Camarinha-Matos, I. Paraskakis, and H. Afsarmanesh, Eds. Springer Berlin Heidelberg, 2009, pp. 651–658.

[3] L. A. Dabbish, R. E. Kraut, S. Fussell, and S. Kiesler, "Understanding email use," in Proc. Conf. on Human Factors in Computing Systems - CHI, 2005, pp. 691–700.

[4] T. Lovejoy and J. Grudin, "Messaging And Formality: Will IM Follow in the Footsteps of Email?," in Proc. Intl. Conf. on Human-Computer Interaction - INTERACT, Zurich, Switzerland, 2003, pp. 817–820.

[5] S. Whittaker, T. Matthews, J. Cerruti, H. Badenes, and J. Tang, "Am I wasting my time organizing email?," in Proc. Conf. Human Factors in Computing Systems - CHI, 2011, p. 3449.

[6] T. Catarci, L. Dong, A. Halevy, and A. Poggi, "Structure Everything," in Personal Information Management, 2007, p. 108.

[7] D. M. Russell and S. Lawrence, "Search everything," in Personal Information Management, W. Jones and J. Teevan, Eds. 2007, pp. 153–166.

[8] V. Bellotti, N. Ducheneaut, M. Howard, and I. Smith, "Taking email to task," in Proc. Conf. Human Factors in Computing Systems - CHI, 2003, p. 345.

[9] S. Whittaker and C. Sidner, "Email overload: exploring personal information management of email," in Proc. Conf. Human Factors in Computing Systems - CHI, 1996, pp. 276–283.

[10] D. Fisher, A. J. Brush, E. Gleave, and M. A. Smith, "Revisiting Whittaker & Sidner's 'email overload' ten years later," in Proc. Conf. on Computer Supported Cooperative Work - CSCW, 2006, p. 309.

[11] G. D. Venolia, L. Dabbish, J. J. Cadiz, and A. Gupta, Supporting Email Workflow: Technical Report MSR-TR-2001-88. Redmond, WA: Microsoft Research, 2001.

[12] S. Whittaker, V. Bellotti, and J. Gwizdka, "Everything through Email," in Personal Information Management, J. Teevan and W. Jones, Eds. University of Washington Press, 2007, pp. 167–189.

[13] T. Winograd, "A language/action perspective on the design of cooperative work," in Proc. Conf. Computer-Supported Cooperative Work - CSCW, 1986, p. 203.

[14] W. W. Cohen, V. R. Carvalho, and T. M. Mitchell, "Learning to Classify Email into 'Speech Acts'," in Proc. Conf. on Empirical Methods in Natural Language Processing - EMNLP, Barcelona, Spain, 2004, pp. 309–316.

[15] A. Matysiak Szóstek, "`Dealing with My Emails': Latent user needs in email management," Comput. Hum. Behav., vol. 27, no. 2, pp. 723–729, 2011.

[16] L. A. Dabbish and R. E. Kraut, "Email overload at work: an analysis of factors associated with email strain," in Proc. Conf. on Computer Supported Cooperative Work - CSCW, New York, NY, USA, 2006, pp. 431–440.

[17] V. Bellotti, N. Ducheneaut, M. Howard, I. Smith, and R. Grinter, "Quality Versus Quantity: E-Mail-Centric Task Management and Its Relation With Overload," Hum.-Comput. Interact., vol. 20, no. 1, pp. 89–138, 2005.

[18] N. Kushmerick, T. Lau, M. Dredze, and R. Khoussainov, "Activity-Centric Email: A Machine Learning Approach," in Proc. National Conf. on Artificial Intelligence - AAAI, Menlo Park, USA, 2006, vol. 2, pp. 1634–1637.

[19] N. Kushmerick and T. Lau, "Automated email activity management," in Proc. Intl. Conf. on Intelligent User Interfaces - IUI, 2005, pp. 67–74.

[20] R. Khoussainov and N. Kushmerick, "Email Task Management: An Iterative Relational Learning Approach," in Proc. Conf. on Email and Anti-Spam - CEAS, Stanford, USA, 2005.

[21] G. D. Venolia and C. Neustaedter, "Understanding sequence and reply relationships within email conversations," in Proc. Conf. Human Factors in Computing Systems - CHI, 2003, p. 361.

