Department of Informatics
University of Fribourg (Switzerland)

Supporting Personal Information Management With Visual Facets

THESIS

presented to the Faculty of Science of the University of Fribourg (Switzerland)
in consideration for the award of the academic grade of
Doctor scientiarum informaticarum

by

Florian EVÉQUOZ

from
Conthey / Vétroz (Switzerland)

Thesis No. 1694
Uniprint, Fribourg
2010

Accepted by the Faculty of Science of the University of Fribourg (Switzerland) upon the recommendation of:

Prof. Beat Hirsbrunner, DIUF, University of Fribourg, Switzerland, President of the Committee;

Dr. Denis Lalanne, DIUF, University of Fribourg, Switzerland, Thesis Director;

Prof. Deborah Barreau, SILS, University of North Carolina at Chapel Hill, USA, Examiner;

Prof. Keith Andrews, IICM, Graz University of Technology, Austria, Examiner;

Prof. Rolf Ingold, DIUF, University of Fribourg, Switzerland, Examiner.

Fribourg, November 8, 2010.

The Thesis Director:

Dr. Denis Lalanne

The Dean:

Prof. Rolf Ingold

Remerciements

Peu de choses, au final. Des lettres jetées sur le papier, quelques dessins, des chiffres. Et la conscience éclairée d'une vérité toujours plus complexe à mesure que l'on s'en approche. Mais ce n'est pas la place d'en débattre. Puisqu'ici est le lieu où la gratitude le dispute à la nostalgie.

La lumière se fait progressivement sur la scène, éclairant les deux personnages qui s'y trouvent.

LA NOSTALGIE
(les yeux au ciel)

Ah ! C'était le bon temps.

LA GRATITUDE
(lisant)

Denis Lalanne, mon directeur de thèse, pour son soutien indéfectible et son optimisme souvent salvateur. Rolf Ingold, pour la confiance qu'il m'a accordée et le partage de son expérience. Deborah Barreau and Keith Andrews who honored me by accepting to be part of my doctoral committee. À eux tous merci, for their availability, kindness, insightful comments and advice.

LA NOSTALGIE

Les membres du groupe DIVA et du département, post-docs, assistants, secrétaires, autour d'une table, avec un crayon et une feuille de papier ou une simple tasse de café, pour les échanges enrichissants et la bonne humeur omniprésente. Une pensée particulière pour Catherine et Bruno avec qui nous avons emprunté le même chemin. Ah ! C'était le bon temps.

LA GRATITUDE

Julien Thomet, pour son important travail dans le développement du prototype et l'évaluation discutés aux chapitres 5 et 6, pour son dynamisme, son efficacité, son sérieux. Agnès Lisowska Masson, pour sa relecture attentive et généreuse. Ainsi que tous les participants aux études réalisées dans le cadre de cette thèse, à qui je dois beaucoup.

LA NOSTALGIE

Les amis de tous horizons, les comédiens, les musiciens, les poètes, les fous, qui ont su m'éviter de me prendre au sérieux. Ah ! C'était le bon temps.

LA GRATITUDE

Mes parents, qui m'ont apporté leur appui dans tout ce que j'entreprenais. Emmanuelle enfin, qui s'est montrée d'un soutien, d'une compréhension et d'une bienveillance à toute épreuve, particulièrement ces derniers mois.

LA NOSTALGIE

Oui. C'était le bon t…

LA GRATITUDE

Bon. On a compris. Viens.

Sort la Gratitude, entraînant la Nostalgie. Entre d'un pas hésitant un nouveau personnage.

LA RECHERCHE

Ça y est ? C'est à moi ?

For they that write learnedly to the understanding of a few Scholers [. . .], seem to me rather to be pitied than happy, as persons that are ever tormenting themselves; adding, changing, putting in, blotting out, revising, reprinting, showing it to friends, and nine years in correcting, yet never fully satisfied; at so great a rate do they purchase this vain reward, to wit, praise, and that too of a very few, with so many watchings, so much sweat, so much vexation and loss of sleep, the most precious of all things. Add to this the waste of health, spoil of complexion, weakness of eyes or rather blindness, poverty, envy, abstinence from pleasure, over-hasty old age, untimely death, and the like; so highly does this wise man value the approbation of one or two blear-eyed fellows.

Erasmus, *In Praise of Folly*, L.

Abstract

Personal Information (PI) refers to digital information owned by an individual and placed under their direct control. Personal Information Management (PIM) describes the processes involved in the acquisition, organization and retrieval of personal information. PIM involves a range of different tasks done by the user, most of which are performed to facilitate the re-finding of needed information. This thesis focuses on the re-finding task of PIM and proposes supporting it with the help of an interface paradigm called faceted navigation.

Previous research highlighted the fact that people remember contextual cues of information (e.g. time, type, people related), but easily forget details about it (e.g. precise content, keywords). Few PIM systems to date make use of contextual cues for re-finding tasks. Moreover, no evaluation has explored how people use these cues in their re-finding tasks.

The work done in the course of this thesis leverages relevant research in cognitive psychology and PIM to identify cues of personal information that could potentially support re-finding. A system supporting the re-finding of files, emails and calendar events is presented. It is based on the faceted navigation paradigm and uses the identified cues as facets. A controlled evaluation of the developed system confirms the potential of combining several cues in a single interface, as it allows people to re-find information based on what they remember. Results show that participants in the study mainly used textual cues (i.e. textual search) and cues related to the type of document (e.g. text, image, music) in their re-finding tasks, and predominantly used social cues (i.e. people's names) to re-find emails.

Additional contributions of this doctoral work include a model to understand PIM strategy changes which is based on an ethnographically-inspired study of strategy changes, and a taxonomy of evaluation in PIM.

The findings of this thesis suggest that future research in PIM could concentrate on integrating a faceted approach into the traditional desktop metaphor, and evaluate how seeing others' PIM practices affects a user's own practices.

Quant à ceux qui soumettent leur érudition au jugement d'un petit nombre de savants [. . .], ils me semblent beaucoup plus misérables qu'heureux, vu la torture sans fin qu'ils s'imposent. Ils ajoutent, changent, suppriment, abandonnent, reprennent, reforgent, consultent sur leur travail, le gardent neuf ans, ne se satisfont jamais; et la gloire, futile récompense que peu reçoivent, ils la payent singulièrement aux dépens du sommeil, ce bien suprême, et par tant de sacrifices, de sueurs et de tracas. Ajoutons la perte de la santé et de la beauté, l'ophtalmie et même la cécité, la pauvreté, les envieux, la privation de tout plaisir, la précoce vieillesse, la mort prématurée et beaucoup d'autres misères. Par cette continuité de sacrifices, notre savant ne croit pas acheter trop cher l'approbation que lui marchande tel ou tel cacochyme.

Erasme, *L'Eloge de la Folie*, L

Résumé

L'information numérique que possède un individu et qui est placée sous son contrôle direct est dite information personnelle. Le terme de "gestion d'information personnelle" (GIP, ou Personal Information Management – PIM) recouvre les processus d'acquisition, d'organisation et de récupération de cette information personnelle. La GIP nécessite que l'utilisateur réalise une grande variété de tâches, la plupart en vue de faciliter la récupération d'informations particulières dont il pourrait avoir besoin. Cette thèse s'intéresse principalement aux tâches de récupération d'information personnelle et propose de soutenir ces tâches au moyen du paradigme de la navigation par facettes.

Les recherches précédentes montrent que l'on se souvient facilement d'indices contextuels de l'information (p. ex. temps, lieu, personnes en rapport) mais que l'on en oublie facilement les détails (p. ex. contenu exact, mots-clefs). Peu de systèmes de GIP tirent parti de ces indices pour faciliter la récupération d'information. En outre, la manière dont les gens utilisent ces indices dans leurs tâches de récupération d'information n'a jamais été évaluée.

Le travail accompli dans le cadre de cette thèse prend appui sur la recherche existante en psychologie cognitive et en GIP pour identifier les indices contextuels ayant le meilleur potentiel de soutenir les tâches de récupération de l'information personnelle. Un système de navigation dans une collection de fichiers, courriels et événements de calendrier a été développé. Il implémente le paradigme de la navigation par facettes, les indices contextuels identifiés étant utilisés comme facettes. Une évaluation contrôlée du système confirme que la combinaison de différents indices dans une seule interface permet aux utilisateurs de retrouver l'information en fonction de ce dont ils se souviennent. Les résultats montrent que les participants ont utilisé principalement les indices textuels (recherche textuelle) ou liés au type de document (p. ex. texte, image, musique) dans leurs tâches de récupération d'information, et principalement des indices sociaux (noms de personnes) dans le cas de courriels.

Cette thèse présente en outre deux contributions annexes : un modèle de classification des changements de stratégie de GIP appuyé sur une étude d'inspiration ethnographique, et une taxinomie des évaluations pour la GIP.

Les conclusions de ce travail suggèrent que la recherche future se concentre sur l'intégration de l'approche par facettes à la métaphore traditionnelle du bureau, et évalue l'impact de la mise en commun de pratiques de GIP sur les propres pratiques d'un utilisateur.

Contents

Lasciate ogni speranza, voi ch'entrate.

Dante Alighieri, *Inferno*, III:9

J'ai plus de souvenirs que si j'avais mille ans.
Un gros meuble à tiroirs encombré de bilans,
De vers, de billets doux, de procès, de romances,
Avec de lourds cheveux roulés dans des quittances [. . .]

Charles Baudelaire, *Les Fleurs du Mal* (Spleen)

Order — let all your things have their places; let each part of your business have its time [— . . .], I found extreamly difficult to acquire. I had not been early accustomed to it, and, having an exceeding good memory, I was not so sensible of the inconvenience attending want of method. [. . .] Now I am grown old, and my memory bad, I feel very sensibly the want of it.

Benjamin Franklin, *Autobiography* (1784)

1

Introduction

1.1 The World Out There

Millions of individuals struggle to manage their personal information on a daily basis. Be it at work or in a leisure context, an important part of their activities involves finding or re-finding information items and classifying them for later retrieval. As more and more media become digital, personal information management (PIM) evolves. New devices let us carry and access our personal information everywhere. The democratization of high capacity storage and progress in compression algorithms make it easier to collect and manage multimedia PI. The low cost of storage also discourages any "spring cleaning" of our PI collections. To paraphrase Jones (2007), we prefer buying a new mansion with plenty of rooms for piling up our personal information instead of cleaning our current home and throwing away our old stuff. Moreover, as portable devices' capacity increases, it seems we could in a near future be proud owners of a device comparable to Rincewind's Luggage in Pratchett's Discworld (Pratchett, 1983–2010): a multi-legged trunk with almost infinite storage for items of every kind and that follows us everywhere, all the time, for better or worse.

However, this vast quantity of information may rapidly become overwhelming. Indeed, a 1976 Swiss newspaper article already warned the *"affluent people"* that information overload was just around the corner and that it may eventually induce feelings of guilt, stress, frustration and aggressiveness and be harmful to social relations (Blakeslee, 1976). When digital data comes into play, the overload is even worse. An internal audit in a big company reported that each worker loses 8 hours a week, on average, coping with information overload (Zeldes, 2009). Consequences include loss of job satisfaction, tension with colleagues, decrease of productivity and, ultimately, of well-being.

Perhaps the most critical and potentially frustrating activity of PIM is re-finding information items. Technological progress in the context of web search engines has still not made its way to the personal desktop. In fact, the main reason for this gap is that personal information retrieval is inherently different from general information retrieval. Unlike a search on the web or in public collections (e.g. libraries), the retrieval process in our own collection entails personal aspects of information. Re-finding does not involve the same cognitive processes as finding. The information we want to re-find is known in advance and can therefore be recognized.

Not only does information pile up, but our needs evolve as well. As a consequence, we often have no other choice than to change the way we manage information, even if constantly changing our PIM habits has drawbacks. When we change our job for instance, we generally have to implement new strategies to adapt to the new context. We usually cannot hold on to a single strategy of PIM because the context constantly evolves. Instead, we sporadically adapt our strategy to fit with the new technical possibilities we have or simply because our current strategy has been pushed to the limit and does not respond to our needs any more.

These are some facets of PIM experienced by everyone. Indeed, PIM is a real-world issue, impacting the productivity and well-being of millions of people. It is not a mere academic concern. I felt it was important to recall it in the preface of this thesis. In a more traditional fashion, the following section shortly examines how academic research has analysed the domain and attempted to bring solutions to people's PIM issues. The section ends by presenting unexplored areas of PIM research that motivate this thesis.

1.2 Research Context

Academic research was interested in personal information management even before the rise of computer-supported PIM. Starting from the definition of broad PIM strategies like the *neat* and *messy* offices (Malone, 1983) in the physical domain, it progressed to the creation of frameworks for characterizing PIM activities in the digital domain (Barreau, 1995; Boardman, 2004). Barreau's still widely acknowledged framework states that PIM comprises the following activities:

1. the acquisition of items (saving, naming, grouping)

2. the organization of items (filing into folders)

3. the maintenance of the collection (updating, archiving, deleting)
4. the retrieval (or re-finding) of items
5. the presentation of retrieved information

Among those activities, re-finding appears to be particularly critical. Indeed, the main reason why people do manage their personal information is to be able to re-find items when they are needed. Thus, most other PIM activities seem to be mainly performed in order to make re-finding easier.

Studies of user behaviour have been conducted to understand how people manage information both in the physical and the digital domains. Those studies have shed light on people's classification schemes and, more recently, on people's re-finding strategies. A common denominator of that research seems to be the role played by human memory. Indeed, it appears that re-finding items is limited by the capacities of people's memory and by what they are able to recall about the sought-after information. In this context, the role of contextual cues of information (people, time, related places, and so forth) seems to be particularly important (Elsweiler, 2007; Bergman et al., 2008c). Still, many studies reveal general user behaviour but fail to provide an understanding of PIM at the task level. However, this finer-grained understanding of PIM would be desirable in order to provide clues for future PIM tool development (Capra and Pérez-Quiñones, 2005; Elsweiler, 2007).

Both academic and industrial research have produced tools to support PIM. Traditional means to manage personal information respects the venerable desktop metaphor, which has been around for 50 years. This metaphor considers the digital world to be so similar to the physical world that it deserves the same way of thinking about and managing information: manipulating files and folders around the desktop and trash bin, and browsing folders to re-find items. Acknowledging the limits of this approach and the particular nature of digital information, tools developed in the recent years try to go beyond the metaphor. Nevertheless, many tools from the state of the art are not grounded on empirical evidence of user behaviour. Worse, most of them are not evaluated or are evaluated in a way that does not assess their efficiency as PIM tools nor provide validations of their approaches that could be built upon. Indeed, the difficulty of conducting evaluations of PIM tools is well known: as information and information needs are personal, participants in a PIM evaluation have to perform real tasks in their own collections for the evaluation to be valid. Obviously, this complicates the experimentation setting and lowers the expectations of the possible outcomes of evaluations.

Acknowledging previous research, this thesis specifically addresses the re-finding task in personal information management. Its main motivation is to evaluate the potential of contextual cues to support the re-finding task. Additionally, the evaluation will increase the understanding of the use of cues with respect to the re-finding task, which can guide the design of future tools. Collateral outcomes of this doctoral work include contributions to the theoretical framework surrounding evaluation in PIM, and to the understanding of PIM strategy changes.

1.3 Research Agenda

Although several approaches to personal information management, and in particular personal information re-finding, have been proposed in recent years, few have been empirically grounded and have had their validity assessed. In this thesis, I explicitly leverage previous research on PIM behaviour and the role of memory in PIM to identify the contextual cues of personal information that are most suited to facilitate re-finding. In two different user studies, I further investigate the preferred means users employ to re-find items in their personal collections. The results of these studies also contribute to the selection of potentially useful cues to facilitate re-finding. In particular, the social and temporal aspects of personal information appear to be promising cues for re-finding. Previous work related to navigating in large public and personal datasets leads me to considering faceted navigation as an appropriate user interface paradigm to leverage contextual cues of personal information. I further adapt the faceted navigation paradigm by proposing to represent the social and temporal facets by means of information visualization techniques. Then, I develop a faceted browser aimed at supporting re-finding which makes use of the contextual cues of personal information as facets. In particular, this requires the development of an original method for extracting the social network and linking personal information to it. Finally, I evaluate the use of cues and the overall potential of the approach by having users performing real re-finding tasks with the developed system.

1.4 Objectives

The objectives of this thesis are the following:

1. *Identify contextual cues* of personal information which support re-finding and can be obtained without a need for user annotation.

2. Identify or develop a *user interface paradigm* suited to support the re-finding task in PIM and which benefits from the cues.

3. *Design, implement and evaluate a PIM re-finding system* based on the cues and interface paradigm.

4. Improve the *theoretical framework underlying evaluation* in the field of PIM.

5. Develop an increased understanding of *PIM strategy changes*.

1.5 Thesis Outline

This thesis is organized as follows:

Chapter 2: Conceptual Background defines the key concepts used throughout this thesis. In particular, it provides precise definitions of personal information management and faceted navigation and introduces a taxonomy of PIM tasks that will be applied during the evaluation presented in Chapter 6.

Chapter 3: State of the Art reviews the scientific literature in connection with PIM, with an emphasis on studies of user behaviour in PIM and previous PIM tools. It also highlights the empirically-supported potential of contextual cues to support personal information re-finding.

Chapter 4: Surveys of PIM behaviour describes the outcomes of preliminary studies of user behaviour in PIM. Those studies are meant to complement the findings from Chapter 3. Additionally, the ethnographically-inspired study presented in this chapter leads to the definition of a framework describing PIM strategy changes.

Chapter 5: A Faceted Browser for PIM presents our approach, based on the findings from Chapters 3 and 4 and the faceted browser used as an evaluation vehicle in Chapter 6.

Chapter 6: Evaluation starts by detailing a formative evaluation which led to the improvement of the faceted browser of Chapter 5. A large part of the chapter is devoted to presenting the main evaluation, performed using the faceted browser introduced in Chapter 5, which aims at assessing the potential of faceted navigation for PIM and studies how people use facets with respect to the specific re-finding tasks they have to perform.

Chapter 7: Discussion provides a summary discussion on the main outcomes of this thesis and comments on the methodology.

Chapter 8: Conclusion closes this thesis.

1.6 Publications Relating To This Thesis

Several publications relate to this thesis. Early publications (Evéquoz and Lalanne, 2007a,b) and a general publication (Mugellini et al., 2009) detail the broad approach and its integration into a larger-scale project for supporting human memory in PIM. Evéquoz and Lalanne (2009) present the outcomes of the ethnographically-inspired study of PIM strategy changes and the resulting framework of PIM strategy changes (see also Chapter 4). It also introduces an early version of the faceted browser developed in the context of this thesis. Evéquoz et al. (2010) is a major publication presenting the whole approach and selected outcomes of the evaluation (see also Chapters 5 and 6). Furthermore, at the time of publishing this thesis, a journal article is under review.

"[Writing], O king, [. . .] will make the Egyptians wiser and will improve their memories; for it is an elixir of memory and wisdom that I have discovered."
But [the king] replied:
"[. . .] This invention will produce forgetfulness in the minds of those who learn to use it, because they will not practice their memory. [. . .] You have invented an elixir not of memory, but of reminding; and you offer your pupils the appearance of wisdom, not true wisdom."

Plato, *Phaedrus* (274e-275a)

Yo no busco. Yo encuentro.

Pablo Picasso

2

Conceptual Background

Before diving into the core of the matter, this chapter is devoted to providing definitions of the key concepts that will be referred to throughout this thesis. The first part of the chapter, in Section 2.1, introduces and defines Personal Information Management (PIM) step by step, starting from a definition of information itself, then restricting its extent to personal information, before defining PIM. A taxonomy of personal information management tasks is then discussed in Section 2.1.5, with a particular focus on re-finding tasks. Complementary discussions are also provided about the role of memory in PIM and the relation between PIM and other domains of research. The second part of the chapter focuses on the interfaces used to manage personal information. Section 2.2.1 discusses the desktop metaphor and its implications regarding PIM. Sections 2.2.2 and 2.2.3 introduce faceted navigation and information visualization, two concepts that will be used in the approach presented in subsequent chapters. Finally, Section 2.3 wraps up the presented topics and introduces the rest of this thesis.

2.1 Personal Information Management

2.1.1 Introduction

How much information is there in the world? When trying to answer this question in 1997, Michael Lesk distinguishes between the quantity of data available in a traditional form, meaning written on paper, and in a digital form, meaning encoded on magnetic storage (Lesk, 1997). His estimations acknowledge the supremacy of traditional information versus its digital counterpart: 12,000 petabytes of traditional information versus about 22,500 terabytes of available storage around the world for digital information. That is an advantage of 500 to one in favor of the traditional form. Nevertheless, Lesk also points out that the amount of digital data seems to be multiplied by ten every year. This naturally leads to the prediction that the amount of digital information would probably overtake traditional information by the year 2000. In 2000, precisely, Varian and Lyman (2000) attempt their own estimation, although with a slightly different focus. Taking the constant increase in information for granted, their primary goal is now to evaluate how much original and unique information is *produced* each year. Their figures are striking. Published information (books, newspapers, office documents, films and music) represent about 285 terabytes of information per year. Individual information created each year (home photographs and videos, X-rays and all the digital information on personal hard disks) is over 600 times larger than this. It means that for every human being on the planet, 250 megabytes of unique and original information is produced each year. An even more recent estimation by Gantz and Reinsel (2010) takes into account the trend towards digital of all major media — print, radio, TV and voice. Their study states that digital information has increased by 62% between 2008 and 2009 to reach 800 billion gigabytes. This means that 300 billion gigabytes of digital information has been created (and possibly replicated) in one single year, that is 40 gigabytes of information produced in one single year for each of the 7 billion persons in the world. The same study forecasts a steady grow for the upcoming years, meaning that the 'digital universe' may be as big as 35 trillion gigabytes, or 35 zettabytes in 2020. Forecasting a world population of 8 billion by 2020, this would be more than 4 terabytes of data per person. Moreover, the number of information containers (files, images, videos, and so on) would be 25 quintillion in 2020, which is more than 3 billion items per person on Earth.

Though the estimation methodologies can be criticized (in particular because they do not account clearly for the uniqueness of information), the glimpse they give is probably not that far from the reality. Let us stop for a moment and consider the following parallel. Victor Hugo, widely acknowledged to be one of the most prolific French authors of the nineteenth century, wrote almost thirty poetry books, a dozen plays, nine long novels, tens of other miscellaneous texts and a huge correspondence. Envious fellow writers and urban legends — refuted since the time — pretended he employed several ghostwriters to achieve his work. Nevertheless, he managed to produce "only" about 60 megabytes of original written text in his whole 83-years lifetime. Needless to say, he did not own a computer.

How are we supposed to handle this vast quantity of information? What is the strength of our cognitive capacities in comparison? Landauer (1986) interestingly provides a snapshot of the capacity of human memory for holding information. He estimates that our brain can memorize 200 megabytes of information on average. He points out that our brain probably accepts some storage inefficiency in order to make effective use of information. Thus the only role of certain neurons might be to maintain links between the information items that we remember.

If 40 gigabytes of original data is created around the world each year for each human being, it does not mean that each human being's personal information does increase by 40 gigabytes a year. Indeed, to our knowledge not any estimation has been published addressing this very question: how much *personal information* is there in the world? However and despite the fact that *"there is no unambiguous way to measure the size of digital information"* (Varian and Lyman, 2000), we can try to figure it out by examining the partial output of the MyLifeBits project[1]. One aim of this project is to record the memories of an entire human life. Gordon Bell offered himself as the subject of this experiment, letting cameras and microphones record his worldly activities, while every digital information he encounters on his computers is carefully captured as well. In 2005, 71 years-old Gordon Bell's digital memories (excluding videos of his real-world activities) totaled up to more than 40 GB (Gemmell et al., 2006). As a matter of comparison, the total size of all the *home* folders and e-mails archives on the several computers that the author of this thesis has been using for the last 15 years sum up to about 150 GB of heterogeneous personal data. Moreover, as storage becomes cheaper[2] and the quantity and quality (and therefore size) of information increase, the amount of personal data that each of us produces is likely to continue its expansion in the future. Storage will not be an issue, as forecast by Bush (1945) describing the storage capacity of his Memex device:

> *If the user inserted 5000 pages of material a day it would take him hundreds of years to fill the repository, so he can be profligate and enter material freely.*

There might be zettabytes of digital information in the world, none of us has the responsibility to manage it alone. However, this very information we consider valuable and call "personal" seems to deserve a particular care.

Following sections define more formally what is meant by information, personal information and personal information management in the context of this thesis.

[1]The *MyLifeBits* project will be presented in the state of the art in Chapter 3.

[2]The cost of one gigabyte was $10 in 2000 and estimated to drop down to $1 by 2005 (Varian and Lyman, 2000). As of January 2010, an internal terabyte hard drive costs about $100, so a gigabyte costs about 10 cents. Source: `http://wiki.answers.com`, retrieved on April 6th, 2010.

2.1.2 Definition of Information

The first meaning of latin word *"informatio"* is the process of putting something into a form (from verb *"in-formare"*), and, by extension, the product resulting from this process. This initial meaning would translate into English as "drawing" or "sketch". In a more abstract sense, *"informatio"* later means the process of putting an *idea* into a certain form and the resulting form of the idea. This form carries the meaning of the original idea while making it suitable for use and communication between people. The form can be words of a natural language, numbers or images. It can be handwritten, printed or pronounced. In the context of this thesis, only digital information is considered. The most basic support of digital information is bits. Information has been encoded so it can be handled by both computers and humans.

From the theoretical point of view undertaken by Shannon (1948), information can be measured by its capacity to reduce uncertainty. In a broader perspective, we can understand from his definition that the value of information to us is relative to our current state of knowledge. The less we know, the more we benefit from information. In the context of PIM, we could restate this as: the less we remember, the more we would benefit from re-finding information.

2.1.3 Definition of Personal Information

The seminal work by Bush (1945) does not give any formal definition, but provides examples of personal information (PI) which can be *"books, records and communication"*. Lansdale (1988) similarly specifies that it *"may be books, notes, folders, diaries, personal records, files or whatever"*. He adds that we *"keep"* it *"for our own use"* and *"would feel deprived if it were taken away"*. The primary reason for keeping this information is *"to be able to retrieve and use it in the future"*. In the following years, the term *personal information* seemed probably so self-explanatory that is was not formally defined, although being used in different contexts. A shift of meaning starting in the 1990s although tends to consider personal information as primarily digital (Barreau, 1995; Rosenberg, 1999). Even in this restricted context, Boardman (2004) recognizes that the term is ambiguous and may be interpreted in different ways. He mentions at least two common interpretations of it:

- Information about an individual, not directly managed by the individual concerned.
- Information managed and stored within a custom personal organizer software (Rosenberg, 1999).

However, he also points out that a proper definition of digital personal information should be *"independent of the subject matter of the information"* and of the software application and digital device on which this information is stored and managed. He therefore proposes to define personal information as *"information owned by an individual and under its direct control"*.

Jones and Teevan (2007) leverage previous definition attempts and summarize the several senses of personal information (digital or not). Thus, personal information may be:

1. Information a person keeps for personal use, under the person's control (e.g. files on a local computer, owned books, CDs, etc.).

2. Information about a person kept by and under the control of others (e.g. health information kept by a doctor).

3. Information experienced by a person but not in the person's control (e.g. books rented in a library, web pages visited, etc.).

4. Information directed to a person (e.g. emails).

All those senses are relevant, though in different contexts. The standard tasks of PIM mainly relate to the first category, but other tasks may relate to the other categories. For instance, some tasks of PIM involve protecting the person from being interrupted by information directed to them, by disabling email notifications, shutting down their cell phone, closing their office door, and so on. Moreover, the distinction between categories can be hard to set in some situations. Consider web browsing, for example. The history kept by browsers surely belongs to the first category of PI, but personal login name and passwords are under the host website administrator's control (second category).

The first sense of personal information, matching Boardman's previous definition, is relevant to this thesis. This is the one we will retain, narrowing it a little more to consider only personal information in a digital form.

Moreover personal information has different possible granularity levels:

- An *information item* (Boardman, 2004) is a *"self-contained unit of information"*. Information items can exist in a range of different formats. Such formats include (but are not restricted to) email, documents, bookmarks, calendar appointments, to-do items and so on. It is the smaller division of information that is considered in this thesis. Information could be subdivided further, though. For example a textual document, which is considered an information item, usually consists of a title, several paragraphs, sentences, etc. These elements also bear information by themselves, nonetheless they would rather be called sub-items of information in this sense. Information items usually have attributes or metadata attached to them, whether by the operating system (or PIM application) itself (e.g. date of creation, size, author), or set by the user as personal annotations (e.g. file name, folder, tags).

- *Personal information collections* (PICs) (Boardman, 2004) are self-contained set of items. Boardman (2004) defines that *"the members of a collection share a particular technological format and* *are accessed through a particular application"*. However, Jones and Teevan (2007) prefer to release the technological format and application constraints, in order to consider PICs as *"personally managed subsets"* of PI. They are a set of items that people have made conscious effort to build and maintain, no matter the format or application involved. Examples of such PICs in the digital domain include: project-related information items organized into a hierarchy of folders, possibly in several different applications (e.g. the filesystem and an email client); a maintained collection of web bookmarks; a database of *bibtex* references. According to the definition by Jones and Teevan (2007), PICs also bear attributes. Such attributes may be their representation, spatial layout, properties, name, and so on. Items in a PICs often share a same format or are managed using the same application. However, this is because current tools do not easily mix different formats in the same application. It does not mean that people would not do it if it were possible. Integrative organization of PI could allow this in the future. The work done in the course of this thesis ought to be one step towards this goal. One last thing that Jones and Teevan add to the definition of PICs is that they are the preferred way of thinking of one's PI with respect to its maintenance and organization. When people struggle to manage their PI better, they target goals such as "cleaning their mail inbox" or "getting their desktop organized". Maintaining personal information collections seems to be a tractable task, whereas maintaining one's whole PI often seems impossible.

- The *personal space of information* (PSI) (Jones and Teevan, 2007) is the union of all information items and collections managed by an individual, as well as all the means used to manage it. A PSI includes, for instance, all the documents, the emails of all the accounts under the person's control, bookmarks, files on servers and web bookmarks. It also includes the applications used to manage them, like email clients, calendars, to-do managers or desktop search engines. Finally it includes all the collections consciously created by the user, e.g. folders in filesystems or on email accounts, custom categories for organizing music, databases of article references, packages of java code and so on. According to Jones and Teevan (2007), the PSI includes, from its heart to its periphery :
  - PI under the person's control (e.g. books, paper documents as well as email messages, e-documents, bookmarks)
  - Applications and tools used to manage this PI
  - User-created constructs (e.g. real folders, piles, or other virtual collections)
  - PI kept by others (e.g. facebook)
  - Public information that is relevant to the person (e.g. public libraries)

Furthermore, an important aspect of digital PI has not been explicitly mentioned in previous definitions, as far as the author knows. In the digital world indeed, data about digital PI, or meta-PI almost always exists. I propose here a terminology characterizing this aspect of digital information items:

- Document: I call "document" the concrete form of an information item, that can be seen as "support". It is in a particular format, be it a file, a calendar appointment, an email, and so on.
- Metadata: represent what we could call the "context" of a document, or additional information about the document. Metadata may belong to two types:
  - Properties: standard operating systems properties, common for all documents and on which the final user has little (if any) control, like owner of the document, creation date, etc.
  - Annotations: information that is added, manually or automatically, to a document. This includes first-level annotations, like folder name or spatial layout on the desktop. Related to PIC, one can consider that the organization of a PIC stems from annotations on documents.

2.1.4 Definition of Personal Information Management

As noticed by Whittaker et al. (2000) and highlighted by Boardman (2004), much research in the field of Human-Computer Interaction has been carried out before researchers agree on common definitions of the key terms in the domain, even the terms forming the very name of the domains. This is true also for personal information management.

Lansdale (1988) defines PIM as *"the methods and procedures by which we handle, categorize, and retrieve"* personal information. Categorization and retrieval seem to be key features of PIM. However, "handling" is probably a little too vague to appear in a formal definition. Barreau (1995) describes a PIM system as an *"information system developed by, or created for an individual in a work environment"*. A PIM system lets the individual build PICs in their PSI. Drawing inspiration from the model of Soergel (1985), she further details this initial definition by representing PIM as a process involving different stages:

1. the acquisition of items (saving, naming, grouping)
2. the organization of items (filing into folders)
3. the maintenance of the collection (updating, archiving, deleting)
4. the retrieval of items
5. the presentation of retrieved information

Her definition of PIM has since been widely acknowledged in the domain. Some later researchers further detailed certain stages. For example, Elsweiler (2007) builds upon this definition and suggests adding a stage between the acquisition of items and their organization, namely the *decision* to keep an information item or not. This decision stage is also defended by Jones and Teevan (2007). Boardman (2004) remarks that the definition of items maintenance (stage 3) is perhaps too broad. Updating items, e.g editing a textual document or changing figures in a spreadsheet, is beyond the scope of PIM in his opinion. He also states that the presentation of retrieved information items is outside of the scope of PIM systems. It can be considered as a sub-task of retrieval or as a task external to PIM, for which the operating system is in charge. He finally notes that although Barreau seems to restrict the role of PIM to a professional context, it certainly is meaningful to consider PIM also in a leisure context.

The most recently acknowledged definition by Jones and Teevan (2007, p. 3) and Jones (2007) is also inspired from Barreau's framework and takes into account previous criticisms as it states:

> *Personal information management or PIM is both the practice and the study of the activities a person performs to* acquire, store, organize, maintain, retrieve, use, and distribute information items *needed to fulfill his or her various roles (as parent, employee, friend, member of a community, etc.).*

New to this definition is the notion of *roles* which formalizes what Boardman felt was missing in Barreau's initial definition: PIM is not just a professional activity, it spreads over all different aspects of an individual's life as far as they involve information, which is the case of most human activities.

As a final note, the definitions of PIM provided above are general enough to be independent on the subject of the information and the software and device used to manage it.

The following section describes the different kinds of tasks belonging to PIM.

2.1.5 Taxonomy of PIM Tasks

PIM tasks are numerous. On a typical day, each computer user has to perform dozens of PIM tasks, ranging from marking an email as read to giving a name to a newly created document. PIM tasks have different granularities: setting a new appointment is a small task compared to re-organizing a whole branch of the filesystem for archiving. They also have different consequences. Taking the decision to keep an email, even if you are almost sure you will never have to access it in the future has few consequences. It may contribute to the creeping overloading of your mailbox, at worst. On the other hand, deleting the whole archive of a web system developed some time ago in order to free up disk space might be more critical. However, providing a taxonomy of PIM tasks is difficult. First, there exist many types and forms of PI, and tasks related to every one of those are different even if some may share similarities. Second and more important, as PI is *personal*, PIM involves creative tasks devised by users and corresponding to the way their own PSI look like. Devising an exhaustive list of possible PIM tasks is thus obviously untractable. The first step would be to categorize them.

Though Capra and Pérez-Quiñones (2005) advocated for the need to understand PIM at the task level, few categorizations of PIM tasks have been proposed. Drawing inspiration from Boardman (2004) who cites the traditional input-storage-output breakdown inherited from information retrieval, Jones and Teevan (2007) propose three broad categories of PIM activities:

Finding / Re-finding activities. Such activities form the output of a PSI and are motivated by a need on the part of the user.

Keeping activities. Such activities input more information into a PSI. People keep information items because they consider that they could need them in the future.

Meta-level activities. Such activities involve higher-level actions, like organizing PI collections, in an effort to make needed items easier to re-find later.

Barreau's framework (see Section 2.1.4) describes activities of PIM that can be included in this broad scheme, giving a finer-grained level of user tasks. In practical terms, the acquisition and organization of items belong to the keeping activity. The retrieval of items and presentation of retrieved items form the re-finding activity. The maintenance of the collection (and, in a certain sense, the re-organization of items) is a meta-level activity.

As can be seen from their definitions, these activities are closely related to the concept of information *need*. This is indeed the very reason why we do PIM. Because we need information all the time, we take measures to make this information easy to access when it is needed. The most immediate activity with respect to this need is therefore the re-finding activity. Keeping activities and meta-level activities are motivated by the final goal of making information items easier to re-find when they are needed. Taking it into account, the next paragraph proposes to look more closely at re-finding tasks, and defines a taxonomy of them.

Taxonomy of re-finding tasks

Re-finding tasks can be categorized on different levels, or along different dimensions. Some of those dimensions have already been proposed by previous researchers in the field, some others are contributions of the author. Note that the taxonomy proposed here is not a hierarchical taxonomy. It consists of independent classifications that are applied to a task. Thus, the classifications of a re-finding task are:

Target. Without actually giving it a proper name, Elsweiler and Ruthven (2007) distinguishes four types of re-finding tasks targets:

- *lookup*. Those tasks *"involve searching for specific information from within a resource, for example an email or a web page, where the resource may or may not be known."* Example tasks include: finding the date of a meeting contained in an email, finding references cited by author X in one of her papers.
- *item*. They involve looking for a particular information item *"when the entire contents are needed to complete the task"*. Looking for an email in order to forward it to someone is a typical example of such a task.
- *multi-items*. Those tasks involve finding many information items in order to complete the tasks. Typical examples are tasks for which several documents need to be accessed in order to collate their contents.
- *unclassified*. Tasks for which the target does not belong to aforementioned categories is set as unclassified. It can be the case when the description of the task is too vague to classify it.

Resources. The user evaluation described in Chapter 6 distinguishes between the resources mentioned in the task. Resources are closely related to formats of PI (like email, textual documents, calendar appointments, etc.). Categories of resources include, but are not limited to:

- *doc*. A task is categorized as *doc* if it explicitly refers to a document within the file system or on a server (e.g. *Office* documents, *Adobe CS* documents).
- *mail*. A task is categorized as *mail* if its description contains an explicit reference to emails.
- *web*. References of web pages or Internet classify a task as a *web* task.
- *cal*. A *cal* task makes an explicit reference to calendar items or a calendar system.
- Combinations of the mentioned resources. For example, tasks related to attachments of emails are classified as both *mail* and *doc* tasks, thus called *mail-doc*.

Temperature. This is a subjective dimension of a task used by Elsweiler and Ruthven (2007) aimed at rating the approximate period of time when the information related to a re-finding task was last seen. During their study, participants logging tasks in a diary were asked to rate their tasks along this dimension. Scales of temperature may be adapted, but they propose to use three categories:

- *hot*: involves information seen less than a week before.
- *warm*: involves information seen less than a month before.
- *cold*: involves older information.

Difficulty. A user involved in a re-finding task on her PIC can rate its *a posteriori* difficulty, answering the question "How difficult was it to re-find the information you have just been looking for ?". Such a task can be rated, e.g. on a scale from 1 (very easy) to 5 (very hard). This rating was used in the evaluations (see Chapter 6).

Cues. The concept of cues is closely related to faceted navigation as we implemented and evaluated it (see Chapters 5 and 6). For now, it is sufficient to consider cues as specific indications found in the descriptions of a task, which can be:

- *search cue*: This type of cue consists in an explicit mention of terms that can be used to issue a textual search to re-find information (e.g. "The document I am looking for contains the term 'absenteeism'").
- *filesystem cue*: A filesystem cue is an explicit mention of a file path (e.g. "A file that is on my Desktop").
- *type cue*: A task contains a type cue if the type of the sought-after information item is mentioned (e.g. "an email" or "a pdf file").
- *social cue*: As soon as a task description contains the name of a related person, the task contains a social cue (e.g. "an email sent by John Doe").
- *temporal cue*: If the task contains an absolute or relative time context, it is said to contain a temporal cue (e.g "document created 6 months ago" or "an email received around Christmas").

Obviously, descriptions of tasks may contain several different types of cues, so these categories are not mutually exclusive.

The dimension of cues point out the importance of what is remembered about an information item and actually used when performing a re-finding task. Indeed, memory plays a crucial role in PIM. This is the subject of the following section.

2.1.6 Role of Memory in PIM

Human memory is a complex and still poorly understood process. The aim of this section is to outline the main findings and theories about human memory that are relevant to PIM. Elsweiler (2007) extensively explored the cognitive psychology research about the role memory plays in PIM. The rest of this section summarizes its main explanations and contributions without arguing or citing first-hand psychological and medical research works. Details and references may be found directly in Elsweiler (2007), as mentioned in parenthesis in the text.

Memory Systems

Memory consists of two main systems, the semantic and episodic memories. The episodic memory *"refers to the storage and retrieval of specific events or episodes occurring in a particular place at a particular time"*, whereas the semantic memory is *"a mental thesaurus, organized knowledge a person possesses about words and other verbal symbols, their meanings and referents, about relations among them, and about rules, formulas, and algorithms for the manipulation of these symbols, concepts and relations"* (Elsweiler, 2007, p. 19).

Furthermore, studies point out that what distinguishes the two memory systems is not the type of information they handle, but the experience that accompanies the system when encoding the information and retrieving it. In other words, the system used to remember information is determined by the context at the time of learning. It seems indeed that episodic memory involves the *"subjective experience of consciously recollecting events from the past, whereas semantic memory does not"*. The cognitive process associated with episodic memory is thus of high-level, involving self-references and resembles a conscious journey into the past (Elsweiler, 2007, p. 20). Moreover, the process of encoding information into memory is more efficient when the information is autobiographical (Elsweiler, 2007, p. 27).

Memory Representations

Several factors affect the quality and persistence of memories. One of them is the categorization of information items. People naturally categorize information items and benefit from the categories to recall the individual items (Elsweiler, 2007, p. 27).

Another important memory representation aid is schemata, that are *"generic knowledge or structures used to represent object, events or knowledge"* (Elsweiler, 2007, p. 28). An alternative term for schemata is "world model". Evidence shows that the precise content of a memory decays over time. However, its relation to a schemata remains. In an experiment quoted by Elsweiler (2007, p. 29), participants were presented with recorded sentences after varying delays of 0 to 50 seconds. They were asked whether the second sentence used exactly the same wording as the first one. After 50 seconds, participants were unable to correctly determine it, whereas after 0 second they could do it easily. This indicates that precise details of textual information rapidly decays, but the general meaning remains. Further experiments show that schemata facilitate both encoding and retrieval of information (Elsweiler, 2007, p. 29).

Memory and Context

Context influences the recollection of events, as the context is encoded along the object in the episodic memory. Everyone has had the experience of failing to recognize their barber in the street, where they are not used to meet him. On the contrary, going to a place may evoke long forgotten memories. Being presented with a context similar to the one occurring during encoding helps recollecting memories. Moreover, certain memories can only be accessed if presented with the appropriate contextual cue, and the amount of available cues positively influences the probability of recall (Elsweiler, 2007, p. 31).

On a side note, the presented studies also show that it is easier to recognize an object in a list than to recall it out of the box (Elsweiler, 2007, p. 31).

Conclusions

The most salient conclusions that can be drawn from this short summary are the following:

1. Auto-biographical information is more easily remembered than other information.
2. Contextual cues of information are naturally encoded in memory.
3. Categorization is a natural process for encoding information in memory and helps recollection.
4. Precise details of information are forgotten, whereas the general meaning and schemata remain.
5. Contextual cues help recollection.

Let us conclude this initial discussion on PIM by presenting different areas of research which are concerned by PIM.

2.1.7 PIM in Relation to Other Domains

PIM is inherently multidisciplinary (Boardman, 2004; Jones and Teevan, 2007). Among other less represented fields, it is an area of inquiry for researchers in cognitive psychology, human-computer interaction and information retrieval.

Cognitive Psychology. Cognitive psychology studies and models mental processes such as memory, thinking or problem solving. The role of memory is specifically of interest for PIM. The way we remember information and its context influence the way we are able to re-find information in personal collections and, by extension, to do PIM efficiently. Elsweiler (2007) extensively reviewed the relevant work in cognitive psychology and explored the role of memory for PIM in detail.

Human-Computer Interaction. Much research in the field of PIM (including this very one) originate from the HCI community. PIM tools developed are numerous, as is presented in Chapter 3. However, as observed by Jones and Teevan (2007), many of those tools remain focused on specific forms of information or applications, thus following a more technology-driven approach to PIM.

Information Retrieval. IR is traditionally defined as the science of searching for documents or textual information in general. As such, it is used by PIM. The re-finding task of PIM can inherently benefit from powerful IR mechanisms. However, IR in PIM differs from traditional IR as it involves re-finding known and personal information and not just finding the most relevant information out of a possibly unknown corpus of textual data.

2.2 Interfaces For Personal Information Management

Managing personal information in a digital form supposes that users interact with a digital system where their personal information is stored. This section reviews the main types of interfaces that can be used to manage personal information. Interfaces following the desktop metaphor are by far the most common. They are the focus of Section 2.2.1. Alternative interfaces are beginning to emerge, though. Faceted navigation is described in Section 2.2.2 and a brief presentation of information visualization is given in Section 2.2.3.

2.2.1 The Desktop Metaphor

The desktop metaphor is usually defined as *"a set of unifying concepts used by graphical user interfaces to help users interact more easily with the computer"*[3]. In the desktop metaphor, the computer's monitor mimics the real physical desktop of the user in allowing virtual objects to be placed on it. Those virtual objects may be digital documents or folders of documents.

The metaphor, though, practically ends here. Physical documents can be piled when digital documents cannot. Physical documents can be read directly when digital documents need to be "opened" in virtual "windows". It is relatively easy to estimate how many documents a physical folder contains, but this is not the case for virtual folders. Moreover, real trash bins usually don't stay on the desktop. Finding the right balance between the purity of the metaphor and the usability of the system has always been a major challenge for UI developers. Initial desktop-like interfaces like Magic Desk on the Commodore (see Figure 2.1) followed the metaphor so closely that the mouse pointer had the shape of an hand interacting with drawers and phones icons on a virtual

[3]Source: `http://en.wikipedia.org/wiki/Desktop_metaphor` (retrieved April 12th, 2010).

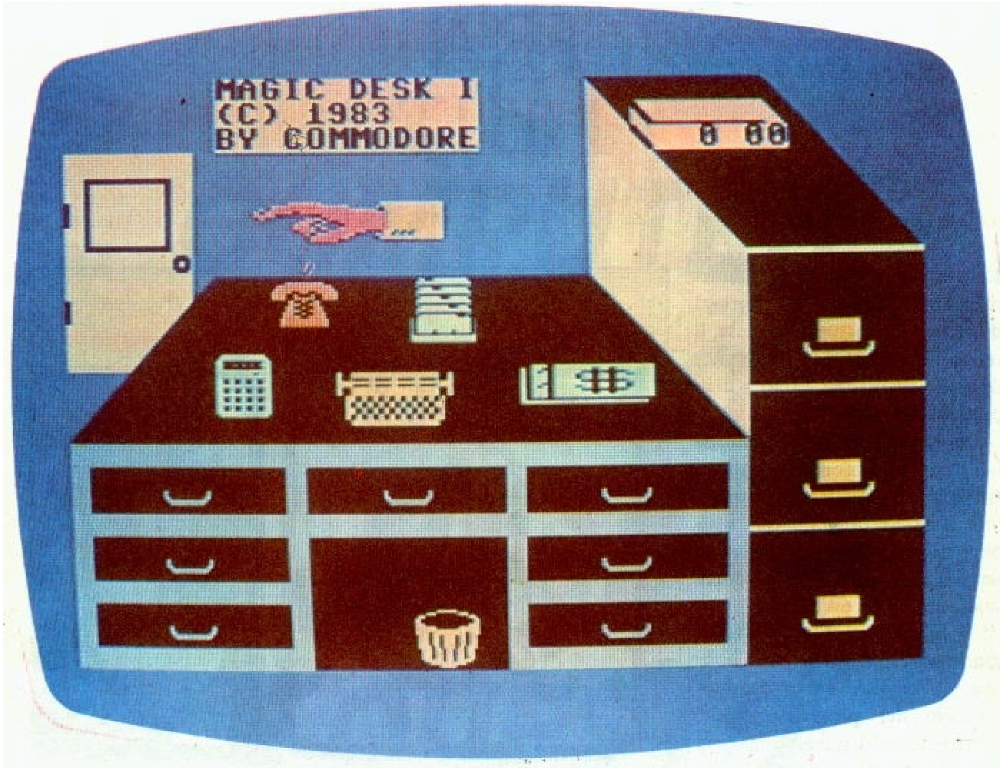

Figure 2.1: The Magic Desk, desktop metaphor implementation on the Commodore 64. Source: http://www.museo8bits.com/anuncios/c64_it4.jpg (Retrieved May 13th, 2010).

desktop. Modern desktops embed objects which have no equivalent in the real world, still "making the computer easier to interact with", like menu bars, task bars and other gadgets.

However, the desktop metaphor as it is implemented in nowerdays systems seems to confuse some users. The desktop can indeed be seen in two different perspectives. From the first perspective, the objective perspective, the desktop is just a folder in the computer's filesystem. According to this perspective, the entry point into the system is the computer itself (symbolized by the 'My Computer' icon on Microsoft Windows systems). On the other hand, from the subjective perspective of the user, the entry point is the desktop. Yet the computer itself may be accessed from the desktop, and the computer contains a folder that is the actual desktop. This inconsistency of the metaphor may confuse some users (Kaptelinin and Czerwinski, 2007). Other inconsistencies may prevent some users from being at ease when using the interface. Features, like shortcuts to files or hypertext links, which have no equivalent in the real world may put a similar burden on casual users. Furthermore, the desktop metaphor has other limitations inherited from the hierarchical categorization of documents. Visualizing the whole hierarchy at a glance is generally impossible[4],

[4]At least, commonly available operating systems implementing the desktop metaphor do not offer means for visualizing the whole hierarchy. Nevertheless, there are several specific approaches or tools which address this issue, e.g.,

and maintaining a coherent hierarchy in the long term is challenging.

The primary means to access information within the desktop metaphor is *browsing*, like in a real desktop environment. Leveraging the capabilities of the computer, automated *search* has emerged as an alternative way of looking for information. Let us present shortly those two means of re-finding information and their limitations.

Browsing. (also called Navigation). In a way similar to what we do with real desktop drawers, browsing on a computer implementing the desktop metaphor involves opening a folder and scanning the list of documents contained within this folder to re-find the one that is needed. Still, there are a number of differences from a real desktop context. First, real folders generally do not exceed two levels of hierarchy (e.g. a binder with folders inside it). File systems hierarchies are usually deeper and more complex. Second, browsing a real folder gives an overview of the documents inside it. Browsing a computer folder usually leads to seeing a list of file names and metadata but not any content. Some file managers enable thumbnails of documents instead of their filenames for certain types of documents, like images. A thumbnail view for textual documents is not realistic, though. Third, computer file systems hierarchies are not one-way like real ones. Shortcuts to other places in the file system hierarchy can be provided, which may further complicate the browsing experience and worsen the "lost in the infospace" effect.

Browsing dates back from the first desktop-like interfaces, as it is the way of re-finding information which is the closest one to the desktop metaphor itself.

Search. Search in a computer system typically involves inputting keywords in order to issue a query to a database, which will return files containing the given keywords.

With the capabilities of computer for quickly matching strings of characters to one another, automated search has been embedded early in operating systems tools, even before the desktop metaphor was put into use. The famous *grep* command-line tool for *Unix* was written in 1973. It can be used to look for regular expressions within file content. Of course, this capability is unmatched in a real desktop environment. As the famous quote states: *"You can't grep dead trees"*: searching digital (textual) media is an easy task but it can't be done with documents printed on paper (dead trees). Searching in file systems became more efficient when desktop search engines began to include indexes of information in the 2000s, which speeded up the search process. The first commercial desktop search systems were released in 2004 as stand-alone applications (Boutin, 2004). Popular operating systems now include this kind of search engines by default (Mac OS since

WinDirStat which is based on treemaps (see `http://windirstat.info`, retrieved October 10, 2010), the hyperbolic browser (Lamping and Rao, 1996) or information pyramids (Wolte, 1998).

2005[5], Windows since 2007[6]).

Searching can be fast. If users remember words contained within a document, chances are that searching will direct them to the document, without having to navigate in the filesystem hierarchy. In that sense, whereas browsing is a navigational process, searching can be seen as teleporting (Teevan et al., 2004). However, searching has a number of limitations, too. First, it is primarily textual. This means that it does work well if and only if the sought-after documents contains text that has been indexed. Searching for images containing a "red truck" will yield no result if the images have not been textually annotated. Second, it forces the user to remember exact keywords, as synonyms will not be understood by the system. Lastly, most searches will yield several possible documents. The re-finding of the needed document will therefore still involve visually scanning a list of filenames.

Accounting the limitations of both access means, efforts have been put into finding a synergy between search and browse (Mackinlay and Zellweger, 1995), even if some of the constraints inherited from the desktop metaphor should be released. Such an alternative way of finding or re-finding information items in information systems which is independent of the desktop metaphor is faceted navigation. The next section describes it.

2.2.2 Faceted Navigation

Faceted classification

Faceted classification has been invented in 1933 by the Indian librarian and mathematician S.R. Ranganathan, who gave it the name of "Colon Classification" (Ranganathan, 2006). It allows the assignment of multiple independent classifications, called *facets*, to an object, rather than a single predefined taxonomic order. Each facet comprises *"clearly defined, mutually exclusive, and collectively exhaustive aspects, properties or characteristics of a class or specific subject"* (Taylor, 1992). These aspects or properties are called *categories* of the facet. For example, a collection of books can be classified along an "author" facet, a "type" facet, a "subject" facet, a "year" facet and so on. Facets can thus be compared to independent taxonomies applied to an initial set of data. The categories of a facet may be:

Flat or hierarchical. An author facet will contain only flat categories (by "Charles Baudelaire"), whereas a location facet will contain hierarchical categories ("Europe" > "France" > "Paris") (Yee et al., 2003).

[5]Source: `http://en.wikipedia.org/wiki/Spotlight_%28software%29` (retrieved April 13th, 2010)
[6]Source: `http://en.wikipedia.org/wiki/Windows_Search` (retrieved April 13th, 2010)

Singlevalued or multi-valued. That is, some facets will only admit one value ("printed in 1998") or can allow multiple values ("edited by D. Johnson and M. Smith") (Yee et al., 2003).

Implicitely ordered or not. A facet related to time will have categories that have an implicit chronological order, whereas a location facet will not. This property is introduced by Lee et al. (2009) under the name "linear facet".

Faceted classification has a number of advantages over classification using a unique taxonomy. Each facet represents a different aspect of the data, which renders the classification itself very expressive. Moreover, it becomes possible to add new facets and to extend each existing facet without having to modify the others (Papa, 2006; Tunkelang, 2009).

Faceted Navigation and Faceted Search

Faceted navigation is an interface paradigm for accessing information structured according to a faceted classification scheme. With this approach, any facet can be used alternatively to navigate in the collection. Each selection in a facet acts as a filter on the original collection, and updates other facets accordingly. Another feature of this approach is that the users may get to the same items while following different paths, or applying different filters, thus enriching the browsing experience (Papa, 2006).

Some data corpora are made of text. Textual data by itself is unstructured, but it can often be augmented with metadata that lend themselves to faceted classification (using facets such as author, date of creation, topic, etc.). Thus, textual data can be considered to form a semi-structured dataset. Faceted navigation can be applied to browse the dataset by following its structured part, whereas text search can be applied to filter the dataset according to the raw text. The combination of both paradigms in a single interface is called "faceted search" by Tunkelang (2009) to differentiate it from strict faceted navigation that does not use text search. Other scholars do not make this distinction and use the term of "faceted navigation" even if the system enables text search (Yee et al., 2003)[7]. This navigation paradigm can thus be seen as a synergy between search and browse as it lets the user combine text search with a progressive narrowing of choices in each facet.

Faceted navigation systems additionally provide a bird's eyes view of the dataset and may be used as an analysis tool of the properties of the dataset.

The interface of faceted navigation systems respect the following main guidelines (Yee et al., 2003):

Dynamic queries. Each selection in a facet should generate a result set immediately, without a need to click on a separate "start search" button (Shneiderman, 1994).

[7] Although the work presented later in this thesis (see Chapters 5 and 6) actually makes use of "faceted search" in Tunkelang's terminology, I only used the designation of "faceted navigation" to refer to it.

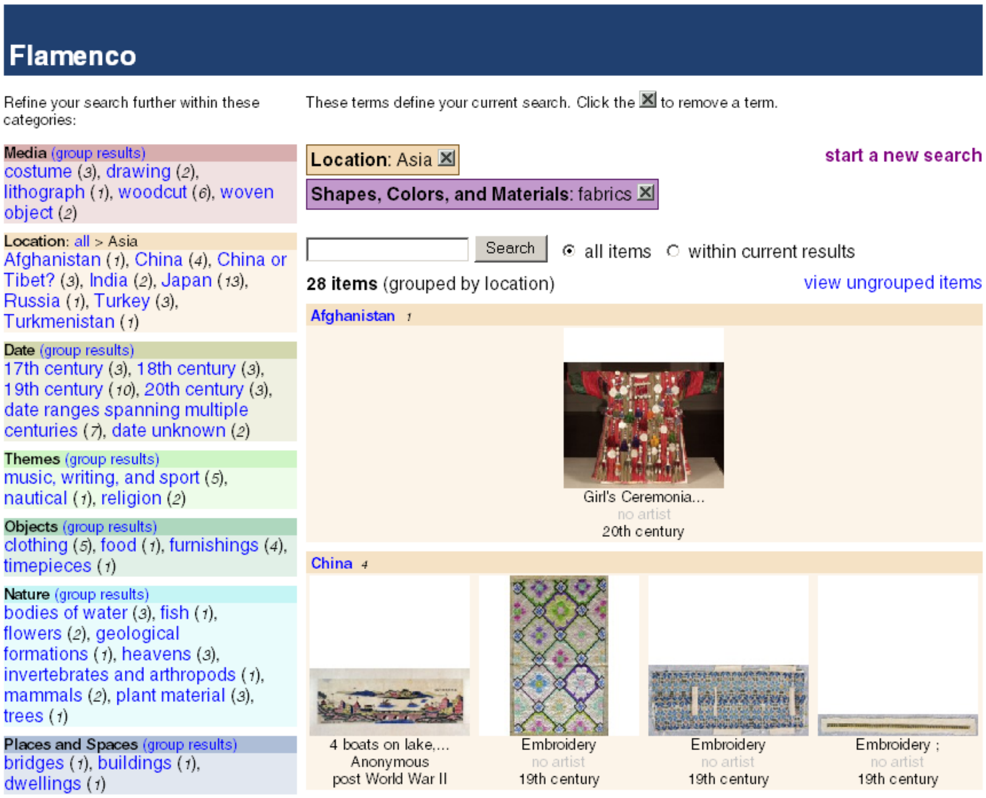

Figure 2.2: Flamenco, a faceted navigation interface by Yee et al. (2003).

Query preview. The number of results to expect from a selection should be shown alongside the interface widget triggering the selection.

No empty result set. Selections that would lead to zero results should not be allowed by the interface.

Enforce feeling of control Several visual measures can help make the user feel in control of the browsing experience. Currently applied filters should be salient and easy to cancel. Each facet should be assigned a particular hue throughout the interface.

Seamless integration of search and browse. The interface should allow keyword searching as well as browsing using pre-assigned metadata terms (facet categories). Both means may be used independently or in combination. Systems allowing textual search in the content of items may be called "faceted search" systems (Tunkelang, 2009).

Figure 2.2 shows the faceted navigation system *Flamenco* (FLexible information Access using MEtadata in Novel COmbinations) by Yee et al. (2003). Its interface is used to browse an archive of images. Note that textual search is present in this interface, but it only supports searching through the facets categories. According to Tunkelang's terminology, this interface is an example of strict faceted navigation. Other academic and industrial work putting faceted navigation and search into use, as well as a thorough presentation of the paradigms can be found in the book by Tunkelang (2009).

Relation to PIM

The application of faceted classification depends on the presence of metadata categorizing the information. If the metadata is sufficient, a faceted classification of it may be automated, and faceted navigation can be applied on top of this classification. Personal information usually bears metadata and can thus be classified along facets. Section 3.3.1 further covers the work related to faceted classification and navigation in the context of PIM. The work presented in Chapter 5 follows a faceted navigation approach to personal information re-finding.

2.2.3 Information Visualization

As the nineteenth century French general and emperor Napoléon Bonaparte claimed, *"un bon croquis vaut mieux qu'un long discours"*[8]. To make sense of complex situations or data, the use of visual representations is often more helpful and effective than long and verbose descriptions. As PIM involves dealing with vast quantities of information, visualization thus appeals as a promethean solution to personal information overload. This section briefly defines information visualization and sketches its relations to PIM.

Definition. Card et al. (1999, p. 6) define information visualization as follows:

> *Information visualization is the use of computer-supported interactive visual representations of abstract data to amplify cognition.*

Several key concepts are made clear by this definition. First, visualization is based on abstract data. This data has to be manipulated in order to be represented visually. Second, interactivity is considered as a needed feature. Manipulation of the visual representation helps the user grasp the meaning of the data[9]. Third, this definition clarifies the fact that visualization itself is a *"cognitive process undertaken by the user"* (Spence, 2006). Users indeed benefit from the visual representation because it lets them form a mental model of the abstract data more easily than with the raw data only. Human perception is therefore a critical stage in the visualization process (Ware, 2004). Thus, the essence of visualization is the formation of a mental model of the data by the user.

[8] Litteraly: *"A good sketch is better than a long speech"*, i.e. *"A picture is worth a thousand words"*.

[9] Interactivity with the visual representation and and computer-support are not considered as needed feature of information visualization for some researchers. Visualizations printed on paper, although providing no interactivity, still deserve to belong to the information visualization realm. Painted graphs and maps created before computers even exist also do (Spence, 2006; Tufte, 2001; Ware, 2004; Few, 2007). However, in the context of this thesis, only computer-

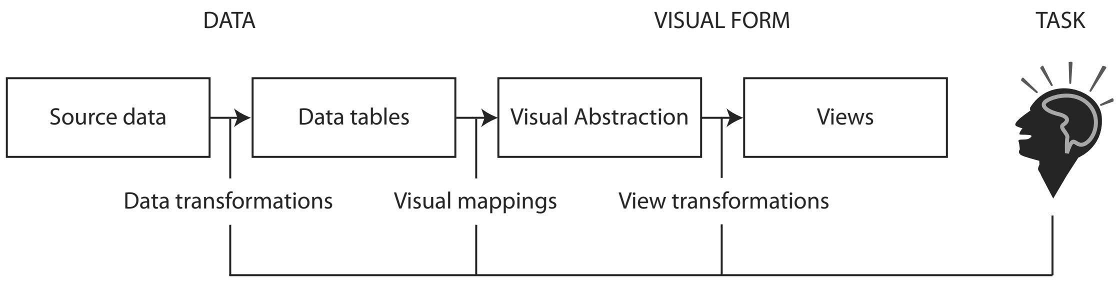

Figure 2.3: The information visualization reference model (Redrawn from Card et al. (1999)).

Visualization Reference Model. All instances of information visualizations can be related to a single reference model, introduced by Card et al. (1999) and further detailed by Chi (1999) (see Figure 2.3). The mapping from data to visualization can be modeled using a series a steps. First, the source data is transformed into data tables, that represent it under a canonical format including metadata relevant for the next steps of the process. Then, the elements contained within the data tables are mapped to corresponding visual structures that form a visual abstraction of the data. This step involves choosing the appropriate graphical properties that will be used to present properties of the inital data. For example, certain properties of the data may be used to determine the shape, spatial positioning or color of the corresponding visual substrates. Finally, view transformations are applied before the visual result is actually presented to the user. View transformations include zooming and panning the visual representation space, for example. User interaction can be applied at any step, as shown in Figure 2.3, to modify the visual output. The computer support lets the user adjust the different steps and see the visual result with enough speed and flexibility for it to be called "interactive". Interactivity is the key to make for the overall appeal of information visualizations.

Aims. The high-level goal of information visualization is to generate *"insight"* in the data (North, 2006). Card refers to this process as "knowledge cristallization" about the data, insisting that the mental process entailed by this term is difficult to apprehend. However, generic high-level goals of information visualization are unclear. Indeed, user tasks for which information visualization systems are conceived are not well defined, if at all, which leads to information visualizations being sometimes presented as *"solutions in search of problems"*[10]. This obviously makes evaluations of information visualization systems challenging. Only recent work by Andrews (2008), Ellis and Dix (2008) or Munzner (2009) propose evaluation guidelines for different stages of the design of information visualization systems which can validate higher-level user tasks and aims. The aims of

supported interactive visualization will be considered.

[10]Mentioned by John Stastko in a 2006 interview given to `www.mentagrafica.it`, relayed by `http://visuale.bertini.me/?p=10` (retrieved May 24th, 2010).

information visualization applied to PIM are tentatively described in the following paragraph.

Relation to PIM. I believe that the application of information visualization techniques can provide several benefits to PIM that current tools do not widely support:

Provide an overview of the PSI and details-on-demand. The different granularities offered by information visualization techniques are able to offer different levels of detail on a personal space of information (PSI), thus facilitating the overall understanding (synthetic view) of the PSI and the navigation in it on the part of the user.

Lower the cost of access to information. Presenting data visually may enhance the way users are able to recognize it, providing ways of accessing the data that does not only involve the visual scanning of textual lists.

Foster the self-analysis of personal information management practices. Information visualization is about making sense of the data. Therefore, it is well suited to foster the self-analysis of PIM practices on the part of the user (e.g. determining how many documents are contained in a folder to assess if creating new subfolders is needed, or quickly evaluating on which project worktime has been assigned in a certain period of time).

2.3 Chapter Summary

This chapter has defined the key concepts that will be used in the remainder of this thesis. In a first part, formal definitions of information and personal information were given. Then, several definitions of personal information management were examined. They led to a summarizing definition. As personal information management involves different categories of tasks, a taxonomy of them was provided. Furthermore, the role of memory in PIM was investigated in relevant psychological and medical work and the crucial role of contextual cues in the recollection process was identified. In a second part, interfaces to personal information were discussed. Beyond the traditional desktop metaphor, novel approaches like faceted navigation and information visualization emerge as promising alternatives to support particular tasks of PIM. Building on this knowledge, the next chapter will present the current state of research in PIM with a specific focus on the psychological aspects of it and on how interfaces for PIM relate to those aspects.

Le seul véritable voyage [. . .], ce ne serait pas d'aller vers de nouveaux paysages, mais d'avoir d'autres yeux.

Marcel Proust, *A la recherche du temps perdu* (1918)

If you leave the smallest corner of your head vacant for a moment, other people's opinions will rush in from all quarters.

George Bernard Shaw

3

State of the Art

This chapter presents the main issues that motivate the research in personal information management (PIM) in Section 3.1. It reviews, in Section 3.2, the empirical contributions that help to understand the users' behaviours and preferences in PIM. Finally, Section 3.3 presents a selected set of PIM tools, the approach of which is relevant for this thesis.

Besides presenting the work related to PIM, one of the goals of this chapter is to leverage previous research in order to identify which contextual cues and interface paradigms may potentially support the re-finding of personal information, according to the first and second objectives of this thesis, presented in Section 1.4. Chapter 4 will present user studies that I conducted and that complement the findings of this chapter. Chapter 5 will leverage the conclusions of the two previous chapters to propose the design of a PIM tool based on the relevant cues and interface paradigm.

3.1 Main Issues in PIM

The issues that motivate research in PIM are:

Information overload. Digitalization of media, popularization of electronic communication means and growth of storage capacities imply that we have more and more information to manage, which may become overwhelming and stressful (Whittaker and Sidner, 1996; Boardman, 2004; Elsweiler, 2007; Bergman et al., 2008b; Jones and Teevan, 2007).

Information fragmentation. As information comes in different formats and on different devices, the information related to a particular project or needed to accomplish a task is inherently fragmented (Dumais et al., 2003; Jones and Teevan, 2007; Jones, 2007). A related, more specific issue is *project fragmentation* (Bergman et al., 2008b) which moves the focus of information fragmentation to the project level.

Loss of context. A related problem addressed by several research works is the loss of context associated with information in electronic environments (Lamming and Flynn, 1994; Fertig et al., 1996; Elsweiler et al., 2007; Bergman et al., 2008b). This is particularly harmful because context is a strong cue for recall.

Lamming and Flynn (1994) and Elsweiler (2007) state that PIM systems do not take advantage of the way human memory works. Fertig et al. (1996) suggest that the problems encountered in PIM merely reflect the limitations of the desktop environment. The following section of this chapter reviews empirical contributions that explain user behaviour in PIM and highlights particular areas of interest for PIM tool development. The subsequent section presents PIM tools that go beyond the traditional desktop metaphor.

3.2 Review of Empirical Studies of User Behaviour

This section is organized according to questions that match the framework of PIM tasks discussed in Section 2.1.4.

3.2.1 Why Do People Organize Their Information?

The primary reason that motivate people to organize their information is to be able to retrieve it later (Lansdale, 1988). Not organizing it, or managing it "badly" (in a subjective sense) leads to *"feelings of guilt, stress, and lack of control"*, to quote participants of a study by Boardman et al. (2003). Moreover, Malone (1983) also discovered that documents have a crucial *reminding* function. People arrange documents in their offices so that they are reminded of tasks to perform with those documents. The final goals of organizing information are thus: (1) to retrieve it later with ease (Jones and Teevan, 2007), (2) to have a feeling of control over it and (3) to remind of tasks to do (Allen, 2002).

3.2.2 What Categories of Information Do People Manage?

Several research works attempted to classify the types of information people manage. Barreau and Nardi (1995) propose to categorize information into three classes according to the frequency with which information items are accessed: ephemeral, working and archived. Sellen and Harper (2003) introduce a temperature metaphor (hot, warm, cold) which is related to the last time information was accessed. The same metaphor was used by Elsweiler (2007). Boardman and Sasse (2004) further refine this conceptual basis by proposing to assess information usefulness, which can belong to one of four categories: active (ephemeral and working), dormant (inactive but potentially useful), not useful, and not assessed (e.g. unread emails) information.

3.2.3 How Do People Organize Their Information?

I review here research in the physical and in the digital domain which sheds light on factors influencing classificatory decisions, and types of classification strategies.

Studies of Physical PIM and Implications for Digital PIM. Malone (1983) conducted an ethnographic study with ten participants, aimed at uncovering the different management strategies people employ to manage their physical desks. From the interviews and observations he did, he was able to devise two broad types of management strategies: the *"neat"* office and the *"messy"* office. In a neat office, the categorization scheme is well determined and the filing system is structured. In a messy office, piles of documents and a generally unstructured layout stand out. Malone explains that the tendency to have a neat or a messy office can be a consequence of the type of job. People having routine jobs with clear flows of documents tend to have neater offices than people having non-routine jobs. Of course, individual differences and personal style also influence the management strategy[1]. He also defines the two main organizational strategies for documents: *filing* and *piling*. Three conclusions of Malone's study are relevant to digital PIM: computer systems can help to (1) create classifications (including hierarchical classifications deeper than real ones) (2) classify documents (in multiple folders at once or automatically), and (3) retrieve documents (using several dimensions at once). Interestingly, with his arguments, Malone advocates for the use of a kind of automated faceted classification:

[1]As a matter of trivia, consider the following excerpt of Benjamin Franklin's Autobiography, that describes the virtue of Order: *"ORDER gave me the most trouble; and I found that, tho' it might be practicable where a man's business was such as to leave him the disposition of his time, that of a journeyman printer, for instance, it was not possible to be exactly observed by a master, who must mix with the world, and often receive people of business at their own hours"*.

> *Some documents, such as electronic mail messages and on-line forms, contain explicit fields for information like title, author, and so forth. These documents can be automatically classified according to these fields, with no effort on the user's part. [...] Another simple way of automatically classifying documents that is potentially very useful is based on when the documents were accessed. [...] One can even imagine a system where users search for a document by a kind of simulated time-lapse photography of the history of their electronic desktop. They could "rewind" and "fast forward" the desktop to locate the last time the desired document was on the desk.* (Malone, 1983)

Indeed, classification itself is seen as a cognitively difficult task (Malone, 1983; Lansdale, 1988), demanding time and effort on the part of users, which they are often not ready to invest. Discussing previous research on the topic, Lansdale (1988) further notes the following dilemma: evidence shows that the more investment is put into organizing information, the less difficult it will be to retrieve it later; on the other hand, the more the user is asked to do at the storage level, the less likely he is to do it.

Regarding the factors influencing the organization of documents, Cole (1982) studied the document collections of 30 office workers. She identified six aspects of documents that influence the filing location: *"type, form, volume, complexity, functions, and levels of information"*. Kwasnik (1989a) reports similar results from her interview and observation of 8 academic staff workers. She organized the factors influencing classification into seven dimensions: Situation attributes (e.g. source, use, circumstance, and access); Document attributes (e.g. author, topic, and form); Disposition (e.g. discard, keep, postpone); Order/Scheme (e.g. group, separate, and arrange); Time (e.g. continuation, duration, and currency); Value (e.g. importance, interest, and confidentiality); and Cognitive state (e.g. "don't know" and "want to remember"). Moreover, she provides a ranking of those dimensions according to how many times the users mentioned them during the interviews. It turns out that "Use" is the most often mentioned attribute (i.e. what an item has been used or will be used for). This means that the main dimension influencing classification of the item seems to be the actual use of the item. Moreover, document-related attributes were mentioned in 30% of the cases, whereas attributes related to interaction with documents and context are mentioned in 70% of the cases.

The studies presented above were all based on the management of physical desks. However, Cole, Malone and Lansdale were explicitly interested in providing design principles for digital PIM systems that were emerging. The next paragraph reviews relevant studies of how people organize their personal information in the digital domain.

Studies of Digital PIM. Barreau (1995) and Barreau and Nardi (1995) pioneered the studies about digital PIM behaviours. Barreau (1995) extended the study by Kwasnik (1991) and tried to assess whether the dimensions devised by Kwasnik for classifying physical documents also apply to digital documents. Her conclusions were that most of the dimensions are the same. Nevertheless, in electronic environments, people also often rely on system defaults to store documents and *"rely upon software features for sorting and reordering information on-the-fly"* (Barreau, 1995). The tools used to manage PI thus become critical parts of the management process. Barreau and Nardi (1995) further observe that people generally avoid too elaborate filing schemes (a conclusion also reached by Boardman et al. (2003)) and archive relatively little information because archiving takes too much time with respect to the perceived value of the information. Reminding is confirmed as a critical function of PIM and one of the main reasons that motivate people to classify information items, as was the case in the physical domain (Malone, 1983).

Whittaker and Sidner (1996) dived into people's emerging email habits and proposed three categories of email management strategies: *no-filers*, *frequent filers* and *spring-cleaners*, the former category being further extended by Bälter (1997) into *folderless cleaner* and *folderless spring-cleaner*. Abrams et al. (1998) did similar work for bookmark management strategies. However, despite the elegance of these categorizations, the authors note that people generally fail to strictly rely on one single PIM strategy. In the study by Boardman and Sasse (2004), an attempt was made to define strategies that span multiple tools and information formats. Interviewing and observing 31 university staff members and students, they found different behaviours across email, file and bookmark management. For files, "total filers" classified almost all files upon creation, "extensive filers" had a majority of classified items, but still many items unclassified, and "occasional filers" had only a few folders and left most files undefined. For emails, "frequent filers" filed or deleted most incoming messages on a daily basis; "extensive filers" filed a large number of messages; "partial filers" filed less than 5 messages per day; "no filers" filed nothing. Bookmark management strategies were "extensive filers", "partial filers" and "no filers". The comparison of strategies across tools provides interesting figures: 26% of participants were pro-organizing in all three tools, 45% were pro-organizing in files and emails only, 23% were pro-organizing in files only and 6% were organizing-neutral in all tools (i.e. low organizing effort involved). Furthermore, they discovered an overlap in folder names for files and emails. Folders named after a project or a role tended to appear in both hierarchies. Bergman et al. (2008b) conducted a similar study of file, email and bookmark management with 34 participants. They report a similar overlap between hierarchies: 20% of all folders had a folder relating to the same project in at least one other hierarchy and 56% of all information items belonged to a folder that overlaps a folder in another hierarchy.

Finally, regarding the way people classify, Bergman et al. (2008b) found that people rather organize their information according to the project it belongs to than its format. He notes that people mostly talk about projects when referring to the way they organize their information: 71% of paragraphs extracted from the interviews mentioned projects, whereas only 28% mentioned format. Moreover, 80% of folders had a name reflecting the name of the project they belonged to; 6% of them had format folder names; the remaining 14% had names of persons or other names. Barreau (2008) conducted a study of the PIM behaviour of managers she already interviewed ten years before (Barreau, 1995). This let her define five filing strategies observed in her study that can be applied to any type of information: (1) task-related filing (applied to working information), (2) topic-related filing (applied to active and archived information), (3) time-related filing (applied in particular for ephemeral information, e.g. email), (4) provenance-related filing (applied for emails, filed by senders) and (5) form-related filing (e.g. all images in the same folder or all correspondence in the same folder).

3.2.4 Do People Change Their Management Strategies and Why?

I distinguish in this paragraph between short-term and long-term strategy changes.

Concerning the short term, Boardman and Sasse (2004) note that two participants of their study reported a strategy change that occurred during on average 286 days their information management strategies were tracked. One moved all active items onto the desktop, using it as the predominant working area. Factors influencing this change were: (1) the need to separate active files for synchronization with a laptop, and (2) the influence of the interviews with the experimenter that made the participant "much more aware of all [her] directory structure". Another participant moved completed project folders under an "old" folder in files and emails hierarchies. The reasons invoked were: (1) the need to have less top-level folders and (2) the influence of the participation in the study. In both cases, it seems that participation in the study induced an increased reflection on PIM and was the main factor causing the change (Boardman and Sasse, 2004). Concerning emails, Bälter (1997) proposed that people who receive high quantities of email are likely to change their email management strategy. He also defines two broad types of strategy changes: towards more structure or towards less structure.

Concerning the long term, the study by Barreau (2008) tried to assess whether the management behaviour of participants in a study conducted 10 years before had changed. She interestingly found that their behaviours *"changed little despite technological improvement"* and *"the expansion of the electronic environment"*.

However, PIM strategy changes are under-studied. Besides anecdotal evidence collected e.g. by Boardman and Sasse (2004) and Bälter (1997), only Barreau explicitly focused on them. No study focused on PIM strategy changes in the short term. They thus remain poorly understood. The study presented in Section 4.3 attempts to shed more light on PIM strategy changes. Additional details on studies reporting strategy changes are further examined in Section 4.3.5.

3.2.5 What Do People Remember About Their Personal Information?

The review of cognitive psychology research by Elsweiler (2007), covered in Section 2.1.6, discusses the role that memory plays in PIM. In particular, the following conclusions were reached, with respect to the functioning of memory in general:

1. Auto-biographical information is more easily remembered than other information.
2. Contextual cues of information are naturally encoded in memory.
3. Categorization is a natural process for encoding information in memory and helps recollection.
4. Precise details of information are forgotten, whereas the general meaning and schemata remain.
5. Contextual cues help recollection.

With respect to PIM, the importance of contextual cues for re-finding information was foreseen by Lansdale (1988), inspired by the concept of *"cue enrichment"* (Cole, 1982). However, he insists on the need to test his hypothesis by experiment. Several experiments have since confirmed it. They are presented hereafter.

The study by Malone (1983), introduced above, collected anecdotal evidence about the importance of contextual cues, which got mentioned often when participants described a re-finding task:

> *Many actual information needs seem to be more naturally specified by using more than one dimension at a time (e.g., "a message from M.A. Smith, last week, about the meeting in Palo Alto"). For instance, about one-third of the descriptions co-workers gave of documents for interviewees to find involved more than one dimension (e.g., author and title of a paper, or title of a form and the name of the person it was about). Furthermore, for about two-thirds of all the retrieval probes, the documents were not filed under the dimension(s) used in the description. For example, a document described by title and author was actually filed under the project to which it pertained.*

In a questionnaire administered in the context of their initial study, Ringel et al. (2003) found that people tend to remember the theme (average 5.5/7), people (average 5.3/7) and time (average 4.4/7) related to a document.

The interviews conducted by Bergman et al. (2008b) consisted in a guided tour of participants' PSI, in order to assess several research questions. The interviews were recorded, transcribed and divided into paragraphs related to one information item or collection at a time. Each paragraph was then annotated by the experimenter according to its content. Annotations belonged to three categories of contextual attributes (cues) of information items:

- External context, i.e. the participant mentioned the environment in which they have worked with this information item (e.g. other applications used, other information items opened)
- Internal context, i.e. the participant mentioned their thoughts about this information item (e.g. what they thought when working with it).
- Item's use context[2], i.e. the participant mentioned what they did with the information item in the past, what is its current state and what are their plans for this information item.
- Social context, i.e. the participant mentioned other persons relating to the information item.

Bergman et al. classified each paragraph extracted from the interviews of participants. Half the paragraphs contained at least one attribute related to the context: 29% mentioned the item's use context; 25% mentioned the social context (mainly when the participants described their email organization); finally, 7% mentioned the external context and 1%, the internal context.

On the other hand, categorization (i.e. in most PIM systems, filing into a hierarchy) is also supposed to help recollecting. In a study by Barreau and Nardi (1995), users were reported to remember well the locations of information items in the hierarchy. Nevertheless, some studies also tend to indicate that it may not be a sufficient cue for re-finding certain items. The results of a small-scale study conducted by Golemati et al. (2007) report significant difficulties on the part of the users with finding information items in hierarchies at depth ≥ 4. Moreover, participants did not remember the exact folder containing the target item in 17% of retrieval tasks, on average. For 13% of the tasks they actually did not know at all in which folder to look, whereas for the remaining 4% they knew a parent folder and had to browse inside it.

3.2.6 How Do People Retrieve Their Personal Information?

As noted by Elsweiler (2007) and Capra and Pérez-Quiñones (2003), there have been significantly more studies of re-finding behaviours concerning general information than personal information. Moreover, studies of re-finding behaviour in PIM has focused on webpage re-finding, which I will not review exhaustively here.

Lansdale (1988) first proposed a framework for information retrieval consisting of two steps: recall-directed search and recognition-based scanning. In the first step, memory is used to recall cues about the sought-after item to get as close as possible to it (e.g. its location in the folder

[2] Bergman calls it the "temporal context", but I prefer the terminology of item's use context which is less confusing in the context of this thesis.

structure). If the item is not reached thanks to an exact recall, the user is forced to scan a list of items (e.g. files in a folder) in order to recognize the relevant one.

In addition, the dichotomy between browsing and searching (discussed in Section 2.2.1) is still a matter of debate (Mackinlay and Zellweger, 1995). Teevan et al. (2004) describe two ways of retrieving information which can be compared to browsing and searching: (1) orienteering, which *"involves using contextual information to narrow in on the actual information target, often in a series of steps"*, and (2) teleporting, where users try *"to take themselves directly to the information they are looking for"*. Barreau and Nardi (1995) observe that users prefer location-based search (i.e. browsing) because it *"more actively engages the mind and the body and imparts a greater sense of control"*. Browsing is also seen as the preferred way for users to re-find personal information by Marchionini (1997), Boardman and Sasse (2004) and Bergman et al. (2008c). Freeman and Gelernter (1996), on the other hand, advocate for the use of advanced search mechanisms, using keywords as well as other attributes to retrieve information items using queries, arguing that operating systems' default search tools are poorly designed.

Bergman et al. (2008c) conducted a large-scale study that aimed at assessing the preference of users between navigation and search for re-finding files. It turned out that 56% to 68% of the retrieval events were navigation-based, depending on the conditions (e.g. Mac or PC, and different phases of the study). Mean search percentages were much lower (between 4% and 15% depending on the conditions). The preference for navigation over search is significant in all conditions. Other means of retrieval were also noted: 16%-20% of retrievals used desktop shortcuts, 5%-12% used recent documents lists and 1%-4% used other means (including "Smart Folders" in the Mac studies). Qualitative results from questionnaires indicate that the inability to remember a file location is the dominant reason for search (75% of responses). A secondary goal of the study was to assess whether advanced search engines (e.g. Spotlight, Google Desktop Search) lead to an increase in searches, compared to traditional search engines (e.g. Sherlock, Windows Search Companion), as was foreseen by Lansdale (1988) and Freeman and Gelernter (1996). Empirical evidence from the study show that those advanced search engines do not make people rely more on search and less on navigation. Another study comparing the search and navigation habits of Windows and Linux users led to similar results (Bergman et al., 2008a).

Dumais et al. (2003) report on the use of a search tool by 234 people during six weeks. Among other interesting results, they discovered that queries used to re-find personal information averaged 1.59 terms, which is shorter than web queries (> 2). Moreover, queries rarely included explicit boolean operators, phrases, or field restrictions. Instead, participants preferred to iteratively filter through result sets by using interface widgets that would dynamically provide the query refinement. This is consistent with the orienteering behaviour observed by Capra and Pérez-Quiñones (2003) and Teevan et al. (2004). In addition, Dumais et al. (2003) noticed the importance of people and time in helping people re-find: people's names were used in 25% of the queries logged and dates were frequently used to sort results.

Finally, Dumais et al. (2003) give some clues concerning the format of accessed items using their personal search tool. The most sought-after items were emails (76%), web pages (14%) and files (10%). With respect to the age of the items, 6.6% of the items were first seen that day, 21.9% within the last week, 45.9% within the last month, and 89.4% during the last year. People searched mainly for recent items, but also for items up to eight years old .

3.2.7 Summary of Empirical Contributions

The empirical studies reviewed in this section show that people organize their information to retrieve it later, to have the feeling of being in control of it and as a reminder of things to do.

People manage different categories of information, with respect to their usefulness, use and frequency of access.

Several factors influence the classification of information, mainly its *use*. Classification strategies can be devised, with respect to the frequency of filing and the dimensions used for filing items. The strategy employed for a particular format of information or in a specific PIM tool is not necessarily reproduced in other tools, though a majority of people seem to be pro-organizing in files and emails. Overlap between different hierarchies may also be important, in particular with files and emails.

Organizing strategy changes in the short-term seem to be motivated by an increased reflection about one's PIM practices. High-level behaviours, though, seem to vary little over the years. However, the process of strategy changes has been under-studied.

Users seem to remember well contextual cues of information items, in particular the item's use, the social context and the temporal context. Categorization of items, though helpful in general, is partially forgotten when it gets too complicated (e.g. hierarchy too deep).

When it comes to information re-finding, people generally prefer browsing over searching. Searching is mainly used when the location of an item has been forgotten and especially for emails. People use few keywords in search queries and prefer refining the results using interface widgets, a process known as orienteering.

Though empirical evidence show a preference for browsing over searching and the importance of contextual cues, existing PIM tools do not explicitly support it, as will be described in the upcoming part of this review.

3.3 Review of Relevant Tools

Vannevar Bush first laid out the foundations of modern PIM tools when he devised the *Memex*, *"a device in which an individual stores all his books, records, and communications"* (Bush, 1945). The device he imagined would be mechanical so that the user could consult it *"with exceeding speed and flexibility"* in order to find the information he is looking for. Now magnetic storage and integrated circuits have replaced microfilms and mechanical devices but the need for such an ideal device remains. In recent years, the exponential growth in the volume of personal information that computer users have to handle, and the lack of tools to efficiently manage it, has pushed the area of PIM research forward. Thus, PIM tool design is an active area. Consequently, this section does not attempt to be exhaustive. It rather provides a representative summary of PIM tools and research prototypes relevant to this work. The focus is put on tools that support (1) the re-finding task of PIM and (2) the self-analysis task of PIM, that may influence PIM strategy changes.

3.3.1 Tools That Support the Re-finding Task

Widely available PIM tools are presented first, along with some improvements or alternatives that have been suggested to improve their performance. Other PIM tools, which have mostly been developed in the context of academic research, are presented afterwards. Building on the conclusions of the previous section, PIM tools are presented based on how they relate to the use of contextual cues for re-finding. Tools that apply faceted navigation are presented last. Table 3.1 summarizes the features of tools presented in this section.

Widely Available Tools

The most widely available PIM tools are those offered by operating systems and standard office programs. They may be navigation-based or search-based tools.

Navigation-based. Navigation tools firstly include filesystem browsers (Windows Explorer, Mac OS Finder, and so on) that support browsing a hierarchy of files. They are not restricted to the re-finding task and enable other PIM tasks, like the creation of collections and their maintenance. Many other tools enabling the management of hierarchical personal information of a particular format also exist. For example, email clients and bookmark managers generally support hierarchies. Classification in these tools is generally done by the user, though it can sometimes be automated (e.g. rules for email clients which put incoming emails into a specific folder based on their content or attributes). Re-finding is done through several steps of navigation in the hierarchy. The main advantage of these systems is the fact that browsing is preferred to searching by users. Browsing gives them a feeling of control over their PSI. Furthermore, those systems support the reminding

	Cues						Interface		
	User-defined cue	Temporal cue	Social cue	Type cue	Keyword search cue	Other cue	Search-based	Browse-based	Faceted navigation
Forget Me Not (Lamming and Flynn, 1994)		×	×	×				×	
LifeStreams (Fertig et al., 1996)		×			×			×	
TimeScape (Rekimoto, 1999)		×			×			×	
Placeless Documents, *Presto* (Dourish et al., 2000)	×						×		
MyLifeBits (Gemmell et al., 2002)	×	×			×		×		
Milestones in time (Ringel et al., 2003)		×			×			×	
UMEA Kaptelinin (2003)	×							×	
Stuff I've Seen (Dumais et al., 2003)	×	×	×	×	×	×	×		
ContactMap (Whittaker et al., 2004)			×					×	
Haystack (Karger et al., 2005)	×				×			×	
TimeSpace (Krishnan and Jones, 2005)	×	×						×	
FacetMap (Smith et al., 2006)		×	×	×	×	×			×
Phlat (Cutrell et al., 2006b)	×	×	×	×	×	×			×
PhotoMemory (Elsweiler et al., 2007)	×	×							×
MemoMail (Elsweiler, 2007)	×	×	×	×	×				×
Project Planner (Planz) Jones et al. (2008)	×							×	
Project contact list Bergman et al. (2008b)	×		×					×	
Feldspar (Chau et al., 2008b)		×	×	×		×		×	
Facet Folders (Weiland and Dachselt, 2008)	×	×							×

Table 3.1: Summary of the cues and interfaces available for re-finding personal items in PIM tools. Systems are presented in the chronological order.

function. The limitations of these systems include the cognitive difficulty of classifying items and, as far as re-finding is concerned, the risk of forgetting where items were classified if the hierarchy is too complex.

Alternatively to hierarchical navigation, another paradigm of navigation is hypertext. In the *XanaduSpace* project by hypertext visionary Ted Nelson[3], navigation in a dummy textual PI space is done by using hyperlinks. Unlike *"one-way ever-breaking"* hyperlinks found on the web, *XanaduSpace* enables multiple categories of links, as well as management of versions and content (e.g. indicating the source of the content). Nevertheless, the author does not know any actual PIM system implementing these concepts.

Search-based. Search tools have become more and more widespread. The first search tools used pattern matching algorithms on file names or the textual content of files (*find* and *grep* Unix tools, Sherlock, Windows Search Companion). Now, common operating systems include more advanced search tools which index personal information beforehand (Google Desktop Search, Windows Search, Copernic Desktop Search, Spotlight and several desktop search engines for Linux). They are capable of indexing a whole range of formats, as well as including many file attributes into the index. With the help of the index, queries are faster and can be parametrized more easily. Specific PIM tools, like email clients or web browsers, generally include search mechanisms as well. Re-finding using search tools can be seen as an attempt to "teleport" directly to the needed information item. The main advantage of these systems is that they can be fast if specific keywords of the sought-after items are remembered. Their main disadvantage is that they are useless if specific keywords are not remembered. And, the previous section of this review shows that this might often be the case.

Given the difficulty of recalling the precise keywords needed to perform textual search, there have been many research works on query expansion. These works used the Wordnet thesaurus to expand queries with synonyms or hyperonyms (Kim et al., 2004). However, other studies (Kluev, 2002) show that query expansions using thesauri do not necessarily improve the results, as the expansion itself is likely to generate mainly noise. Furthermore, the author does not know of any application of query expansion that has been done in the particular context of search in PIM.

Tools that Focus on User-defined Cues

MyLifeBits (Gemmell et al., 2002) is a database of resources and links, providing a primarily search-based re-finding mechanism. Links indeed represent either user-created collections of resources or so-called transclusion (i.e. a resource cited or used by another one, e.g. a spreadsheet embedded in a textual document). Results of queries in MyLifeBits can be viewed using traditional detailed

[3]See `http://www.xanadu.net/`.

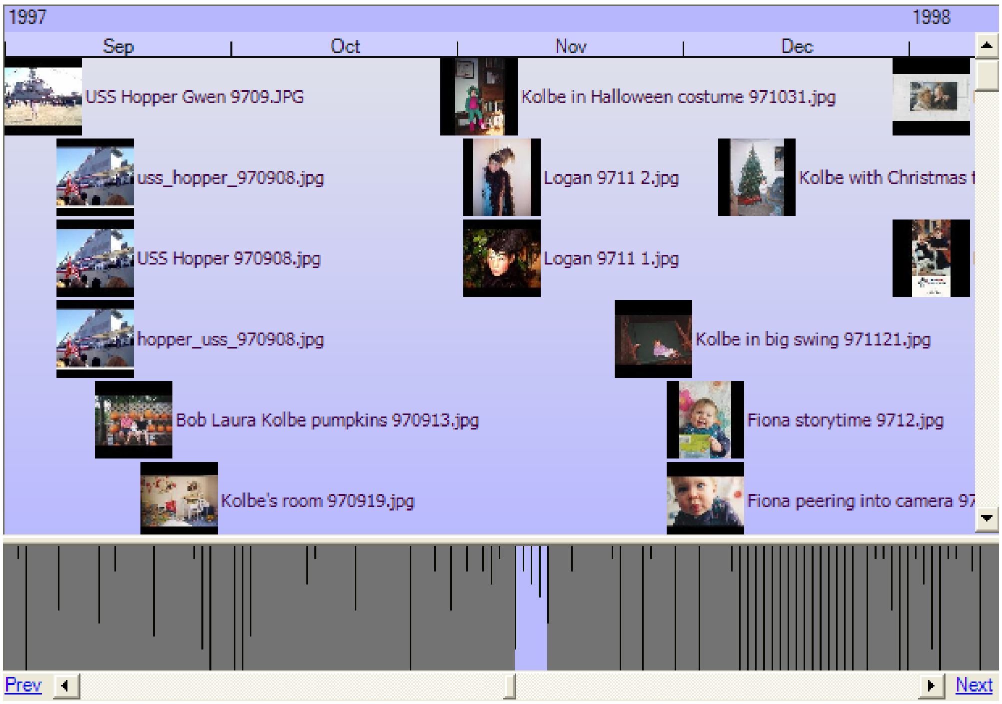

Figure 3.1: *MyLifeBits*'s time-based view, by Gemmell et al. (2002).

or thumbnail views, or using more original and flexible time-based visualizations which support navigation (see Figure 3.1). Searching on personal annotations is also possible.

Placeless Documents and the related *Presto* (Dourish et al., 2000) systems are based on user-defined collections and system-level annotations. Collections can be defined by several complementary means: (1) a query, (2) an inclusion list, (3) an exclusion list. Moreover, system-level annotations like document use patterns in time are automatically collected by the system, which allows to re-find documents based on other documents that were opened at the same time. Both apply searching for re-finding.

Haystack (Karger et al., 2005) lets the user define annotations and link related PI items together, no matter their format, thus enabling a rich semantics to be put on top of PI. However flexible, the presentation of PI for re-finding is generally conservative, although ways of benefitting from links between related data are provided. No specific improvement on re-finding over traditional search or navigation systems has been devised. Xiao and Cruz (2005) also focus on the semantic aspect of personal information, without putting an emphasis on the re-finding task.

The *Project Planner (Planz)* by Jones et al. (2008) allows the creation of project plans. Plans are lists of tasks related to a project. Plans are hierarchical and external documents can be attached to them. This allows integration of information items coming from different sources into the planner, which acts both as a repository of links to documents and as a task manager. Re-finding items related to a project created in the planner is facilitated.

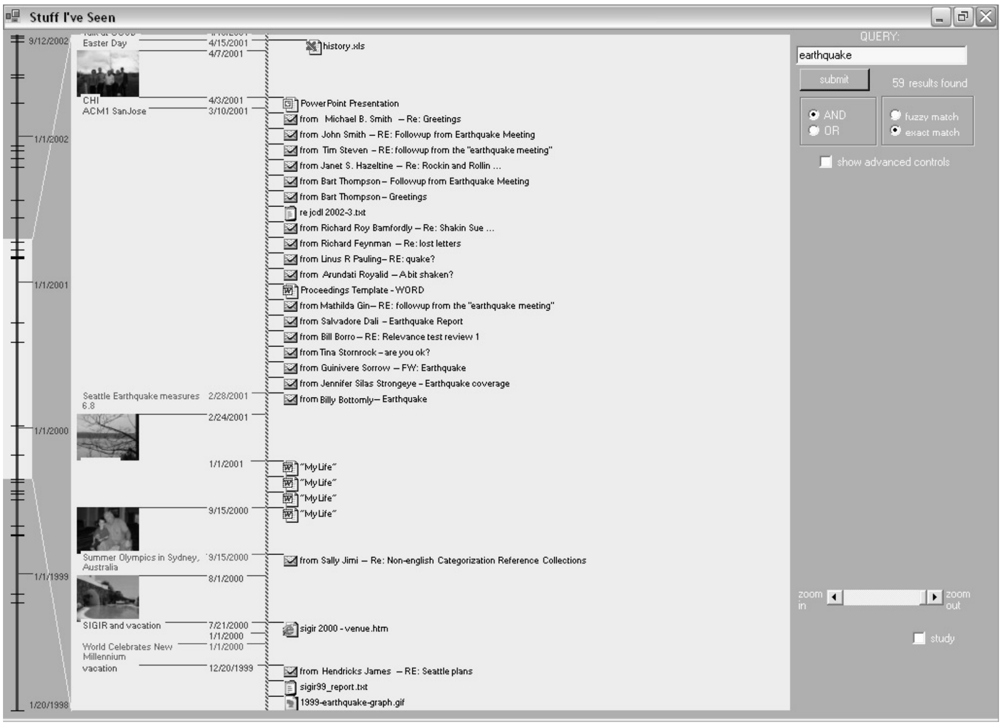

Figure 3.2: *Milestones in time*, by Ringel et al. (2003).

Finally, Google's *Gmail* and other commercial email managers enable multiple concurrent categorizations of emails (labels), and enable search-based re-finding.

All in all, the main issue with the systems presented in this section is the effort needed on the part of users to maintain annotation collections on their PI. Evidence presented in the preceeding section show that people are unlikely to invest time and effort into manually defining annotations and relationships between PI items.

Tools that Focus on Temporal Cues

Milestones in time (Ringel et al., 2003) is a search-based interface extending *Stuff I've Seen* (Dumais et al., 2003) that tries to replace information in its temporal context, making use of episodes and temporal landmarks. The evaluation of two variants of the system shows that search time is significantly reduced when the user has access to the episodic content (landmarks) instead of having only date information (without landmarks). It is depicted in Figure 3.2.

Forget Me Not by Lamming and Flynn (1994) (see Figure 3.3) suggests exploiting episodic memory. The developed prototype automatically stores encounters with people, phone calls, emails and file exchanges along with their timestamp, the place where they happened and the people involved. Re-finding events or documents follows the orienteering approach. Indeed, it involves applying filters on people, places or kinds of communication means (encounter, phone, email,

Figure 3.3: *Forget Me Not*, by Lamming and Flynn (1994).

file). Unfortunately, this system only allows the management of personal information related to communication (emails, exchanged files, phone calls).

LifeStreams, by Fertig et al. (1996), set time as the main dimension for organizing PI items, in order to exploit our ability to recall the approximate temporal period when a document was created or modified. *TimeScape*, by Rekimoto (1999), extends on the previous work by allowing both temporal and spatial organization of items. Both those tools use navigation as the main re-finding means, though searching is also possible to a certain extent. However, no evaluation sufficiently validates their approach. The commercial Time Machine environment featured in the latest MacOS X operating system leverages the seminal work by Rekimoto (1999). But, it is used rather as a backup tool, complementary to the more traditional desktop metaphor which remains the standard for organizing and re-finding items.

TimeSpace (Krishnan and Jones, 2005) further extends the work by Rekimoto (1999). It organizes information items based on the activities (projects) they are related to and the time when they are accessed. Categorization of information items is done manually, and interaction with the environment is automatically recorded so that the state of the environment when the user was working on a specific activity can be restored. Re-finding proceeds by selecting the appropriate activity and then using either the spatial layout or the timeline to re-find the sought-after item or work environment. Search or filters are not available. *UMEA* by Kaptelinin (2003) is close to *TimeSpace* and also proposes both activity-based and time-based organization of information items. Re-finding mechanisms are similar. The definition of activities is nonetheless left to the user, thus suffering from the same limitations as the previously discussed systems that focus on user-defined cues.