[22] B. Kerr and E. Wilcox, "Designing remail: reinventing the email client through innovation and integration," in Extended Abstracts Conf. Human Factors in Computing Systems - CHI, 2004, p. 837.

[23] B. Kerr, "Thread Arcs: an email thread visualization," in Symposium on Information Visualization, 2003, pp. 211–218.

[24] R. Segal and J. O. Kephart, "Incremental Learning in SwiftFile," in Proc. Intl. Conf. Machine Learning - ICML, P. Langley, Ed. Morgan Kaufmann Publishers Inc, 2000, pp. 863–870.

[25] R. B. Segal and J. O. Kephart, "SwiftFile: An Intelligent Assistant for Organizing E-Mail," in Proc. AAAI Spring Symposium - Adaptive User Interfaces, S. Rogers and W. Iba, Eds. Menlo Park, CA, 2001.

[26] E. Crawford, J. Kay, and E. McCreath, "IEMS-the intelligent email sorter," in Proc. Intl. Conf. Machine Learning - ICML, 2002, pp. 83–90.

[27] E. Crawford, J. Kay, and E. McCreath, "An intelligent interface for sorting electronic mail," in Proc. Intl. Conf. Intelligent User Interfaces - IUI, 2002, pp. 182–183.

[28] O. Bälter, "Keystroke level analysis of email message organization," in Proc. Conf. Human Factors in Computing Systems - CHI, New York, NY, USA, 2000, pp. 105–112.

[29] C. Neustaedter, A. J. B. Brush, and M. A. Smith, "Beyond 'from' and 'received': exploring the dynamics of email triage," in Exended Abstracts Human Factors in Computing Systems - CHI, New York, NY, USA, 2005, pp. 1977–1980.

[30] N. Ducheneaut and V. Bellotti, "E-mail as habitat: an exploration of embedded personal information management," Interactions, vol. 8, no. 5, pp. 30–38, Sep. 2001.

[31] A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design Science in Information Systems Research," Mis Q., vol. 28, no. 1, pp. 75–105, 2004.

[32] S. Gregor, "A theory of theories in information systems," Inf. Syst. Found. Build. Theor. Base, pp. 1–20, 2002.

[33] C. J. Fillmore, C. J. Johnson, and M. R. L. Petruk, "Background to Framenet," Intl J Lexicogr., vol. 16, no. 3, pp. 235–250, 2003.

[34] K. Toutanova, D. Klein, C. D. Manning, and Y. Singer, "Feature-rich part-of-speech tagging with a cyclic dependency network," in Proc. Conf. North American Chapter of the Association for Computational Linguistics on Human Language Technology - NAACL, 2003, pp. 173–180.

[35] B. M. Sundheim, "Overview of results of the MUC-6 evaluation," in Proc. Message Understanding Conf. - MUC, Columbia, USA, 1995, p. 13.

[36] J. R. Finkel, T. Grenager, and C. Manning, "Incorporating non-local information into information extraction systems by Gibbs sampling," in Proc. Conf. Association for Computational Linguistics - ACL, 2005, pp. 363–370.

[37] O. Medelyan and I. H. Witten, "Domain-independent automatic keyphrase indexing with small training sets," J Am. Soc. Inf. Sci. Technol., vol. 59, no. 7, pp. 1026–1040, 2008.

[38] V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, "User Acceptance of Information Technology: Toward a Unified View," Mis Q., vol. 27, no. 3, pp. 425–478, 2003.

[39] M. Hassenzahl, M. Burmester, and F. Koller, "AttrakDiff: Ein Fragebogen zur Messung wahrgenommener hedonischer und pragmatischer Qualität," Mensch Comput., pp. 187–196, 2003.

[40] V. Kaptelinin and R. Boardman, "Toward Integrated Work Environments: Application-centric Versus Workspace-level Design.," in Beyond the Desktop Metaphor: Designing Integrated Digital Work Environments, Cambridge, MA, USA: MIT Press, 2007, pp. 295–331.

[41] D. R. Karger, "Unify Everything," in Personal Information Management, 2007, pp. 127–152.