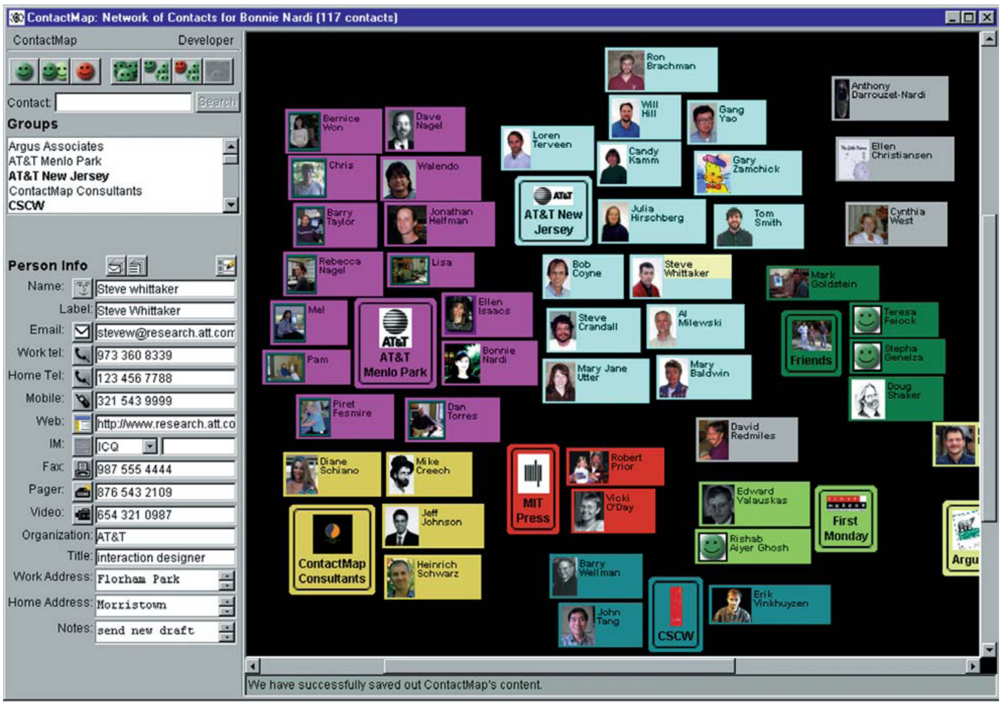

Figure 3.4: *ContactMap*, by Whittaker et al. (2004).

Tools that Focus on Social Cues

ContactMap, by Whittaker et al. (2004), is an interface using the social network of the user as a means to access their PI (see Figure 3.4). Interacting with the visual representation of their social network, the user can access emails and chats related to their contacts. The authors present the system as a means to access information based on people rather than messages. Their evaluations show a great interest from the part of the participants for the approach. As with the system developed by Lamming and Flynn (1994), however, it is limited to a subset of PI relating to communication. Furthermore, the social network in *ContactMap* is built by manual annotation.

Bergman et al. (2008b) introduces a design principle for PIM systems they call the *Project contact list*. It proposes to assign contact lists to folders, that would map the information items contained within the folder to a subset of the user's contacts. No implementation of the system has been proposed, though it seems implementations are ongoing work[4].

Tools that Focus on Multiple Cues

A precursor of desktop search engines in its philosophy, *Stuff I've Seen* or SIS (Dumais et al., 2003), depicted in Figure 3.5, is a universal search system working over several formats of PI (files,

[4] `http://www.user-subjective.com/` demonstrates prototypes putting some of Bergman et al.'s design principles in use (retrieved May 17th, 2010)

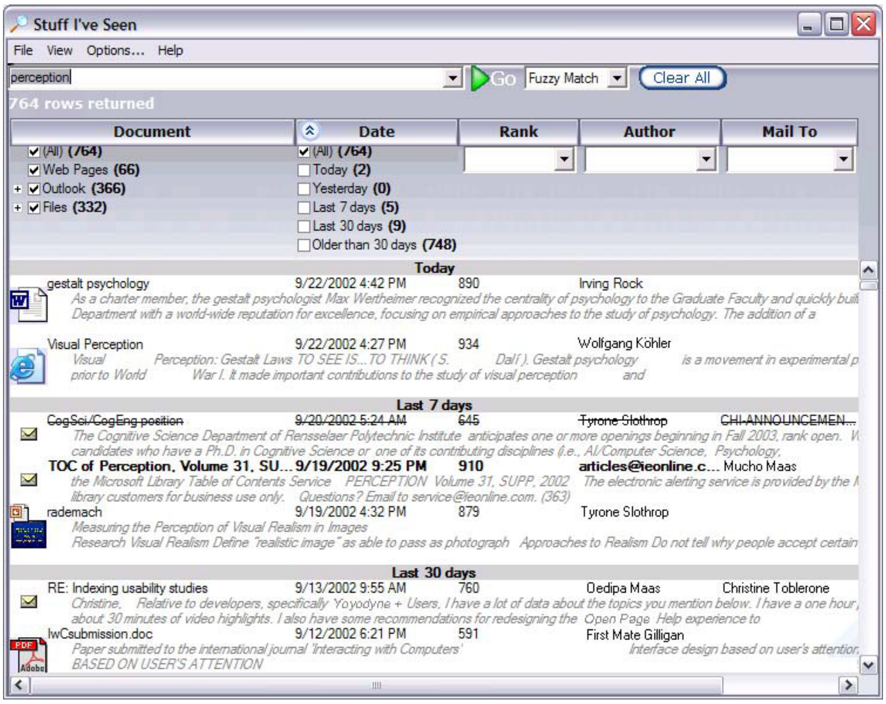

Figure 3.5: *Stuff I've Seen*, by Dumais et al. (2003).

emails, web history, etc.). It builds a single common index for all types of personal information, as advanced search systems do (e.g. Windows Search, Spotlight, Google Desktop Search and so on). SIS enables free-text queries as well as filters on date, author, type of document and other available metadata. Search results are presented in a textual list that may be ranked by relevance or sorted chronologically. Acting on interface widgets directly filters the result set without having to resubmit a query, thus supporting an orienteering approach to information re-finding. The developers conducted a user evaluation with more than 200 users that confirmed the importance of social and temporal cues for retrieving information.

In the same vein, *Feldspar* (Chau et al., 2008b) provides a visual interface to help build complex search queries by association of multiple cues.

The open-source mail client *Mozilla Thunderbird 3*[5], released in December 2009, offers a search interface augmented with filters that is very similar to SIS.

The tools that enable multiple cues as filters for search results are very close in spirit to faceted navigation tools presented in the next section.

[5]See `http://www.mozillamessaging.com/en-US/thunderbird/features/` (retrieved June 8th, 2010).

Tools Enabling Faceted Navigation

As defined in Section 2.2.2, faceted navigation is a paradigm which supports browsing and searching a collection using so-called facets. Facets can be seen as independent dimensions of the data, consisting of several categories. In the navigation process, each selection in a facet acts as a filter on the original collection. Several aspects differentiate faceted navigation from traditional browsing and orienteering with filters applied on search results. I present here a set of features that faceted navigation systems may have, depending on how much they conform to the faceted approach:

1. The categories within a facet are collectively exhaustive, i.e. each information item belongs to at least one category for this facet.

2. The categories within a facet are mutually exclusive, i.e. each information item belongs to at most one category for this facet.

3. Representations of categories show aggregative metadata about the data, e.g. the category of Charles Baudelaire in the author facet shows how many books by Charles Baudelaire are available.

4. A selection in a facet updates the representations of the other facets so that their categories reflect aggregative metadata about the subset of the initial collection currently represented, e.g. selecting Charles Baudelaire in the author facet updates the language facet which will only show the French language (unless Baudelaire did write books in another language).

5. Categories which do not appear in the collection are hidden. This implies that one cannot navigate to a "dead-end".

Since the work by Daniel Tunkelang[6] and Marti Hearst (Yee et al., 2003) made faceted navigation popular, numerous websites apply it (*Yahoo!*, *Flamenco*, countless e-commerce websites and so on), but local applications are still rare, with the notable exception of music library managers (*iTunes* and the like). Evaluations of the faceted interface developed by Yee et al. (2003) against a standard image search system nevertheless show a clear benefit for the faceted approach: *"90% of the participants preferred the faceted approach overall, 97% said that it helped them learn more about the collection, 75% found it more flexible, and 72% found it easier to use"* (Yee et al., 2003), despite the fact that it was often much slower than the search-based interface.

Most of the work on faceted navigation has focused on collections that are not personal, in particular collections of images and media (Yee et al., 2003; Tvarozek and Bielikova, 2007) or generic

[6]His *Endeca* e-commerce platform applied faceted search as early as 2001, see http://web.archive.org/web/20010202153500/endeca.com/app/ecommerce.shtml (retrieved October 12, 2010).

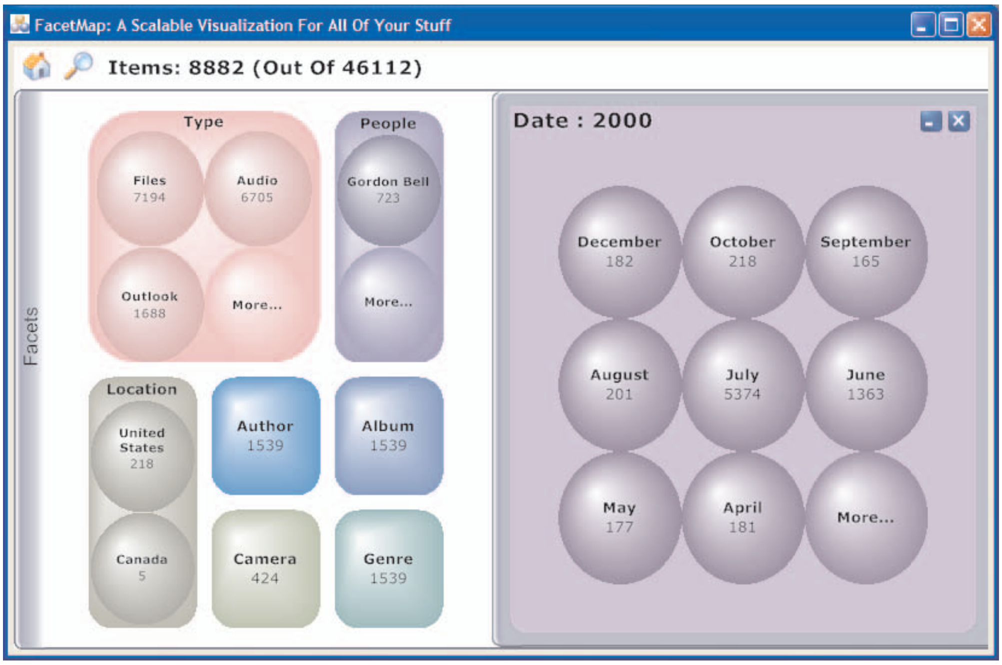

Figure 3.6: *FacetMap*, by Smith et al. (2006).

datasets (Schraefel et al., 2005; Capra and Marchionini, 2008). In PIM, Malone (1983) already suggested that electronic documents may be automatically classified along facets, thus advocating for faceted navigation: *"some documents [...] contain explicit fields for information like title, author, and so forth. These documents can be automatically classified according to these fields, with no effort on the user's part"* (Malone, 1983). Several developed systems use faceted navigation at least to a certain extent. They are presented next.

FacetMap (Smith et al., 2006) (see Figure 3.6) is an extension of *MyLifeBits*. Its interface is built on top of the *MyLifeBits* data store and it offers a browsing paradigm based on facets. Facets are represented as bubbles laid out hierarchically in a treemap fashion. However, a formative evaluation of the system fails to reveal a clear preference for faceted browsing. The chosen graphical representation of categories with bubbles is also criticized as being less effective than textual lists in some cases. The ecological validity of this evaluation is void, however. The data used comes from Jim Gemell's own personal information, that he agreed to share as part of the *MyLifeBits* project. This was not the participants' real PI.

FacetLens (Lee et al., 2009) is an extension of the previous work. It is not specifically targeted on PIM but nevertheless introduces an interesting new facet layout, coined the "linear facet". In this facet, categories are not presented as bubbles, but ordered along a line, which conveys their relative ordering. This is a natural way of organizing dates, for example. It also has the increased benefit of providing a support for displaying barcharts which easily convey aggregative values. A

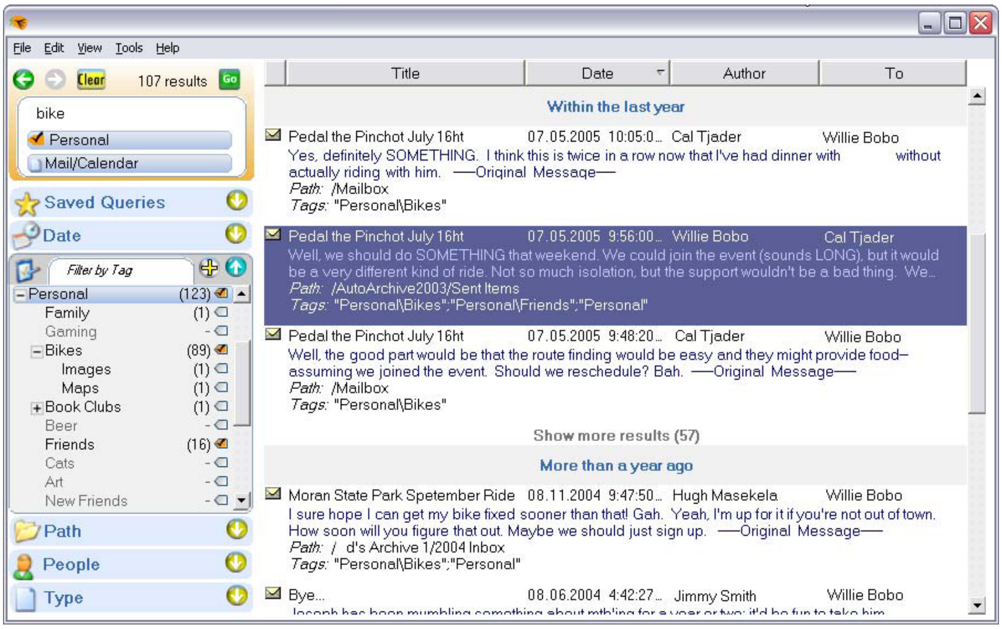

Figure 3.7: *Phlat*, by Cutrell et al. (2006b).

formative study of the system indicates a potential of the faceted approach for gaining insight into the dataset (which was not personal information for that matter).

PhotoMemory (Elsweiler et al., 2007) applies faceted navigation on a store of personal pictures, where the pictures are annotated by hand. The evaluation of the system shows a faster retrieval time with this approach than with traditional browsing in a hierarchical file system. *MemoMail*, by the same author, applies faceted navigation to browse an email archive. Several contextual cues can be used to trigger recollection: social cues, temporal cues, folder and keywords.

Phlat (Cutrell et al., 2006b) (see Figure 3.7) (Prototype for Helpful Lookup and Tagging) improves over *Stuff I've Seen* in order to facilitate the dynamic building of queries. The system is an interface-level on top of *Stuff I've Seen*. Browsing PI is achieved using queries and filters, which are nothing more than facets. Date, path, people and type "filters" can be used in addition to a traditional search box. Moreover, the user can tag PI items and a tag filter is also provided in the re-finding interface, although the relevance of this tagging facility in a re-finding-oriented tool is questionable. With respect to the capabilities of human memory presented in the first part of this review, *Phlat* improves significantly over previous research. Its evaluation confirms an interest for filtering by people and data type. The interface of *Phlat*, however, reflects the will of its developers to make it look primarily like an advanced *search* tool. Before the user explicitly issues a query, no documents are displayed in the result set. Furthermore, a facet is considered to be a filter on the search result set. The authors indeed emphasize that *"a filter is a query"*. Moreover, as the result panel takes up the most space on the screen, there is little space for the facets, despite a design guideline claiming that the *"current search criteria must be salient and visible all the time"*. Additionally, facets are hidden by default. When they are shown, most of their categories are hidden and the user needs to scroll to see them all. All in all, *Phlat* is a good approach as it provides contextual cues in the retrieval process. However, it does not implement all the features expected of a full faceted navigation tool. It therefore fails to give a big picture view of the PSI and does not encourage a navigation experience.

Weiland and Dachselt (2008) introduce a system called *Facet Folders* which embed faceted metadata into the traditional folders hierarchy. The authors claim that this approach provides both the flexibility of the faceted navigation and the benefits of spatial encoding into memory provided by the traditional hierarchy. The system seems to be at an early design stage, though.

In summary, in recent years, tools applying a faceted approach to the re-finding task of PIM have emerged. Nevertheless, only *Phlat* by Cutrell et al. (2006b) was mature enough to be applied in an ecologically valid context and explicitly supports the unification of different data formats (emails, files, web history and possibly other sources). However, it does not fully implement the faceted navigation approach and thus fails to take all the benefit it could from it.

3.3.2 Tools That Support the Self-analysis Task

Self-analysis in PIM is a task that occurs infrequently. It is a high-level and time-consuming task and is therefore not performed regularly. On the one hand, it may be triggered by situations where a management strategy employed becomes obsolete or does not meet needs anymore. Time then needs to be taken to observe strategies from a critical point of view in order to adapt them or to devise new ones. In this sense, the self-analysis task is preliminary to changes in PIM strategies. On the other hand, it can be triggered by simple curiosity. Getting back in contact with an old friend may lead to tracking the previous activities with him, and how they evolved over time. Self-analysis in PIM is not an activity as critical as re-finding needed items. It is an exploratory activity, belonging to the set of meta-level activities (Jones and Teevan, 2007). It is about getting insight into one's personal space of information and examining not its precise content, but its general shape, particular aspects of its evolution over time, or relationships between different facets of it.

Work in PIM that explicitly supports the self-analysis task was mainly developed in the context of emails. Viegas et al. (2006) propose to visualize the "conversational history" between the mailbox owner and a chosen contact during a certain period of time. Their *Themail* system is thus tailored for psycho-social practices analysis. Perer et al. (2006) and Perer and Smith (2006) visually explore relationships through past emails. They use visualization to gain insight into the email practice of the mailbox owner, answering questions such as: how many emails do I exchange with each of my contacts, which contacts are similar according to the pattern of email exchanged with them? The commercial system *Xobni* (Xobni Corporation, 2010) also allows similar analysis tasks on an

Outlook mailbox.

In other academic work on PIM tools, the self-analysis task is never explicitly targeted. However, among the research work presented in the previous section, some mention they improved the ability of users to spot and analyse trends (Smith et al., 2006; Lee et al., 2009) or to help users manage and track the progress of long-term work. Krishnan and Jones (2005) mention that users of *TimeScape "readily related activity histograms to work intensity over time, noticing variations in patterns and work styles of activities carried out and supported within visualizations. These indications of work intensity helped in tracking progress, comparing activities and quickly scanning activities for a particular time-frame of interest."*

On the other hand, faceted navigation tools are intrinsically exploratory (Marchionini, 2006) and help users get a better understanding of the underlying dataset (Yee et al., 2003). Indeed, 97% of the participants in the study by Yee et al. (2003) agreed that the faceted approach helped them learn more about the collection than a search-based interface.

3.4 Chapter Summary

This chapter reviewed the research literature related to personal information management. The main issues in PIM identified in the literature are information fragmentation, information overload and loss of context. To understand how users cope with those difficulties, a number of studies of user behaviour have been conducted. As far as classification is concerned, studies point out that users apply various strategies which may depend on several factors. Users seem to change their strategy as a reaction to a self-analysis of their PIM practices, too, but these changes have been little studied. Surveys of re-finding strategies highlight a preference for browsing over search, and the importance of contextual cues which are easily remembered. System-wise, few PIM tools supporting the re-finding task of PIM exploit several contextual cues simultaneously (see table 3.1). Even fewer explicitly support the self-analysis of PIM practices.

The findings of this chapter guide the work presented in the remainder of this thesis. In particular, previous research pointed out some contextual cues of personal information that could support re-finding. Moreover, recent approaches identified faceted navigation as a promising interface paradigm for supporting the re-finding of information items and the self-analysis of practices. The next chapter presents studies of user behaviour that were conducted to complement those findings and give directions for the development of the PIM tool presented in Chapter 5. In accordance with the objectives of the thesis presented in Section 1.4, the tool exploits faceted navigation and contextual cues of information to facilitate re-finding and self-analysis. This tool is further used in an evaluation presented in Chapter 6 which mainly aims at understanding how users benefit from contextual cues in PIM.

It has been said that man is a rational animal. All my life I have been searching for evidence which could support this.

Bertrand Russell

Wandel und Wechsel liebt wer lebt.

Richard Wagner, *Das Rheingold*, 1:2 (Wotan)

4

Surveys of PIM Behaviour

This chapter describes surveys of personal information management (PIM) behaviour conducted in the course of this thesis work. These surveys aim at increasing the understanding of PIM behaviours and complementing the findings of Chapter 3. At the end of this chapter, the two first objectives of the thesis, presented in Section 1.4, will therefore be fulfilled: (1) identify contextual cues of personal information which support re-finding, (2) identify a user interface paradigm suited to support the re-finding task in PIM and which benefits from the cues. Additionally, this chapter tackles the fifth objective of the thesis: develop an increased understanding of PIM strategy changes. In addition to providing a complement to what could be learned by previous research, the studies presented here were also a way for me to "get my hands dirty" by confronting with real users' PIM problems.

The chapter describes three separate studies. Section 4.1 introduces a preliminary survey which gave directions for the future studies. Section 4.2 presents a questionnaire-based survey of PIM practices which outputs frequency distributions of the use of different strategies for classifying and re-finding files. Then Section 4.3 describes an ethnographically-inspired study of PIM practices with a particular focus on strategy changes. An original model for strategy changes in PIM is presented along the results of this latter study. Finally, Section 4.4 covers the implications of the presented studies for the design of the system introduced in Chapter 5.

4.1 Preliminary Survey and Directions

The first user survey that was conducted targeted the use of technology in relation to meetings. The dedicated online questionnaire received 118 responses from people from various countries and domains. Most of the conclusions extracted from this survey are of little interest in the context of this thesis. However, one particular result relates to the re-finding task of PIM. I will present it briefly here, along with other observations made in the course of this survey. The full report of the survey has been published internally by Bertini and Lalanne (2007).

Table 4.1 presents the answers to the question: "How often do you use the following methods to find files on your PC?". The most used means are: (1) remember the location where the file was stored, and (2) search for the email in which the file was attached. Interestingly, people tend to use less the search facilities and more to remember where to look for. Quite often an email is used to find a file. This supports the idea that browsing has still a great importance over search, even though many current systems concentrate on the latter. Only few participants in the study use desktop search engines to find files on their computer. Other methods suggested by respondents are: "ask colleagues to send them by e-mail (when I don't find them)"; "I have a local html file (maintained by me) with links to files not stored locally + an Excel file with links to locally stored documents"; "different view sorting (e.g. list files by date)"; "I store documents common to a project online in a Yahoo or Google group as do other members of the project"; "We use a specific application design for information management. It is called CrowdTrust"; "I keep working folders on my desktop for easy access to current projects"; "Copernic Desktop Search".

Apart from this, collateral observations confirmed conclusions suggested by related research presented in Chapter 3, in particular the fact that humans have a good memory for certain things (people, places) but not for dates or filenames.

Despite the fact that the results of this survey are clearly not sufficient to improve significantly our understanding of PIM practices nor to guide the design of our system, they give interesting

	Never	Sometimes	Often	Always
I remember where the information is stored in the file system (e.g., the specific containing folder)	0.0%	12.0%	**65.0%**	23.1%
I use the operating system's search functions	17.2%	**65.5%**	17.2%	0.0%
I search for the email having the file attached	10.5%	43.0%	**45.6%**	0.9%
I use other search functions (e.g., Google Search)	**64.9%**	24.3%	9.9%	0.9%

Table 4.1: Responses to the question : *How often do you use the following methods to find files on your PC?*

clues on directions to follow in the upcoming study, presented in the next section.

4.2 Study of Classification and Re-finding Strategies

The study presented in this section was conducted by an undergrad student as part of his M.Sc. work done under my supervision (Thomet, 2010).

4.2.1 Goal

The goal of the study presented here was to get an overview of people's PIM practices relating to files on their personal computer. We wanted to know how often people use the different classification and re-finding strategies identified in previous research (see Section 3.2.3) in their daily practice. In particular, we were interested in hierarchical and temporal aspects of their classification and re-finding strategies that could guide the design of the faceted navigation tool presented in Chapter 5. The focus was explicitly put on files rather than emails and bookmarks or cross-tools as was the case with other studies (Whittaker and Sidner, 1996; Bälter, 1997; Abrams et al., 1998; Boardman and Sasse, 2004). Moreover, the study presented here originally outputs reported frequencies in the use of different means of classification and re-finding.

4.2.2 Questionnaire

We created a web questionnaire using the infrastructure provided by an online tool[1]. The questionnaire comprised 23 questions organized in 4 different parts:

Generalities. Questions about the general profile of the user (age, sex, profession, operating system, frequency of computer use, experience in the use of computers).

Strategies for classifying documents. Questions in this part focused on the following topics: (1) how often do people use filing and piling, (2) how often do they use specific types of

[1]`http://www.surveymonkey.com`.

	Almost never	Rarely	Sometimes	Often	Almost always
Filing	0.7%	4.1%	8.2%	34.7%	**52.4%**
Piling	17.6%	**28.2%**	**28.2%**	20.4%	5.6%
Desktop	22.4%	20.3%	**29.4%**	23.8%	4.2%

Table 4.2: Frequency of strategy use for classifying documents.

folder names, (3) what kind of information items are stored on their desktop. The answers to the questions were based on a 5-level Likert scale (1: almost never, 2: rarely, 3: sometimes, 4: often, 5: almost always).

Strategies for re-finding items. Those questions were centered on the type of cues used in search queries, and other strategies for re-finding items with respect to the age of the sought-after information item. The answers were based on the same Likert scale as in the previous part.

Use of dedicated tools and features. Those last questions aimed at spotting the actual use of and interest in specialized dedicated features of widespread PIM tools.

The full questionnaire and a summary of its results are available in Appendix A. Relevant results are presented next.

4.2.3 Results

Generalities

The questionnaire was completed by 147 persons. The ages of participants ranged from 18 to 62 (mean: 34.63, stdev: 11.04). Fifty-five (55) were females and 92 were males. The major field of work represented was IT (49.4%). Important minorities included teaching (11.9%), internet-multimedia (8.3%), communication (4.2%) and audiovisual (3.6%). Most participants used a computer for more than 6 hours a day (66.1%) and 20.2% used it between 3 and 6 hours a day. Operating systems used included Windows XP (50.6%), Mac OS X (21.4%), Windows Vista (13.7%), Linux (7.1%) and others (7.1%). The vast majority of users (81.5%) rated their expertise with computers as high, and the remaining rated their expertise as knowledgeable.

Strategies for Classifying Documents

Filing Versus Piling. In the questionnaire, we defined "filing" as: "organization into a hierarchy: you define a folder hierarchy into which you put your documents". Piling was defined as: "piling: you put your documents into a unique folder or in folders having less than 2 subfolders". Finally, the last strategy we proposed was desktop piling, labeled as: "you put your documents on the

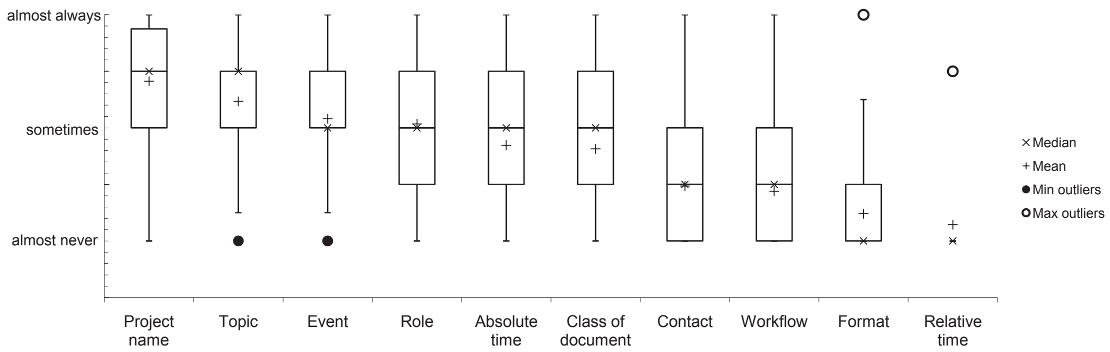

Figure 4.1: Responses to the question: To name a folder, how often do you use this category of labels?

	Almost never	Rarely	Sometimes	Often	Almost always
Recent Documents	17.4%	11.1%	27.1%	**28.5%**	16.0%
Files (any date)	**45.8%**	23.9%	18.3%	7.8%	4.2%
Information to be processed	18.2%	11.9%	**29.4%**	26.6%	14.0%
Important information to remember	19.7%	16.2%	22.5%	**31.7%**	9.9%

Table 4.3: Frequency of use of the desktop for storing particular kinds of items.

desktop". Results are presented in table 4.2. Filing is used more often than piling for classifying documents in general. The majority of participants (52.4%) used filing "almost always" and 34.7% "often". Piling is used "sometimes" or "rarely" with an equivalent 28.2% of answers. Nevertheless, 20.4% of participants reported using piling "often". Desktop piling is used a bit more often on average. The majority of participants (53.2%) indeed use this strategy whether "sometimes" (29.4%) or "often" (23.8%).

Naming Strategies for Folders. Using the taxonomy defined by Boardman (2004), we asked participants to rate the frequency with which they use specific categories of labels to name their folders. The distribution of participants' responses are plotted in Figure 4.1. "Project name" and "topic" are the two most used dimensions, a result that is consistent with Boardman's observations. "Format" and "relative time" are "almost never" used by 70%, respectively 80%, of participants. Additional comments provided by participants seem to indicate that they may also use a combination of categories to name folders: "project name or event with date"; "role + absolute time or topic"; "project name, document type specifications and version".

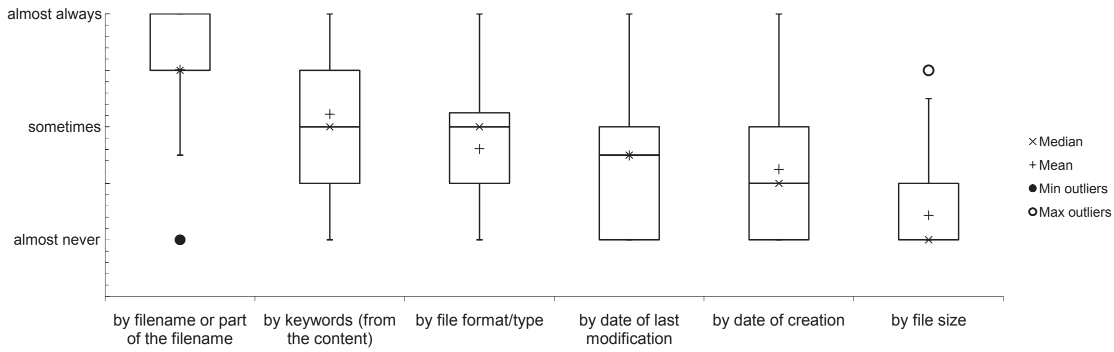

Figure 4.2: Frequency of use of given cues when searching for a file.

Kinds of Information Items Stored on the Desktop. Participants use their desktop to store mainly "recent documents", "information to be processed" and "important information to remember", as shown in Table 4.3. "Files (any date)" are less likely to be found on the desktop. Comments given by the participants also indicate that the desktop is used to store shortcuts to "folders used frequently" or "tools used every day". Additionally, several participants described their desktop as a temporary location for information items being processed.

Strategies for Re-finding Items

When using search to re-find files, participants primarily search by filename or part of the filename (see Figure 4.2). Keywords and format/type are the two most frequently used of the remaining cues for searching. Dates are used less often despite many search engine facilities for querying by absolute or relative dates (e.g. Windows Search). File size is almost never used on average, but some outliers do employ it. Interestingly, 4 comments out of 6 point out that the participant does not use search engines but directly browses to the sought-after information item using the file hierarchy.

Figure 4.3 compares the strategies used by participants to re-find documents based on the age of the document. The questions were of the form: "How often do you find documents last modified x days ago ?". Possible answers were:

- Using direct browsing (you know exactly where the document is)
- Using step-by-step browsing (you don't know exactly where it is but browse the file hierarchy iteratively)
- Using a search engine

Most people use direct browsing "almost always" for 2-days old documents. As time passes, they tend to use more and more step-by-step browsing and search engines. For documents modified

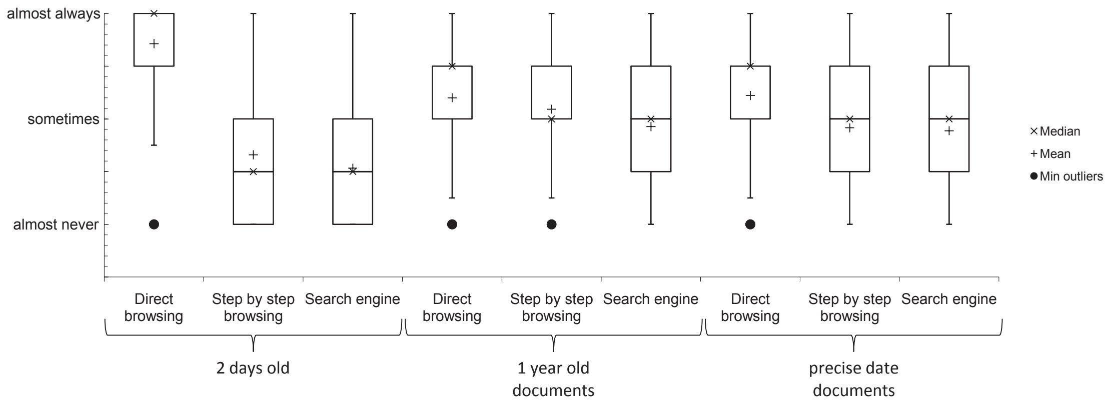

Figure 4.3: Frequency of use of given methods for refinding information items.

at a precise date, they still say to use direct browsing more often than step-by-step browsing or search engines. Additionally, for 2-days old documents, two other answers were available:

- Using "Start > Recent Documents"
- Using the "Recent Documents" option in applications

Those two answers received a median value of 2.15 (stdev 1.21), respectively 2.64 (stdev 1.29) (i.e. between "rarely" — 2 — and "sometimes" — 3).

The last question of this part of the study asked people to rate how often they use temporal cues in search queries. The mean value of use for relative dates (e.g. "yesterday") is rather low (1.84, stdev 1.07), as is the mean value for absolute dates (1.96, stdev 1.14). People therefore seem to use those features very rarely. Interestingly, 4 participants commented that they rather use time intervals in search queries.

Use of Dedicated Tools and Features

A majority of participants in our study used a calendar software (77.5%). However, a large proportion of them (38.8%) did not know that they could attach information items (files, emails) to calendar items. Amongst the participants that knew the feature, the majority found it rather not useful (33.8%), whereas only 27.4% found it rather useful. Commenting this particular question, people mainly complained that it makes information fragmentation even worse ("Too many repositories with files", "often too much information in the wrong place", "Using more than 1 calendar - prefer referring to files where stored").

4.2.4 Discussion

In summary, the results of our questionnaire are mostly consistent with previous research, while adding reported frequencies in the use of different means for classifying and re-finding information. Classification-wise, people still use more filing than piling strategies, which confirms observations made in particular by Boardman and Sasse (2004). Names of folders rarely refer to format and time, but often relate to project names, as already noticed by Bergman et al. (2008b), or topics. The desktop often endorses the role of a reminder of documents to process, as was identified by Malone (1983). Re-finding-wise, users prefer browsing over searching, which was already noted in several research works (Barreau and Nardi, 1995; Marchionini, 1997; Boardman and Sasse, 2004; Bergman et al., 2008c). People seem to rely more on search and step-by-step browsing when looking for older information. When searching, people report to use filenames most often, which is a result not outlined by previous studies (Dumais et al., 2003). The sparse use of temporal cues in search queries matches Dumais et al.'s observations. However, comments point out that time intervals seem to be a potentially interesting cue for re-finding.

On a more general note, this study only reports what people say about their strategies, which is a similar approach to the one endorsed by Barreau (1995) or Bergman et al. (2008b). Instead, other studies like those by Dumais et al. (2003) and Boardman and Sasse (2004) report on real use of PIM tools to devise classification and re-finding strategies. To complete the understanding of PIM strategies, I also explored people's PIM strategies from an ethnographic point of view, with an interest on what people really do rather than what they say. This is presented in the following section.

4.3 Study of PIM Strategy Changes

4.3.1 Presentation of the Study

There have been a number of empirical studies of PIM behaviour (see Section 3.2). Most of them are naturalistic in the tradition of ethnographic field work, some of them are controlled experiments. Ethnographic studies are qualitative methods used in social sciences to assess how people act in their real-life, uncontrolled context. Several tools are used by ethnographers, in particular questionnaires, interviews, and observation. It has been argued by Forsythe (1999) and Fitzpatrick (2009) that observation is necessary, as much of the insight of ethnography lies in the "invisible", the things that people take for granted and tend to apply unconsciously, without even feeling the need to report it. In PIM, there are a few research works that involve observation of users (Kwasnik, 1989b; Barreau, 1995; Hartel, 2006). Naturalistic studies are considered a complementary approach to quantitative studies in PIM, because they can provide a more holistic and contextual understanding of the actual practices and needs of people. Their final goal is to apprehend how they make sense of their personal information and what are their actual management strategies (Naumer and Fisher, 2007). Naturalistic studies rely on fieldwork and contextual observations, which leads to the formulation of hypothesis and knowledge modelling. As computer scientists, I of course have a lesser experience of the fieldwork than the experience a trained ethnographer would have. However, I am a domain expert for digital PIM, which has the advantage of providing another insight into the matter. The following sections explain the goals, context and results of the study I conducted and provides a discussion of them.

4.3.2 Goals

My initial goal was to get an overview of the strategies people employ to manage their personal information in the context of their work. I planned to cover more precisely the areas of task and calendar management as well as email and files management, but remained opened to other opportunities that I could have not envisioned, according to an holistic approach. During the fieldwork, I consolidated my understanding of PIM strategies. Among the many promising questions that arose, one caught my attention: how do people analyse their own PIM practices and why do they change their strategies? This is a topic not yet well covered in the literature and I saw my fieldwork as an opportunity to gain insight into it. I therefore reframed my goals in order to elicit the means people employ to analyse their own PIM strategies and the motivations behind their strategy changes.

4.3.3 Context

Twelve participants (coded P1 to P12) were recruited by convenience among the circle of friends and acquaintances of the author, as is often the case with naturalistic studies. Moreover, as this study deals with personal matters, I believe this also helps the participants to know and trust the interviewer in order to feel comfortable while discussing their PIM strategies. The participants were aged 26 to 49 (average 36), 7 of them were female, 5 male. Half of the participants had a computer science background, whereas the other half came from diverse fields like psychology, humanities or business. Their professions ranged from full-time university professor to secretary, including research assistants and high-school teachers. Most of them had used a computer at work for more than 10 years, only two of them used one for less than 5 years.

The time and resources at disposal for the fieldwork were limited. In particular, I could not perform long-term observations of user behaviour, but would focus on a snapshot of their practice, combined with a semi-structured interview. I decided not to use recording devices that would be intrusive for the participants and require much time to transcribe and analyse. Instead, I used a notebook for taking live notes during the interviews. The raw field notes were taken on the right page while left pages were left empty. They would be used later to identify possible themes, add comments, notes and sketches as a first effort of synthesis, as advised by Fitzpatrick (2009). While note taking, I took care to write down not only what the participant was saying, but also relevant details I could notice, or what the participant was actually doing while talking.

The interviews took place at each participant's workplace, behind the computer they usually work with. The interviews lasted between one and two hours on average, depending on the time the participant could invest. I used a "backpocket guide" to help me frame the interview. But, as it is the rule with naturalistic studies, open discussion was encouraged in order to discover unsuspected areas of interest. The backpocket guide was inspired from the questions set proposed by Barreau (2008) and included specific directions towards the understanding of strategy changes:

- What do you do on a typical work day?
- What information do you have and use in your personal workspace?
- How do you organize information in your personal workspace?
- How do you typically go about finding information when you need it?
- What features do you wish were available for organizing and retrieving information from your workspace that you do not already have?
- What are the main problems you have encountered managing your personal information?
- How have your organizing and retrieving strategies evolved over time?
- What motivates you to change a strategy?

The raw notes were then reworked in order to build an organized collection and interpretation of the participants' responses. These were published internally (Evéquoz, 2009). In the following I present selected results and discuss them.

4.3.4 Results

All 12 participants were able to talk about their PIM strategies and to explain at least one PIM strategy change. Moreover, some participants were able to say that the strategy they disclosed orally was ideal, and that exceptions happen in practice. Exceptions can be due to "lack of time" (P5) or "laziness" (P1) for example. Others tried to formulate a strategy, but evidence shows that it was inconsistent with their actual practice. Indeed, the interview time was an opportunity for the participants to reflect on their PIM practices, which they had seldom done before. Some also became aware of strategy flaws or forgotten items that needed to be handled (P1-P3-P9-P10). Moreover, after the interview, one participant (P4) wanted to give precisions on her strategies and notify the interviewer that she noticed something new in her practice. This seems to emphasize the fact that the interview time made her think about her strategy and analyse it more precisely.

However, strategy changes are not always conscious and seem to "happen almost naturally" (P3). After people have explained their strategy, I could sometimes notice evidence that they used another strategy in the past. For example, P1 explained that she manages her tasks with her calendar. While browsing through her emails, I noticed a folder called *Todo*. She could explain she planned to use it for managing tasks related to emails but forgot it in the end. Similar situations happened with other participants (P3-P10). Therefore, it seems that some people do not have a good mental model of their personal space of information (PSI) as a whole, and that they happen to forget portions of their personal collections and related strategies.

Other interesting results of the study show that concurrent management strategies can coexist. In particular, archived items often adhere to a classifying scheme which is not used any more. For example, P12 mail archive contains past emails classified in directories. However, his current management strategy is that of a folderless cleaner in the taxonomy introduced by Bälter (1997).

The different strategies exhibited by participants during the interviews let us define a model for characterizing PIM strategy changes. I present this model in the following paragraph and classify some strategy changes that I observed during the study according to it.

Model of Strategy Changes

From my empirical observations of PIM strategy changes, I attempt to formulate a simple model for characterizing changes. In this model, a PIM strategy change consist in a general *scope*, a *cause* and an *action*:

Scope. The scope of a strategy change is either *specific* to a PIM tool or *longitudinal*, across tools but related to one meaningful activity. For example, using more folders to classify emails is a strategy change specific to the email client. Centralizing the task management that was scattered across many tools to a single one is a decision taken from a longitudinal scope.

Cause. The cause of a strategy change is either *external* or *intentional*. External causes involve conformance to organizational requirements, for example rules to use collaboratively a filesystem or calendar. Intentional causes involve the user wanting to make their management better without external constraints, for example because they happen to have trouble re-finding emails or they take too much time putting files into folders.

Action. The action may be a decision to *simplify*, or *complete* the existing strategy. Depending on the scope, simplify and complete take different meanings. In a tool-specific scope, simplify means using *less categorization* (e.g. less folders, less labels, etc.), while complete means using *more categorization*. In a longitudinal scope, simplify maps to *centralize* (e.g. use one single tool for managing all tasks) and complete to *fragment* (e.g. keep copies of email attachments both in the mail archive and the filesystem). Note that I consider a change of strategy towards

	Scope		Cause		Action	
Strategy change	Specific	Longitudinal	External	Intentional	Simplify	Complete
Stop filing emails into folders	×			×	×	
Stop deleting all emails, keep them in the inbox	×			×		×
Use naming conventions for files	×			×		×
Abandon the use of todo folder and do task management through emails		×		×	×	
Keep several copies of documents in different places (including attachments)		×	×			×
Stop using a folder for "source" documents	×			×	×	
Use a common filing strategy for documents on a collaborative platform	×		×			×
Stop tagging items in del.icio.us	×			×	×	
Add folders and subfolders in a systematic way for all projects		×		×		×
Use a single tool for task management		×		×	×	
Centralize collaborative tasks management		×	×		×	
Use two calendars for private and professional matters	×			×		×
Use less folders		×		×	×	
Centralize all tasks management in a specific tool		×		×	×	
Use a common calendar	×		×		×	

Table 4.4: Types of some observed strategy changes

deletion of items a simplifying action, since it is a kind of simplification or centralization. As well, a strategy which is abandoned altogether is considered being simplified to the extreme.

Table 4.4 details a selected sample of strategy changes that were disclosed or exhibited by participants during interviews. Some changes were mentioned by participants directly, others were asked by the interviewer when he noticed something unexpected according to a strategy disclosed by the participant.

Motivations Behind PIM Strategy Changes

Motivations to Use a New Strategy. People adapt their strategies. But the concrete reasons that make them change their strategies are many. For example, they may have discovered a new PIM tool that they want to give a try (P8). They have had a hard time re-asking for emails they had deleted and do not want to experience that situation again (P2). Their strategy for tasks management failed as they forgot appointments and tasks (P3-P8). They feel overwhelmed by tasks, because they use the same calendar for private and professional tasks (P10). They face collaboration problems for task management (P9). They have new job responsibilities (P8). They must respond to organizational requirements (P4-P6-P11-P12). They have learned a new feature in a PIM program that they never used before (P6). They find that some features are missing in their current strategy (P12). They cannot cope any more with fragmentation (P9). Or they simply want to try a new idea inspired by a colleague's practice (P8).

Motivations to Retain Existing Strategy. Moreover, several people admitted being not satisfied with their strategy and consider it a makeshift solution, waiting for a better way of managing their information. Thus, they invoke reasons for not changing their strategy. Some participants explained that changing would take too much time: "I prefer losing some time searching for a document than classifying everything perfectly" (P5). Other people also mentioned a lack of education for justifying their strategy. For example, P10 feels overloaded by email, but she does not use any mail client, because "no one ever told [her] how to use them". The standard email web interface she uses is poor, slow and does not provide search tools. Thus it takes her much time and energy to manage emails. When asked to describe the strategies she was using for managing her personal information in her workspace, P6 expressed her disappointment in these terms: "I thought you were going to show me how to do it". She further described herself as an "old-school secretary", educated to the use of real papers and binders and with a very basic knowledge of computer-supported information management tools, which somewhat frustrates her. Other participants also asked the interviewer for advice in their PIM or complained that they never have received lessons or hints for doing PIM.

Collateral Observations

During the study, some side problems were raised by participants. It seems that they contributed to the decision of changing or retaining their strategy.

- *Classification*: One participant clearly said that keeping things organised takes too much time (P5). On the other hand, several people said that they are afraid of losing things which encourages them to classify information or even duplicate it (P4-P5). On a related note, one participant said that even if time-consuming, he considers classification calming and reassuring (P11).

- *Information fragmentation*: At least two participants (P5-P10) complained about the fragmented nature of their personal information, either because it stands on various materials (disks, servers, USB sticks, etc.) or because it is distributed over various applications (emails, calendar, documents, etc.).

- *Collaborative PIM*: Although shared information is not personal, various participants considered that it enters in the personal information sphere since information is sometimes shared among small groups of people. Various people said that they have problems to find things in collaborative settings (P4-P5-P6-P8-P9). Strangely, even when they consider that they have messy shared directories, they admit they do not have the will to change their organization because they cannot take decisions about shared data on their own (P4-P6).

In addition, participants explicitly mentioned the following needs:

- *Browsing using memory-supporting cues*: Various people use facets implicitly to organize and re-find information. For example, P2, P4 and P5 classify some documents and emails in directories having the name of people from their social network. Others classify and/or access documents or emails by temporal indexes (P1-P3-P4-P8). Some participants told me that they would like to have their communication items (emails, chats) automatically clustered by people and further to have their broader PI automatically linked to the people concerned. Still, related to information linking, several persons expressed the need for tools to help them organize specific types of information. For instance, a system administrator (P7) said she would like to be able to keep track of software and hardware problems-solutions pairs.

- *Ubiquity versus fragmentation*: Although several people expressed their problems with the fragmented nature of information, on different disks or applications, as discussed in the previous section, they reversely expressed their desire to have information available and accessible from everywhere, for instance to have access to appointments online or to have information available on the field (in a mobile task) (P3-P5-P9-P12).

- *Usage history*: Some participants expressed their will to visualize the usage history of their PI items and collections. For instance, P2 said that he would like to see the usage frequency and other attributes of folders without having to open them. P5 explicitly said that he would appreciate to have the operating system taking care of the versioning of documents in order to be able to jump back to a previous version, without having to manually use a naming convention. Finally, some people said they would like to have the history of the tasks they performed, to observe how they invested time on specific tasks or projects (P2-P11, indeed P11 did it by hand at the time of the interview).

- *Education and support*: Various people (P6-P10) said that they would like to be more educated on the way to use the existing tools and would like to have support.

4.3.5 Related work

Section 3.2.4 shortly presented previous studies where strategy changes were reported. Here, we discuss them in greater detail. In addition to the two anecdotal strategy changes reported in Boardman and Sasse (2004) and already presented in Section 3.2.4, Boardman (2004) reports further strategy changes occurring during a later study. Three participants reported an increase in their organizing tendency (two in files and one in emails), four participants reported a decrease (three in bookmarks and one in files). However, the data monitored during the study did not reveal those strategy changes. Boardman therefore claims that such small-scale strategy changes do not modify the general orientation of users. Therefore they can only be identified by the users themselves. As for the reasons of strategy changes, Boardman only states that the study itself had a "self-auditing" influence on the participants who decided to adapt their strategies. On the other hand, Bälter (1997) did a survey of email use in a company where employees used at least one of three available email clients. He found out several factors possibly influencing the email management strategy (which can be one of *frequent filer*, *spring-cleaner*, *folderless cleaner* or *folderless spring-cleaner*). He shows that the strategy is related to the experience with computers and the amount of data to manage, and that the cleaning and filing habits are related to the email tool. He also points out that the work tasks and position seem to have no influence on the choice of a strategy. He identifies two main reasons for strategy changes. A change may whether be determined by the user or be a reaction to outer pressure. He also identifies two broad kinds of changes: towards less structure or towards more structure. Furthermore he proposes a model of possible strategy changes between the four aforementioned strategies for emails. He nevertheless admits that his model needs to be confirmed by empirical evidence. Interestingly, the two reasons identified by Bälter resemble my model's internal and external cause and the two actual changes pointed out by both Boardman and Bälter closely map the actions of "simplifying" or "completing" the existing strategy. However, neither Boardman nor Bälter explicitly characterized the scope of a strategy change like I did in my model.

4.3.6 Discussion

As Barreau (2008) stated recently, *"people manage their work in unique and creative ways"*. I found this assertion particularly true while performing that naturalistic study of PIM. Each participant seemed to have developed a different way to organize their information, in some way expressing not only their roles but also their personality. Despite this great variety, I found some common patterns related to changes of PIM strategies. Each change has a cause: external, or intentional. Each change has a scope: tool-specific or longitudinal (cross-tools). And finally, each change implies an action over the strategy: either to make it more simple, or to enrich it. Of course, stronger empirical evidence (e.g. alternate studies) would be necessary to validate my model attempt.

What I found particularly singular in my study is that people are not aware of their strategies or often not applying it as they would like. Various people even asked me to help them better organize or train them to do better. But what is the best way, since each way seems suited to a particular person, job or task? This is a challenging question yet unanswered. This study also suggests that what people say about their PI and strategies often differs from the reality of their data and practice.

4.4 Implications for the System Design

Some outputs of the studies presented in this chapter provide guidelines or clues that guide or confirm the design choices for the faceted personal information browser introduced in Chapter 5. The strongest evidence from these studies is that browsing is used more often than searching, which confirms the potential benefit of a browse-based re-finding interface. Alternatively, empirical evidence show that certain contextual cues may be more promising than others for re-finding information:

- File format/type is used relatively often on average and could be integrated as a facet.
- Dates are rarely used for re-finding information, but comments point out that time intervals could be more interesting.
- Textual searches using keywords contained in the sought-after document or parts of its filename seem to be used often.
- On the contrary, file sizes are almost never used as re-finding means and may not be a good choice of facet.

Lastly, the study of PIM strategies revealed that people may have a poor overview of their personal collections and of the strategies they actually use to manage them. Facets and information visualization techniques are known to provide support for gaining insight into datasets (see Chapter 3) which is another argument in favour of their use.

4.5 Chapter Summary

This chapter has presented preliminary user surveys that were conducted with three main goals in mind, which were also three objectives of this thesis: (1) identify contextual cues of personal information which support re-finding, (2) identify a user interface paradigm suited to support the re-finding task in PIM and which benefits from the cues, and (3) develop an increased understanding of PIM strategy changes.

The contributions of the first two studies presented in Sections 4.1 and 4.2 are twofold. First, they complement existing studies of PIM classification and re-finding behaviours by providing the reported frequency of the use of certain means of classifying and re-finding personal information amongst a panel of users. Doing this, they give further arguments in favour of certain contextual cues to support re-finding. Second, they provide evidence that a faceted approach to personal information browsing is promising. The naturalistic study presented in Section 4.3 highlights an under-studied domain of PIM: strategy changes. Empirical evidence collected during this study allows the definition of a model for strategy changes which complement previous attempts by fellow researchers. Moreover, side observations collected in the course of this study suggest that broadly used PIM tools fail to provide a satisfying overview of the personal collections they help to manage. Faceted navigation may be a promising approach to help people gain more insight into their own personal collections. The next chapter is devoted to the presentation of the faceted navigation tool that was developed in accordance with the findings of this chapter and Chapter 3. Chapter 6 will then present evaluations conducted with the help of this tool that assesses the validity of the approach and studies the use of facets to re-find information items and gain insight into one's own personal information space.

But do you know that, although I have kept the diary for months past, it never once struck me how I was going to find any particular part of it in case I wanted to look it up?

Bram Stoker, *Dracula*

J'ai tendu des cordes de clocher à clocher; des guirlandes de fenêtre à fenêtre; des chaînes d'or d'étoile à étoile et je danse.

Arthur Rimbaud, *Illuminations*

All we have to decide is what to do with the time that is given to us.

J.R.R. Tolkien, *The Fellowship of the Ring* (Gandalf)

5

A Faceted Browser For PIM

Previous research in personal information management (PIM) exposed in Chapter 3, as well as the preliminary studies presented in Chapter 4, identified the potential of faceted navigation and contextual cues to support re-finding and self-analysis in PIM. This chapter presents the experimental system that was designed and developed to tackle the third objective of this thesis (see Section 1.4): design, implement and evaluate a PIM re-finding system based on the cues and interface paradigm identified in previous chapters. This system consists in a faceted browser of personal information using contextual cues, that exploits the findings of Chapters 3 and 4. The system will be used in Chapter 6 to understand better how people use facets of personal information and to assess the validity of a faceted approach to personal information management.

This chapter is divided into two main parts. Section 5.1 motivates the design of the system and replaces it in context. Section 5.2 describes the system extensively and gives relevant implementation details.

5.1 System Design

5.1.1 Aims and Guidelines of the Experimental System

The system introduced in this chapter is aimed at experimenting faceted navigation in a realistic PIM context and at assessing its usefulness to support the re-finding and self-analysis tasks of PIM. As the developed tool is meant to be used in an experimental setting, a main goal was that it be reliable and realistic. It was designed with the following guidelines in mind:

Unified access. Unify access to (at least) three main types of personal information: textual files, e-mails, calendar entries. One of the main issues with PIM is that personal information is inherently fragmented. The documents needed to perform a task may be scattered between folders or network repositories, some information may have been received as an e-mail attachment and details might be found in a calendar entry only. A realistic re-finding tool must be able to provide a transparent access to the three main kinds of textual personal information: files, e-mails and calendar events.

Combined search/browsing. Combine search and browsing within a single interface. Search interfaces provide speed and require a low cognitive load at the cost of indexing difficulty and strong dependency on the keywords used in queries. Browsing might cause cognitive overload, but it is not tied to particular keywords, often helps to form a mental structure of information and may lead to information finding by serendipity. Conforming to the aims of faceted navigation, we consider that search and browse should be seen as complementary. Moreover, a PIM interface should take care of the presentation clarity so that the user does not feel lost in his personal space of information (PSI). Specifically, the interface should emphasize what queries, filters or navigation paths are active to generate the current view of personal information. Faceted navigation explicitly combines search and browse in a single interface.

Promote recognition. Whenever possible, inspire recognition of information rather than requiring the user to recall. By presenting the user with a list of choices, the tool should decrease the cognitive load that is required to recall the exact name of a person, for example. Recognizing and picking an item out of a list of explicit choices is faster and cognitively easier than expressing it *ab nihilo*. Again, the faceted navigation paradigm explicitly supports this.

Provide contextual information. Expose contextual cues of personal information to the user in a way that he can benefit from his memory. This aspect is also closely tight to faceted navigation.

Do not rely on user annotation. As users are unlikely to maintain annotations consistently in the long-term on their personal collections, we do not want our system to be dependent on user annotations.

The faceted navigation paradigm explicitly follows most of those guidelines. Though, there are design challenges related to it. First, determining the most appropriate facets of PI to support the re-finding and self-analysis tasks is a double-edged challenge. On one side, those facets should be meaningful for the user and map their memory representation of the sought-after information and the mental model of their PSI (user-centered aspect). On the other side, faceted metadata should be obtained by automatic means, without requiring a manual annotation on the part of the user — as is often the case with faceted classifications of generic datasets (Yee et al., 2003) (technology-driven aspect). Another challenge of faceted navigation is to find a meaningful and usable way of representing faceted metadata in the interface, possibly extending the traditional list-based representations of facet categories.

The following paragraphs briefly remind the aims and means of faceted navigation in general and discuss previous approaches of PIM that exploit this paradigm. The contribution of our faceted browser is presented afterwards. The section ends by presenting our actual design choices and methodology.

5.1.2 Reminder of Faceted Navigation

Faceted navigation, shortly introduced in Section 2.2.2, is an interface paradigm for accessing information using contextual cues. According to the seminal work by Yee et al. (2003), the main goal of faceted navigation is to improve over traditional searching by avoiding two negative consequences: empty result sets and feeling of being lost. To achieve this, the underlying data which is to be navigated through should be classified using faceted classification and the navigation interface should respect the following main guidelines (Yee et al., 2003):

Seamless integration of search and browse. The interface should allow keyword searching as well as browsing using pre-assigned metadata terms (facet categories). Both means may be used independently or in combination.

Dynamic queries. Each selection in a facet should generate a result set immediately, without a need to click on a separate "start search" button (Shneiderman, 1994).

Query preview. The number of results to expect from a selection should be shown alongside the interface widget triggering the selection.

No empty result set. Selections that would lead to zero results should not be allowed by the interface.

Enforce feeling of control. Several visual measures can help make the user feel in control of the browsing experience. Currently applied filters should be salient and easy to cancel. Each facet should be assigned a particular hue throughout the interface.

5.1.3 Reminder of Previous Faceted Navigation Approaches to PIM

Previous research presented in Chapter 3 highlighted the potential of faceted navigation for (1) exploring large datasets in order to (re-)find items and (2) learn about the datasets. The two main systems from the state of the art that apply faceted navigation to PIM are *FacetMap* (Smith et al., 2006) and *Phlat* (Cutrell et al., 2006b). A closer study of them however reveals that they do not completely implement the faceted navigation paradigm. Indeed, *FacetMap* does not allow to see a result set while navigating. When there are too many results, they are aggregated and the user is required to continue navigating using facets until there remains only a small enough amount of results for them to be displayed. Additionally, the layout of facets using bubbles seems to be primarily motivated by the fact that it is visually appealing, not because it is suited to represent the facet categories. *Phlat* distinguishes between textual search (*"query area"*) and the filters (*"filter area"*), which is consistent with the faceted navigation paradigm, as textual search is not a real facet with mutually exclusive categories. However, in *Phlat*, filters are not visible at all times which may hinder the navigation experience and increase the feeling of being lost. Moreover, the visual representation of facets based on lists or tree widgets may not be optimal for certain facets, like the social and temporal facets.

5.1.4 Contribution and Novelty of Our System

Our system is a contribution to existing faceted navigation approaches to PIM, as it:

- Uses facets of personal information that have been chosen based on studies or acknowledged empirical results (user-centered approach).
- Uses an automatic faceted classification without the need for user annotation (technology-driven approach).
- Improves over list-based facet categories for certain facets which may be better represented using information visualization-inspired techniques.
- Is developed towards an evaluation.

5.1.5 Choosing the Right Facets of Personal Information

Elsweiler (2007) summarized previous works in psychological research and confirms that contextual cues of information help recollecting (see Section 2.1.6). However, these research works do not tell which particular cues are better suited to trigger recollection. The surveys of user behaviour presented in Chapter 4 (both the questionnaire and user observation) have shed light on some contextual cues that people use for re-finding tasks:

- File format/type seems to be used often in re-finding tasks.
- Exact dates are rarely used for re-finding information, but time intervals may be.
- Textual searches using keywords contained in the sought-after document or parts of its file-name are used often.
- On the contrary, file sizes are almost never used as re-finding means.

Additionally, other empirical data give directions for the choice of efficient facets. Some of this data comes from questionnaires where people were asked to describe what they remember from information (Ringel et al., 2003). Some other data comes from log analysis of users performing real re-finding tasks (Dumais et al., 2003; Teevan et al., 2004; Jones et al., 2005; Cutrell et al., 2006b). These studies confirm the potential interest of the following facets:

- the type of item (Dumais et al., 2003; Teevan et al., 2004; Cutrell et al., 2006b)
- the people (social network) related to the item (Ringel et al., 2003; Lamming and Flynn, 1994; Whittaker et al., 2004; Cutrell et al., 2006b)
- the (approximate) date of the item (Fertig et al., 1996; Rekimoto, 1999; Ringel et al., 2003; Cutrell et al., 2006b)
- the path of the item in the user-defined hierarchy (Jones et al., 2005; Cutrell et al., 2006b)

The categorisation of an item according to the aforementioned facets has the additional advantage that it can be automated, without the need for the user to manually annotate its PI, which is another constraint that we have set.

Moreover, Capra and Pérez-Quiñones (2005) have shown that users performing re-finding tasks like to be able to use textual search, a conclusion also reached by our studies introduced in Chapter 4. Nevertheless, textual search cannot be considered a facet *stricto sensu*, as the set of all possible search queries is not a set of mutually exclusive categories as required by the facets definition. Despite this fact, we call it "textual facet" in this work for simplifying reasons. Note that this textual facet is slightly different from the textual search fields appearing in traditional faceted navigation systems. In traditional systems indeed, textual search is meant to be used to search amongst existing *categories of facets*. For example, in the work by Yee et al. (2003), a textual search query produces a list of categories of existing facets corresponding to the keywords entered. In our implementation

— as in Cutrell et al. (2006a), textual search does not output matching categories in the existing facets. Instead, it acts as a facet by itself as it filters the current result set. This behaviour is also more natural for users who are accustomed to textual searches which output actual information items and not categories of a faceted classification of their PI.

Our system will therefore enable the following facets of personal information:

File explorer facet (FE). This is the user-defined filesystem hierarchy. The categories of this facet are hierarchical, single-valued and not implicitly ordered (see Section 2.2.2 for a reminder of facets taxonomy).

Textual search facet (TS). This textual search field allows to issue search queries for words contained within a document or its filename. As said above, textual search is not a real facet as it does not comprise categories.

Type facet (T). It represents the type or format of document (e.g. email, MS-Word document, image, PDF). We have defined the categories of this facet as flat, single-valued and not implicitly ordered.

Social network facet (SN). This facet comprises the people relating to information items. The categories are thus flat, multi-valued and not implicitly ordered.

Date facets (D). We define three separate date facets which correspond to the three temporal metadata available in most operating systems, namely: date of creation (DC), date of last modification (DM) and date of last access (DA). The categories for each of these facets are flat, single-valued and implicitly ordered.

The next part of this chapter describes the developed system in greater detail.

5.2 System Presentation

5.2.1 Early Prototypes

Appendix B describes some early software prototypes developed in the context of this thesis but that were abandoned. Architecture and scalability issues prevented us from including them in the final system, which is presented hereafter.

5.2.2 The *Weena* Faceted Browser

Technical Considerations

The system presented in this section has been developed in Java[1] using the Netbeans[2] and Eclipse[3] IDE. Two persons were involved in the development, a M.Sc. student supervised by the author, and the author himself. Various software libraries were used: *Windows Desktop Search* (WDS)[4] as the main indexer and data source, *Jacob*[5] as a bridge between COM and Java, the *prefuse* information visualization toolkit[6], SwingLabs *SwingX*[7] for particular interface widgets, *Swing GUI Builder*[8], Apache *log4j*[9] for logging purposes.

Backend and Constraints

The first step towards developing a faceted navigation interface to PI is to build a faceted classification of it. Indexing the whole PSI to get relevant metadata is clearly out of the scope of this thesis, though. Therefore, I chose to leverage an existing indexing engine. Some previous PIM systems followed a similar approach: *FacetMap* (Smith et al., 2006) works on top of the *MyLifeBits* (Gemmell et al., 2002) data store; *Phlat* (Cutrell et al., 2006b) is on top of *Windows Desktop Search*, that is the successor of *Stuff I've Seen* (Dumais et al., 2003); *Feldspar* uses the facilities offered by *Google Desktop Search* (Chau et al., 2008b). Before the development of the tool, I compared two indexing systems: *Google Desktop Search* (GDS) and *Windows (Desktop) Search* (WDS). Both are reliable and offer a wide set of features. I finally chose to use *Windows Search* for several reasons, mainly because:

- It gathers a greater amount of metadata for each information item than GDS.
- It offers powerful SQL-like query mechanisms, versus HTTP queries for GDS.
- It is embedded in MS *Vista* and later whereas GDS needs to be installed separately and allowed time to index PI which would make the user evaluation more complicated.

[1] `http://www.java.com/`, retrieved July 24th 2010.
[2] `http://www.netbeans.org/`, retrieved July 24th 2010.
[3] `http://www.eclipse.org/`, retrieved July 24th 2010.
[4] `http://www.microsoft.com/windows/products/winfamily/desktopsearch/default.mspx`, retrieved July 24th 2010.
[5] `http://sourceforge.net/projects/jacob-project/`, retrieved July 24th 2010.
[6] `http://prefuse.org/`, retrieved July 24th 2010.
[7] `https://swingx.dev.java.net/`, retrieved July 24th 2010.
[8] `http://netbeans.org/features/java/swing.html`, retrieved July 24th 2010.
[9] `http://logging.apache.org/log4j/`, retrieved July 24th 2010.

WDS offers enough metadata to build the faceted classification[10]. Though, as described in detail in Section 5.2.3, I complemented the building of the social network facet by using another source of metadata.

A critical aspect of the tool is that it needs to be efficient enough to be used in a user evaluation. It was therefore developed incrementally, assuring that the performances were sufficient throughout the process. This forced us to put some restrictions on the quantity of data that could be made available in our faceted browser. In particular, we only allowed access to the data that was last modified during the last year. We thought it was a reasonable limitation since the study by Dumais et al. (2003) reports that 89.4% of re-finding tasks concern information created during the last year. Moreover, we limited the amount of contacts in the social network facet to 150 at most. Again, this is motivated by observations by Whittaker et al. (2004) that *"most people actively communicate with fewer than 150 contacts"*.

General Software Architecture

The system uses WDS as its data source. Upon launch, a query is issued to WDS to retrieve all information items last modified during the last year. Every information item comes with its relevant faceted classification obtained from the WDS index and additional metadata needed by the system (e.g. its name, its URL). All the information items are kept in a backing table loaded in memory. Table 5.1 depicts a slightly simplified example of what this backing table may contain. Each information item occupies one row. Facets appear in columns (except the special columns containing the name and URL of the item). The values in a column are the categories of the corresponding information item for the facet represented by the column.

Each selection of categories in a particular facet acts as a filter on the corresponding column in the backing table. Applying the filter actually generates a cascaded table which has the same structure as the original table but contains only the matching information items. This cascaded table is used as the new backing table. All facets in the interface update their display on this cascaded table. Filters may be implemented very easily with predicates which return *true* or *false* for each information item depending on the category it has for the corresponding facet.

[10] As of March 2010, each file in the Windows operating system can be assigned 540 different properties organized in 28 categories. Certain properties are common for all files, certain are suited to particular formats (e.g. music, pictures, textual documents). All those properties can be queried from WDS. The list of properties is available at `http://msdn.microsoft.com/en-us/library/dd561977(VS.85).aspx`.

SYSTEM.ITEMNAMEDISPLAY	SYSTEM.ITEMURL	SYSTEM.ITEMTYPE	SYSTEM.ITEMPARTICIPANTS	SYSTEM.DATEMODIFIED	SYSTEM.DATEACCESSED	SYSTEM.DATECREATED
documents and settings	file:c:/…	Directory	null	21.11.2007 08:35:52	28.05.2010 08:29:06	18.09.2007
/	mapi://…	MAPI/Folder	null	23.07.2010 08:31:10		
EightQueens.class	file:c:/…	.class	Florian Evéquoz	04.04.2007 09:04:40	06.01.2009 14:17:47	04.12.2007
RE: TR: [Diuf-g-ingold] Programme so…	mapi://…	MAPI/IPM.Note.Read	florian.evequoz@unifr.ch, anonymized@unifr.ch	16.06.2010 12:59:32		16.06.2010
Master Thesis Presentation "Perva…	mapi://…	MAPI/IPM.Schedule.Meeting	evequozf	12.03.2010 13:46:45		06.03.2010
GDQueryResult.txt	file:c:/…	.txt	null	21/11/2007 14:14	06/01/2009 14:20	04/12/2007 11:14
Fin de service_'EVEQUOZ F…	otfs://…	.eml	null	29/10/2009 15:00	15/06/2007 07:11	12/06/2007 07:45
Fw: rapport sur windows d…	mapi://…	MAPI/IPM.Note.Read	EVEQUOZ Florian, ANO Nymous	10/06/2009 08:04	01/01/1970 01:00	21/08/2008 14:46
Fwd: Re: Fwd: Roadshow…	mapi://…	MAPI/IPM.Note.Read	EVEQUOZ Florian, ANO Nymous	10/06/2009 08:03	01/01/1970 01:00	27/10/2008 08:32
Articles pour état de l'a…	mapi://…	MAPI/IPM.Note	LALANNE Denis, ano.nymous@smth.com	10/06/2009 07:26	01/01/1970 01:00	08/06/2006 15:23
FW: Fwd: Re: Fwd: Roads…	mapi://…	MAPI/IPM.Note.Read	Denis Lalanne', Florian Evéquoz	29/10/2008 10:19	01/01/1970 01:00	29/10/2008 10:19
Survey emails.doc	file:c:/…	.doc	evequozf	02/05/2007 07:43	06/04/2010 07:25	05/12/2007 17:27
MLMI06_MRFEDLRI_final.pdf	file:c:/…	.pdf	evequozf	05/12/2007 14:05	06/01/2009 14:41	05/12/2007 17:31

Table 5.1: Illustrative sample of a backing data table. Each row is an information item along with its faceted metadata. Metadata relating to facets appear in columns. Labels of the columns are the actual properties names in WDS. The SYSTEM.ITEMTYPE column contains the type facet categories. The SYSTEM.ITEMPARTICIPANTS column contains the social facet categories and the three SYSTEM.ITEMDATE* columns contain the temporal facet categories. Some metadata lack, the WDS index being incomplete.

Overall Interface

The overall interface of our faceted browser is shown in Figure 5.1. Each facet is depicted in a panel with a coloured border. In this example, the user first selected 'pdf' and 'java' documents in the type facet. Then he selected the folder '3 - Current Projects' in the file explorer facet. A breadcrumb trail is on top (with label 'Path'). On the right-hand side lays the result panel. The button labeled 'Start' at the bottom of the window was only used during the evaluation to start a task. The following sections detail the different parts of the interface.

Breadcrumb trail

Similar to breadcrumbs often found on the web, *Weena*'s breadcrumb trail shows the path that leads to the current selection. Each time a selection is done in a facet, a box corresponding to this selection is added to the breadcrumbs. This box has the same colour as the facet border, and is labeled according to the facet name. It furthermore indicates in parenthesis how many items remain after this selection. When the mouse hovers over the box, a tooltip gives further details about the selection (e.g. which categories of the facet were actually selected). A cross-shaped button inside each box in the breadcrumbs cancels this selection.

For efficiency reasons, cancelling a selection works in a particular way. Cancelling any selection in the middle of the breadcrumbs is not permitted. The cancelling of a selection actually cancels all selections that have been done after this selection. This behaviour is suggested visually by the interface, as shown in Figure 5.2. This behaviour was chosen because it allows a previously saved version of the backing data table to be restored quickly in the case of the user cancelling a selection.

Result Panel

The result panel depicted in Figure 5.3 shows the information items resulting from the current filtering in a tabular view. Five columns may be displayed: type (icon), name, size, last modification date, last access date, creation date. Results may be sorted by any of those columns. By default, as in most systems, results are sorted by last access date, most recent first. Tooltips additionally show the URL of the item and a preview of it, in the case of an image. A contextual menu offers the following actions on each item: open, open its containing folder, copy the item, copy the path of the item, delete the item. At the start of the application, this panel shows all the information items available, ordered by date of last access. At each selection step, the result set is filtered. Above the result panel, right to the breadcrumbs, the count of currently available items is displayed.

Furthermore, the result panel is paginated. Moving between pages of results is performed using the dedicated buttons at the bottom of the panel. We have done this for several reasons. First, it is easier to scroll down with a middle-sized scrollbar. With several hundreds of results on the same

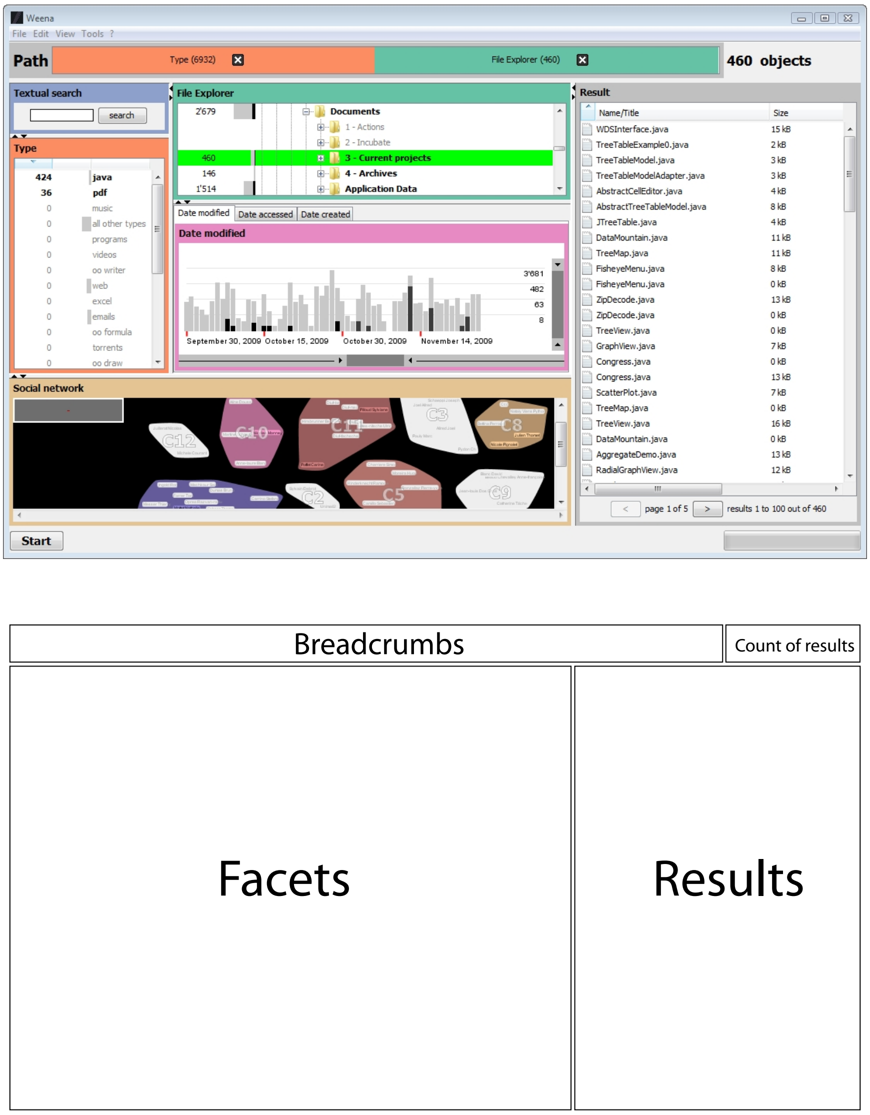

Figure 5.1: The *Weena* faceted browser interface and its different parts.

Figure 5.2: The breadcumb trail and its use to cancel a selection.

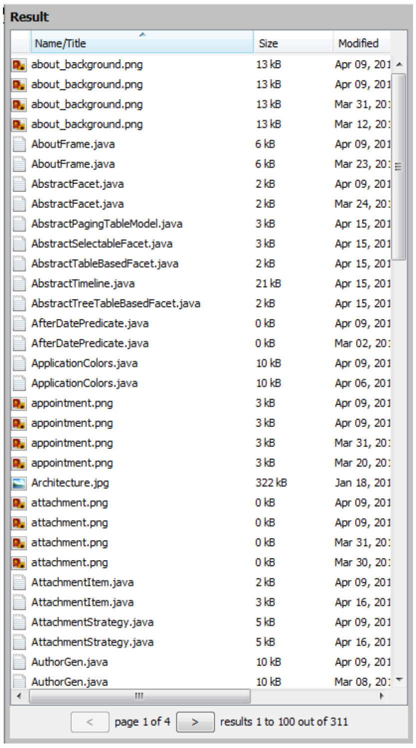

Figure 5.3: The result panel.

page, the scrollbar is so minimized that it makes scrolling challenging. Second, pagination may encourage the user to apply further selections in facets in order to reduce the amount of results shown, and consequently the amount of pages. Anyway, if the user prefers not to have paginated results, the amount of results shown per page can be configured or the pagination may be cancelled altogether.

Facets: Common Features

Weena proposes seven facets. Each facet, except the textual search facet, displays a set of constituting categories, the properties of which have been presented in Section 5.1.5. When the user selects a category, the result set is filtered in order to keep only the information items matching this category. At the same time, the other facets are updated to reflect the current state of the result set: the amount of items related to each category is re-computed and categories which do not appear any more in the result set are disabled. In order to visually differentiate selected categories, enabled categories and disabled categories, colour conventions have been set. Selected categories have a green background. Categories which can be selected have a black foreground and, in the case of textual categories, a bold font. Categories which cannot be selected have a gray foreground and, in the case of textual categories, a plain font. The file explorer, type and dates facets contain bars that

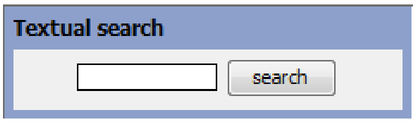

Figure 5.4: The textual search facet.

represent the amount of documents for each category. Again, black is used to show the amount of items that match the currently applied filters, while gray is used to represent the total amount of items when no filter was applied. Figure 5.1 illustrates this. The textual search does not display categories. However, when the user issues a search query, the background of the search field turns green to enforce visual consistency with the rest of the interface.

The following sections further present the different facets.

Textual Search Facet

The textual search facet depicted in Figure 5.4 is not a real facet as the set of categories that constitutes it is not a set of mutually exclusive categories, as was already discussed. The search field accepts single words, phrases put into quotation marks, or several words separated by the boolean operators 'AND' or 'OR'. Issuing a search query actually looks for the given keywords in the information items' names (or titles in the case of emails) and content. It does not look into other properties. This was done to prevent this facet to overlap with other facets and compromise the consistency of the faceted navigation paradigm: each facet is indeed set to be applied on a specific dimension of information items, which is isolated from the other facets.

Type Facet

The type facet shows a list of item types. We defined the type of an item to be based on its file extension. Mapping the extension to a type may be configured in an external XML file as the combination of an inclusion list of extensions (i.e. which file extensions should included as belonging to this type) and an exclusion list of extensions for each category[11]. Tooltips show the inclusion and exclusion list of extensions for each type. Several types may be selected at once in the interface. Note also the bars which show the amount of items that correspond to the currently applied filters.

Figure 5.5: The type facet.

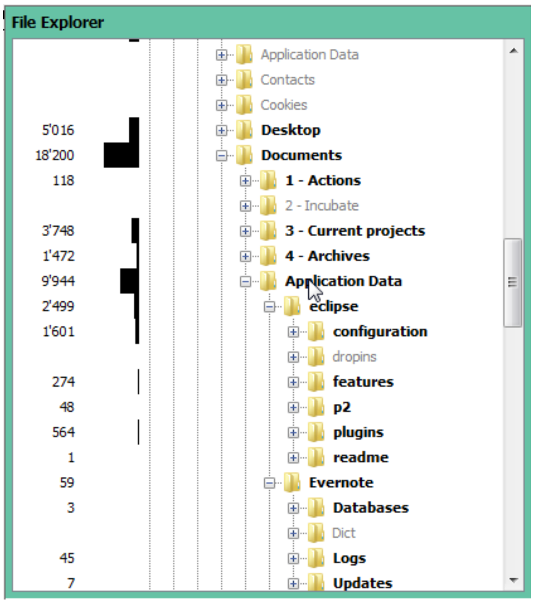

Figure 5.6: The file explorer facet.

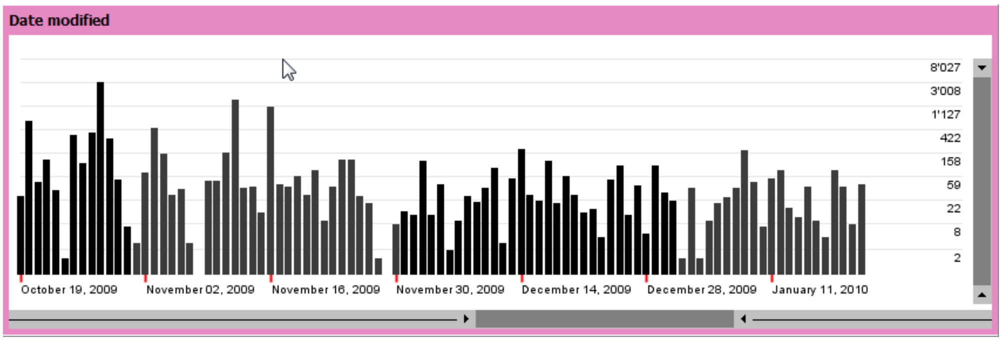

Figure 5.7: The date modified facet.

File Explorer Facet

The file explorer facet depicted in Figure 5.6 shows the user's filesystem hierarchy. We have chosen to keep the visual tree hierarchy which is familiar to the user. Nevertheless, this view has been augmented by the bars on the left which show how many items corresponding to the current filters are contained within each folder. The user may select more than one folder in this facet. Additionally, shortcuts to specific folders in the hierarchy may be added and appear as buttons at the bottom of the panel. These shortcuts can be configured in an external XML file.

Date Facets

Three date facets similar to the one shown in Figure 5.7 are available in *Weena*. By default, only one is visible, the other ones being under tabs (see Figure 5.1), but they can be set to be visible all at the same time, on top of one another. The date facets all have the same layout. Similar to the work by Lee et al. (2009), categories of this facet are solely represented as bars on a timeline. The dates, seen as facet categories, are indeed implicitly ordered and therefore well suited to such a representation. Each bar represents the amount of items for one day. Tooltips give the exact date and amount of items for each bar. Custom scrollbars allow movements along time while changing the portion of the timeline that is visible (zoom in or out), as suggested by the arrows. Two shades of black (one being actually a dark gray) are used to facilitate the differentiation between months. The scale on the horizontal axis is automatically set according to the current zoom level. Selection of time intervals can be achieved by dragging a selection box around the bars. Discontinuous selections can also be made by holding a key.

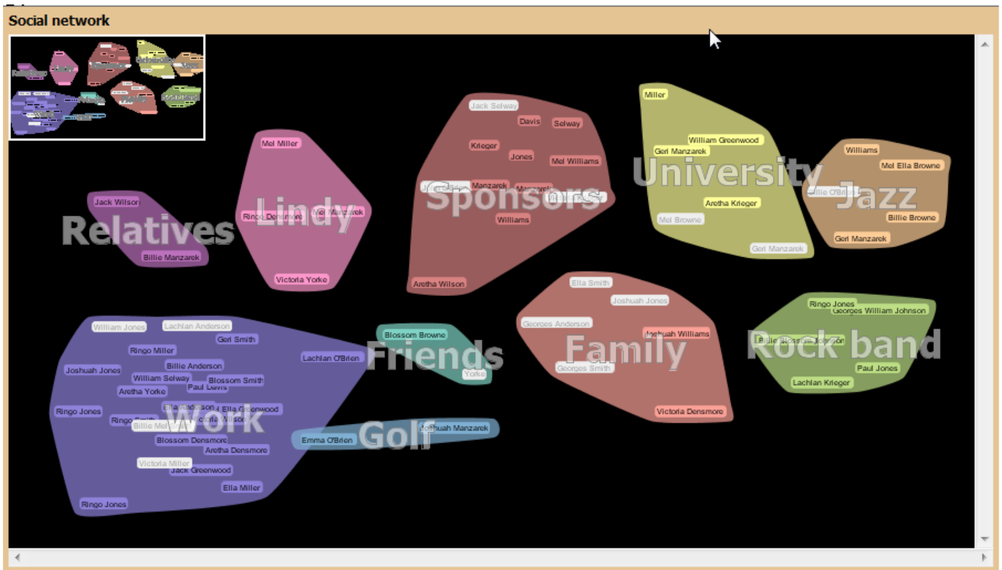

Figure 5.8: The social facet.

Social Facet

Figure 5.8 illustrates the social facet. Contacts of the user are laid out in a two-dimensional plane and grouped by communities represented as coloured bubbles around contacts. The labels of communities appear as overlaid text. Consistent with the rest of the interface, grayed out contacts cannot be selected but black-foreground contacts can. Communities are coloured to facilitate recognition. Selecting contacts highlights them in green. Selecting contacts can be achieved by clicking on a contact directly, or clicking on several contacts while holding a key down, or drawing a rectangular selection around desired contacts. The user can zoom and pan the view. Scrollbars are provided, as well as an overview of the social network in the upper left corner. Hovering over a contact shows links that go from this contact to related contacts. This illustrates who is related to this contact by exposing the relevant parts of the underlying social graph (see Section 5.2.3).

Use Case

Figure 5.9 shows the aspect of *Weena* upon launch. Note the blue color scheme used for the borders of the different facets. Color schemes may be user-defined in menus. In this example, I want to find textual files containing diary studies of participants in a study. I first try to unroll the hierarchy in the file explorer facet but cannot remember which is the containing folder (Figure 5.10). I remember the file is textual, so I select 'word', 'text' and 'pdf' in the type facet (Figure 5.11). The display of the categories in other facets change to reflect this selection. The result panel also shows

[11] Special types, like emails or calendar items, can also be distinguished thanks to their extension in the Windows Properties scheme.

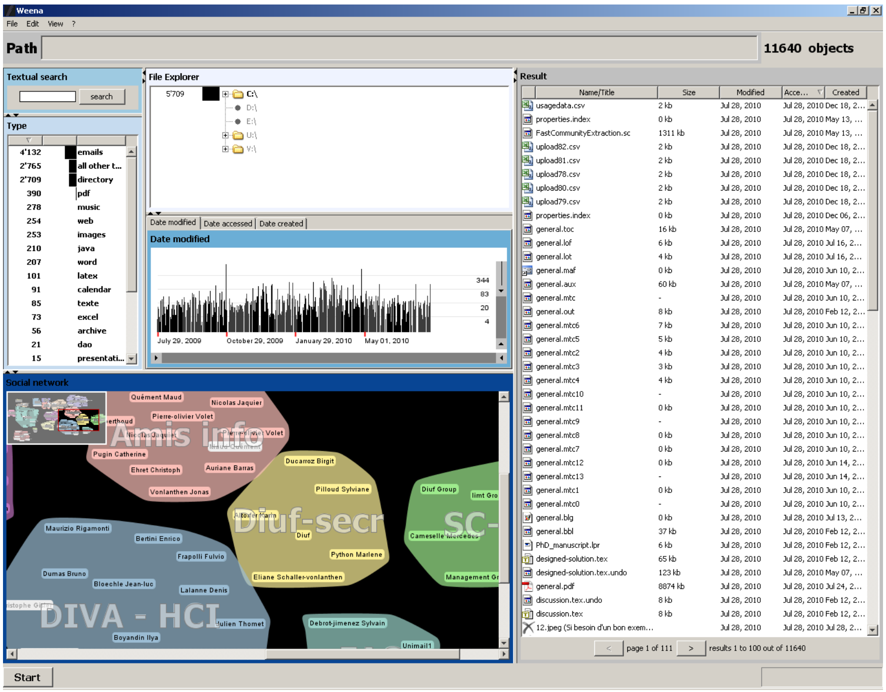

Figure 5.9: Use case, step 1.

now only the information items of the selected types. A box appears in the breadcrumbs indicating that 682 objects (items) match the criteria. As the file is related to the 'DIVA-HCI' community of my social network, I select the members of this community in the social network (Figure 5.12). The other facets' categories are updated and now only 167 objects remain. In the date modified facet, I select the month of March 2010 (Figures 5.13 and 5.14), further filtering the result set. As a last step, I select the 'diaries' folder in the file explorer facet, which outputs the files I was looking for (Figure 5.15). Note that the categories in other facets have also been updated to reflect only the properties of the remaining files.

5.2.3 Social Network: *WotanEye*

This section takes a closer look at the generation of the social facet view which is part of the *Weena* faceted browser presented above.

The social facet is different from all the other facets of *Weena* as the data used to generate it does not come from the WDS index. Indeed, limitations appeared with social data in the WDS index: no more than three participants can be registered with an information item and participants names vary a lot (e.g. the author of this thesis could appear as 'florian.evequoz@unifr.ch' in some emails, 'Florian Evéquoz' in some others, 'evequozf' in some textual documents, 'FE' in others and so on). This social data was not sufficient to be leveraged on automatically. Thus, we exploited the outcomes of a previously developed project, initially centered on the management of personal

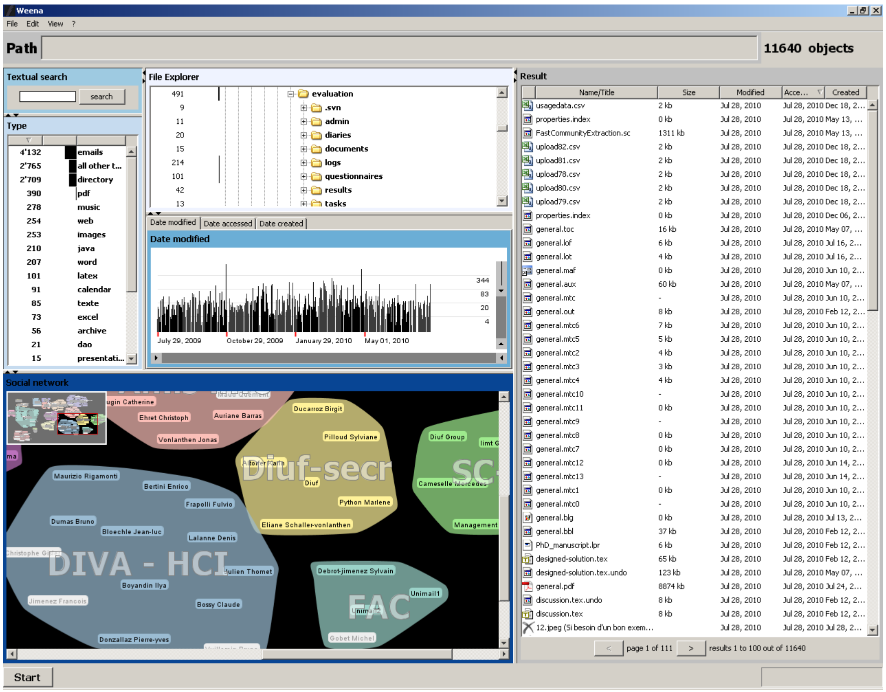

Figure 5.10: Use case, step 2.

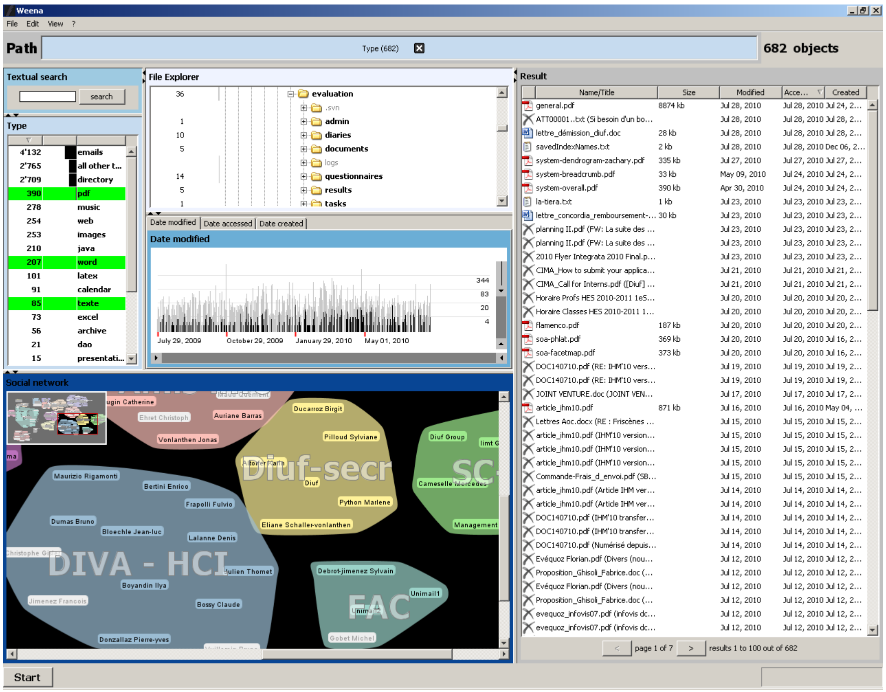

Figure 5.11: Use case, step 3.

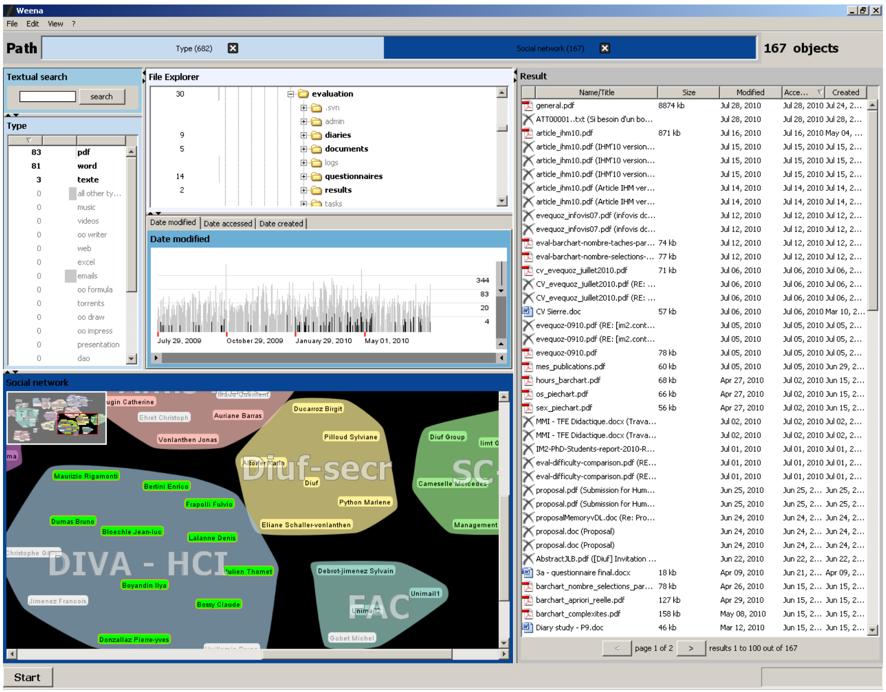

Figure 5.12: Use case, step 4.

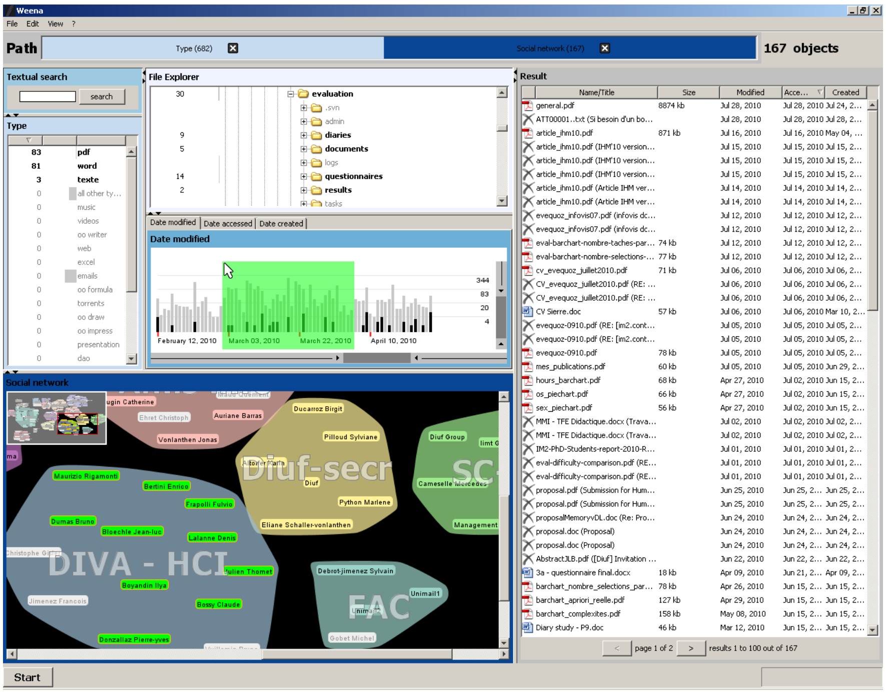

Figure 5.13: Use case, step 5.

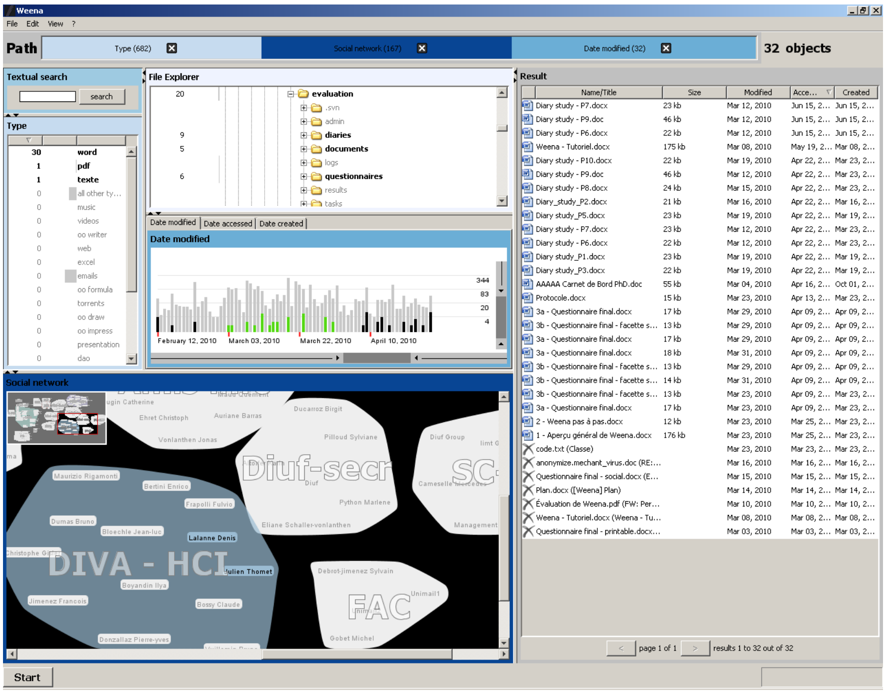

Figure 5.14: Use case, step 6.

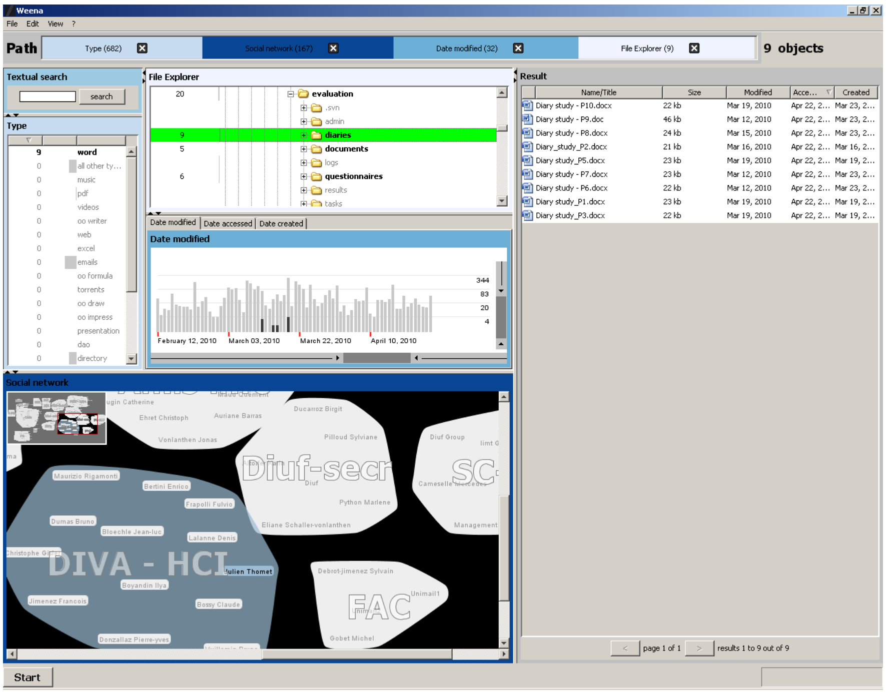

Figure 5.15: Use case, step 7.

emails, which was called *WotanEye*. A subproject of *WotanEye*, coined *MailViz*, involved generating visualizations of email archives, in particular visualizations of the social network extracted from emails. *Weena*'s social facet thus adapts relevant parts of the *WotanEye* code.

The whole process of generating the social facet is covered step by step in this section, starting from the network building, continuing with the community extraction algorithm and the visual representation of the network, introducing briefly the use of *WotanEye*'s interface to activate the social facet in *Weena*, and ending with a description of how contacts in the social network are mapped to personal information items in *Weena*.

Network Building

We extract a personal social network from the main email archive of the user by examining the senders and recipients of emails. Each distinct address in the user's social network is considered to be one contact (i.e. we do not consider that several email addresses may refer to the same person). We define a closeness measure between two contact's addresses a_m and a_n that is proportional to the number of times these addresses appear together in the headers of emails, yet depends on the number of recipients of those common emails. Indeed if Georges and Nick belong to the same group of people and know each other well, there are chances that they often appear together on emails envelopes, including envelopes with few other recipients. Though, if Georges and Jack systematically appear together on envelopes along with many other recipients, chances are that they do not know each other well but are both members of some mailing-list. Our closeness measure takes these observations into account. More precisely, let us define the set $E = \{e_1, e_2, \ldots, e_n\}$ of all emails e_i in the archive. Let $A = \{a_1, a_2, \ldots, a_m\}$ be the set of all email addresses that appear in the archive. We call $envelope(e_i)$ the set of contacts appearing whether in the `From:`, `To:` or `Cc:` headers of an email e_i. We then define the function:

$$on\_envelope(a_i, e_j) = \begin{cases} 1 & \text{if } a_i \in envelope(e_j) \\ 0 & \text{if } a_i \notin envelope(e_j) \end{cases}$$

Using this function, we define the closeness measure $closeness(a_i, a_j)$ between two contacts a_i and a_j as follows:

$$closeness(a_i, a_j) = \sum_{n=1}^{|E|} \frac{on\_envelope(a_i, e_n) \cdot on\_envelope(a_j, e_n)}{|envelope(e_n)|}$$

The weight of the relationship between a_m and a_n thus depends on the number of recipients of an email. This closeness measure has been inspired by previous work on automatic email classification by activities (Dredze et al., 2006; Cselle et al., 2007). With this closeness measure between all pairs of contacts, we can build a social network graph. Vertices of the graph are contacts, and

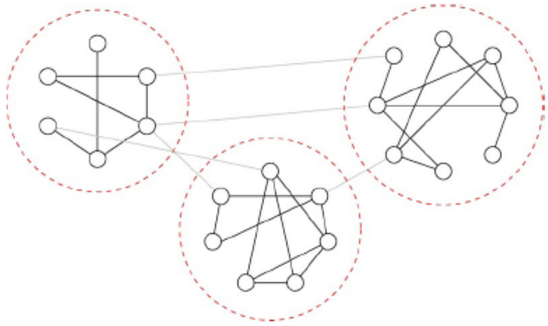

Figure 5.16: Communities are areas of a graph which have many internal edges and few external edges (Source of the picture: Newman and Girvan (2004)).

edges are relationships between contacts, weighted by their closeness. This social network is user-centered as it reflects the view the user has on their contacts. This kind of network is obviously different from an "objective" social network that could be obtained by combining people's own networks together, e.g. by digging popular social network websites.

Communities Extraction

Communities, also called clusters, cohesive groups or modules, are regions of a network within which the vertices are highly connected and between which the connections are sparser as depicted by Figure 5.16. There has been a recent interest on community extraction in networks with applications to different fields like the World Wide Web, biology, scientific citations, actors networks, characters interactions in novels and many others (Newman and Girvan, 2004; Palla et al., 2005; Pujol et al., 2006). Extracting and representing communities indeed facilitate the understanding of networks by revealing their hierarchical nature. We used Newman's fast community extraction algorithm for weighted graphs to extract communities out of the social networks generated by *WotanEye* as was described above[12]. Its main advantages over other algorithms is its speed and its ability to be applied on weighted graphs (whereas other algorithms are suited for unweighted graphs only). Its main drawback is that communities cannot overlap, i.e. a vertex can only belong to one community. This algorithm has been applied in recent research on social networks visualization (Heer and Boyd, 2005; Perer and Shneiderman, 2006), though in its unweighted graph version.

The aforementioned community extraction method actually outputs a full hierarchical structure of the original network. This structure may be represented by a dendrogram (see Figure 5.17). Cutting the dendrogram at a certain depth gives a specific splitting of the graph into communities.

[12]We used the C++ version of this algorithm, introduced by Clauset et al. (2004). The source code has been released at the end of 2008 under `http://www.cs.unm.edu/~aaron/research/fastmodularity.htm`.

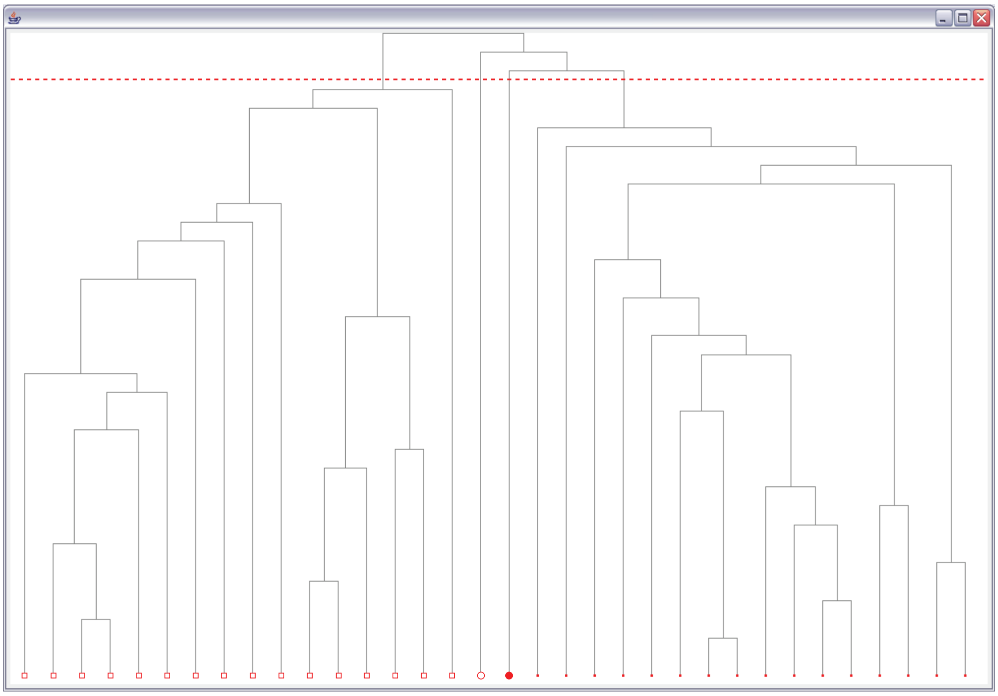

Figure 5.17: Dendrogram depicting the hierarchical structure of a network. Leaves of the tree (red shapes at the bottom) represent actual vertices in the network. Moving up in the tree join vertices that form larger and larger communities until the root is reached. Cutting the dendrogram (red line near the root) at a certain point gives a particular splitting of the network into communities. The communities are depicted here by different shapes of the vertices.

The strength of the community structure obtained by cutting the dendrogram can be measured by its modularity Q (Newman and Girvan, 2004). The highest Q means the "best" splitting of the graph into communities, i.e. the splitting which maximizes edges internal to communities and minimizes edges external to communities (Newman and Girvan, 2004).

During the development of our prototype, we noticed that the "best" community splitting happened to be suboptimal in some cases. Indeed, it often resulted in few communities, each containing many contacts that the user would actually have further split into subcommunities. Therefore, we slightly adapted the community extraction method by applying the algorithm in several recursive passes as follows:

1. Let the original graph be $G = (V, E)$ with a set of vertices $V = \{v_1, v_2, \ldots\}$ and a set of edges $E \subseteq \{\{v_m, v_n\} \mid v_m, v_n \in V\}$.
2. Run the community extraction algorithm and get the best communities splitting $C(G) = \{C_1(G), C_2(G), \ldots\}$ with $C_i(G) \subseteq V$ and $\bigcup_{i=1}^{|C(G)|} C_i(G) = V$.
3. For each community $C_i(G)$
   (a) If $|C_i(G)| < C_{threshold}$ then leave $C(G)$ unchanged.
   (b) Else
      i. Create a subgraph $G_i = (V_i, E_i)$ with $V_i = C_i(G)$ and $E_i = \{\{v_m, v_n\} \mid v_m, v_n \in V_i\}$.
      ii. Get the best communities splitting $C(G_i) = \{C_1(G_i), C_2(G_i), \ldots\}$ and its respective Q value as $Q(C(G_i))$.
      iii. If $Q(C(G_i)) < Q_{threshold}$ then leave $C(G)$ unchanged.
      iv. Else
         A. Replace the community $C_i(G)$ in $C(G)$ by the communities in $C(G_i)$.
         B. Recursively apply step 3 on each community $C_k(G_i) \in C(G_i)$, *mutatis mutandis*.

This method further splits a community with enough vertices ($|C_i(G)| >= C_{threshold}$) into smaller ones. However if the resulting splitting is bad ($Q(C(G_i)) < Q_{threshold}$), the original community is kept. Empirical data suggested that values of $Q_{threshold} = 0.3$ and $C_{threshold} = 5$ might work well for our needs. We used these values during the evaluations presented in Chapter 6.

Visual Representations

Visualization of social networks is an area of research on its own. Social networks may be represented by adjacency matrix, node-link diagrams or hybrid methods (Henry et al., 2007), the most popular view being probably the node-link diagram, also called visual graph (Heer and Boyd, 2005; Perer and Shneiderman, 2006). Attempting to develop a novel visualization of social data was clearly out of the scope of this thesis. To represent the personal social network graph along with its communities, we tried two alternatives: a visual graph and a simple tabular view.

Visual Graph. There are many different algorithms used to layout graphs so that vertices are well separated from one another and the amount of edge crossings is minimized. In our case, the graph that is to be visualized is a weighted graph, where the weight of edges convey information about the closeness of the vertices they connect. Furthermore, communities of vertices are known and should be shown as well. Experiments with a weighted graph layout algorithm (Rodgers and Mutton, 2003) were disappointing, causing the social network to be hardly readable. We ended up using a classical spring-embedder algorithm (Fruchterman and Reingold, 1991) to layout the graph. This algorithm models edges as springs. Springs act as forces which draw connected vertices towards one another. To prevent vertices from collapsing, repulsive forces are also put into the model. Thus, all vertices tend to repulse one another like if they were particles having the same electrical charge. The spring-embedder layout algorithm actually simulates those antagonist forces and leads to a particular, generally unstable, graph topology. At first, the basic spring-embedder used to layout the graph proved to be quite efficient to separate main clusters from one another. However, it did not take into account the edges weights and we had to set a closeness threshold above which two people get connected by an edge. Furthermore, it appeared that communities (represented as bubbles around vertices) may overlap in some situations, making the graph difficult to read. To account the facts that our social network graph is weighted and split into communities, we actually adapted the forces in the basic spring-embedder as follows:

- Springs lengths are inversely proportional to edges weight.
- Springs connecting vertices in the same community have a stronger force and a smaller length than springs connecting vertices in separate communities.

This generates a decent view of social networks as can be seen in the screenshot of an early prototype of *WotanEye* in Figure 5.18.

Tabular View. We also implemented an alternate view using standard GUI widgets, depicted on the right-hand side of Figure 5.19. Each community is assigned a tab bearing its name and colour. Contacts which are members of the community are shown in a table inside this tab.

Visualization Choosen for the Social Facet in *Weena*. To avoid unnecessary interface complexity in *Weena*, we did not want to include two variants of the social facet interface. Arguments in

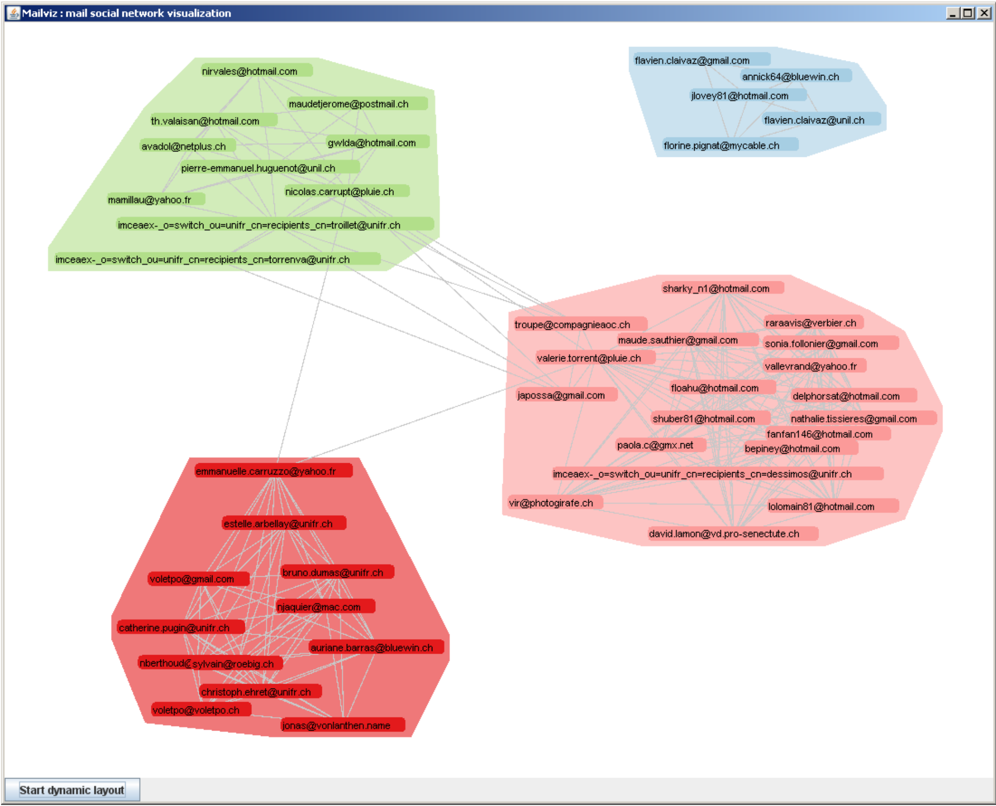

Figure 5.18: View of a small social network as a visual graph laid out using our modified spring-embedder algorithm.

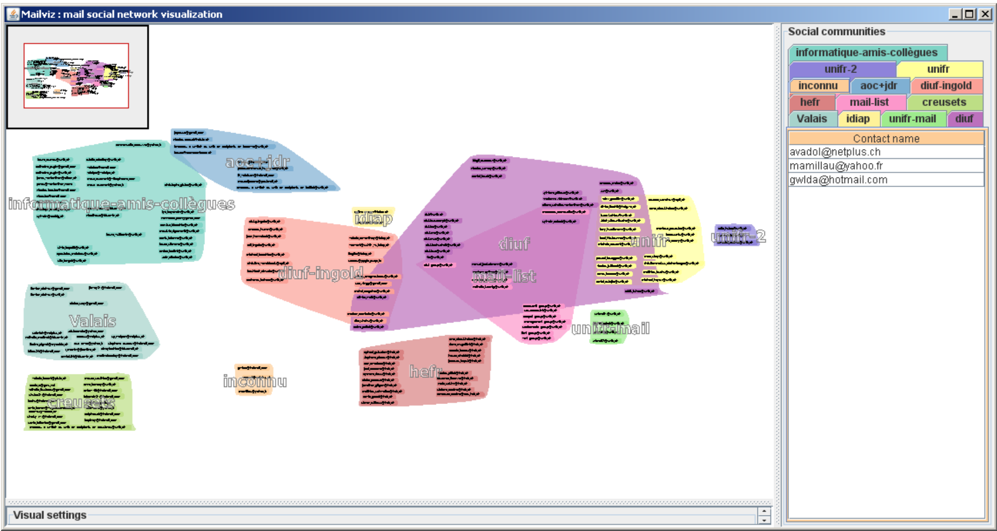

Figure 5.19: Two views of a social network. On the left, as a visual graph partially laid out manually with edges hidden. On the right, as a tabular view with each community under a tab.

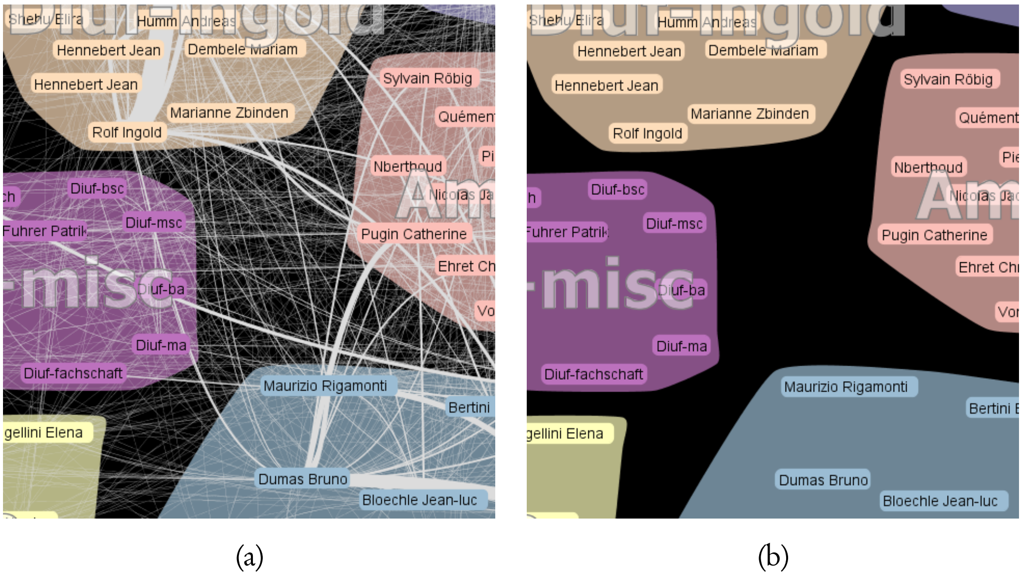

Figure 5.20: Close view of a social network with edges (a) and without (b). Hiding edges reduces visual cluttering.

favour of the graph-based version of the social facet were found in Elsweiler et al. (2007): (1) a non-linear organisation helps retrieval from a memory point of view, and (2) users learn the location of interface objects as a by-product of interaction. Whittaker et al. (2004) also notes that users like to have their contacts organized in a two-dimensional plane.

Some slight visual improvements were made to the basic *WotanEye* social network visualization in order to integrate it as the social facet of *Weena*. Names of contacts were displayed instead of their email addresses. The weight of edges was made proportional to the closeness of contacts. Edges were hidden to avoid visual cluttering (see Figure 5.20). Hovering over a contact shows the edges incident to it and also highlights connected contacts. A tooltip shows additional information about the hovered contact: their name, the community they belong to, and the contacts they are related to, sorted by decreasing closeness (see Figure 5.21).

WotanEye: Preamble to *Weena*

We formatively evaluated the social facet used in isolation (see Section 6.2). Observations made in the course of this evaluation encouraged us to extract users' social networks and configure the social facet prior to involve the users in our faceted browser's evaluation (presented in Section 6.3). We felt it would be important for users to get accustomed to the visual representation of their social network before having to employ it as a facet of their PI. Furthermore, as the automatic social graph extraction, the communities building and the graph layout are not flawless, we wanted to give users a chance to correct it in order to suit their needs. However, as we explicitly excluded

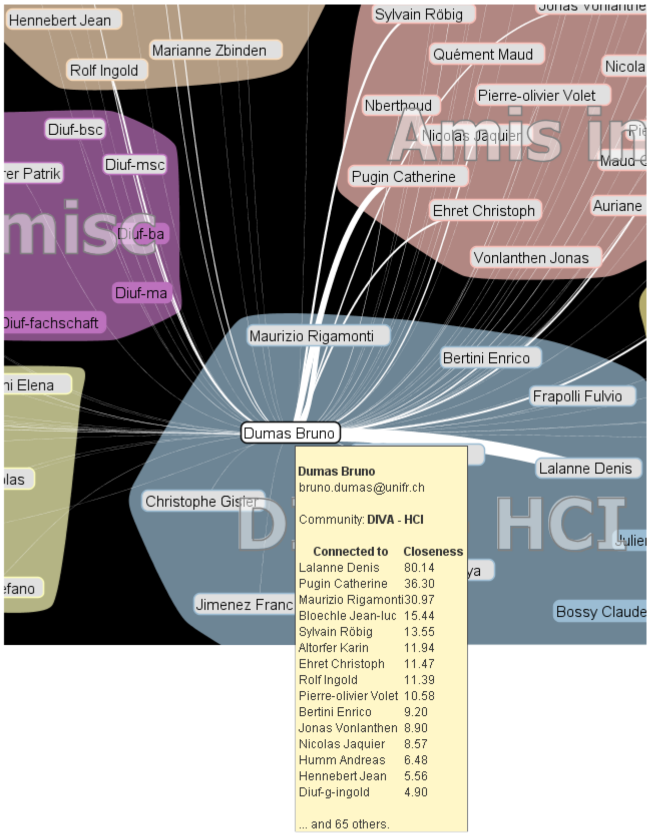

Figure 5.21: Hovering over a contact.

to rely on user annotations in the design phase of our system, we limited the extent to which users can impact on the automatically generated social network. Indeed, users can do two things: give names to communities and manually move contacts around in the visual graph. They cannot change the underlying graph structure, nor modify the communities.

During the evaluation of our faceted browser presented in Section 6.3, *WotanEye* was executed first in order to generate users' social networks. Then, users were asked to interact with their network, naming communities and moving contacts until they were satisfied with the layout of their network. At that point, the visual state of the social network was stored. Upon launch, *Weena* would restore this visual state of the social facet. We found this way of doing as a tradeoff between letting the users discover their automatically-generated social network directly in *Weena* and risking potential frustration on their part, and letting them construct their social network manually which would have been against our design principles[13].

Mapping Contacts in the Social Facet to Information Items

As the social facet is built from the email archive of the user, it does not share the same backing data table as the other *Weena* facets which get theirs from the WDS index. Three different strategies have been implemented to map contacts in the social facet to information items obtained from the WDS index. They are presented hereafter.

Attachments Names. The first strategy compares the names of files attached to emails to the names of files indexed by WDS. If file A in the WDS index appears in an email, then all contacts who are on the envelope of this email will be related to file A.

Participants Names. The participants strategy maps email addresses from the social facet to items in the WDS index based on the value of their SYSTEM.ITEMPARTICIPANTS property. The mapping from addresses to names and *vice versa* is based on rules and string patterns. For example, it maps the address 'florian.evequoz@unifr.ch' to the participant 'Florian Evéquoz', and marks items having 'Florian Evéquoz' as participant as related to the contact labeled 'florian.evequoz@unifr.ch' in the social facet.

Query. The last strategy relies on a WDS query to find further information items that may be related to contacts. Actually, this strategy consists in issuing a WDS fulltext query with the

[13]While developing *ContactMap*, Whittaker et al. (2004) noted the following paradox related to how users perceive their social network: *"On one hand, users found tedious to build their social network by themselves. On the other hand, when presented with automatic ways of classifying their social network, or helping them in this task, users said they would prefer to do it all by themselves"*.

displayed name of the contact in the social facet as keyword. All items resulting from the query are marked as related to this contact.

The actual mapping strategy used in *Weena* is indeed a combination of the three strategies presented above. Information items related to a contact are thus obtained by taking the union of the items related to this contact according to each of the strategies.

5.2.4 Implementation Challenges and Limitations

There were a number of challenges that appeared during the implementation of the system. First, the WDS index is sometimes incomplete, some metadata missing. We had to take this into account to be able to recover if such a case happens. Then, we had to make WDS COM-based interface work flawlessly with Java. The development of an *ad hoc* bridge led us to face memory leaks, losses of data and performance issues. To prevent those performance issues, the final version of the tool was as much disconnected from WDS as possible. On the interface side, as filtering the backing data table may be time-consuming (especially in the case of the textual search and social facets which involve sending a query to WDS and waiting for the answer), synchronization mechanisms had to be developed to prevent the interface from freezing and hindering the user experience.

Additionally, the mapping between contacts and information items obtained from the WDS index is suboptimal as implemented. Ideally, this link should be done in a preliminary phase, so that the relations are fixed once and for all during the execution of *Weena*. However, this could not be achieved in the version of the system used during the evaluation. Instead, the linking of contacts and information items was done dynamically on-the-fly when needed. Apart from slowing down the interface, it also has a few drawbacks: (1) the mapping is not symmetrical (i.e. a contact C might get connected to a file F after C was selected in the social facet, but C might not get connected to F if F was kept in the result set after a selection in another facet), (2) it is impossible to know how many items are related to a contact, thus impoverishing the information that the social facet can convey (no query preview).

5.3 Chapter Summary

This chapter described *Weena*, the experimental system developed according to the findings of Chapters 3 and 4. The system consists in a faceted browser of personal information enabling the use of seven facets of PI to navigate in a personal space of information. The system improves over previous applications of faceted navigation to PIM, in particular because it does not require user annotation of the data and presents the categories of facets using information visualization techniques (date and social facets). As the main research goal of this thesis is to evaluate the potential of faceted navigation using contextual cues to support the re-finding and self-analysis tasks of PIM, the following chapter is dedicated to the evaluations conducted using this system.

Not everything that can be counted counts and not everything that counts can be counted.

Albert Einstein

Les trois quarts de l'univers peuvent trouver délicieuse l'odeur d'une rose, sans que cela puisse servir de preuve, ni pour condamner le quart qui pourrait la trouver mauvaise, ni pour démontrer que cette odeur soit véritablement agréable.

Sade, *Justine* (1791)

Experience has shown, and a true philosophy will always show, that a vast, perhaps the larger portion of the truth arises from the seemingly irrelevant.

Edgar Allan Poe, *The Mystery of Marie Rogêt* (1842)

6

Evaluation

This chapter presents the evaluations that were conducted using the system introduced in Chapter 5. As the preliminary studies discussed in Chapter 4 outlined, what people say about how they manage their personal information may be very different from what they actually do. My motivation was therefore to observe their behaviour using a real system and performing real tasks. The main objectives of this evaluation were: (1) to assess the usefulness of faceted navigation for re-finding personal information, and (2) to explore how people actually perform real re-finding tasks with the help of facets. Secondary goals included an informal exploratory evaluation of the potential of faceted navigation for gaining insight into one's PSI and a summative evaluation of the system.

This chapter is divided into three parts. The first part, Section 6.1, offers a general discussion of evaluation in personal information management (PIM). It presents the different types of evaluation that can be conducted, discusses previous evaluations in PIM and uses them to identify pitfalls in PIM evaluation. It concludes by presenting an evaluation methodology aimed at overcoming those difficulties. This general discussion also contributes to improve the theoretical framework underlying evaluation in the field of PIM, which is a side objective of this thesis (see Section 1.4). The second part, Section 6.2, presents a formative evaluation of the social facet interface introduced in Section 5.2.3. The results of this evaluation have been incorporated into the development of the *Weena* faceted browser described in Chapter 5. The third part, Section 6.3, presents the main evaluation conducted with the *Weena* faceted browser and using the methodology introduced in Section 6.1.

6.1 Evaluation in PIM

This section first covers the different types of evaluation in the field of human-computer interaction (HCI). Then, it describes different methods available for evaluations involving PIM tools. Previous evaluations of PIM tools are then discussed and used to identify the difficulties that may arise in such evaluations. Finally, evaluation guidelines for overcoming those difficulties are examined.

6.1.1 Evaluation Types and Levels in HCI

Andrews (2008) and Munzner (2009) discuss the different kinds of evaluation performed in the field of information visualization. I believe that this discussion is also relevant for HCI at large, and in particular for PIM tools. This section briefly presents the specificities of evaluations and discusses them in the context of my own evaluation presented in Section 6.3.

Evaluation can be divided into exploratory (i.e. how do people use the interface and what for?), predictive (i.e. estimate user performance with the interface based on a specialist's expertise), formative (i.e. provide design feedback and issue recommendations) and summative (i.e. overall assessment of the interface, often consisting of statistical measures). Only summative evaluations are suited to demonstrate the utility or superiority of an interface (Andrews, 2008). Summative evaluations can use different concrete methods. *Guideline scoring* is a method where an expert scores an interface against a list of guidelines. *Questionnaires* filled in by testers after they performed tasks on the interface are another summative method. *Formal experiments* are also summative. They allow the collection of objective measurements on how testers perform with the interface. Lastly, a special type of *observational studies* can be considered summative. Observing participants who use the interface and collecting evidence may lead to an informal validation of the interface. Nonetheless, observational studies are primarily meant for exploratory evaluations (Andrews, 2008).

Furthermore, evaluations of software interfaces can be applied at different levels. Munzner (2009) discusses these levels in the context of information visualization. She distinguishes 4 nested levels of evaluation that are, from outermost to innermost:

Domain-level evaluations aim at answering the question: "does the interface help target users in fulfilling their real tasks?" Such evaluations assess whether users like the interface and whether is it adopted.

Abstraction-level evaluations aim at answering the question: "does the abstract form of the data, and operations associated with it, match the users' mental models?" Such evaluations assess whether the abstraction of the data is comprehensible and usable.

Encoding/interaction-level evaluations aim at answering the question: "are visual encoding of the data in the interface, and interaction with them good design choices?" Such evaluations assess whether the actual interface and the widgets it offers are usable.

Algorithm-level evaluations assess whether algorithms used at different stages in the interface are efficient in terms of time, memory, or precision and recall measures.

In the context of PIM and the evaluation of the tool introduced in Chapter 5, those levels would match the following. Domain-level evaluations would assess if real users like and adopt the system as a support for their everyday re-finding tasks. Abstraction-level evaluations would assess if faceted navigation in general is a way of re-finding PI that is comprehensible and usable. Encoding/interaction-level evaluations would assess if the particular design choices that were made in the faceted navigation approach are comprehensible and usable. Algorithm-level evaluations would assess for instance whether algorithms dedicated to generating the faceted metadata are efficient in terms of time, memory, precision/recall, etc.

The evaluations presented in this chapter are of different types and address different levels. The usability evaluation of the social facet presented in Section 6.2 is a primarily formative evaluation addressing the encoding and the abstraction level for the particular case of the social facet interface. The evaluation of the tool presented in Section 6.3 comprises different phases, some of which are summative, while some others are exploratory. It addresses the encoding and abstraction levels, as well as the domain level, to a certain extent.

6.1.2 Evaluation Methods for PIM Tools

Elsweiler (2007) extensively reviewed the different methods of evaluation of PIM tools that can be used. Table 6.1 summarizes them and additionally provides the main advantages and drawbacks of these methods.

On the one hand, there are naturalistic studies, which aim to observe users in their environment, with the observer trying to disturb the user as little as possible. Advantages of these techniques are that both the corpus of information (the users' PI) and the tasks are real and occur in real work contexts. This guarantees the ecological validity of the evaluations. Fieldwork and ethnographic methods, however, require the presence of an observer during long periods of time to collect evidence and are very expensive when it comes to analysing the notes and recordings from the fieldwork. Log-based studies can be spread more easily over many users and allow automatic gathering of many low-level features of the users' interactions with the system. However, nothing is known about the tasks that make people actually use the system the way they do. Moreover, the evaluated system needs to be at a stage of development that allows it to be on a person's critical path. This criteria is typically difficult to satisfy in many academic contexts. Furthermore, PIM tasks done with legacy tools would not be logged. Parts of the users' behaviour would therefore not be understood.

On the other hand, laboratory-based studies occur in a controlled setting. The most controlled setting uses a common corpus and a predefined set of tasks for all participants. While this permits a precise analysis and comparison of tools, it lacks the ecological validity of working with information that is personal. Nevertheless, it may be used to test precise features of interfaces or to compare alternative interfaces one against another, when applied to the same tasks. The last evaluation technique involves using participants' PI in a controlled setting. The main challenge of this approach resides in the selection of tasks that must be comparable between participants, while still remaining realistic.

6.1.3 Discussion of Previous PIM Tool Evaluations

As was already noted in Chapter 3, most PIM tools have not been evaluated at all, or only informally (Freeman and Gelernter, 1996; Fertig et al., 1996; Rekimoto, 1999; Dourish et al., 2000; Gemmell et al., 2002; Karger et al., 2005). Some tools have been evaluated with a formative goal in mind (Kaptelinin, 2003; Whittaker et al., 2004; Krishnan and Jones, 2005; Smith et al., 2006; Jones

<table>
<tr><th>Evaluation methods</th><th></th><th>Advantages</th><th>Drawbacks</th></tr>
<tr><td rowspan="2">Naturalistic studies</td><td>Fieldwork-ethnographic studies</td><td>Real corpus
Real tasks</td><td>Experimenter's presence
Few participants
Analysis of results costly</td></tr>
<tr><td>Log-file analysis</td><td>Real corpus
Real tasks
No experimenter's presence
Many participants
Much data about interaction</td><td>Tasks are unknown
Tool must be deployable
Legacy tools not integrated</td></tr>
<tr><td rowspan="2">Lab studies</td><td>With common corpus</td><td>Controlled setting</td><td>Corpus is not personal</td></tr>
<tr><td>With personal information</td><td>Controlled setting
Real corpus</td><td>Tasks definition challenging</td></tr>
</table>

Table 6.1: Methods for summative evaluation of PIM tools, with their advantages and drawbacks.

et al., 2008). Some tools were evaluated summatively on aspects not related to the support of the re-finding task (Whittaker et al., 2004). Here, I describe and discuss previous evaluations of PIM tools that are of interest and classify the methodology employed according to the evaluation types and levels presented above in Section 6.1.1. Table 6.2 summarizes this discussion.

The *TimeSpace* system (Krishnan and Jones, 2005) was evaluated with a qualitative observational study, addressing the abstraction and domain levels. However, the outputs of the study suggest a formative rather than a summative goal. Moreover, the re-finding task received little interest. The development of *ContactMap* involved formative evaluations (Whittaker et al., 2004) where representative users provided feedback that could be incorporated into the next stages of development. A summative evaluation at the abstraction level was also conducted. It compares the use of the tool with a regular email client on 4 specific tasks, two of them supporting the reminding function of PIM and the other two being social data mining tasks. The corpus of information was the participants' real PSI. The tasks were simulated and common to all participants. The capacities of the tool to support re-finding tasks was evaluated on one task which involved re-finding the last 5 emails sent and received which pertained to a particular project, which was a real project of the participant. *Milestones in Time* (Ringel et al., 2003) was evaluated summatively at the abstraction level for re-finding tasks. The evaluation method was a laboratory study in a controlled setting. The corpus of information was semi-personal: it consisted of a single corpus of announcement emails sent to the participants. All participants had received those emails, which were broadcasted to the whole staff. Nevertheless, in their real PSI participants may have deleted or classified those emails, a fact that the simulated corpus did not take into account. The tasks were artificially created and were the same for all participants. Tasks were performed on two variants of the systems using different abstractions: in the first system, only a timeline was provided, whereas in the second system landmarks were added to the timeline. Additionally, a questionnaire addressed the domain level in a summative way.

Stuff I've Seen (Dumais et al., 2003) and *Phlat* (Cutrell et al., 2006b) are two projects done by the same team at Microsoft Research. They have been evaluated very similarly. Their evaluations are examples of naturalistic evaluation, as they involved real corpora of PI and real tasks. *SIS* was used by 234 participants during six weeks. *Phlat* was used by 225 participants during eight months. In both cases, the system was installed on the participants' own personal computer. Participants were free to use it for their PI re-finding tasks. During this period, user interactions with the system were logged. Properties of items opened with the help of the system, such as age and file types, were also logged. Some information was not kept, for privacy reasons, e.g. file names. Users could provide feedback to the developers by sending them emails. Their feedback constituted a formative evaluation of the system, and led to defining recommendations for the system's improvement. The logs analysis provided very useful general observations about how the systems were used. The adoption of the system could have been assessed from the logs. For instance, this would have answered the question: did the participants use the system more often as time passes or not? This would have constituted a summative evaluation at the domain level but the authors did not report on this. Instead, they presented statistics about the age of items accessed, the types of words appearing in queries, the types of files opened, etc. Thus, the evaluations qualify as exploratory. Furthermore, this evaluation method exhibits an important weakness: the users' tasks are unknown. Thus, the behaviour of the users obtained from the logs cannot be mapped to the specific task they were performing at the time. The evaluation, presented in Section 6.3, maps user behaviour to specific tasks and is a contribution with respect to this aspect.

Feldspar (Chau et al., 2008b) exemplifies the case of a controlled laboratory study. The evaluation was summative and addressed the abstraction level. It was done under a controlled setting, with a common corpus and common tasks. The corpus consisted of: (1) emails artificially created for a research project, (2) a filesystem hierarchy built randomly and populated with files obtained from wildcard searches in Google and (3) calendar events artificially created by the authors. Fourteen tasks were created for this corpus, and assigned a difficulty rating by authors. The results show that people took significantly less time to perform the tasks with *Feldspar* than with *Windows Explorer*, *Outlook*, and *Google Desktop Search* used freely. However, I claim that the chosen methodology only allows to evaluate *Feldspar* as a general information retrieval tool, and not as a personal information retrieval tool. Indeed, it seems reasonable that finding the files required by the task *"open the file folder containing all the email attachments from Spence"*, for instance, is difficult without the help of their prototype, because the participants did not know the filesystem hierarchy, nor who is Spence. One can imagine a real situation where a participant would be much more efficient navigating to the target folder using the file explorer (because she knows which folder is concerned), than using the prototype. I similarly claim that the summative questionnaire filled

	Evaluation level				Evaluation type			
	Domain-level	Abstraction-level	Encoding/interaction level	Algorithm level	Exploratory	Predictive	Formative	Summative
Stuff I've Seen (Dumais et al., 2003)	×				×		×	
Phlat (Cutrell et al., 2006b)	×				×		×	
PhotoMemory, *MemoMail* (Elsweiler, 2007)	×	×	×		×			×
TimeSpace (Krishnan and Jones, 2005)	×	×					×	
ContactMap (Whittaker et al., 2004)		×	×				×	×
Milestones in Time (Ringel et al., 2003)	×	×						×
Feldspar (Chau et al., 2008b)		×						×

Table 6.2: Types and levels of previous PIM tools evaluations.

in by participants in their study rates the system as if it was a general, not personal, information retrieval tool. Consequently, this evaluation, though well conducted, fails to assess the superiority of the system as a *personal* information retrieval tool.

PhotoMemory and *MemoMail*, by Elsweiler (2007) were evaluated using their authors' methodology described in Section 6.1.5. Both were evaluated on participants' real corpus and with real tasks — or simulated tasks based on real tasks, as described in Elsweiler et al. (2007). However, the focus of the studies were not on the tools themselves, but on the recollection of participants and the role of their memory in re-finding activities. Indeed, the main outputs of the study relate to the recollection of attributes of sought-after information items. Thus, it qualifies the study as exploratory with respect to the tools developed. Satisfaction questionnaires filled in by participants nevertheless provide some summative validation of the abstraction and encoding levels of the implemented tools.

6.1.4 Pitfalls of PIM Evaluations

The above discussion of PIM tool evaluations has shown that conducting ecologically valid evaluations of PIM tools involves accounting for the following facts:

- PI consists of information collected and maintained individually. It is linked to an individual's own experience. The only realistic PI corpus for an individual is this individual's real PI corpus.

- PI may contain areas that are and should stay private. An individual shall be the master of what part of their PSI they agree to share with an experimenter.

- Reasons that make people re-find PI are not known in advance. They happen in the course of their work.

Evaluations of PIM tools supporting the re-finding task which do not account for those facts lack validity and reliability. Consequently, it is difficult to accumulate, compare and integrate results across such evaluations (Kelly, 2006). In particular, evaluations of PIM tools supporting the re-finding task cannot be conducted using the traditional information retrieval (IR) model, which corresponds to a particular kind of controlled laboratory studies. The evaluation of *Feldspar* presented above is an example of a PIM tool evaluation based on the IR model. Standard evaluation methods in IR require common corpora, where experts have classified documents into a so-called ground truth. The idea of a "common corpus" is a completely unrealistic information collection with respect to PIM. Therefore, traditional IR evaluation methods based on precision/recall measures cannot be applied to PIM because it would not involve the personal aspects of it. For the same reason, there cannot be common tasks in PIM, as required by the IR model, as the corpus is not common for all participants. To assess the adequacy of a PIM tool for users' needs, the users themselves should handle the expert's work of devising tasks and creating ground-truths.

Indeed, the most natural way of obtaining realistic user tasks would be to let the participants themselves provide tasks. However, user-defined tasks may be vague. Kelly (2006) notes for example that *"'doing email' might be subdivided into at least four separate tasks: searching for a specific piece of email; managing and filing email; setting up and accessing an address book; and reading email"*. Moreover, without some similarity between tasks, results between different evaluations cannot be compared. Thus, tasks need to be characterized precisely enough to allow comparison between studies. In other terms, taxonomies of PIM tasks are needed. There have been previous attempts at providing such taxonomies. Some of them classify tasks according to properties of the sought-after information item. For example, Elsweiler et al. (2007) used a taxonomy of re-finding tasks pertaining to personal photographs, based on how photographs relate to events. Capra and Pérez-Quiñones (2005) classified tasks based on how familiar users are with the sought-after information. Barreau and Nardi (1995) classified tasks according to the frequency of access to the information. Elsweiler and Ruthven (2007) classified tasks according to the temperature of the target item, i.e. the last time the sought-after information item was accessed. Some other scholars provided taxonomies of tasks that do not focus on the sought-after item but more on the task itself. Byström and Järvelin (1995) suggested classifying tasks based on how users try to find information. Capra and Pérez-Quiñones (2005) classifed tasks according to how often users perform the task. Bell and Ruthven (2004) additionnaly proposed to use the subjective complexity of tasks to classify them, a categorization that is also used by Elsweiler and Ruthven (2007). The chosen taxonomy of tasks obviously influences the conclusions that can be made when examining the results of the study.

Thus, it should be determined carefully and its validity should be assessed as well. The next section discusses guidelines for evaluations of PIM tools that are valid with respect to the issues presented in this section.

6.1.5 Guidelines for Evaluations of PIM Tools

Elsweiler and Ruthven (2007) introduce an evaluation methodology that takes the issues mentioned in the previous section into account and thus aims at ensuring the validity and reliability of evaluations of PIM tools supporting the re-finding task. More specifically, their evaluation methodology comprises two distinct phases. In the first phase, participants are asked to fill in a diary of their real re-finding tasks during a one-week period. After this first phase, an experimenter gathers the different tasks and classifies them into an *ad hoc* taxonomy of tasks. He then selects a subset of relevant tasks. In the second phase, participants are asked to perform the relevant tasks with the PIM tool to be evaluated. The performance and behaviour of the participants related to different kinds of tasks can then be assessed. Section 6.3 describes this methodology in greater detail in the context of my own evaluation.

The proposed methodology has several benefits. First and most important, it is realistic in terms of corpus and tasks. Participants work on their own PI and the tasks they perform are real tasks, devised by themselves. Moreover, they of course know the targets of their re-finding tasks (ground-truth). Second, it allows to evaluate aspects of the re-finding behaviour at the *task* level, which is particularly interesting (Capra and Pérez-Quiñones, 2003) and still under-evaluated. Third, the controlled experiment is suitable for different kinds of qualitative and quantitative observations that can be obtained concurrently, e.g. time taken to fulfill tasks, low-level use of the interface, participants thinking aloud, etc.

There are a few drawbacks, though. First, as the evaluation relies on tasks provided by the participants, the tasks obtained from the diary study should be of sufficient quality to be re-used. It is therefore important that: (1) participants log enough tasks so that the experimenter can make a selection from them and (2) the participants understand clearly the need to be precise enough in their tasks descriptions. A second drawback of this evaluation is that is is costly in terms of setting up. Collecting tasks from diaries, devising an *ad hoc* taxonomy of them, selecting relevant ones and conducting the controlled experiment itself requires a lot of resources. Therefore, it is inherently limited to a small sample of participants, which may not be representative. Given this, results may not lead to general interpretations, but rather give directions for future research. In the evaluation presented in Section 6.3, we attempt to take measures aimed at limiting the effects of these drawbacks.

6.2 Usability Evaluation of the Social Facet Interface

This section describes a usability evaluation of the social facet interface used in isolation. This evaluation was proposed in the framework of a collaboration with the Department for Psychology of the University of Fribourg as a topic for a research seminar in usability evaluation of human-computer interfaces. The full study has been published internally by the two psychology students involved in the evaluation process (Berlowitz and Gautschi, 2010).

6.2.1 Description and Goals

This evaluation was both exploratory and formative. The goal of the exploratory study was to investigate the way users think of their social network. In particular, we wanted to assess whether the automatically constructed social network was considered meaningful by participants. It is therefore an evaluation at the abstraction level according to the model by Munzner (2009). The goal of the formative evaluation was to evaluate the encoding/interaction level of the interface to assess its usability, to devise possible improvements and find missing features.

A questionnaire was also filled in by participants. Even though it provides summative data about the perceived satisfaction of participants, the number of participants was not enough to assess the actual efficiency of the system. Its results therefore have a merely informative value.

6.2.2 Evaluated System

The evaluated system was *WotanEye*'s social network interface, coined *Mailviz*. This system aims at displaying a user's social network of her most important contacts, automatically clustered in communities. As described in detail in Section 5.2.3, this system reads the emails from a user's mailbox and builds a social network out of them. The social network is represented as an undirected weighted graph of email addresses of contacts. The weight of edges represent the connectivity strength between two contacts. An algorithm automatically extracts communities from the graph, based on its topology. Communities consist of contacts which have strong links with one another inside the community (i.e. many edges having high-weighing values) and weak links to other contacts in the graph (i.e. few or low-weight edges).

A screenshot of *Mailviz* is shown in Figure 6.1. The right-hand side of the window represents the actual social network as a graph. Vertices represent email addresses and are labeled accordingly. In this anonymized screenshot, email addresses have been replaced by numbers. Edges are gray curves connecting vertices. The weight of edges is mapped to their width. Communities are colored bubbles grouping contacts together. Vertices and their surrounding community are painted in the same colour, the community being lightly desaturated. In the middle of each community, a semi-transparent text overlay shows the community's label. The graph is laid out according to a

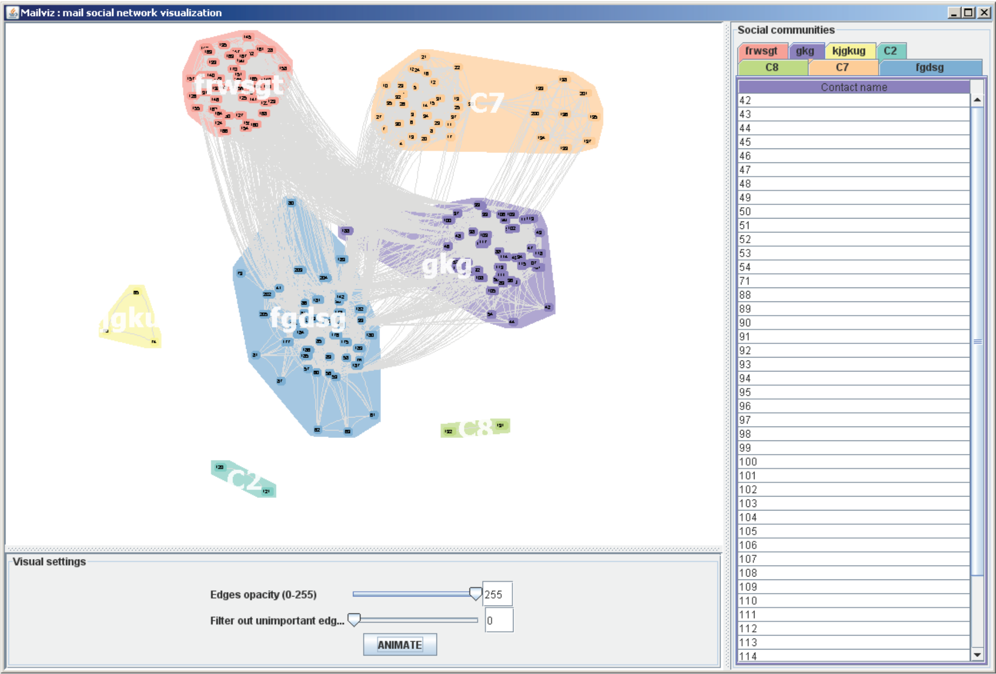

Figure 6.1: An anonymized screenshot of *Mailviz* displaying a participant's social network.

modified spring-embedder graph layout algorithm. The default layout can be overriden, as vertices and communities may be moved freely by the user. Below the social network view, a panel is dedicated to visual settings. A slider lets the user set the opacity of edges, from fully opaque to transparent. Another slider lets the user filter out the edges based on their weight, to keep only the most weighted edges. Finally, a button lets the user start and stop the automatic graph layout. On the left-hand side of the window, each community takes up a tabbed panel. The colors on tabs identify communities. When a community is selected, its members are displayed in a tabular fashion.

Several interactions are enabled in the interface. The visual settings may be adapted using the provided widgets. Communities can be renamed by double-clicking on the corresponding tab. In the social network, the mouse wheel enables zooming; a right-click drag-and-drop allows to pan the view; hovering over a contact with the mouse highlights the connected contacts; double-clicking a community zooms in on it.

6.2.3 Methodology

The exploratory goal of the study consisted in assessing whether the representation of the social network made sense for the participants. To attain this, evaluators asked each participant to sketch her social network on a sheet of paper (see Figure 6.2) at the beginning of the interview. When the automatically generated social network was revealed, evaluators further asked the participants to comment on it and to describe the differences from their previous sketch.

To attain the formative goal, evaluators asked the participants to perform specific tasks with the interface, without giving them any tutorial on how to use it. Those tasks were:

1. Stop the animation of the social network.
2. Zoom in and out the social network view.
3. Look at a section currently out of sight.
4. Hide links.
5. Change the opacity of links.
6. Name the communities.
7. Explain the meaning of colors.

If the participant was unable to perform the task after a certain amount of time, the evaluators would explain how to proceed.

After the participant had gained some familiarity with the interface, the evaluators proposed a scenario. The participant supposedly has a birthday party to organize. She is under great time pressure and wants to send an invitation email to her friends as soon as possible. She should decide what interface she would use to select her friends in the quickest way possible from all her contacts.

While they performed the task, participants were asked to think aloud. Additionally the evaluators observed the participants interacting with the system and took notes. The evaluation session was also tape-recorded.

At the end of the study, the participants filled in a satisfaction questionnaire created according to the System Usability Scale (Brooke, 1996).

6.2.4 Setting

Five participants (3 female) were recruited for this evaluation. Two of them were Ph.D. students in computer science and the remaining three were grad students in psychology. Their ages ranged from 23 to 31. The study took place in a laboratory where a computer was available. Two evaluators were always present, one acting as a moderator and the other one in the background observing the participant using the system while taking notes. Interviews were also tape-recorded. The tapes were later examined in detail, and observations were extracted from them.

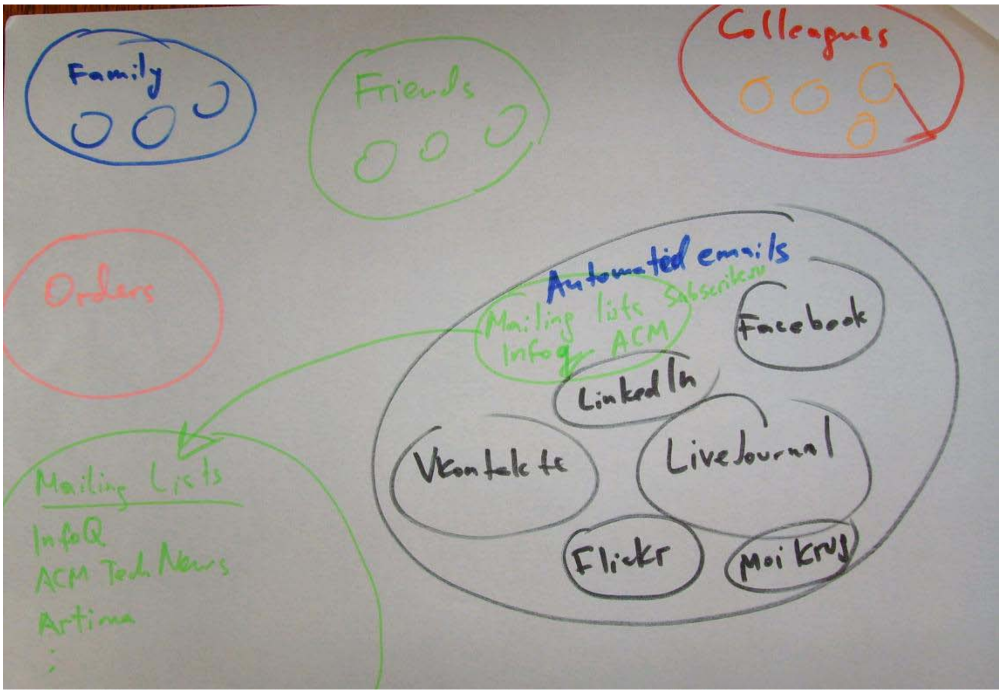

Figure 6.2: Sketch of a participant's social network.

6.2.5 Results

Results of the Exploratory Evaluation

When they were asked to draw their social network, the participants drew two main types of sketches. In the first type, participants drew themselves in the middle, and the groups were connected to them. In the second one, participants did not sketch themselves and only drew their communities of contacts. *Mailviz* uses the latter type of view. All participants also used colors in their sketches to differentiate communities, as is done automatically in *Mailviz*. Comparing their sketches with the automatically extracted network, the participants managed to recognize partial similarities between them. They found at least some meaningful communities (e.g. family, working colleagues, close friends). However, unrelated contacts also often appeared in such communities. Noticing inconsistencies between their expected view and the actual result of the community building led some participants to lose trust in the system. One of the participants wanted to have communities further subdivided. The largest group often consisted of people unfamiliar to the participant. This can be explained as follows. Announcement emails are very often sent to many recipients. Most of those recipients are unknown to the participant. As those unfamiliar people only appeared in announcement emails, they become grouped together. Interestingly, the two computer science participants were able to explain this fact by themselves. Additionally, one participant interpreted communities with respect to time. Some groups of people were indeed past colleagues with whom he had no more contact.

Results of the Formative Evaluation

All users could easily stop the animation, zoom, and understand the meaning of colors in the social network view. Most participants failed to pan the view in order to visualize another portion of the network without the assistance of the evaluator. Changing the opacity of links and hiding unimportant edges was also seen as a difficult task. The main issue explaining this is that most participants did not understand the labels of the dedicated controls.

When presented with the birthday scenario, most participants said they would use the tabular lists to select relevant contacts rather than the social network view. The main reason invoked by participants was the difficulty to pan, zoom and spot specific contacts in the visualization. However, most participants liked the fact that the *"social map"* allows them to have a bird's eye view on their contacts and said that they would like to have this feature included in their regular email client.

General remarks about particular features of the interface were also given by participants. One complained about zooming which does not work if the cursor hovers over anything other than blank space: *"it seems as if I can zoom only sometimes, just before I could do it but now it won't work, I don't get it"*. Two participants explicitly mentioned the labels of widgets which were not clear: *"I don't see what 'hide edges' means"* and *"I'm not quite sure what these lines mean, there should be an explanation for that somewhere"*. A participant was surprised that zooming did not focus on the cursor (focused zooming) but on the center of the screen (central zooming). A more general remark specifically addressed the flaws of the automatic communities extraction: *"if I could remove individuals or place them in different groups, or if I could create new groups myself, then I think the grouping would make sense"*[1].

Other remarks of the participants were in favor of the interface. One participant liked the highlighting of connected contacts when the mouse hovers over a specific contact. Another one liked the possibility to hide unimportant edges.

In addition to user feedback, the evaluators took notes that further described some unexpected situations. For example, a participant's social network exhibited a very wide edge which hid parts of his social network. Some networks also contained different addresses for the same person, which disturbed the participants. Furthermore, some bugs and algorithmic inefficiencies were also discovered and listed.

The overall satisfaction score from the questionnaire (on a scale from 0 to 100), based on Brooke (1996), ranges from 35 to 87.5 (average 61). While it does not meet the requirements for a "usable" system (above 80), the score was understandable for a system such as *Mailviz* which is at a de-

[1] Some participants' remarks have been translated from German to English by Berlowitz and Gautschi (2010).

velopmental stage. The most often mentioned reasons for this low score concerned the need for technical support or explanations of how to use the interface, and the fact that some interaction mechanisms were not intuitive enough.

Requirements

Based on the results of the formative study, the following list of recommendations for enhancing the system was issued:

- The contact's mail addresses should be replaced by real names. These real names can be automatically extracted from the email headers.
- Users should be able to set the maximum width of edges with a slider.
- Scrolling using a right-click drag-and-drop is not sufficient. Scrollbars should be provided. Additionally, an overview would help users see the big picture of their social network at all times.
- Zooming should work even if the cursor does not hover over blank space.
- The labels of the two filters (opacity and number of edges) should be made clearer.
- Tooltips with explanatory descriptions when hovering over interface elements and visual items could help rule out ambiguities and provide additional information.
- A contact who possesses several email addresses should be recognized as a single person and displayed only once.
- Users should be able to modify the communities produced by the system.

The above requirements were integrated into the system's development. All but the last two ones were successfully satisfied for the final version of the prototype that would be used as the social facet of the *Weena* faceted browser presented in Section 5.2.2. It was decided not to address the last two requirements because they would require a significant programming effort while at the same time being non-critical as far as system use was concerned.

6.2.6 Conclusion

This section presented a usability study of *Mailviz*, a visualization of a social network extracted from emails. The exploratory part of this study showed that the presentation of contacts as a social graph of communities is related to the way people sketch their social network. The visualization itself was perceived as interesting, even if the grouping of contacts was not flawless. The formative part of the study led to the formulation of a list of requirements. Most of these requirements were met in the next development stages of the system. Participants said they would prefer the tabular view to quickly re-find contacts. However, their reasons for doing so were mainly related to the poor interaction mechanisms with the visual social network. These interaction mechanisms having been corrected, the overall results of the study encouraged me to keep the social network view and integrate it into the *Weena* faceted browser whose evaluation is presented in the following section.

6.3 Evaluation of Facets Use

This section presents an evaluation of the faceted browser presented in Chapter 5. Participants' own personal information and tasks are used to make the evaluation as ecologically valid as possible, using the browser *in situ*. This section describes the evaluation setting in detail, its results, and a discussion of them.

6.3.1 Objectives

The main objectives of this evaluation are:

1. Assess the usefulness of faceted navigation for re-finding personal information (summative evaluation at the abstraction and domain levels)

2. Explore how people actually perform real re-finding tasks with the help of facets (exploratory evaluation)

Secondary goals include (1) informally evaluating the potential of faceted navigation for gaining insight into the personal space of information (PSI) and how faceted navigation can help users analyse their own PIM strategies and (2) evaluate the usability of the system (summative evaluation at the encoding/interaction levels).

6.3.2 Setting

The evaluation process we used was based on the methodology proposed by Elsweiler and Ruthven (2007) to evaluate PIM systems based on real tasks. The evaluation process consisted of two successive phases. In the first phase, participants were asked to fill in a diary documenting which kind of personal information re-finding tasks they performed during a certain period of time. In a second phase, they were presented with a subset of the tasks they actually recorded and asked to perform them using our system. Of course, when filling in the diary in the first phase, participants did not know that the facts they reported were going to be used again in the second phase, as this would have biased the evaluation process.

Experimenters

Two experimenters worked on this evaluation. They were a grad student involved in the development of the *Weena* faceted browser in the context of his MSc thesis (Thomet, 2010) and the author.

Participants

Ten participants (P1 to P10), including 5 females, were recruited for this evaluation. Their ages ranged from 18 to 46 (average 28). They came from very diverse fields, but mainly computer science (3) and psychology (2). Their professions include administrative jobs (secretary, project leader), research, high-school studies and arts (graphical, musical). We chose people from different backgrounds in order to get as representative a sample as possible. All participants spoke French except P2. Thus, the instructions were provided in French and the participants wrote their diaries in French. P2 could read French, so she could use the same instructions and documents as the others. However, she wrote her diary in English and we conducted the interview with her in English as well.

As is often the case with PIM studies (Boardman, 2004), these participants were chosen among acquaintances of the experimenters. The main advantage of doing this is that it reduces the privacy issues of working with strangers' personal information. It also facilitates informal and open discussion during the evaluation itself. However, participants knew that the program they were evaluating was (at least partly) a contribution of the experimenters, which could of course influence their feedback towards a more positive assessment. To minimize this possible methodological bias, we took two measures: (1) we asked the participants to be sincere in their feedback and report their real impressions without fear of offending the experimenters; (2) as two experimenters were involved, we chose the "best" experimenter for each participant, i.e. the experimenter with which the participant would hopefully be the most honest in their feedback. As we received a broad range of feedback for most questions asked, we conclude that this goal was achieved. Moreover, the evaluation conducted here was qualitative. The small sample of participants also prevents us from making conclusions that are too general. Therefore, we are confident that even with this choice of participants the results provide good directions for future research on the faceted navigation of personal information.

6.3.3 Diary Study

Benefits and Drawbacks. Diary studies belong to the category of naturalistic studies. They have a number of advantages. Notably, they are marginally intrusive. They allow to record real user tasks without the participants being disturbed by the presence of an observer. Nevertheless, diary studies also have some limitations. For example, it may be hard to maintain participants' dedication to the logging task. Furthermore, participants may not record facts that seem uninteresting to them, even if those facts would be useful and profitable to the experimenter. The quality of their fact reporting can also be very uneven and they may provide too little context to allow the re-use of the recorded task in a second evaluation phase. Another drawback of diary studies is that participants may log tasks unrelated to what is expected from them.

We tried to minimize the negative impacts of the diary study in several ways. First, we proposed that participants report only about 10 tasks. This amount was chosen because we thought it would not be a burden for participants, while providing enough usable tasks for the second phase of the evaluation. However, we put no pressure on participants, telling them that any amount of tasks they would agree to report would be useful for us. Furthermore, we provided participants with both a paper-based and a digital form. This allowed them to choose the one they preferred for reporting facts. In order to maintain their dedication, we also wanted the form to be as easy as possible to fill out. We did not want to force participants into evaluating several contextual cues about the re-finding task they were reporting. Thus, we only asked them to provide a written description of the task and to rate its subjective difficulty.

However, the quality of the written description of the task was critical for the task to be re-usable in the second phase. To achieve this quality, we provided examples of different recorded tasks. Example tasks contained descriptions of everything that was remembered about the information being looked for. This included contextual cues relevant for the second phase of the evaluation. A more precise description of the instructions provided for the diary study is given below.

Instructions for the Diary. We provided participants with a document explaining what we expected from them during the diary study (cf. Appendix C). More specifically, the document described that the diary is a journal of their re-finding tasks, and that they should log some of them, each day of the week when the study takes place. We specified that they should also write down the context they remembered about the information (for example its type, approximate time, or people concerned) and the reason why they needed it. Examples of logged tasks were provided. They were also asked to also rate the *a posteriori* difficulty of the task, i.e. the subjective difficulty of the task after it has been carried out, on a scale from 1 (very easy) to 5 (very hard). Of course, no mention was made in this document that the tasks were going to be re-used in the second phase of evaluation, to avoid possible bias when the participants decided whether or not to write down a particular task.

Difference from Elsweiler's Diary Protocol. Our protocol for the diary study differs from the one used by (Elsweiler and Ruthven, 2007) in that we do not ask the participants to rate the

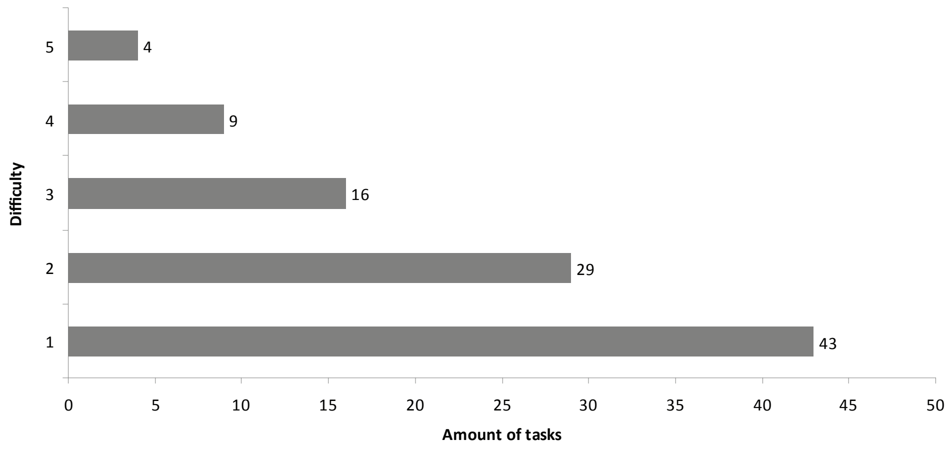

Figure 6.3: Rated difficulties of tasks issued from the diary study.

"temperature" of the sought-after information, i.e. the last time the information item was accessed. With the help this measure, Elsweiler shows that: (1) people mainly search for recent information, and need older information less and less as time passes and (2) the length of time between accessing and re-accessing the sought-after information influences the perceived difficulty of the task. We take these results as trustworthy. However, we chose not to include this measure, because we thought that its usefulness for our particular goals is eclipsed by the extra burden put on the participants when logging their re-finding tasks.

6.3.4 Results from the Diary Study

101 tasks (coded T1 to T101) were recorded by participants in their diaries. Their difficulty covered the full range of possible values, from 1 ("very easy") to 5 ("very difficult"). Figure 6.3 shows the distribution of difficulties. Most reported tasks were rated as easy.

However, the classification of tasks according to their subjective difficulty was insufficient to make sense of participant behaviour in the second phase of the evaluation. We therefore classified tasks along other dimensions. The first dimension we used was inspired by Elsweiler and Ruthven (2007). It classifies tasks according to their **target**. Possible targets are:

- *lookup* (L). The sought-after information is contained within a specific ressource (e.g. "Looking for the exact title of a song in a file"[2]).
- *item* (I). The sought-after information consists of a specific document (e.g. "I'm looking for the initial presentation of project XYZ"[3]).

[2] "Recherche du titre exact d'un chant dans un fichier" - P7.
[3] "Je cherche la présentation initiale pour le Projet XYZ" - P3.

- *multi-items* (M). Task that involves finding many information items in order to complete the task (e.g. "I'm looking for all the documents concerning the project ABC from past years"[4]).
- *unclassified* (U). Tasks for which the target does not belong to the aforementioned categories is set as unclassified. It can be the case when the description of the task is too vague to classify (e.g. "Opening and searching the opened document"[5]) or when the task itself is unrelated to re-finding (e.g. "I had to change my outlook password: I clicked on options, change password" - P2). Tasks that belong to this category are excluded from the future stages of the evaluation.

Tasks have been coded independently by the two experimenters along this dimension. The intercoder reliability (ICR) was 93% percent (94 tasks out of 101) according to Stempel's Percentage Agreement Index, which is more that the 80% generally acceptable for studies in social sciences. Among the 7 tasks for which the coders did not agree are such examples:

- "I had to search in the literature saved on my computer if there were articles concerning the subject of 'absenteism'" (P2): the task can be classified as a lookup task if we consider that the user wants to answer the question "do such articles exist?" or multi-items if we consider she needs the documents.
- "Looking for a PDF file containg images of molecules computed on a cluster"[6]: the task can be considered an item task if the target of the search is the whole PDF file, or a lookup task if he needs particular images contained within the file.

A second taxonomic dimension classifies tasks according to the **resource** they use:

- *mail*: tasks related to emails.
- *doc*: tasks related to files or documents.
- *cal*: tasks related to the calendar.
- *web*: tasks involving the web.
- *unknown*: none of the above.

Tasks can use several resources. Some tasks were thus classified as *mail-doc* (e.g. looking for a document attached to an email) or *web-doc* (e.g. looking for a document on the web). The two

[4] "Je cherche tous les documents concernant le projet ABC des années précédentes" - P3.
[5] "Ouverture et recherche du document ouvert" - P6.
[6] "Recherche d'un fichier PDF contenant des images de molécules calculées sur un cluster" - P9.

experimenters have assessed together the resources used by the tasks. They discussed conflicts and solved them.

The third and last taxonomic dimension we used to classify tasks was the dimension of **cues**. Cues are closely related to the choice of facets devised for the system to be evaluated in the second phase. For example, terms related to time in the description of a task (e.g. "May 2009") would classify the task as containing a *cue related to the temporal facet*. Categories of this dimension are not exclusive, thus a task can contain multiple cues. Possible cues are:

- *file explorer (FE)*: the description of the task contains an explicit mention of the location of the information item in the file hierarchy (including the desktop).
- *textual search (TS)*: the description of the task contains a mention of a term contained within the sought-after information item, or the name of the document.
- *type (T)*: the description of the task mentions the type of document (e.g. email, presentation, pdf, etc.).
- *social network (SN)*: the description of the task mentions the name of a person.
- *date (D)*: the description of the task mentions an absolute (e.g. "May 2009") or relative (e.g. "last week") date.

Cues have been assessed independently by the two evaluators on a sample of 45 tasks (ICR 80.1%). Disagreement was discussed and resolved between the coders, which led to further detailing of the coding rules for the textual search cue.

6.3.5 Tasks Selection

The selection of tasks for the second phase of the evaluation was done by the two evaluators together. They proceeded as follows. First, as the evaluated system only allows to re-find emails and documents, only tasks that had *mail*, *doc* or *mail-doc* resources were kept. They still constituted the vast majority of tasks (93 out of 101 tasks). From this pool of 93 tasks, tasks having an *unknown* target (8) were also eliminated. The remaining pool thus contained 85 tasks. Another 22 tasks were finally removed from the pool because the evaluators agreed that their descriptions were too vague for the task to be repeated. Such problematic descriptions of tasks included:

- "Looking for an email attachment but I don't remember in which mailbox I have received it. Outlook and Firefox."[7]

[7] "Recherche d'une pièce jointe d'un mail mais je ne sais pas dans quelle boîte mail je l'ai reçue. Outlook + Firefox." - P6.

Category	# tasks
I	30
LU	22
M	11
doc	43
mail	13
mail-doc	7
Diff. 1	28
Diff. 2	20
Diff. 3	9
Diff. 4	4
Diff. 5	2

Table 6.3: Count of usable tasks by category.

- "Looking for an email of which I know the sender, in Outlook, by date."[8]
- "I needed to reread an email I received about a month ago to my outlook: I typed the name of the sender into the search field and chose the inbox search option." - P2.

Sixty-three (63) tasks remained in the pool of usable tasks for the second phase of the study. Among them, 30 were item tasks, 22 lookup tasks and 11 multi-items tasks. The distribution of resources used was 43 doc, 13 mail and 7 mail-doc. Table 6.3 summarizes the usable tasks.

Among those 63 remaining tasks, we selected tasks by trying to optimize the following rules, in order of importance, for each participant:

1. Balance targets of tasks.
2. Balance resources of tasks.
3. Favour higher difficulties.
4. Have an equal amount of tasks per participant.

The application of these rules led to the selection of 41 tasks. Each participant would have 4 or 5 tasks to perform. Additionally, a task was artificially created for P6 who had only 3 tasks available. We used our knowledge of the participant's PSI to devise this task, following the guidelines presented by Elsweiler and Ruthven (2007). Artificially created tasks were coded T100x to differentiate them from real tasks. They were not assigned a difficulty rating. Moreover, an issue during the evaluation led us to create 3 other artificial tasks for P2 to replace 3 previously selected tasks for

[8] "Recherche d'un mail dont je connais l'expéditeur, dans Outlook, par date." - P6.

	doc			mail			mail-doc			
	I	LU	M	I	LU	M	I	LU	M	**Total**
P1	1 (1)	1 (1)			3 (2)				1 (1)	6 (5)
P2	1 (0+1)	1 (1)		0 (0+1)			1 (1)		1 (0)	4 (2+2)
P3	2 (1)	1 (1)	1 (1)		5 (2)					9 (5)
P4	5 (3)			1 (1)						6 (4)
P5	2 (1+1)	2 (2)	1 (1)							5 (4+1)
P6	1 (1+1)		1 (1)		1 (1)					3 (3+1)
P7	4 (3)	3 (1)								7 (4)
P8	3 (3)		4 (1)						1 (1)	8 (5)
P9	4 (3)							1 (1)		5 (4)
P10	3 (2)	1 (0)	1 (1)	1 (1)	2 (1)		1 (0)	1 (0)		10 (5)
Total	26 (18+3)	9 (6)	8 (5)	2 (2+1)	11 (6)		2 (1)	2 (1)	3 (2)	63 (41+4)

Table 6.4: Summary of task selection per participant and category of task. Regular numbers are usable tasks; numbers inside parenthesis are selected tasks. Artificially created tasks are preceeded by a plus sign (+).

her. Indeed, P2 filled in the diary study with tasks related to personal documents. However, she came to the evaluation with her professional laptop. Most of the personal documents mentioned in the diary study were not available on this laptop. Consequently we could use only 2 tasks from the diary study and had to use 3 created tasks. All in all, there were 45 tasks actually performed by the participants in the second phase of the evaluation. Table 6.4 presents the tasks selected for each participant. The numbers without parenthesis are the count of usable tasks. The numbers in parenthesis are the count of selected tasks. Artificially created tasks are preceeded by a plus sign (+). Table 6.5 summarizes the count of tasks per target, resource and difficulty. Artificially created tasks have no difficulty rating.

The last step was the reformulating of tasks. Indeed, many tasks from the diaries contained descriptions of the way people performed the task. These descriptions were removed altogether. An example of such tasks is:

> "I look for an old email from a professor who has sent us literature for the exam. I browse the names in my mailbox and I look according to the approximate date in order to find the right email of this professor concerning the sought-after documents. Then, I order the literature in my personal folder concerning this class."[9]

[9]"Je cherche un ancien e-mail d'un professeur qui nous a envoyé de la littérature concernant l'examen. Je parcours les noms de ma boîte e-mail et je cherche selon la date approximative afin de trouver le bon mail de ce professeur contenant les documents recherchés. Je classe ensuite la littérature dans mon dossier personnel concernant ce cours." - P1.

Category	# tasks
I	25
LU	13
M	7
doc	31
mail	9
mail-doc	5
Diff. 1	14
Diff. 2	12
Diff. 3	8
Diff. 4	4
Diff. 5	2
Diff N/A	4

Table 6.5: Count of selected tasks by category. Artificially created tasks have no difficulty rating.

This task — which was rated as a *multi-items* task involving the *mail* and *doc* resources — was reformulated as follows:

"Find the literature that a professor sent you by email for the exam."

The experimenters took care to keep the reformulated task consistent with the classification of the original task (i.e. having the same target, resources and cues).

6.3.6 Field Evaluation

Procedure

The field evaluation took place at each participant's office, on average 4 weeks after the distributions of the diaries. Only the participant and one evaluator were present. Each session lasted between 1 hour and 1 hour and a half depending on the time taken by the system to build the social network and by the time taken by the participant to fulfill the tasks and the questionnaires. The evaluation protocol comprised several steps:

1. Environment checking: the evaluator checks that all required software and virtual machines are available and properly configured (Java, Windows Search, Outlook or a similar email client, etc.)
2. *WotanEye* configuration: the evaluator sets the address of the IMAP server containing the participant's mailbox.
3. *WotanEye* execution: the system extracts the social network of the participant from her IMAP mail archive.

4. *WotanEye* communities handling: the participant uses the *Mailviz* interface to reorganize her contacts and name her communities. The interaction is not logged by the system. The participant is asked to think aloud. The evaluator takes notes and offers support if needed.

5. *Weena* configuration: the evaluator configures (1) additional file types that the participant works with (e.g. differentiate between several formats of images for a art designer) (used within the type facet) and (2) shortcuts to specific folders on participant's demand (used within the filesystem facet). The configuration is done in two separate XML files.

6. *Weena* execution: the evaluator launches *Weena*.

7. Overview of *Weena*: the participant reads the document presenting the interface (see Appendix C.2).

8. *Weena* guided tour: the evaluator guides the participant in using the interface, simulating a re-finding task. All facets and features of the interface are presented according to a defined protocol.

9. *Weena* free use: the participant is free to use the system for a couple of minutes. The evaluator guides her in case of problems.

10. *Weena* actual evaluation: the participant is asked to perform the tasks selected for her from her diary. Task order is randomized according to the difficulty rating obtained from the diary study. A widget included in the interface is used to trigger the beginning and end of each task. The interaction of the participant with the system is logged during this phase. Additionally, the participant is asked to think aloud and the evaluator takes notes, without providing any further information. At the end of the task, a dialog box asks for the success and perceived difficulty of the task.

11. Additional questions: the evaluator then asks the participant four questions related to the self-analysis of PIM (see Section 6.3.8, on page 142). These are:

    - "Please summarize what you did last week."
    - "Please summarize what you did last month."
    - "With whom have you had the most to do in October?"
    - "Now that you know the system, do you have any idea of tasks that you could do with it? If yes, please formulate them and perform them."

12. Anonymization check: the evaluator presents the logs and automatic screenshots to the participant. The participant gives her consent to keep the data.

13. Questionnaires: the participant is asked to fill in two questionnaires (see Appendices C.3 and C.4).

A few more configuration steps were also planned to optimize performance if needed. For example, it was possible to select a subset of email folders instead of the full mailbox if the social network computation took too much time. The visual settings could also be adjusted. The full protocol is available as an appendix (see Appendix C.5).

6.3.7 Collected Data

Data was logged automatically while the participant was performing the task. Figure 6.4 gives an overview of the data collected. Besides general task information (identifier, completion time, difficulty from the diary study and perceived difficulty after performing the task with the system), the sequence of facet selection was also logged in detail along with the amount of remaining documents after each selection. At the end of the task, a screenshot was automatically taken in an anonymized fashion (Figure 6.5). In particular, keywords in the textual search box, folder names and file names in the results panel were replaced by ###. Names of contacts and communities in the social network facet were replaced by random names obtained from a seed-randomizer.

In addition to the automatic logging of user's interactions with the system, the evaluator also took notes of what the participant was saying and how she interacted with the system to reach her goal. Furthermore, the interaction of participants with the system for the additional questions (point 11 of the protocol above) could not be logged for technical reasons. Only the notes of the evaluator are thus available for this step.

The last data available to make sense of the field study are the questionnaires filled in by participants. The first questionnaire was about the general use of *Weena* and faceted navigation in general (see Appendix C.3) and the second questionnaire was centered more particularly on the social facet (see Appendix C.4).

6.3.8 Results

Table 6.6 presents general information about the participants of the evaluation and the amount of items and contacts that appeared in *Weena* at the time of the evaluation. The detailed results of the evaluation are given in Table 6.7. The rest of this section gives a thorough interpretation of them.

```
+++++++++++++++++++++++++++++++++++++++++++++++++++++++
> Summary for task: 3
+++++++++++++++++++++++++++++++++++++++++++++++++++++++
--------- Task infos ---------
ID                                  3
Diary Complexity                    5
Complexity                          3
Completion time                     119313 ms
Successfully done                   yes
----------- Facets -----------
First selected facet                Type
Last selected facet                 Social Network
Max selected facet                  Type (selected 1 times)
--------- Selections ---------
Max item selected                   1
Min item selected                   1
Mean item selected                  1.0
Variation item selected             0.0
----------- Results ----------
First results filtered              990
Last results filtered               9
--------- Operations ---------
Total Number of operations          2
Number of operations to goal        2
  (Undos: 0 | Redos: 0 | Selections: 2)

********************** Details ************************
Facets selected
 Type - selections: 1 / results: 990
 Social Network - selections: 1 / results: 9
Operations
 selection
 selection
+++++++++++++++++++++++++++++++++++++++++++++++++++++++
```

Figure 6.4: Usage data for task T3 by P1, generated by the automatic logger.

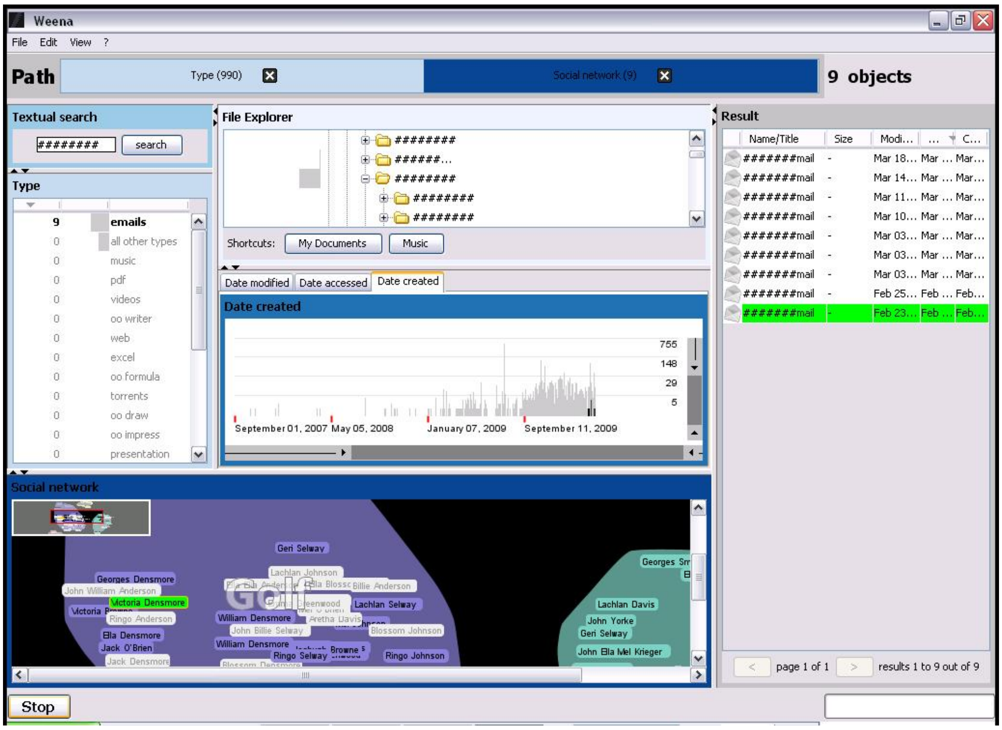

Figure 6.5: Anonymous screenshot automatically taken at the end of task T3 by P1.

Id	Gender	Age	Profession	Items	Contacts
P1	F	23	MA Student (Psychology)	2552	70
P2	F	26	MA Student (Psychology)	1024	37
P3	M	29	PhD Student (CS)	100078	138
P4	F	46	Secretary	22427	95
P5	M	32	PhD Student (CS)	2979	144
P6	F	29	Polygraph	10796	92
P7	M	28	Musician	20742	73
P8	F	18	High-School Student	5050	104
P9	M	24	MSc student (Chemistry)	8031	116
P10	M	29	Project leader (Admin)	22567	197
Mean		28.40		19624.60	106.60
Median		28.50		9413.50	99.50
Stdev		7.35		29495.42	45.11
IQR		4.50		18509.00	54.75

Table 6.6: General information about the participants and count of their items and contacts available in *Weena* at the time of the field evaluation.

Participant	Task ID	Diff. diary	Category	Resource	TS cue	T cue	SN cue	D cue	FS cue	Diff. field	Completion time (s)	Operations
P1	T2	2	I	doc		1				1	82	F
P1	T3	5	LU	mail		1	1			3	119	TN
P1	T4	3	M	mail-doc		1	1			5	364	TSUNSUC
P1	T5	3	LU	mail			1			1	49	N
P1	T8	4	LU	doc		1			1	1	69	F
P2	T13	2	LU	doc	1				1	1	94	TFU
P3	T22	2	LU	mail		1				1	213	NNUN
P3	T24	2	LU	doc						1	148	FFT
P3	T27	2	LU	mail		1	1			3	408	NUNNNNNNUN
P3	T28	3	M	doc		1				1	161	FF
P3	T29	1	I	doc		1	1	1		3	343	NSUMMMMSUNNUUNNN
P4	T34	4	I	doc	1	1			1	2	142	FTS
P4	T35	1	I	mail		1	1			2	149	TUTS
P4	T36	4	I	doc		1	1	1	1	2	148	FSFM
P4	T39	1	I	doc		1				2	93	TFS
P6	T42	3	LU	mail	1	1				3	354	SNNNUUS
P6	T44	3	M	doc		1				1	118	TTS
P6	T47	1	I	doc			1			1	153	TF
P7	T50	1	I	doc		1		1		1	132	ST
P7	T54	2	I	doc						2	186	SFUUFS
P7	T55	2	I	doc						2	75	SUS
P7	T59	1	LU	doc						1	71	S
P8	T63	2	I	doc					1	1	71	S
P8	T64	1	I	doc	1	1			1	1	47	SS
P8	T69	3	I	doc		1				2	187	FTT
P8	T71	2	M	mail-doc		1	1			1	78	N
P9	T72	2	I	doc	1					2	127	F
P9	T73	1	I	doc		1				1	56	TF
P9	T80	2	I	doc		1				1	78	TFS
P10	T82	3	LU	mail	1	1				4	359	STUSUSTM
P10	T84	1	I	mail		1	1			2	82	SST
P10	T89	4	I	doc	1					3	161	SF
P10	T90	1	I	doc		1		1		1	72	SSS
P5	T95	1	LU	doc	1				1	1	46	TTS
P5	T100	5	LU	doc						2	121	TFFFFFFFFF
P5	T101	1	M	doc	1				1	1	141	TSUUS
P6	T1001		M	doc				1		2	86	FT
P2	T1002		I	mail		1				2	187	NTN
P5	T1003		I	doc		1				1	36	TS
P2	T1005		I	mail-doc		1				2	115	TSUUT

Table 6.7: Detailed data for the 40 successful tasks. Operations are coded using one letter per operation: T for the type facet; F for the file explorer facet; S for the textual search facet; N for the social network facet; M for the date modified facet; C for the date created facet; U for undo. Each single letter represents one selection in a facet or one step back using undo.

General Observations

Forty (40) tasks out of 45 were successful. The reasons for the failure of the 5 remaining tasks were external to the system evaluated (e.g. looking for a file mentioned in the diary study but deleted since then, or a file not indexed by Windows Search, etc.). One of the failed task was a simulated task, whereas the 3 remaining simulated task were successfully performed. In the remainder of this section, we consider only the 40 successful tasks and the related collected data. A total of 147 operations (selection in a facet, undo or redo) were performed by participants, which makes an average of 3.675 operations per task (min: 1, max: 16, see Figure 6.6 for a graphical view of the distribution); only 4 tasks used more than 7 selections; 22 undos were logged.

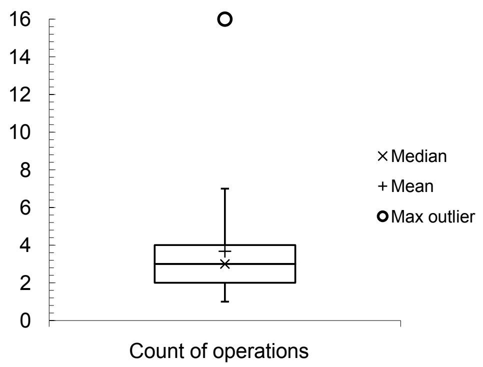

Figure 6.6: Number of operations per task.

A comparison of the difficulty rating given during the diary phase versus the rating given during the evaluation phase is presented in Table 6.8. Out of the 37 real tasks that were successful, 12 tasks received the same difficulty rating during the diary phase and after their realization with the system (yellow background). Seven (7) tasks were rated as more difficult to fulfill with the system (red background). Eighteen (18) tasks were rated as easier with the system (green background). Interestingly, tasks that were perceived as difficult to perform with the system were mostly easy tasks in the diary. On the contrary, tasks perceived as difficult in the diary phase were rated as easier to perform with the system. However it is unclear whether the system is better suited for this kind of task or whether the initial difficulty of the task contributed to a better recollection of the sought-after information, which would have negatively influenced the perceived difficulty of the task during the second phase.

	Difficulty - Field evaluation					
Difficulty - Diary study	very easy	easy	average	difficult	very difficult	**Total**
very easy	8	3	1			12
easy	8	3	1			12
average	3	1	1	1	1	7
difficult	1	2	1			4
very difficult		1	1			2
Total	20	10	5	1	1	37

Table 6.8: Difficulty comparisons of tasks during the diary phase and during the field evaluation phase.

Use of Facets

Figure 6.7 shows the number of selections per facet. The textual search was clearly the most used facet with 35 selections. The file explorer (29), type (28) and social facet (26) were close behind. However, field notes taken by the evaluators reveal that the high amount of selections within the social facet may be a consequence of the issues some participants faced in order to select the proper contact or group of contacts in the social network interface. The date facets were clearly less used (7 selections total). Note that the 'date accessed' facet was never used. This may be due to the fact that the 'date accessed' facet is particularly useful for recent documents and that documents are automatically sorted by date of last access in the result panel, which reduces the usefulness of this facet.

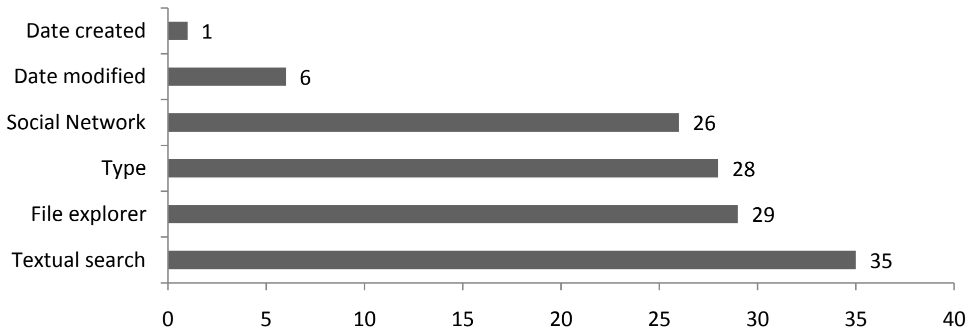

Figure 6.7: Number of selections per facet.

The number of tasks involving each facet is shown in Figure 6.8. The facets are represented on the vertical axis in the same order as in Figure 6.7 to facilitate comparison. The textual search (23), file explorer (17) and type (22) facet were involved in the most tasks. The textual search and type

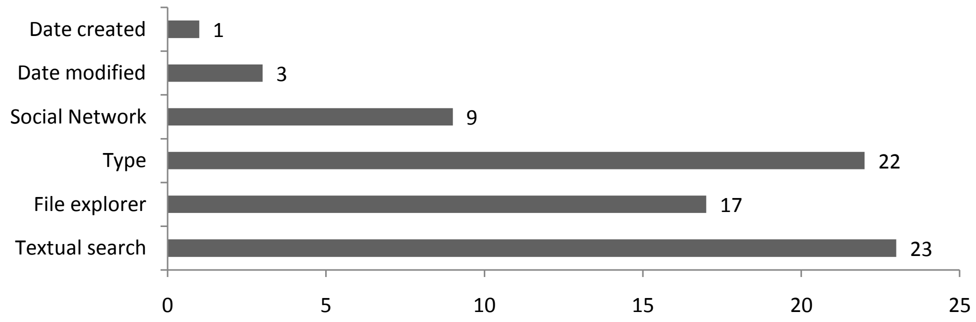

Figure 6.8: Number of tasks involving each facet.

	-	FE	T	TS	SN	DM
File explorer (FE)	4					
Type (T)		11				
Textual search (TS)	5	7	15			
Social network (SN)	4		3	3		
Date modified (DM)		1	1	3	1	
Date created (DC)			1	1	1	

Table 6.9: Number of tasks performed for each pair of facets.

facets were used for more than half of the successful tasks. Moreover, several observations can be made by comparing the barcharts of Figure 6.7 and Figure 6.8. When textual search was used, on average 1.52 textual searches were made. The use of the file explorer required an average of 1.71 selections. The use of the 'type' facet required less selections on average (1.27) which may indicate a better affordance of the facet. The 'date modified' facet required 2 selections on average and the 'date created' was used only once. The most striking comparison involves the social facet. It was used for only 9 tasks, but required 26 selections. This tends to indicate a bad affordance of the facet which requires several selections (average 2.89 per task using it) to reach the goal.

Table 6.9 presents the number of tasks performed for each pair of facets. The first column, labeled '-' indicates the amount of times that the facet was used in isolation. The 'type' facet (T) and the 'date' facets (DM and DC) were never used in isolation to fulfill a task but were always combined with other facets. Both 'date' facets were always used in combination with textual search (TS). Similarly, more than half of the tasks involving the 'type' facet also involved textual search (15 tasks out of 22, i.e. 68.18%). This is the highest number of facet combinations. The 'file explorer' facet (FE) was mainly used in combination with the 'type' facet (T) (11 tasks out of 17 involving it, i.e. 64.7%), but also textual search (TS) (7 tasks out of 17, i.e. 41.17%). The social network (SN) was used mainly in isolation (4 tasks out of 9, i.e. 44.45%) with success.

A statistical correlation analysis of these data also yielded interesting results[10]. The significant negative correlation between the file explorer and the social network ($\Phi = -0.4632$; $\chi^2 = 8.5835$; df$=1$; $p < 0.01$) was expected: as the social network is used mainly in the case of emails and the file explorer mainly in the case of documents other than emails, they are likely not to be used together. The significant negative correlation between the file explorer and text search ($\Phi = -0.3303$; $\chi^2 = 4.3649$; df$=1$; $p < 0.05$) is more interesting. It means that people tend to use either search or the file explorer, but not both of them together. The correlation is low, however. Similarly, the joint use of search and the social network ($\Phi = -0.2933$; $\chi^2 = 3.4409$; df$=1$; $p > 0.05$) could indicate that people would use either search or the social network, but not both of them together. However, further observations would be needed, as the dependency is not significant in this case ($p > 0.05$). The use of the textual search and date facets together may be very lowly correlated ($\Phi = 0.2722$; $\chi^2 = 2.9630$; df$=1$; $p > 0.05$), but again it is not significant. Other correlations of pairwise facet use are not significant.

Use of Facets with Respect to Targets and Resources

Table 6.10 presents the use rate of facets with respect to targets and resources of tasks. As was already presented before, the 'textual search' (TS) and the 'type' (T) facets were the most used, with 58% (23 tasks) and 55% (22 tasks) of all successful tasks using them. The 'file explorer' facet (FE) was used for 43% of the tasks. The social network (SN) was used for one quarter of the tasks and the temporal facets (D) are used for 10% of the tasks.

The file explorer (FE) was used exclusively in the case of file-related tasks. Its use rate in this context was 61% (17 tasks). This result was expected, as people are used to navigating their filesystem hierarchy using a file-explorer approach. 'Textual search' (TS) and 'type' (T) facets were used in an homogeneous way for about half the tasks, no matter which resource was involved. The use of the social facet was predominant for emails. In particular, lookup tasks in emails led to the use of the social facet in 83% of the cases (5 tasks out of 6). It scored better than all other facets in this case. In tasks with a mail resource, the use of the type facet was also higher. This can be explained by the observations of the evaluators. They noted that people often began their re-finding tasks involving emails by selecting the *email* type. There were too few tasks involving the mail-doc resources to formulate reliable conclusions. Nevertheless, the social facet seemed to also be favoured in this context (67% i.e. 2 tasks out of 3). The temporal facets were rarely used. However, notes taken by the evaluators in the field seem to show that their use was always effective

[10] I performed a Pearson χ^2 test of independence for 2×2 contingency tables. When the two variables were significantly dependent (with probability $p < 0.05$ or less), I computed the Φ coefficient of their correlation. The value of Φ may vary between -1 (perfect negative correlation) and 1 (perfect positive correlation). A value of 0 means that the variables are not correlated.

	Facets use rate (%)					
	FE	*TS*	*T*	*SN*	*D*	***Tasks amount***
mail	-	*44*	*56*	*67*	*11*	***9***
LU	-	33	33	83	17	6
I	-	67	100	33	-	3
M	-	-	-	-	-	0
doc	*61*	*61*	*54*	*04*	*07*	***28***
LU	67	33	67	-	-	6
I	61	72	44	06	11	18
M	50	50	75	-	-	4
mail-doc	-	*67*	*67*	*67*	*33*	***3***
LU	-	-	-	-	-	0
I	-	100	100	-	-	1
M	-	50	50	100	50	2
Total	**43**	**58**	**55**	**23**	**10**	***40***

Table 6.10: Use rate of facets according to the targets and resources of tasks. The rightmost column indicates the amount of successful tasks in each category.

when the participant used them. Moreover, the logged data indicates that a selection in one of the temporal facets was never undone.

I additionally performed a statistical correlation test between the targets and resources of tasks and the use of facets. Some significant correlations were observed. *Item* tasks are lowly negatively correlated with the use of the social facet ($\Phi = -0.3550$; $\chi^2 = 5.0411$; df= 1; $p < 0.05$). Thus, *item* tasks probably will not lead to the use of the social facet. However, the correlation value is low, which makes this conclusion uncertain. The significances of *multi-items* and *lookup* tasks' correlations with facet use are not sufficient ($p > 0.05$). As far as resources are concerned, the correlations confirm the observations made in the previous paragraph. The *doc* resource is positively correlated to the use of the file explorer ($\Phi = 0.5628$; $\chi^2 = 12.6708$; df= 1; $p < 0.001$) and is negatively correlated to the use of the social facet ($\Phi = -0.6924$; $\chi^2 = 19.1773$; df= 1; $p < 0.001$). Conversely, the *mail* resource is negatively correlated to the use of the 'file explorer' facet ($\Phi = -0.4632$; $\chi^2 = 8.5835$; df= 1; $p < 0.01$) and positively correlated to the use of the social facet ($\Phi = 0.5699$; $\chi^2 = 12.9911$; df= 1; $p < 0.001$). Those two observations confirm our previous analysis about the use of the 'file explorer' (FE) and 'social network' (SN) facets. Other correlations of resources with facet use are not significant.

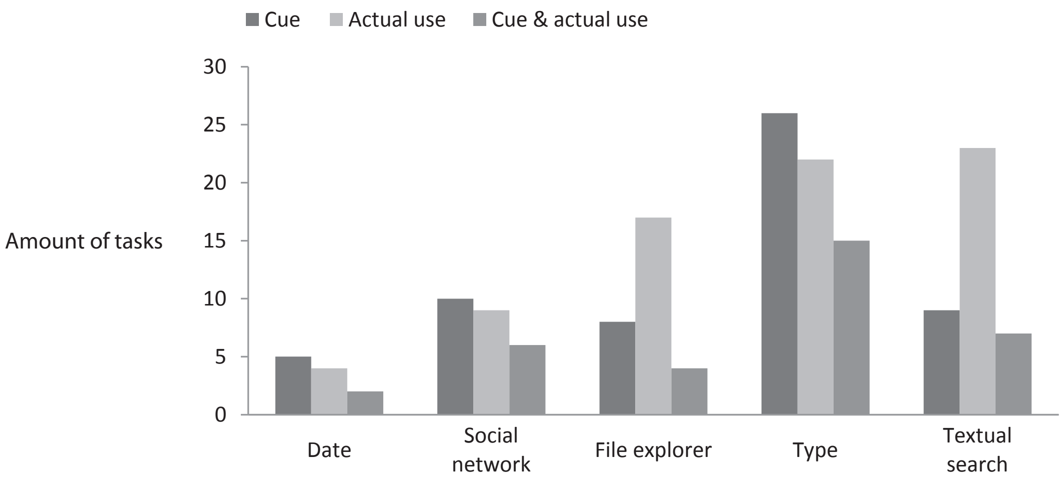

	Date	Social network	File explorer	Type	Textual search
Cue	5	10	8	26	9
Actual use	4	9	17	22	23
Cue & actual use	2	6	4	15	7

Facets

Figure 6.9: Number of tasks : (1) having cues related to each facet, (2) having triggered an actual use of the facet, and (3) having a cue and having triggered an actual use of the facet.

Use of Facets with Respect to Cues

Figure 6.9 compares: (1) the amount of tasks having a cue relating to each facet (dark gray), (2) the amount of tasks leading to an actual use of the facets (light gray), and (3) the amount of tasks that both have a cue and have involved the facet while completing them (middle gray). What stands out first is that the 'file explorer' and 'textual search' facets were used more than what could be expected from the cues in the task descriptions. All other facets were used less often than the cues would have suggested. It is also interesting to examine the proportion of tasks having a cue which led to a use of the facet. The textual search has the highest value: 77.78% of tasks containing a 'textual search' cue were actually performed using the 'textual search' facet (7 tasks out of 9). Tasks having cues related to the social network and the 'type' facet come next: 60% (6 out of 10) of tasks having a social cue were done using the social facet, and 57.69% (15 out of 26) having a 'type' cue were carried out using the 'type' facet. Finally, 50% (4 out of 8) of tasks containing a 'file explorer' cue were executed using the 'file explorer' facet and 40% (2 out of 5) of tasks having a temporal cue led to a use of the temporal facets.

Table 6.11 further compares the use of facets with the presence of cues in the task description. The values indicate the portion of all tasks involving a facet where (1) the task description contained the corresponding cue and (2) the task description did not contain the corresponding cue. The

	FE	*TS*	*T*	*SN*	*D*
Facet use rate with cue (%)	24	30	**68**	**67**	50
Facet use rate w/o cue (%)	**76**	70	32	33	50
Number of tasks	17	23	22	9	4

Table 6.11: Use rate of facets with respect to corresponding cues in the initial description of the task. The final row presents the number of tasks using the facet.

‘type’ (T) and ‘social network’ (SN) facets were mainly used when a cue was present explicitly in the task description. To the contrary, the ‘file explorer’ (FE) and ‘textual search’ (TS) facets were used most often without a corresponding cue in the task description. We tentatively explain this observation as follows. Participants use textual search facilities and the file explorer on a daily basis. It is their usual way of re-finding information and they favour their use. On the other hand, they are less experienced in the use of the ‘type’ and ‘social network’ facets to re-find information. Thus, they do not think of using them without an explicit cue in the task description. The temporal facet was used half of the time with a cue and half the time without. However, the low amount of tasks carried out with the temporal facet makes this conclusion uncertain.

I also carried out a statistical correlation analysis of these data, observing the correlation between the use of the facet and the presence of a corresponding cue in the task description. A very low correlation might exist for ‘type’ ($\Phi = 0.0737$), for ‘file explorer’ ($\Phi = 0.0758$), for ‘date’ ($\Phi = 0.1667$), and for ‘text search’ ($\Phi = 0.2806$) but these values are not significant ($\chi^2 < 4$; df$= 1$; $p > 0.05$). Nevertheless, a significant correlation does exist with the social facet ($\Phi = 0.5184$; $\chi^2 = 10.7527$; df$= 1$; $p < 0.01$). Therefore, the observation made in the previous paragraph that the use of the social facet may be triggered by the presence of a cue in the task description is confirmed by this statistical analysis. The other interpretations we made in the previous paragraph would need further empirical evidence to be generalized with confidence.

Self-analysis and Reporting Tasks

Apart from the tasks coming from their diaries, participants were asked additional general questions by the evaluators which aimed at exploring the possible use of *Weena* for gaining insight into one’s own PSI, PIM strategies and other metalevel activities. Those questions are mentioned in the protocol of the field evaluation described above in Section 6.3.6. They were:

- "Please summarize what you did last week."
- "Please summarize what you did last month."
- "With whom have you had the most to do in October?"

- "Now that you know the system, do you have any idea of tasks that you could do with it? If yes, please formulate them and perform them."

The user activity with the system during the realization of these tasks could not be logged for technical reasons. The participant was asked to talk aloud and comment on her interaction with the system. To answer the first two questions, eight participants selected the relevant period in the 'date modified' facet. One participant answered without using the system (P2) at all, relying only on her memory. Another participant sorted the elements by date in the result panel and used this to remember her activities (P8), navigating through pages of results. Some participants observed the folders that were highlighted in the file explorer after the selection in the temporal facet (P1, P5, P7, P10). Some looked at the results (P1, P4, P8, P9). P4 additionally selected in turn the types of files highlighted and looked at the documents displayed in the results panel.

To answer the third question, all participants but one (P8) selected the month of October in the date modified facet. They then looked at their social network to see who was highlighted, except P9 who selected the email type and looked at the email titles in the result panel to answer, and P3, who looked at the type facet, the file explorer and the result to answer. Furthermore, P4, P6 and P7 additionally selected the email type in the type facet before looking at their social network. P8 wanted to select all contacts in his social network but eventually gave up and skipped the question.

The last question received few answers. P1 and P10 liked the social network view and found that it would be interesting to see how people relate to activities, by selecting folders related to projects, or using keyword searches. They had no precise idea of tasks to perform at the time of the interview, though. P7 mentioned that the system could be useful to help maintain the filesystem hierarchy, by providing an easy way to select files of the same type in order to classify them in a single folder. He also would use the system to create archives of files based on their last modification date and find too large files in order to delete them. P9 would use the amount of documents by type to decide whether to keep or delete a specific software that creates documents of this type. Similarly, he would select a type and see if documents were created recently. If it is not the case, this implies that the software is no longer needed and can be removed. Finally, P10 would see the system being used to help with reporting tasks. He could imagine a scenario where the system is connected to a versioning tool like SVN in order to visualize which parts of the hierarchy were modified with respect to time.

Results from the Questionnaires

This section presents the results of the questionnaires that the participants filled in after the controlled study. The questionnaires contained both quantitative and open questions. They are available in Appendices C.3 and C.4.

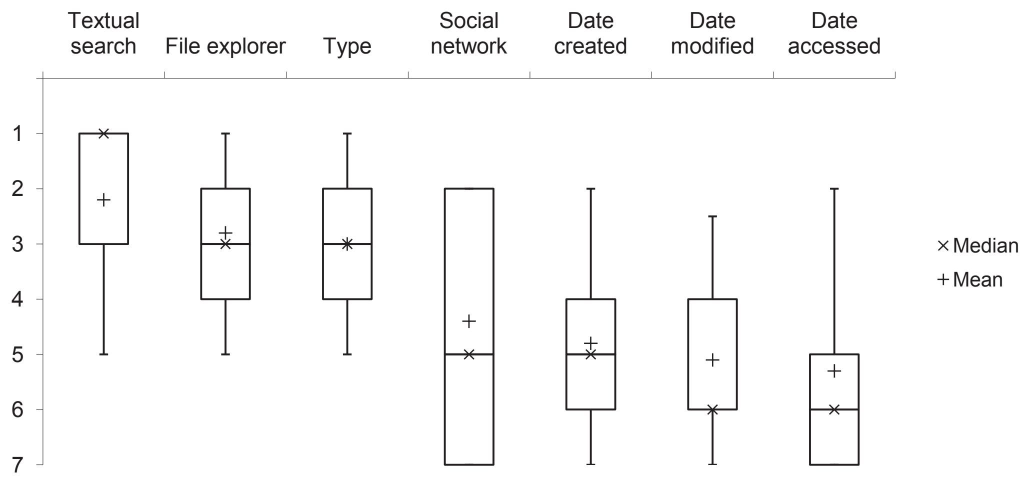

Figure 6.10: Perceived preference of facets.

Perceived Preference and Efficiency of Facets. Participants were asked to rank the facets by subjective preference (1 — favourite — to 7) and perceived efficiency (1 — most efficient — to 7). Figure 6.10 shows a box-and-whisker plot of the perceived preference rating. Facets are ordered by decreasing median value first, and by decreasing mean value second. The most common facets are also the preferred ones, namely the textual search and file explorer. The social facet received the most variable marks as its interquartile range (IQR) of 5 indicates. Indeed, some participants (P6, P7, P10) had troubles using it and thus gave it a low mark whereas others liked it very much (P3, P8). Temporal facets are the least liked ones. Moreover, a comparison of the perceived preference with the actual use of the facets obtained from the automatic logs (see Figure 6.7) reveals that the preference seems to be closely related to the actual use of the facets.

Figure 6.11 shows the distribution of participants' ranking of facets according to their perceived efficiency. The file explorer (median 2, IQR 2), textual search (median 2, IQR 3) and type (median 2.5, IQR 2) facets have been perceived as the most efficient. The social and temporal facets were perceived as generally less efficient.

Overall User Satisfaction. The system was rated as rather "useful" (median 4, IQR 1.75) (scale of 1 to 5) and "usable" (median 4, IQR 0) (scale of 1 to 5) by the participants (see Figure 6.12).

Open Questions. Open questions revealed that participants especially liked to be able to combine research criteria (P4, P6, P7, P8, P9) and the unification of all PI in a single interface (P1, P2, P3, P8). Some perceived the system as faster and more efficient (P1, P8, P10). They also ex-

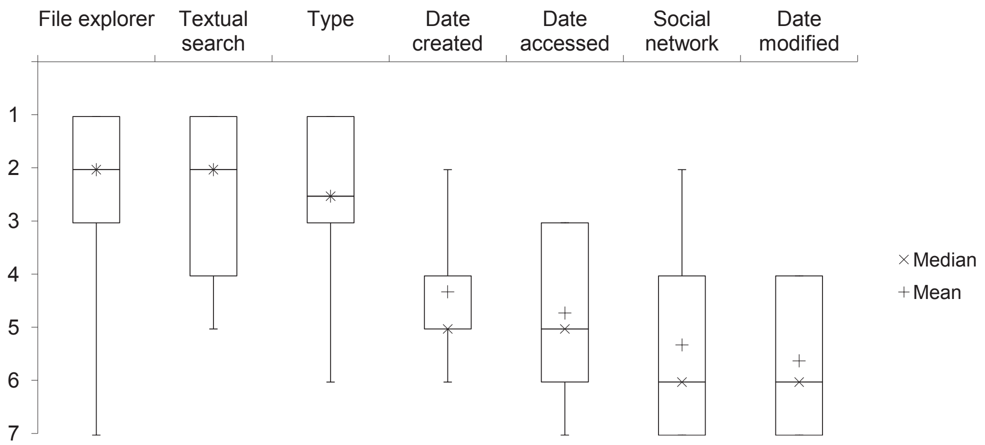

Figure 6.11: Perceived efficiency of facets.

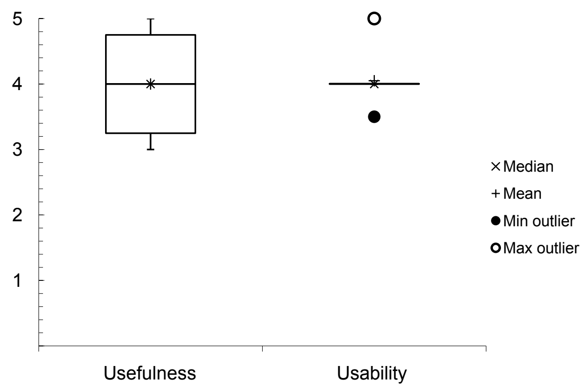

Figure 6.12: User rating of the system's usefulness and usability

plicitly mentioned the 'type' facet as an useful feature (P5, P7). Miscellaneous appreciated features included: having multiple views of the same data (facets) (P2, P5), the fact that it was visually appealing (P2, P4), graphical views and a responsive interface (P7), shortcuts and the ability to sort results (P4) and the ability to open files and containing folders (P10).

Disliked features included: technical issues (installation, speed, screen resolution, image previews bugs, emails not indexed correctly) (P1, P4, P7, P8, P9), an overly restrictive textual search which works only on file content and file names (P3) and needs explicit boolean operators to connect keywords (P10). The interface as a whole was found to be not intuitive enough (P6, P7). Some visual features were also criticized: the whole look of the interface (P6), the layout of facets (P3), the large number of items on the screen (P1), highlight colors which were different from usual ones (P2), the highlight in the file explorer (subfolders and not current folder) (P5). Finally, P3 did not understand the link between the files and social network.

Participants said that they would use *Weena* in order to search for documents (P1, P2, P3, P4, P8, P9) or more specifically documents attached to emails (P10). They would also use it especially if the search is difficult and other means do not work (P5, P6, P7). Additionally, P3 mentioned that he would use it to search for documents from a certain period of time and to organize and tidy up his PI.

Temporal Facet Use. Participants said that they would use the temporal facets mainly when they look for a document from which they remember the date or period (P1, P2, P6, P7, P10) or to report on what they did and what they used during a certain period of time (P3, P8, P9). Interestingly, P6 would use it at the beginning of a search for which she remembers few cues, to do a first filtering, benefitting from the filtering that occurs in the other facets. P3 would use it to find the last version of a document or to do a backup. Furthermore, P4 and P5 pointed out that this facet is only useful for recent files, because the date of older items is often forgotten.

Social Facet Use and Satisfaction. The social facet questionnaire (see Appendix C.4) contained quantitative and open questions. The quantitative data reveals a mixed feeling of the participants regarding the social facet. They found the graphical representation of their contacts in the form of a social graph rather understandable (mean: 3.7, min: 2, max: 5, IQR: 1.75) (scale of 1 to 5), but rated the usefulness of such a representation as rather average (mean: 3.1, min: 1, max: 5, IQR: 1.5). The usefulness of the communities was rather liked (mean: 3.45, min: 1, max: 5, IQR: 1.375).

Two users (P4, P7) were particularly harsh with the social network. They gave the marks 2-2-3 and 3-1-1 to the above questions. They were users with many emails (P7 had 24,000 emails) and addresses, and the extraction of their social network took too much time. The evaluators actually had to select a subset of their emails for the extraction stage to be done in a reasonable amount of time. As the social network view does not scale well, it was also too difficult for them to use it efficiently. P7 however admitted that, even if it was useless for him, it would be useful for other people having a different social network.

Participants would use the social facet when looking for an email (P2, P6, P8, P9), to tell how many people are concerned by a file (P10) or to find people contacted infrequently (P3).

Participants would use the social facet to send an email to a group (P2, P4, P6, P8, P9) e.g. for advertising for an event[11]. Other possible uses of the social facet that were mentioned included: maintain the mailbox (e.g. delete all messages of a group) (P3), manage emails one community at a time (P3).

Some users complained about missing features. P3 would like a search embedded in the social view for people and mail addresses and P10 would like to be able to manually modify the automatic clustering as well as a support for subgroups.

6.3.9 Discussion

Reminder of the Evaluation Goals

The main objectives of this evaluation were: (1) to assess the usefulness of faceted navigation for re-finding personal information (summative evaluation at the abstraction and domain levels), and (2) to explore how people actually perform real re-finding tasks with the help of facets (exploratory evaluation at the domain level). Secondary goals included an informal exploratory evaluation of the potential of faceted navigation for gaining insight into one's PSI and a summative evaluation of the prototype at the encoding/interaction level. We discuss the corresponding results in the following section and end by presenting the possible evaluation bias of the methodology we employed.

Is Faceted Navigation Useful for Re-finding Personal Information?

The results from the questionnaires indicate that the participants found the system useful. Responses to open questions show that two main features of the faceted navigation approach are appreciated: the unification of all personal information in a single interface and the combination of research criteria allowed by the faceted approach. Moreover, the 'type' facet was explicitly mentioned as useful compared to existing systems. Some participants also said they liked having multiple views of their data, which is the core of faceted navigation. Finally, a majority of tasks were rated as easier to perform with the system (during the evaluation phase) than without it (during the diary phase). In particular, tasks rated as difficult became easier to perform with the faceted browser. On the contrary, some easy tasks became more difficult to perform. This tends to indicate that the faceted navigation approach could be used as a complement to more traditional means of re-finding information when the task is particularly difficult.

[11] P9 even mentioned that it would remove the need for Facebook, as it is mainly useful for sending such advertisements, in her opinion.

How Do People Use Facets in Re-finding Tasks?

Participants primarily used the most common facets: the file explorer and the textual search. The 'type' facet was also often used, mainly in combination with the file explorer or the textual search. The 'social network' facet was used predominantly for tasks related to email re-finding. It was used more often than all other facets in this context and often quickly led to the sought-after information without being combined with other facets. As emails are associated to people, it seemed natural for users to re-find emails by selecting the people concerned. Furthermore, the presence of cues in the description of the task related to the social network, the textual search and the type of sought-after information seems to positively influence the use of the related facet. However, without the corresponding cue, participants seemed to favour the use of the 'textual search' and 'file explorer' rather than the 'social network' and 'type' facets. We attribute this to the habit of the participants to re-find information using textual search and the file explorer instead of using the less common 'social' and 'type' facets. On the other hand, the temporal facets were rarely used. We think that this is at least partly due to the low number of tasks containing temporal cues, and the lack of tasks involving calendar events. Indeed, the diaries contained no such usable tasks.

In summary, people preferred and used mainly the file explorer and textual search. However, the strong use of the type facet for all kinds of resources and the predominant use of the social network for re-finding information in emails advocates for their usefulness in personal information re-finding systems.

Does Faceted Navigation Help to Gain Insight into one's PSI?

Evaluating how much an application allows to "gain insight" into a dataset is an inherently vague and difficult task. This is even more difficult to achieve when the dataset is unknown to the evaluator, as is the case with personal information. From our evaluation, only observational data allows to estimate it. Indeed, when asked questions about their activities during the last week, the last month and the people with whom they had the most to do some months ago, participants mainly made use of the system to answer. They usually selected the given period in the temporal facet, as could be expected from the questions, and observed how the other facets were filtered accordingly.

Moreover, it seems that participants could imagine using the system for exploring their PSI, not only for re-finding information. Indeed, some users could imagine using the system: for reporting their activities at a certain period of time; for maintaining the filesystem hierarchy; to assess whether a software license should be renewed after observing if documents using this software were created recently; to see how people relate to activities; to find people contacted infrequently. All these tasks are metalevel activites which qualify faceted navigation as a way to gain insight into a PSI.

Is the *Weena* Faceted Browser Usable?

The usability of the system is a summative question addressing the encoding/interaction level. The whole system was rated as usable in the questionnaires. However, some particular features were disliked, in particular the overly restrictive textual search and some miscellaneous graphical choices. On the other hand, the large number of selections versus low number of tasks concerned by the social facet tends to indicate that the encoding of information and interaction mechanisms with this facet is still suboptimal, despite the improvements implemented after the usability evaluation presented in Section 6.2. Results from the questionnaire indicate that the choice of representing the social network as a graph is understandable, but may not be optimal in terms of usability. Community extraction is rather liked, although a way of manually correcting the automatic community building would be appreciated, which confirms the result of the usability study.

Possible Evaluation Bias

Several issues with the evaluation itself may also have had an influence on the results. Primarily, we were limited by the tasks available from the diaries. The irregular amount of tasks available for each type in the taxonomy prevented us from having a fully balanced experimental setting. This therefore qualifies our evaluations as mainly qualitative and exploratory. On the technical side, our system leverages an existing indexing system which may have some flaws: recent changes in the PSI may not be reflected in the index and names of people and authors of documents are ambiguous (one person can be referred to with different names, two different person may have the same initials, etc.) (similar issues are reported by Cutrell et al. (2006b) who use the *Windows Search* index as well). Moreover, a usability evaluation of each facet independently could not be performed, except for the social facet. Finally, we did not have the resources to perform a long-term evaluation in order to assess the adoption of the system after the people have gotten used to it, which would have constituted a summative evaluation at the domain level.

6.4 Chapter Summary

This chapter presented the evaluations conducted in the context of this thesis. First, based on previous evaluation attempts, it identified the difficulties and pitfalls of performing evaluations in PIM and presented a methodology allowing ecologically valid evaluations to take place. Furthermore, it discussed evaluation types and levels in PIM and, as such, contributed to the theoretical framework underlying PIM evaluation, which was a side objective of this thesis (see Section 1.4). A formative evaluation of the system presented in Chapter 5 allowed to refine specific parts of its interface. The main evaluation introduced a taxonomy of tasks in PIM that completes previous attempts and allows the understanding of re-finding behaviour at the task level. Results pointed out an interest on the part of participants for the faceted navigation approach to personal information re-finding. In particular, the type facet has been used often in all contexts and the social facet has been used often to re-find emails. The most used facets were also the most common: file explorer and textual search. The users' lack of practice with the system might be a factor in this. It was also difficult to evaluate how much the faceted browser allowed the users to gain insight into their own personal information space. Some of their comments nonetheless point out that they can see a use for the system for such kinds of metalevel PIM activities.

The following chapter discusses the findings of this thesis as well as methodological issues.

There is a crack in everything,
That’s how the light gets in.

Leonard Cohen, *Anthem*

Et comme toutes sont entre elles ressemblantes,
Quand il les voit venir, avec leurs gros drapeaux,
Le sage, en hésitant, tourne autour du tombeau:
Mourons pour des idées, d’accord, mais de mort lente.

Georges Brassens, *Mourir pour des idées*

7
Discussion

7.1 Introduction

This chapter looks more closely at three themes of this doctoral work. I believe these themes are original and novel with respect to the current state of personal information management (PIM) research. Therefore, I discuss them in detail in this chapter. Section 7.2 deals with PIM strategy changes and discusses the findings of Chapter 4 in a broader perspective. Section 7.3 analyses faceted classification of personal information (PI) in light of the findings of Chapter 5. Section 7.4 discusses faceted navigation as an interface paradigm for PIM leveraging the experience acquired in Chapters 5 and 6. For each theme, a brief summary of the outcomes of this work is provided first and commented on. Then, limitations of the approach and methodology are considered. Finally, perspectives, future work and open questions related to each theme are discussed.

7.2 On PIM Strategy Changes

7.2.1 Discussion of Outcomes

The ethnographically-inspired study I conducted shed light on strategies employed by people in their daily PIM practices. A number of interesting observations were made in the course of the study.

First, the ethnographic nature of the study highlighted the fact that strategies described by participants often do not match their real practices. It seems that people have ideal strategies that they would like to follow, but often cannot, due to the somewhat unpredictable nature of PIM and the difficulty of foreseeing the future use of information items. Time constraints and motivation also seem to impair how well they are able to follow their ideal strategies. Moreover, for almost every participant, there was evidence that relics of abandoned strategies were still present in their personal space of information (PSI). This points out the fact that observation may be a necessary complement to interviews in order to fully understand PIM behaviour. To my knowledge, most studies of user behaviour in PIM are based on interviews or questionnaires. In other terms, they are based on "what people say they do" and not on "what they actually do". I believe that much insight into their practices can be obtained by observing them from an ethnographic stance. However, long-term ethnographic studies also raise a number of issues, related in particular to the sensitive nature of personal information, and privacy concerns.

Second, I noticed that strategy changes were numerous for every participant. Indeed, people change strategies for multiple specific reasons: a bad organization hindering to re-find information, the discovery of a new technology, an excessive amount of documents making classification time-consuming, the abandon of a useless or ineffective strategy, and so on. Based on the interviews, I proposed a framework for characterizing short-term strategy changes. In my model, each strategy change has a cause (either internal or external), a scope (tool-specific or longitudinal) and require an action (simplify or complete the existing strategy).

Third, evidence shows that at least some participants had a bad overview of their PSI. They had forgotten entire areas of it. Tours of their PSI made them discover things that, in their own terms, "should not be there". I believe that this bad conception of their PSI is at least partly due to the tools at their disposal. Indeed, the desktop metaphor enforces filing which hides items under a hierarchy and provides no way of getting an other overview than this very hierarchy. Thus, "misplaced" items may be very hard to re-find, for example.

Finally, most participants said they did not like doing PIM for itself, because it is a painful and time-consuming task. However, although people participating in my naturalistic study had spent one or two hours of their working day with me, they had all appreciated spending time to share and analyse their strategies. It seems that the occasions to share experience and reflect on one's own practices are rare. Moreover, several participants explicitly asked for advice or complained because they had never been educated to use PIM tools effectively. This points out a more general problem of PIM: people never study how to do PIM, they learn it on the job. Therefore, there seem to be a clear need for supporting the choice of PIM strategies.

7.2.2 Limitations

The evaluation I conducted, by its very methodology, is only able to highlight general trends. Indeed, observation and open discussion with participants are not suited to explore their behaviour systematically. Moreover, the study was done in a short time scale. Traditional ethnographic work demands repeated observations in order to formulate interpretations of behaviour, whereas my work is based on single observation sessions. Indeed, many observations made in the course of the study were founded on people's reports, not on actual observation. For example, strategy changes were not observed, but described and explained by participants afterwards. Finally, the use of a recording device would have helped minimize the bias and disambiguate hand notes taken by the interviewer. The results of this study shall be considered with respect to these points.

7.2.3 Perspectives and Future Work

PIM behaviours have been the subject of many studies already. But the subject is particularly difficult: strategies are tailored to individuals, goals and needs of individuals are often not clearly defined even for themselves. Nevertheless, empirical data clearly showed that people need support regarding the choice of strategy to manage their PI efficiently. This raises the question to know whether certain strategies are better suited to certain tasks. If it is the case, it would mean strategies could be taught or suggested with respect to the needs of the individual.

However, the practical outcomes of studying PIM strategies seem to be poor. Those studies may have a *"sociological and psychological value"*, but they offer little *"to inform future information architects, analysts, archivists, and librarians"* (Barreau et al., 2009). Is there a benefit to expect from technological solutions to support the choice of PIM strategy or does the solution lay in human knowledge sharing? This is still an open question.

7.3 On Faceted Classification of PI

7.3.1 Discussion of Outcomes

The work done in the context of this thesis identified contextual cues of personal information that can be used as facets of a faceted classification system. The choice of facets was motivated by results from the state of the art and from my own studies of user behaviour.

The faceted classification of personal information performed in the context of this thesis leveraged an external indexing engine. This has a number of advantages — some of them have been highlighted by Chau et al. (2008a): (1) the indexing engine can handle a broad range of different data types; (2) support for other data types can be added easily through plugins; (3) the indexing engine can notify an external system that the index has changed; (4) many system-level metadata are available in the index; (5) the indexer can be queried using a rich SQL-like syntax; (6) users do not need to change the way they organize their information; (7) user information is left unchanged.

As the social metadata included in the Windows Desktop Search index used by the system was poor, I developed the social facet separately. This occurs in two phases. In a first phase, a social network graph was extracted from the email archive of the user and split into communities. In the second phase, information items were mapped to pertaining contacts. The generation of the social network graph uses an original metric to weight the relationships of people. The communities splitting expands a state-of-the-art community extraction algorithm. Linking information items to contacts in the social network is achieved using rule-based strategies.

7.3.2 Limitations

Using an indexer has some disadvantages — some of them have been discussed by Cutrell et al. (2006b): (1) the creation, update and deletion of items take some time to propagate to the index, so the index does not reflect the current state of the PSI; (2) metadata may be wrong or missing in some fields. Additionally, the social metadata in the index was poor. A single person may be represented by several different names in the index, and those co-references should be resolved.

Regarding the way the social network is created, there are also a number of limitations. First, several addresses could refer to the same person. Thus, co-reference resolution would also be needed at that level. Second, Whittaker et al. (2004) found that 10% of contacts on average belong to several groups. Nonetheless in my approach, a contact can only belong to one single community. Another limitation of my social network building is that it is not dynamic. As it never changes, an advantage is that its topology is always the same. But if new people appear in the life of the user, they will not appear in their social network. It would need to be generated again from scratch, which would change its overall topology.

The connection of contacts in the social network to information items is simplistic in my implementation. Indeed, it is based on: (1) pattern matching of contacts' and document authors' names using transformation rules, (2) comparison of attachments and file names and (3) fulltext queries on the contact's name. Moreover, the classification of items along their social facet could not be achieved in a preliminary phase for technical reasons. In the evaluated system, the link between contacts and items is computed on-the-fly when the user selects contacts in the interface. In consequence, there is so to speak no proper classification of items along their social facet, which implies that it is impossible to know how many information items are related to each contact, thus preventing from taking advantage of the query preview associated with faceted navigation, in the case of the social facet. However, my goal was to enable the social facet as a proof of concept. Even if a proper classification was missing, consequences were rarely noticeable during system use.

Finally, the methodology I employed to select relevant contextual cues of personal information to use them as facets (i.e. by reviewing previous research) may have prevented me from discovering other unexplored cues of personal information that might have interest for particular re-finding tasks.

7.3.3 Perspectives and Future Work

Several improvements can be applied to the social network extraction. First, as far as the graph building is concerned, co-reference resolution techniques can be applied to regroup different email addresses referring to one single person. Moreover, links between contacts could be weighted using other metrics, e.g. based on similar terms appearing in emails or based on the folders into which emails are classified. Furthermore, knowledge from online social networks could be integrated. As social networks evolve over time, solutions to dynamically update them could be investigated as well. Concerning the extraction of communities in the social graph, other algorithms can be used, in particular algorithms that allow overlapping communities — as one contact may belong to several communities in real social networks. Furthermore, the community splitting was flattened to one single level. Future research could investigate the utility and affordance of representing a social network on several levels.

Regarding the connection between the social network and the individual information items, my simplistic approach constituted a proof of concept. Nevertheless, future research could address the issue by considering more advanced approaches (e.g. machine-learning) to classify information items according to the people concerned.

From a more general point of view, a great deal of work is still needed to obtain a complete and reliable faceted classification of personal information, at least along the facets identified in this thesis. In my opinion, future work has two aspects. First, more metadata should be collected to reflect the temporal evolution of information items' context. For example, if I work on a document at several points in time, all the dates of access should be logged and not only the last one. Thanks to this, it would be possible to retrieve a document that I remember having accessed at the same time as another one, no matter if it was the last time I accessed them or not. Moreover, it opens a large field for retrospectively analysing my own past activities. Considering not only desktop computers but also mobile devices, additional metadata could be automatically collected, like those issued from geo-localisation. It would enrich the context of information items and provide even more ways to trigger recollection. A second aspect that would be needed to improve faceted classification is related to the social facet. A recurring issue with systems that deal with social data is to regroup the different digital identities (e.g. several email addresses, user name, initials) into a single "person" construct. To make an effective use of the social facet, social data should be organized in a comprehensive way. Moreover, a systematic social tagging of all information items at the time of authoring or sharing would be needed, i.e. not only chats and emails but also textual documents, spreadsheets, presentations, calendar events, and so on.

This brief discussion shows that building a comprehensive faceted classification of personal information is clearly a low-level task that has to be carried out at the operating system level. All in all, most of the remarks I made in this section point out to a currently *missing infrastructure* in operating systems, that future research should try to design and implement.

At a more general level, further naturalistic studies conducted on the longer term might reveal other contextual cues of personal information that the ones I used in my work, and that may be effective in particular re-finding tasks.

7.4 On Faceted Navigation For PIM

7.4.1 Discussion of Outcomes

Faceted navigation was chosen as an interface paradigm to support re-finding in PIM given its successful application as a way to navigate large datasets. The implementation of faceted navigation for PIM introduced in this thesis differs from previous applications in some aspects. First, textual search was used as a separate facet in the interface, which filters the original dataset according to words contained within documents or their filenames. The traditional application of textual search in faceted browsers uses it to search inside other facets' categories, not as a separate facet. I therefore took some liberty from the faceted navigation paradigm with respect to search[1]. Nevertheless, this use of search made more sense from the point of view of the users which are accustomed to searching in their personal collections. Second, facets in the interface were augmented with visual clues. Indeed, bars were provided next to facet categories. The grayed out portion of those bars indicate how many items of this category exist in the whole dataset, while the black portion represents

[1]With its use of search, the faceted browser presented in this thesis qualifies as a "faceted search" system, according to the terminology used by Tunkelang (2009). See also Section 2.2.2.

how many items would match a selection of this category — which relates to the query preview mechanism of the faceted navigation paradigm. This visual encoding provides more context for categories than in the traditional paradigm which lacks those visual bars and usually only shows the number of items to expect from a selection. Third, visual representations of the temporal and social facet were introduced, rather than more conventional textual lists of categories.

The evaluation conducted using the system led to an increased understanding of how people use facets when they perform particular kinds of re-finding tasks. Results show that participants predominantly use the file hierarchy and textual search. The type facet is used regularly for all kinds of tasks. The social facet is the most used facet in the case of email re-finding. Further evidence point out that the predominant use of textual search and the file explorer might be a clue that people are not used to other facets of personal information and do not think of using them. Indeed, when explicit clues about the social or type facets were present in the description of the task, people tend to use the associated facets. On the other hand, empirical data seem to indicate that faceted navigation helps to gain insight into the PSI, as participants in the evaluation could imagine higher-level tasks to be performed with the help of the system (e.g. maintenance tasks). The temporal facet in particular seemed to facilitate the analysis of past PIM activities and may be useful for reporting one's work occupations during a certain period of time. Finally, from a more general point of view, participants in the evaluation found the approach of faceted navigation both usable and useful for re-finding items in their PSI.

7.4.2 Limitations

The implementation of faceted navigation presented in this thesis has a number of limitations. The main limitation is a consequence of the imperfect faceted classification presented in the previous section. Indeed, the lack of complete faceted metadata for the social facet did not allow the representation of categories (i.e. persons) according to the faceted navigation paradigm: no query preview was available and it could theoretically happen that selecting contacts led to an empty result set which contradicts the faceted navigation requirements. One could also mention the breadcrumb trail which does not allow the removal of a filter without removing all filters that were applied after it. Furthermore, for technical reasons, the 'file explorer' facet could only show the filesystem hierarchy and not the email archive hierarchy.

Regarding the evaluation, perhaps the main limitation was the lack of practice time the participants have had with the system. They did a guided tour of the interface in about 15 minutes and then were free to use it for a couple of additional minutes. It is difficult to assess how much this impacted the results of the evaluation. For example, it is unclear how much the predominant use of the most usual facets was a consequence of participants being unfamiliar with the other possibilities offered by the interface. One could also argue that this lack of practice was beneficial. As the system used in the study was a prototype, participants using it on a longer-term would maybe have experienced flaws and lost trust in the system. This would also have biased the evaluation. In my opinion, the relative lack of practice of the participants may have avoided that they rely only on "flawless" facets (i.e. file explorer). Therefore, it helped to show the true *potential* of other facets (e.g. social facet), as participants used them because they seemed promising to perform the tasks, without anticipating their actual efficiency (or lack thereof) in the implemented system.

Other variables may have influenced the evaluation. In particular, tasks used during the evaluation were real tasks of participants, reported as part of a preliminary diary study. This pool of tasks was unfortunately unbalanced. Therefore, when selecting tasks out of this pool, it was impossible to balance all dimensions (difficulty, cues, targets, resources). Nevertheless I tried to take it into account in the interpretation of the results. Another criticism that can be made is the relatively small number of people involved in the study (10 people). Again, the interpretation of the results acknowledges this fact.

7.4.3 Perspectives and Future Work

Visual encoding choices

The representation of facets included in the interface could be further adapted. In the social facet, the use of pictures of contacts instead of their names could be more evocative (Whittaker et al., 2004; Elsweiler, 2007). This nonetheless raises the issue of obtaining those images automatically. Online social networks could be leveraged on to this aim. The temporal facet could benefit from other representations as well. A timeline view augmented with personal memory landmarks (Ringel et al., 2003) could help relate items with remembered events (personal calendar events, Christmas, birthday, holiday and so on), thus increasing the value of the temporal facet. Moreover, the temporal view chosen in the system is aggregative: single items are not displayed, only aggregates of items for a specific date are shown. Users could benefit from seeing all individual items. It would let them visualize the relationships between individual items based on the time they were accessed or modified. Appendix B.1 describes a prototype of such an temporal interface that was made towards this goal. Such an approach however raises a number of other issues related to information visualization, in particular visual cluttering.

Evaluations

Further evaluations would be needed to assess with more confidence this approach of faceted navigation to PIM. In particular, regarding the evaluation types and levels introduced in Section 6.1.1, visual encoding choices would need to be evaluated separately from the whole system. For example, such an evaluation would take the form of a comparison between two interfaces, each one using a different visual representation for the social or temporal facet. Quantitative data related to success rates or access times would be collected in order to assess the most effective visual encoding. Regarding the benefit of faceted navigation, one could also imagine comparing the system with a variant of it that would not display the query preview or that would not filter out other facets after a selection has been made in a specific facet.

On a more general note, faceted navigation is believed to bring an increased insight into a dataset. Many information visualization systems, in particular those who aim at supporting the analysis of large datasets, share the same purpose. However, evaluating how much interfaces allow to gain insight into a dataset is a vague and complex question and methodological means are still missing to address it.

General perspective

I believe that faceted navigation would develop its full potential if it was combined with current interfaces. One should be able to apply faceted navigation as a basic feature in every PIM tool as a complement to other means of re-finding information. For example, the visual part of the desktop metaphor (files and folders arranged spatially) could be considered another visual encoding of the file explorer facet developed in our system. Re-finding items in an augmented desktop metaphor would be done by selecting categories in type, temporal or social facets, which would for example gray out irrelevant items on the desktop or in folders. Such a faceted navigation interface would have the additional benefit of supporting the spatial memory, which is a strength of the desktop metaphor with respect to re-finding items.

7.5 Chapter Summary

This chapter reviewed what I think are the three most original and novel themes with respect to current research in PIM that this doctoral work has dealt with. The outcomes and limitations of my research were discussed as honestly as possible and future work and perspectives related to those themes were presented. The following chapter summarizes the findings of the thesis according to its objectives and considers general perspectives for this research.

This is the end
My only friend, the end
It hurts to set you free
But you’ll never follow me.

James Douglas Morrison, *The End*

Eilt mit dem Werk: widerlich ist mir’s!

Richard Wagner, *Das Rheingold*, 1:4 (Wotan)

8 Conclusion

This thesis has examined how faceted navigation can help support personal information management (PIM). The most widespread systems for managing personal information implement the so-called desktop metaphor which considers digital information in a way similar to physical information, collected in files and folders and which is accessed mainly through browsing. However, previous research on PIM behaviours and on the role of memory in PIM has highlighted the fact that contextual cues of information are easily remembered and help re-finding, an aspect the desktop metaphor inadequately supports. Recent PIM systems have begun making use of particular cues of personal information to make re-finding easier. The system designed and developed in the context of this thesis combines several contextual cues of personal information in a faceted browser and was evaluated in a ecologically valid context. Section 8.1 presents the contributions of this thesis and Section 8.2 examines perspectives and future work. Both sections are organized according to the objectives of the thesis introduced in Section 1.4:

1. *Identify contextual cues* of personal information which support re-finding and can be obtained without a need for user annotation.
2. Identify or develop a *user interface paradigm* suited to support the re-finding task in PIM and which benefits from the cues.
3. *Design, implement and evaluate a PIM re-finding system* based on the cues and interface paradigm.
4. Improve the *theoretical framework underlying evaluation* in the field of PIM.
5. Develop an increased understanding of *PIM strategy changes*.

Finally, I give a short concluding statement of my research and set out a more general and open conclusion highlighting real-world PIM challenges for which my doctoral work may have implications.

8.1 Contributions

8.1.1 Identify Contextual Cues

The first objective of this thesis was to identify the contextual cues of PI which can support re-finding and be obtained automatically without a need for user annotation. Previous research in PIM reviewed in Chapter 3 highlighted the potential of several cues which seem to be easily remembered and facilitate recollecting. The study presented in the first part of Chapter 4 confirmed and extended the findings by allowing the rejection of other candidate cues. The possibility to obtain metadata related to these cues without annotation on the part of the users further led to selecting the following cues:

- the type of item
- the people (social network) related to the item
- the approximate date of the item
- the path of the item in the user-defined hierarchy
- textual search applied on the item's name and content

8.1.2 User Interface Paradigm

Given its use as a browsing and searching interface on public and personal datasets presented in Chapter 3, faceted navigation using the aforementioned contextual cues appears to be an appropriate interface paradigm. Chapter 5 described how the identified contextual cues can be used to build a faceted classification of personal information. Going beyond the traditional approach of faceted navigation using textual lists of facet categories, I proposed to apply information visualization techniques to represent the social and temporal facets. Furthermore, visual encodings were used to represent the amount of filtered items for each of the facet categories. Empirical data from the evaluation in Chapter 6 revealed that users found the approach of faceted navigation using contextual cues both usable and useful to re-find information items.

8.1.3 Design, Implement and Evaluate the System

A faceted browser of personal information was developed based on the findings described in Chapters 3 and 4. Chapter 5 detailed its design and implementation. Faceted classification of information items was mainly acquired by leveraging an existing indexing engine embedded in the operating system. A technical contribution of this thesis is the extraction of a social network from the email archive of the user. Technical solutions were developed to allow generating a social network graph and splitting it into communities. Furthermore, a simple approach was proposed to connect contacts in the social network to information items. There are limitations to this implementation which were addressed in Chapter 7. It could nevertheless be used as a proof of concept and the social facet could be integrated into the final evaluated system.

The evaluation presented in Chapter 6 gave insight into how people use facets of personal information with respect to the kind of task they have to perform. An important finding is that some facets are used more for certain kinds of tasks and certain kinds of resources. Therefore, combining several facets offers a benefit, as it allows people to re-find items applying what they recollect about the sought-after item. Moreover, the system seemed to be particularly beneficial for tasks that were rated as difficult to perform without it. Another important finding is that people seem to favour facets of personal information that they are accustomed to. For example, the social, type and temporal facets were selected less than the file explorer and textual search facets. Habits and lack of practice with the systems may nonetheless be factors. Some evidence further showed that faceted navigation has a potential to help people gain insight into their personal space of information and perform metalevel activities of PIM, like the maintenance of their personal collections.

8.1.4 Theoretical Framework Underlying Evaluation

Several scholars have discussed the difficulty of performing PIM evaluations. Various reasons explain this difficulty. In particular, the personal aspect of PIM requires that realistic data and tasks be used. Elsweiler et al. (2007) recently proposed a practical methodology that overcomes those difficulties. Building on this work, Chapters 2 and 6 proposed an extended taxonomy of PIM re-finding tasks. The goal of this extended taxonomy was to permit a more complete understanding of PIM practices at the task level. In particular, in Chapter 6, it allowed for the study of how people perform re-finding tasks with respect to the cues mentioned in the description of the task.

The discussion presented at the beginning of Chapter 6 pointed out that PIM evaluation can be carried out at different levels and with different goals in mind. Leveraging previous attempts at classifying the types and levels of evaluation in the particular domain of information visualization, I proposed a classification of evaluation in PIM. I examined some previous evaluations from the state of the art and classified them according to the evaluation type and level they used. I believe that a more careful handling of evaluation in PIM in general is needed and I encourage researchers and developers of PIM systems to refer to this model when they claim contributions at a specific level.

8.1.5 PIM Strategy Changes

PIM strategy changes in the short-term have been under-studied. In particular, most previous attempts only considered tool-specific strategy changes. The ethnographic study presented in Chapter 4 shed more light on strategy changes and resulted in the definition of a framework for strategy changes at a more general level. Another finding of this study is that people are often dissatisfied with their PIM strategies, but do not know how to make them better.

8.2 Perspectives

Based on the findings of the thesis, this section presents future work and research perspectives. It is organized along the thesis objectives.

8.2.1 Identify Contextual Cues

The contextual cues I identified based on previous studies are common to many kinds of information items. Further research could consider taking advantage of cues specific to certain types of items, like camera type or geographical location for pictures, or artist and genre for audio files. Nonetheless, I believe that it is still important for those contextual cues to be obtained automatically, as users are not likely to maintain collections of annotations consistently in the long-term. Future research should not rely too much on user annotations to support PIM.

8.2.2 User Interface Paradigm

This thesis applied faceted navigation to PIM and proposed specific visual encodings of certain facets and of their categories. While faceted navigation has a number of advantages over the desktop metaphor, it falls short of taking advantage of the spatial encoding in memory. Further research could address this by embedding faceted navigation into the desktop metaphor, e.g. by providing ways to select facet categories in a regular desktop environment which would gray out irrelevant items according to the categories currently selected in the facets. From a visual encoding point of view, the choices made in our approach for the social and temporal facets should be further evaluated, or compared to alternate encodings. Possible improvements include providing temporal landmarks in the temporal facet or using photographs of contacts along with their names in the social facet.

8.2.3 Design, Implement and Evaluate the System

Design and Implementation. From a technical point of view, the faceted classification upon which our faceted browser is built is obtained partly from an external indexing engine and partly from a custom extraction of the social network from the personal email archive. Several improvements could be made to the way the social network is generated and represented, as well as to how it is connected to individual information items. The most important point overall is that attempts to automatically extract the social network should be evaluated with users, to see how much the extracted social network match with their own representations. Scholars reported that users are often dissatisfied with automatically produced social networks, while at the same time reluctant to create their social network from scratch or correct the results of automatic generation. Our observations also pointed out that being dissatisfied with the overall look of the social network seems to induce a loss of trust in the whole system.

In a more general perspective related to the faceted classification, I believe that the whole work of classification should be delegated to the indexer. The discussion in Chapter 7 emphasized the need for a more complete backend infrastructure in operating systems which would be able to provide coherent and reliable contextual metadata for all information items.

Evaluation. Regarding the evaluation of the system, Chapter 7 mentioned that several variables may have influenced the outcomes. Further evaluations should try to limit the impact of these variables. In particular, the participants should be allowed a longer time to get used to the system before the actual evaluation takes place. Moreover, further evaluations could focus on specific parts of the interface, comparing the usability of particular visual encodings for example, or be conducted on a long-term basis to assess the adoption of the system.

8.2.4 Theoretical Framework Underlying Evaluation

There is still much fundamental work to be done in this area. In particular, the methodology of evaluation used in this work may not suit all evaluations types and levels presented in Chapter 6. Finding other methodologies to support them should be addressed by future research. For example, the methodology we applied clearly demands using real tasks and real personal information. A future work in the theory of PIM evaluations could be to assess whether these constraints are necessary when summatively evaluating the encoding or algorithm level of a PIM system.

8.2.5 PIM Strategy Changes

The understanding of PIM strategies and changes thereof would benefit from longer-term observation, in the ethnographic tradition. However, such evaluations are very costly in terms of resources. Given the particular nature of personal information, such observations may also raise important privacy concerns. The model I have proposed for characterizing PIM strategy changes would also need further empirical evidence to be validated, e.g. alternate studies. Evidence has shown that people are often dissatisfied with their PIM practices and would appreciate support for choosing and changing their PIM strategies. However, as discussed in Chapter 7, it is still unclear whether technical solutions alone could help them, or if they would benefit more from sharing their experiences with co-workers or possibly receive hints or courses on PIM. Future work on PIM behaviour studies could therefore examine how sharing PIM practices impact one's own practices and satisfaction.

8.3 Concluding Statement

This thesis contributes to PIM research in different ways. In particular, it showed that a faceted navigation approach using contextual cues of information helps users re-find information items and gain insight in their own personal collections. Furthermore, it increased the understanding of people's behaviour when performing specific kinds of re-finding tasks.

8.4 Back to the World Out There

Will the work carried out in this doctoral thesis impact the "real" world out there? Probably not, or at least not directly. The desktop metaphor still has many good days ahead on personal desktop computers. However, PIM is gaining more and more importance on portable devices, as their storage capacity and features are improving. Their interfaces for managing personal information may often look primitive and application-centric, far away from the desktop metaphor. Nevertheless, it does not prevent them from being accepted by a wide audience. As a blog recently noted[1]:

> *The majority of [smartphone and portable media player] owners do not have access to the filesystem of their device and don't lose sleep over it. It's even the opposite: it makes information technology accessible to them.*

Most portable devices include geo-localisation and are tailored for communication, a task that has an inherent social nature. They could therefore be designed to associate even more contextual cues to information items. These contextual cues could then be applied in interfaces inspired by faceted navigation to support re-finding. Imagine being able to re-find the email that you remember having read in the suburb between two stations, by using this single cue.

Personal information management has left the ivory tower of academia. People begin to discuss their PIM problems and practices on online forums[2]. Reading them, one feels relieved to know that others share the same experiences of frustration and dissatisfaction with PIM. Or, that they have some good advice to give. Moreover, books have popularized the topic of *"Keeping Found Things Found"* (Jones, 2007) or offered recipes for coping with information overload and *"Getting Things Done"* (Allen, 2002; Hurst, 2007).

Difficulties associated with information management, like information overload and fragmentation, are not particular to the 21st century. Since the invention of printing, technology has always contributed to the spread of information. So, even if technology could help overcome the difficulties of personal information management, it would probably never be sufficient on its own, as it is part of the problem. Sharing of personal information management practices and education could and should be called to the rescue. The millions of individuals struggling to manage their personal information are not alone. It is time to let them know.

[1] `http://www.gete.net/blog/`, translation by the author.
[2] `http://www.talesofpim.org` or `http://www.addforums.com`.

Il semble que la perfection soit atteinte non quand il n'y a plus rien à ajouter, mais quand il n'y a plus rien à retrancher.

Antoine de St-Exupéry, *Terre des Hommes*

Appendix

A
Online Survey Results

This appendix presents a summary of the results of the online survey introduced in Chapter 4.

Summary of Survey on Personal Documents Management (University of Fribourg)

1. Sex

	Response Percent	Response Count
woman	38.1%	64
man	**61.9%**	**104**
	answered question	**168**
	skipped question	**0**

2. Age

	Response Count
view	168
answered question	**168**
skipped question	**0**

3. Current job

Details removed. Main responses were :

- informatics - IT - computer science : 49.4% (83)
- teaching : 11.9% (20)
- internet - multimedia : 8.3% (14)
- communication : 4.2% (7)
- audiovisual : 3.6% (6)
- health - social : 3.6% (6)
- psychology - sociology : 2.4% (4)
- secretary - assistant : 2.4% (4)
- public service : 1.8% (3)
- answered question: 168
- skipped question: 0

4. Frequency of use of computer

	Response Percent	Response Count
less than 1 hour per day	4.2%	7
between 1 and 3 hours per day	9.5%	16
between 3 and 6 hours per day	20.2%	34
more than 6 hours per day	**66.1%**	**111**
	answered question	**168**
	skipped question	**0**

5. What operating system do you mostly use?

	Response Percent	Response Count
Linux	7.1%	12
Mac OS X	21.4%	36
	answered question	**168**
	skipped question	**0**

5. What operating system do you mostly use?

		Response Percent	Response Count
Windows XP		**50.6%**	**85**
Windows Vista		13.7%	23
view	Other (please specify)	7.1%	12
		answered question	**168**
		skipped question	**0**

6. How do you rate your expertise with computers?

	Response Percent	Response Count
Expert (I know how to install applications, do some programming, configure my operating system)	**81.5%**	**137**
Knowledgeable (I know how to send/receive emails, navigate in internet)	18.5%	31
Passing (I use computers from time to time for limited activities like writing a document and send one email)	0.0%	0
	answered question	**168**
	skipped question	**0**

7. To manage your documents, how often do you use these strategies?

	almost never	rarely	sometimes	often	almost always	Response Count
Organization into a hierarchy: you define a folder hierarchy into which you put your documents	0.7% (1)	4.1% (6)	8.2% (12)	34.7% (51)	**52.4% (77)**	147
Piling: you put your documents into a unique folder or in folders having less than 2 subfolders	17.6% (25)	**28.2% (40)**	**28.2% (40)**	20.4% (29)	5.6% (8)	142
You put your documents on the desktop	22.4% (32)	20.3% (29)	**29.4% (42)**	23.8% (34)	4.2% (6)	143
					Other (please specify) view	9
					answered question	**147**
					skipped question	**21**

8. To name a folder, how often do you use these labellings?

	almost never	rarely	sometimes	often	almost always	Response Count
Project name ("experiment", "agent code",...)	4.2% (6)	3.5% (5)	20.4% (29)	**46.5% (66)**	25.4% (36)	142
Role ("teaching", "personal",...)	13.2% (19)	13.9% (20)	**33.3% (48)**	31.9% (46)	7.6% (11)	144
Topic ("banking", "music", "politics",...)	4.2% (6)	14.6% (21)	23.6% (34)	**45.1% (65)**	12.5% (18)	144
Contact ("Rick", "John Doe",...)	**45.1% (64)**	23.2% (33)	23.2% (33)	5.6% (8)	2.8% (4)	142
					answered question	**147**
					skipped question	**21**

8. To name a folder, how often do you use these labellings?

Absolute time ("1970-01-01","year 2009",...)	22.8% (33)	19.3% (28)	**26.9% (39)**	**26.9% (39)**	4.1% (6)	145
Relative time ("yesterday", "after Christmas",...)	**80.1% (113)**	12.8% (18)	5.7% (8)	1.4% (2)	0.0% (0)	141
Format ("excel sheets", "word document", "mp3s",...)	**71.5% (103)**	16.7% (24)	5.6% (8)	4.9% (7)	1.4% (2)	144
Class of document ("letters", "presentations", "images",...)	23.4% (34)	21.4% (31)	**27.6% (40)**	20.7% (30)	6.9% (10)	145
Workflow ("pending","to do",...)	**48.3% (70)**	24.8% (36)	20.0% (29)	4.1% (6)	2.8% (4)	145
Event ("Christmas 2009", meetings or conferences,...)	11.2% (16)	10.5% (15)	**37.8% (54)**	31.5% (45)	9.1% (13)	143
					Other (please specify) view	7
					answered question	**147**
					skipped question	**21**

9. You use your desktop for storing:

	almost never	rarely	sometimes	often	almost always	Response Count
recent documents	16.0% (23)	12.5% (18)	26.4% (38)	**28.5% (41)**	16.7% (24)	144
files (any date)	**44.8% (64)**	24.5% (35)	18.9% (27)	7.7% (11)	4.2% (6)	143
information to be processed	18.1% (26)	12.5% (18)	**27.8% (40)**	**27.8% (40)**	13.9% (20)	144
important information to remember	20.3% (29)	16.1% (23)	21.7% (31)	**31.5% (45)**	10.5% (15)	143
					other (please specify) view	20
					answered question	**145**
					skipped question	**23**

10. When you search for a file, you search:

	almost never	rarely	sometimes	often	almost always	Response Count
by filename or part of the filename	3.4% (5)	7.6% (11)	9.7% (14)	**40.0% (58)**	39.3% (57)	145
by file size	**67.6% (96)**	21.8% (31)	9.2% (13)	1.4% (2)	0.0% (0)	142
by date of creation	29.0% (42)	**30.3% (44)**	27.6% (40)	10.3% (15)	2.8% (4)	145
by date of last modification	**29.7% (43)**	20.0% (29)	26.9% (39)	17.2% (25)	6.2% (9)	145
by file format/type	20.8% (30)	23.6% (34)	**29.9% (43)**	22.9% (33)	2.8% (4)	144
by keywords (from file content)	18.8% (27)	8.3% (12)	22.9% (33)	**33.3% (48)**	16.7% (24)	144
					Other (please specify) view	6
					answered question	**145**
					skipped question	**23**

11. How do you access documents modified or created 2 days ago?

	almost never	rarely	sometimes	often	almost always	Response
					answered question	**143**
					skipped question	**25**

11. How do you access documents modified or created 2 days ago?

						Count
Direct browsing to desired file (you know exactly where it is)	0.7% (1)	1.4% (2)	5.7% (8)	38.3% (54)	**53.9% (76)**	141
Step by step browsing to desired file (you don't know exactly where it is, i.e. look into each folder you open to find next folder to browse in)	25.5% (36)	**33.3% (47)**	27.7% (39)	9.9% (14)	3.5% (5)	141
Using "Recent documents" option in applications	**25.4% (36)**	21.1% (30)	22.5% (32)	23.9% (34)	7.0% (10)	142
Using menu "Start > Recent documents"	**39.7% (56)**	27.0% (38)	16.3% (23)	12.8% (18)	4.3% (6)	141
Using a desktop search engine	**42.1% (59)**	26.4% (37)	16.4% (23)	11.4% (16)	3.6% (5)	140
					Other (please specify) view	13
					answered question	**143**
					skipped question	**25**

12. How do you access 1-year-old documents?

	almost never	rarely	sometimes	often	almost always	Response Count
Direct browsing to desired file	5.6% (8)	15.5% (22)	28.2% (40)	**33.8% (48)**	16.9% (24)	142
Step by step browsing to desired file	8.5% (12)	10.6% (15)	**41.8% (59)**	31.2% (44)	7.8% (11)	141
Using a Desktop search engine	18.9% (27)	18.9% (27)	**28.0% (40)**	25.2% (36)	9.1% (13)	143
					Other (please specify) view	2
					answered question	**143**
					skipped question	**25**

13. How do you access documents modified or created at a precise date?

	almost never	rarely	sometimes	often	almost always	Response Count
Direct browsing to desired file	8.0% (11)	13.8% (19)	25.4% (35)	**31.2% (43)**	21.7% (30)	138
Step by step browsing to desired file	14.7% (20)	14.7% (20)	**44.9% (61)**	23.5% (32)	2.2% (3)	136
Using a Desktop search engine	23.2% (32)	18.8% (26)	**25.4% (35)**	21.0% (29)	11.6% (16)	138
					Other (please specify) view	5
					answered question	**139**
					skipped question	**29**

14. When using a search engine, how often do you use these kinds of temporal queries?

	almost never	rarely	sometimes	often	almost always	Response Count
Relative dates ("documents of yesterday")	**51.7% (74)**	21.7% (31)	18.9% (27)	4.9% (7)	2.8% (4)	143
Absolute dates ("1970-01-01")	**45.7% (64)**	25.7% (36)	17.1% (24)	7.9% (11)	3.6% (5)	140
					answered question	**143**
					skipped question	**25**

14. When using a search engine, how often do you use these kinds of temporal queries?

	Other (please specify)	view	8
		answered question	**143**
		skipped question	**25**

15. Do you use a calendar software (outlook's calendar, google calendar, ...)?

		Response Percent	Response Count
yes		**77.5%**	**110**
no		22.5%	32
	Can you please specify why? view		108
		answered question	**142**
		skipped question	**26**

16. Do you find useful to attach files to calendar events?

		Response Percent	Response Count
5 (very useful)		6.3%	9
4		21.1%	30
3 (no opinion)		**38.7%**	**55**
2		11.3%	16
1 (useless)		22.5%	32
	Can you please specify why? view		77
		answered question	**142**
		skipped question	**26**

17. Do you use electronic "to-do" lists (to manage your work, remember things to do,...)?

		Response Percent	Response Count
yes		**54.9%**	**78**
no		45.1%	64
	Can you please specify why? view		87
		answered question	**142**
		skipped question	**26**

18. Do you find useful to attach files to "to-do" lists?

		Response Percent	Response Count
5 (very useful)		7.2%	10
4		12.3%	17
		answered question	**138**
		skipped question	**30**

18. Do you find useful to attach files to "to-do" lists?

	Response Percent	Response Count
3 (no opinion)	**37.7%**	**52**
2	11.6%	16
1 (useless)	31.2%	43
Can you please specify why? view		53
answered question		**138**
skipped question		**30**

19. Do you use electronic notes (to remember things like quotes, web pages, movies to see,...)?

	Response Percent	Response Count
yes	**63.8%**	**90**
no	36.2%	51
Can you please specify why? view		69
answered question		**141**
skipped question		**27**

20. Do you find useful to date electronic notes?

	Response Percent	Response Count
5 (very useful)	12.9%	18
4	21.6%	30
3 (no opinion)	**39.6%**	**55**
2	9.4%	13
1 (useless)	16.5%	23
Can you please specify why? view		45
answered question		**139**
skipped question		**29**

21. Do you make web bookmarks?

	Response Percent	Response Count
yes	**88.5%**	**123**
no	11.5%	16
Can you please specify why? view		75
answered question		**139**
skipped question		**29**

22. To manage your web bookmarks, how often do you use these strategies?

	almost never	rarely	sometimes	often	almost always	Response
					answered question	**134**
					skipped question	**34**

22. To manage your web bookmarks, how often do you use these strategies?

						Count
Organization into a hierarchy: you define a category hierarchy into which you put your bookmarks.	16.4% (22)	9.7% (13)	15.7% (21)	26.9% (36)	**31.3% (42)**	134
Piling: you put your bookmarks in the root category	**28.2% (37)**	18.3% (24)	22.9% (30)	19.1% (25)	11.5% (15)	131
					Other (please specify) view	18
					answered question	**134**
					skipped question	**34**

23. How often do use your web bookmarks?

	Response Percent	Response Count
almost always	25.9%	36
often	**38.1%**	**53**
sometimes	19.4%	27
rarely	8.6%	12
almost never	7.9%	11
	Can you please specify why? view	45
	answered question	**139**
	skipped question	**29**

24. General comment about the survey and/or personal documents management

	Response Count
view	46
answered question	**46**
skipped question	**122**

B

Early Prototypes

This appendix presents early prototypes of a faceted browser and a document-reading assistant developed in the context of this thesis. Parts of the findings and technologies related to them were used in the development of the system presented in Chapter 5.

B.1 WotanEye: Calendar View of Temporal Facet

This section presents a preliminary prototype of the faceted browser introduced in Chapter 5. In particular, this prototype was centered on the temporal facet of personal information, in the context of the *WotanEye* project. In this prototype, the temporal facet of personal information is represented as a calendar-like view rather than as an aggregative timeline view like in the final version of the *Weena* faceted browser (see Chapter 5).

Figure B.1 shows the prototype's main window. On the left part of the figure, information items are plotted according to their last modification date in a calendar. The view represents a full month, where each line shows a week from Monday to Sunday. The week line is divided into four horizontal axes which are used to display (from top to bottom) the user's appointments (rectangles, width depending on appointment duration), documents (circles), and emails (triangles pointing up are outgoing emails whereas triangles pointing down are incoming emails). A traditional search box is available on top of this temporal view and allows textual queries. On the right part of the window lays the document-list view. It lists the documents along with further details like the

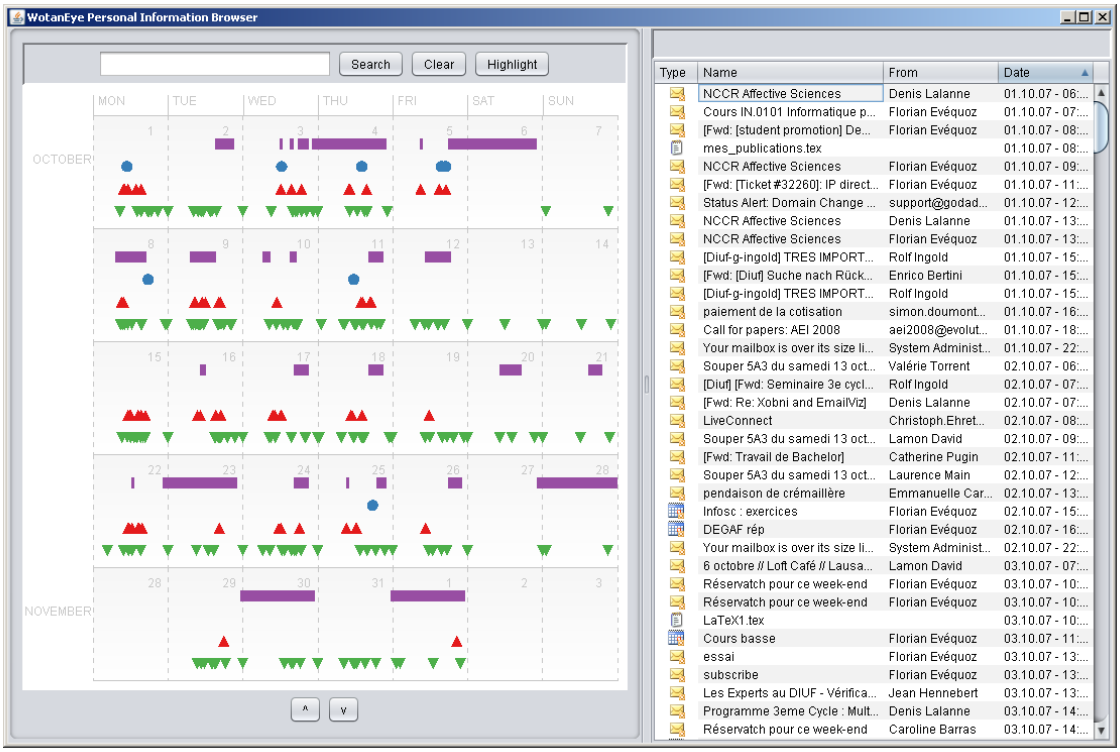

Figure B.1: Global view of *WotanEye* 's calendar view.

document's type, name, its author and timestamp. Issuing a query grays out irrelevant items in the calendar view and removes them from the document-list view. The selection of items displays their name in the calendar view and mark them in the document-list view (see Figure B.2).

Initial plans were to use this calendar view as a basis for the temporal facet of our browser. Figure B.3 depicts a mockup of the attended design. On the middle right is the social facet, in the form of a social network where people are displayed and clustered by communities. Selecting a group of users in the social network highlights pertaining items in the temporal view. Similarly, selecting items in the temporal facet view highlights relevant contacts. On the bottom right is a topic map of the user's personal information, shown as a treemap. Relevant topics are highlighted. The mockup in Figure B.3 shows the result of a query with keywords "Hasler Memodules".

Actually, as the underlying software architecture of this prototype was not flexible enough, adding more facets was unnecessarily difficult. The faceted browser presented in Chapter 5 was thus developed anew. The calendar-like temporal facet could not have been integrated due to software architecture incompatibility and time constraints. Furthermore, visual cluttering was an issue in the calendar view, with possibly overlapping items, thus an aggregative view using a timeline was preferred.

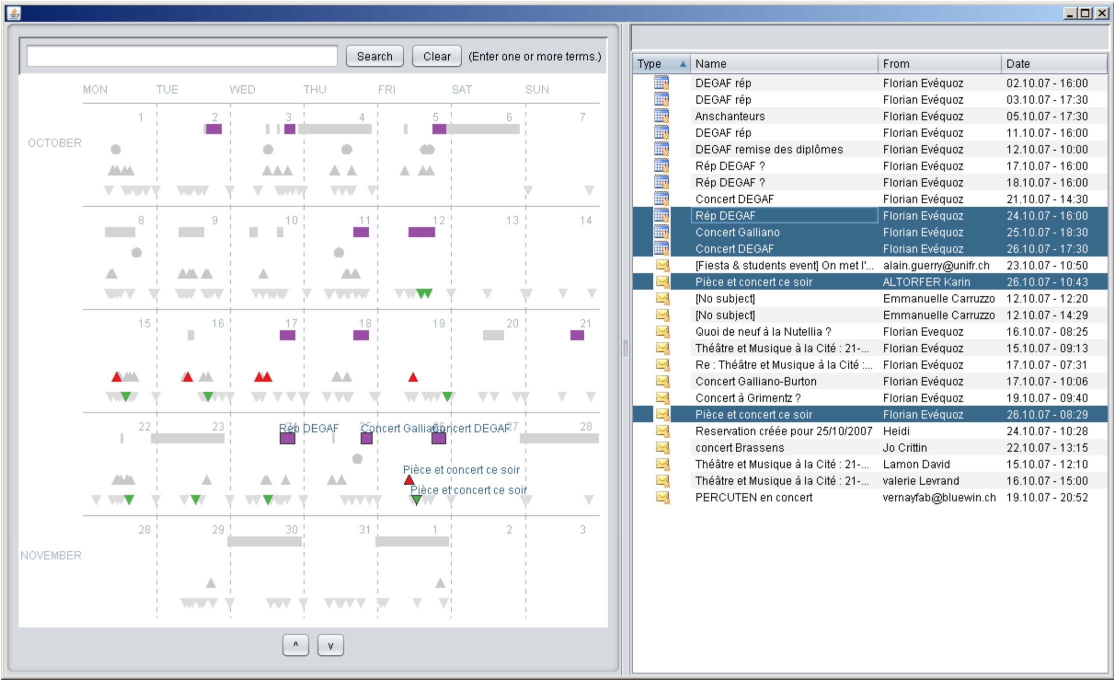

Figure B.2: Selection of items in the calendar view propagates to the document-list view.

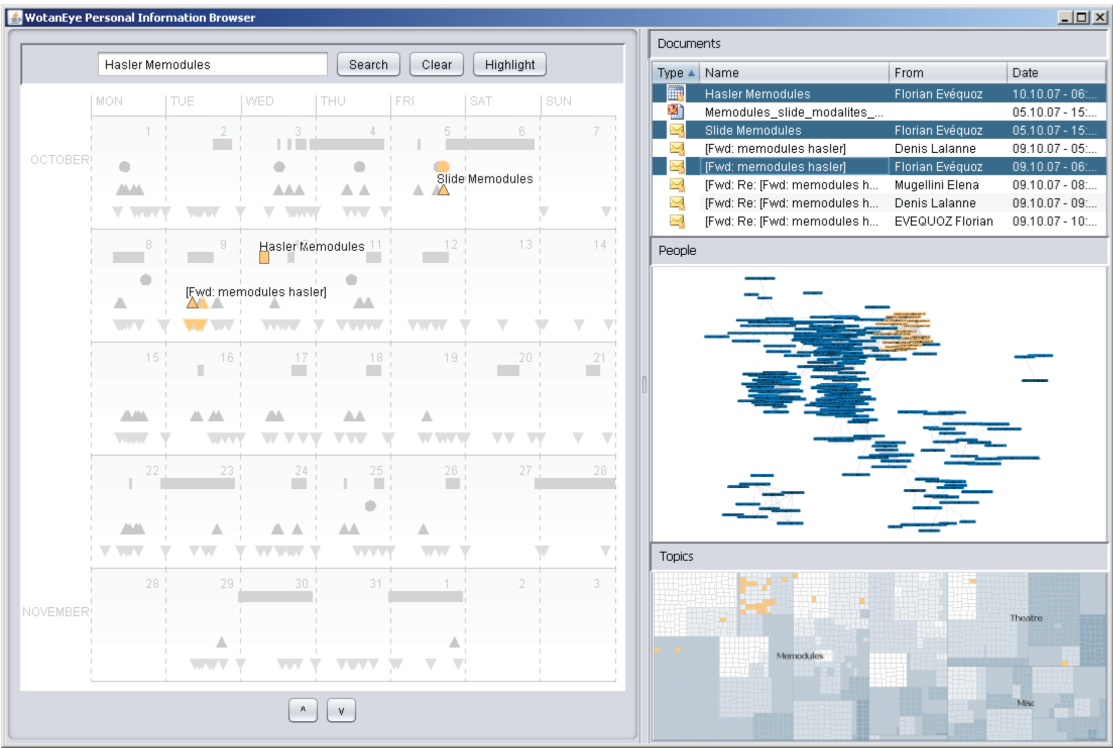

Figure B.3: Mockup of a preliminary faceted navigation interface.

B.2 MyLink: Ego-Centric Document-Reading Assistant

This section describes *MyLink* , a personal document-reading assistant which links document parts with personal information. This software was developed as a BSc thesis by Pierre-Yves Donzallaz under the supervision of the author.In the following, *MyLink* functionalities are illustrated through two scenarios related to meetings: (1) ego-centric meeting browsing and (2) meeting preparation.

B.2.1 Ego-centric meeting browsing

MyLink lets the user open a formatted textual document and see at a glance the parts of this document that will be the most interesting for her. In our scenario, we let the document be related to a meeting (e.g. discussed during the meeting, meeting transcript or minutes). When reviewing a meeting, a user would not be interested in reading the full transcript or minutes. Instead she wants to review in particular (1) the main topic of the meeting and (2) the specific parts concerning her duties.

MyLink works as follows. First, the user selects the document she wants to read. *MyLink* then parses the opened document and separates it down into parts (paragraphs). It then analyses its content, removing stop words (articles, phrasal verbs, etc.) and applies a stemming algorithm on the remaining words in order to count the occurrences of each word stem in every part of the document. The most frequent stems are the salient themes in the document. Once the analysis is done, *MyLink* provides two graphical summaries of the document that were inspired by the works of Eick et al. (1992); Hearst (1995). Those views are displayed next to the document itself.

The "global summary" is a narrow vertical outline of the document made of stacked rectangles. Each rectangle represents a paragraph in the document. The darker the paragraph, the more interesting for the user according to its themes. The global summary thus presents an overview of the document with highlighted zones corresponding to the regions where most salient themes appear. With the help of this view, the user can quickly spot the regions of the document where the main topics of it get discussed. For example, Figure B.4 shows the global summary view of a document discussed during a meeting. This meeting was about the use of emails in personal information management. The prominent theme of this document is "email". Other minor themes also appear, like "management", "organization", "user". Paragraphs with the most occurrences of those themes get darkened in the global summary view. Figure B.5 depicts the view of a meeting transcript opened in the interface.

User's main themes can be either automatically extracted (e.g. by using text retrieval methods), or alternatively, the user can create a reading profile. She inputs the keywords she is interested in, and let the global summary view present these only. Therefore, she can focus on the regions of current interest (combine or not with more persistent interests automatically derived from her personal information using text retrieval methods).

The main document window and the global summary view are synchronized. Clicking on a region in the summary scrolls the document window down to the corresponding paragraph in the document. As well, moving the viewport in the document view moves a marker (white highlighting in the margins) in the global summary view to show which part of the document is currently visible. Thus, one can focus on a specific part of the document in the main view while keeping an eye on the document's outline in the global summary view.

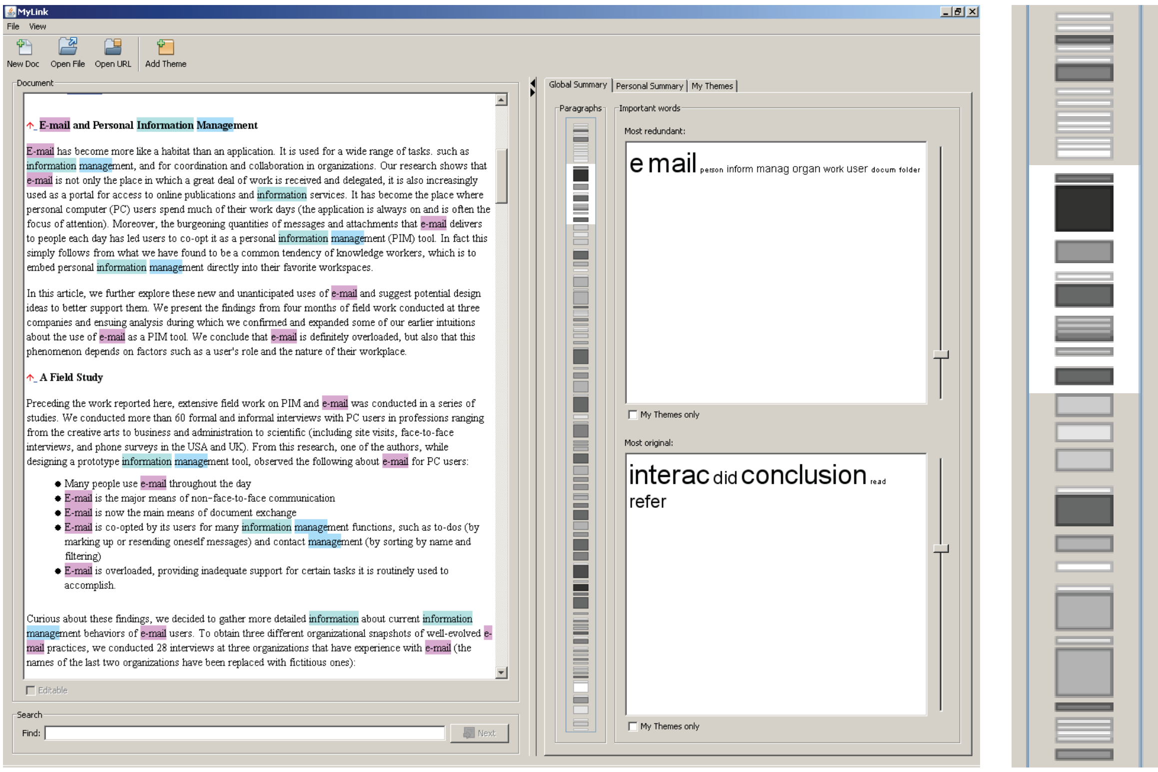

Figure B.4: (a) *MyLink* main window. A document is opened and its global summary, as well as salient themes are displayed. (b) Cropped view of the global summary. The most interesting regions are darker.

The second summary view that *MyLink* proposes is called the "personal summary view". The focus can be put on a particular theme or on a combination of themes. Figure B.6 depicts this view applied to a meeting transcript. The meeting was a movie club deciding which movie to project on their next session. Each colored bar on this figure is associated to a user theme or selected keyword. The bar width is proportional to the importance of the associated theme in the document (i.e. its amount of occurrences). The bar is also a miniature view of the document, focused on this

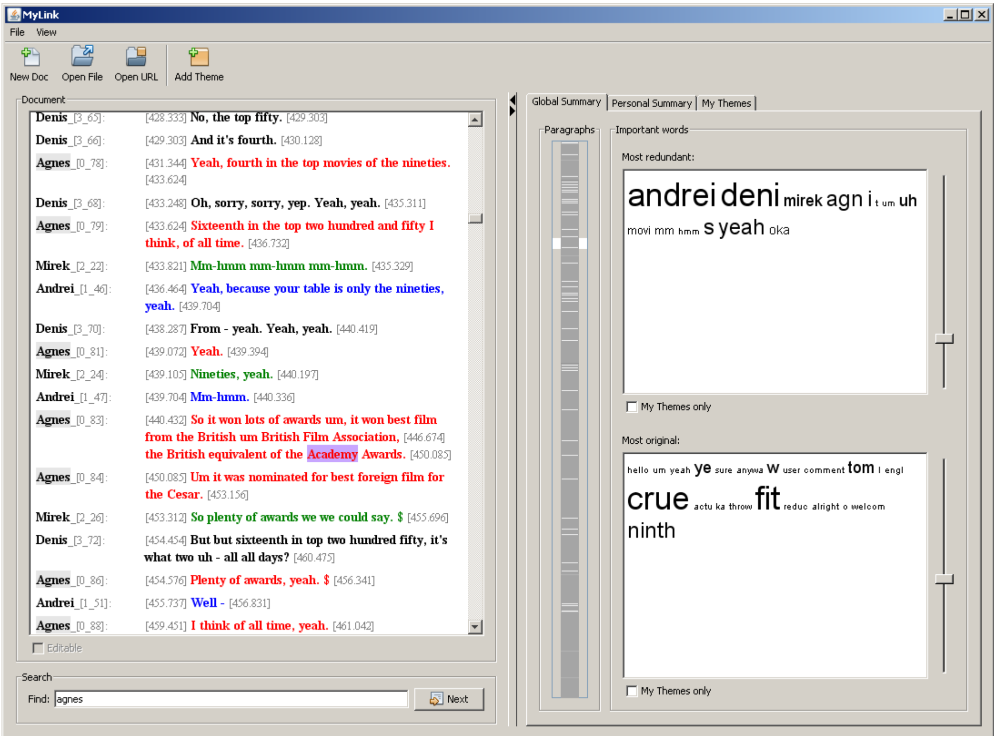

Figure B.5: *MyLink* 's global summary view of a meeting transcript.

keyword. Dark horizontal rectangles on the bar map to the specific positions within the document where this theme is discussed. Clicking on a dark rectangle scrolls to the corresponding part of the document. Alternatively, by putting the bars of two (or more) interesting keywords next to each other, one can quickly spot areas of the document where they appear together. For example, in Figure B.6, the user chose to focus on three keywords that were discussed during the movie club meeting, namely the movies *Private* Ryan and Usual *Suspects* and what was said about *Academy* Awards (emphasis is put on the keyword used to build the personal summary). From the personal summary view, she can quickly go to the interesting part of the meeting and play the associated audio/video recording.

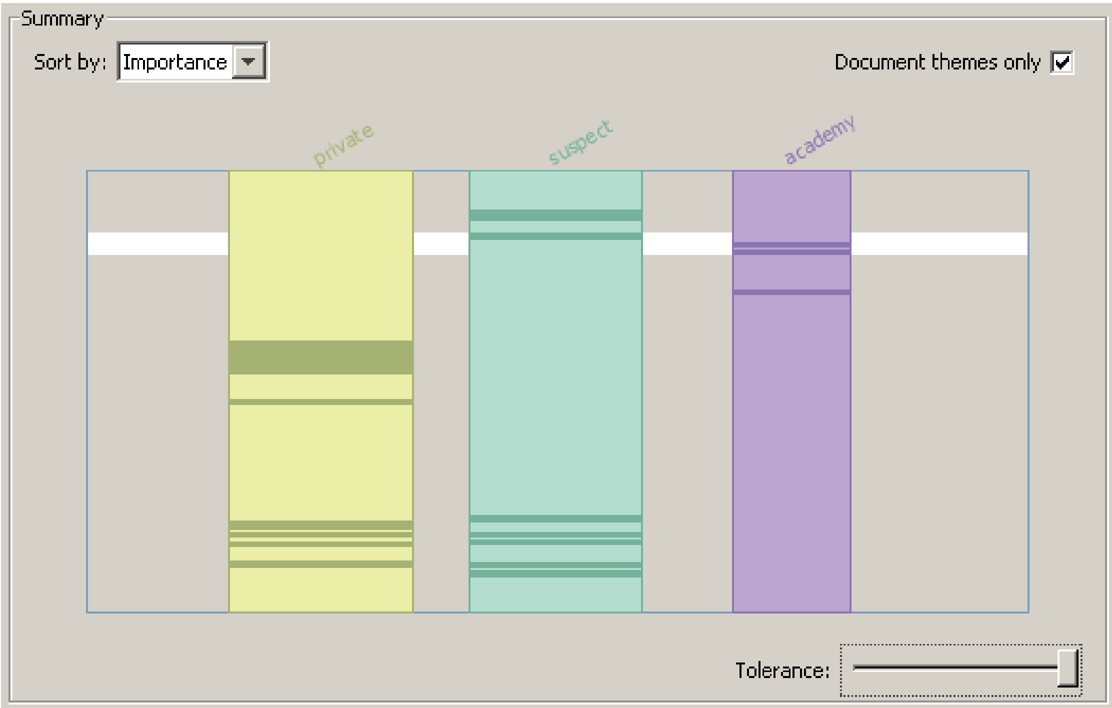

Figure B.6: Cropped view of *MyLink* 's personal summary view of a meeting transcript.

B.2.2 Meeting preparation

In the meeting preparation scenario, a meeting participant wants to gather all necessary documents in order to be ready for an upcoming meeting. Using *MyLink* to open the meeting agenda, this gathering task can be made easier. Indeed, *MyLink* automatically shows a list of all personal documents (emails, files, appointments, etc.) that are related to the themes currently examined within a given document. The actual linking with the personal documents is done leveraging Windows Desktop Search (WDS). When the document viewer scrolls over a new region, *MyLink* builds queries using the salient keywords associated with the paragraph, send them to the WDS engine, and displays results along the currently opened document.

B.2.3 User feedback

MyLink has been tested informally by several users. We asked them to think aloud while using the program. They understood quickly the use of the global and personal summary to spot the most interesting regions of the document, according to one or more keywords. They formulated some criticisms, too. In particular, they wanted to see the global and personal summaries together on the same panel, which is not possible with our current prototype. Moreover, people would have liked to create several "profiles" (set of keywords) and be able to select them according to the document they want to read.

C

Evaluation of the Faceted Browser

This appendix contains the documents related to the evaluation presented in Chapter 6. All documents, except the one in Appendix C.4, have been authored by Julien Thomet as part of his MSc thesis done under my supervision. The evaluation was conducted by Julien Thomet and me.

C.1 Description of Evaluation

Évaluation utilisateur de *Weena*

1 Déroulement

Dans le cadre de mon travail de master, j'ai conçu et développé une application, *Weena*, dont le but est d'améliorer l'efficacité de la recherche et de la navigation de documents personnels.

Vous avez donc été sollicité afin d'évaluer concrètement cette application. La méthode d'évaluation se déroulera en deux étapes :

	Étape	*Durée*
1.	Rédaction d'un *diary study*	quelques minutes par jours pendant une semaine
2.	Évaluation contrôlée de *Weena*	environ 1h

1.1 Rédaction d'un *diary study*

Un *diary study* est un journal (au sens « carnet de bord ») dans lequel l'utilisateur répertorie un certain nombre d'informations spécifiques à l'activité évaluée (par exemple la recherche d'informations personnelles).

Dans cette étude, le *diary study* consistera en une **liste d'au minimum 10 tâches** concernant la recherche et la navigation de documents personnels auxquelles seront associées des **notes de difficulté**. L'idée d'un tel document est de lister les tâches auxquelles vous êtes confronté tous les jours en utilisant vos méthodes de recherche et vos outils habituels. Le tableau ci-dessous expose un exemple de *diary study* dans le cadre de ce projet :

Tâche	Difficulté
Je cherche une lettre au format Word que j'utilise comme modèle pour toutes mes autres lettres professionnelles. Le « modèle » est adressé à l'entreprise X. J'ai besoin de cette lettre pour en créer une similaire pour l'entreprise Y.	☐ 1 ☐ 2 ☐ 3 ☒ 4 ☐ 5
Retrouver tous les articles concernant la reconnaissance de geste téléchargés il y a quelques mois pour un séminaire sur le sujet afin de les imprimer.	☒ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5
Je cherche un logiciel que j'avais utilisé pour faire créer un proxy et dont j'avais peut-être gardé le fichier d'installation.	☐ 1 ☐ 2 ☐ 3 ☐ 4 ☒ 5
Trouver tous les fichiers musicaux téléchargés récemment et qui se trouvent à différents endroits sur mon disque dur afin de les écouter.	☐ 1 ☐ 2 ☒ 3 ☐ 4 ☐ 5
J'ai besoin de récupérer la liste des membres de l'association culturelle A (dont je suis membre) que le secrétaire a rédigé et envoyé à tous les membres en début de saison.	☐ 1 ☒ 2 ☐ 3 ☐ 4 ☐ 5

Tableau 1: Exemple de *diary study*

Les tâches listées doivent être assez précises concernant le type de documents et leurs caractéristiques utiles à la recherche/navigation.

La **note de difficulté** est à considérer **a posteriori** ; idéalement, lorsque vous avez une tâche à réaliser, décrivez-la dans le *diary study*, réalisez-la et ensuite inscrivez une note évaluant la difficulté de la tâche (nombre de clics effectués, temps…).

Pour remplir ce document le meilleur moyen est de l'**avoir à portée de main chaque fois que vous utilisez votre ordinateur** et d'y inscrire les de tâches de recherche et de navigation de document que vous effectuez.

1.2 Évaluation de *Weena*

Dans la seconde partie de l'évaluation, une liste de tâches vous sera proposée. *Weena* aura préalablement été installé et configuré sur votre ordinateur personnel et les tâches seront réalisées **à l'aide de cette application uniquement**.

Lors de l'utilisation de *Weena*, certaines informations seront récupérées pour analyse. Ces informations resteront strictement anonymes et ne contiendront que des statistiques sur les nombres de documents filtrés (en tout, par type et par date) et sur l'interaction avec l'application (type de sélection, nombre d'éléments sélectionnés…). Les données enregistrées seront visibles dans un fichier texte que vous pourrez contrôler à l'issue de l'évaluation pour vous assurer de leur anonymat.

2 Système requis pour l'installation de *Weena*

Système **Windows** (XP, Vista, 7) avec l'indexation de **Windows Desktop Search** activée. Windows Desktop Search est intégré à Windows Vista et 7. Il doit cependant être installé sur des machines avec Windows XP.

Pour pouvoir utiliser pleinement *Weena*, il faut également avoir accès à une messagerie supportant le protocole **IMAP** par exemple, *bluewin.ch*, *unifr.ch*, *gmail.com*…

Julien Thomet
[julien.thomet@unifr.ch]

Annexe : modèle *diary study* à imprimer ou remplir électroniquement

C.2 Overview of *Weena*

Aperçu général de *Weena*

Weena est une application de gestion d'informations personnelles, plus particulièrement de la recherche et de la navigation d'informations personnelles. L'image suivante présente l'interface de *Weena* et ces différentes parties.

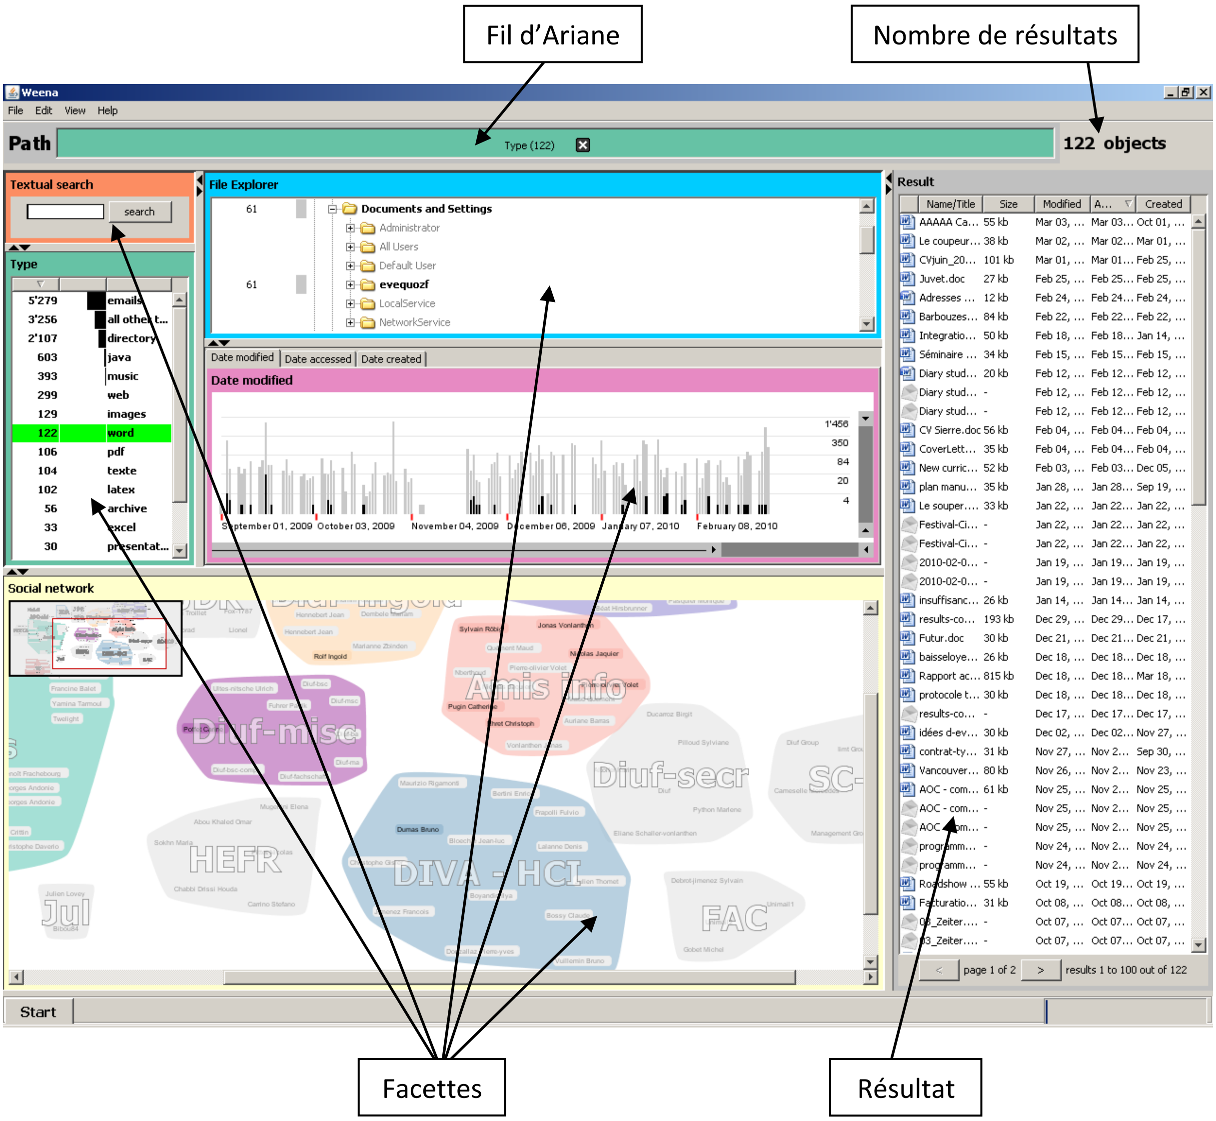

Fil d'Ariane : présente les facettes sélectionnées pour arriver au résultat courant. Il permet également d'annuler les sélections précédentes.

Nombre de résultats : donne une indication sur le nombre d'éléments présent dans le panneau de résultat.

Résultat : présente les éléments filtrés par la sélection de facettes. Ce panneau se comporte comme la vue de droite de l'explorateur de fichiers de *Windows*.

Facettes : chaque facette présente les documents selon une caractéristique différente (type du document, date de création…) et permet de filtrer le résultat en sélectionnant les éléments idoines. Les éléments sélectionnables des facettes sont en **noir** et les éléments désactivés en **gris**. Lorsqu'une sélection est réalisée dans une facette, celle-ci est surlignée en **vert**. Une seule facette est sélectionnée à la fois ; lorsqu'une facette a été sélectionnée et qu'une nouvelle sélection est réalisée, l'ancienne facette est désélectionnée.

Sept facettes sont disponibles :

- La facette ***File explorer*** qui se comporte comme la vue de gauche de l'explorateur de fichiers de *Windows*, c'est-à-dire que l'arborescence des disques durs de l'ordinateur est présentée sous forme d'arbre. Vis-à-vis des dossiers apparaît le nombre de documents présents dans ses sous-dossiers qui correspondentà la requête courante. Le rectangle associé à ce nombre donne une indication supplémentaire sur le nombre initial de documents présents dans les dossiers. Ces indicateurs apparaissent également dans la facette *Type*. Il est à noter que les e-mails issus d'*Outlook* n'apparaissent pas dans cette hiérarchie.
- La facette ***Textual search*** permet de rechercher directement des expressions dans les documents disponibles.
- La facette ***Type*** permet de filtrer par type de fichiers.
- Les facettes temporelles ***Date modified***, ***Date created*** et ***Date accessed*** présentent des histogrammes (semblables aux indicateurs des facettes *File explorer* et *Type*) du nombre de documents respectivement modifiés, créés ou accédés. Les éléments sélectionnables sont les « barres » de l'histogramme.
- La facette ***Social network*** présente votre réseau social. Les différents contacts extraits de votre boîte mail sont organisés par communauté. Les contacts affichés en gris ne sont pas concernés par la sélection courante sur les différentes facettes et ne sont par conséquent pas sélectionnables.

C.3 *Weena* Questionnaire

Questionnaire final

1) Trouvez-vous *Weena* **utile** ?

(1 inutile - 5 très utile) 1 2 3 4 5

2) Trouvez-vous *Weena* **utilisable** ?

(1 inutilisable - 5 très utilisable) 1 2 3 4 5

3) **Classez** les facettes par **préférence**[1].

☐ File explorer ☐ Date accessed ☐ Textual search ☐ Date created

☐ Social network ☐ Date modified ☐ Type

4) **Classez** les facettes par **efficacité**[2].

☐ File explorer ☐ Date accessed ☐ Textual search ☐ Date created

☐ Social network ☐ Date modified ☐ Type

5) Qu'est ce qui vous **plaît** dans *Weena* par rapport à vos outils habituels ?

6) Qu'est-ce qui ne vous **plaît pas** dans *Weena* par rapport à vos outils habituels ?

[1] 1=votre facette préférée
[2] 1=la facette la plus efficace

7) Dans quel(s) cas la **facette sociale** (*Social network*) est-elle utile pour vous ?

8) Dans quel(s) cas les **facettes temporelles** (*Date modified*, *Date accessed*, *Date created*) sont-elles utiles pour vous ?

9) Utiliseriez-vous *Weena* ? Pour quel type de tâches ?

10) Commentaire général sur *Weena* et son évaluation

C.4 Social Facet Questionnaire

Questionnaire 'Facette sociale'

1) Trouvez-vous **compréhensible** la représentation graphique de vos contacts sous forme d'un graphe ou d'une carte sociale ?
(1 incompréhensible - 5 très compréhensible) 1 2 3 4 5

Commentaires

2) Trouvez-vous **utile** la représentation graphique de vos contacts sous forme d'un graphe ou d'une carte sociale ?
(1 inutile - 5 très utile) 1 2 3 4 5

Commentaires

3) Trouvez-vous **utile** que vos contacts soient regroupés en **communautés** ?
(1 inutile - 5 très utile) 1 2 3 4 5

Commentaires

4) Avez-vous découvert quelque chose de nouveau dans votre réseau social grâce à cette représentation graphique ?

5) Quelles autres applications pratiques de cette représentation graphique de votre réseau social pourriez-vous imaginer ?

Merci beaucoup pour votre participation.

C.5 Evaluation Protocol

Weena pas à pas

1. File explorer : expliquer l'histogramme, expliquer les couleurs (gris,noir,vert) (et que c'est partout la même chose), raccourcis de l'explorateur (dans le menu *view*)
2. Panneau de résultat : montrer qu'une sélection dans le FE influence le résultat. Présenter le panneau : trie, dates, tooltip, pagination
3. Mise à jour des facettes lors de la sélection
4. Sélection de type (histogramme)
5. Fil d'Ariane :
    a. Affiche les facettes sélectionnées avec le nombre de résultats avec les couleurs idoines + tooltip
    b. Undo/Redo
6. Undo de tout : montre comment on recommence une nouvelle requête
7. Facette sociale
    a. Récap des couleurs
    b. Sélection : clic sur un nom, clic sur une communauté, sélection d'ensemble, +shift
8. Facette temporelle
    a. Représente le nombre de fichiers
    b. Récap des couleurs (gris clair, gris foncé, noir, vert)
    c. Navigation
    d. Sélection d'intervalles (item, intervalle, + ctrl)
    e. Vue des timelines toutes en même temps (ctrl + t ou *view*)
9. Textual search

Bibliography

Abrams, David, Ron Baecker, and Mark Chignell (1998). *Information Archiving With Bookmarks: Personal Web Space Construction and Organization*. In *CHI '98: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 41–48. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA. ISBN 0201309874. doi:10.1145/274644.274651.

Allen, David (2002). *Getting Things Done*. Penguin Books. ISBN 0142000280.

Andrews, Keith (2008). *Evaluation Comes in Many Guises*. In *AVI Workshop on BELIV, Position Paper*. `http://ftp.iicm.tugraz.at/pub/papers/andrews-beliv08.pdf`.

Bälter, Olle (1997). *Strategies for Organising Email*. In *HCI 97: Proceedings of HCI on People and Computers XII*, pages 21–38. Springer-Verlag, London, UK. ISBN 3540761721.

Barreau, Deborah (2008). *The Persistence of Behavior and Form in the Organization of Personal Information*. *J. Am. Soc. Inf. Sci. Technol.*, 59(2), pages 307–317. ISSN 1532-2882. doi:10.1002/asi.20752.

Barreau, Deborah and Bonnie A. Nardi (1995). *Finding and Reminding: File Organization from the Desktop*. *SIGCHI Bull.*, 27(3), pages 39–43. ISSN 0736-6906. doi:10.1145/221296.221307.

Barreau, Deborah, Laura O'Neill, and Amanda Stevens (2009). *The Relationship Between Research and Practice: What Are We Learning About Teaching Personal Information Management?* In *ASIS&T PIM Workshop, Vancouver, BC, Canada*. `http://pimworkshop.org/2009/papers/barreau-pim2009.pdf`.

Barreau, Deborah K. (1995). *Context as a Factor in Personal Information Management Systems*. *J. Am. Soc. Inf. Sci.*, 46(5), pages 327–339. ISSN 0002-8231. doi:10.1002/(SICI)1097-4571(199506)46:5<327::AID-ASI4>3.0.CO;2-C.

Bell, David and Ian Ruthven (2004). *Searcher's Assessments of Task Complexity for Web Searching*. In *Proceedings of ECIR 04*, pages 57–71. Springer Berlin / Heidelberg, Sunderland. doi:10.1007/978-3-540-24752-4_5.

Bergman, O., R. Beyth-Marom, R. Nachmias, N. Gradovitch, and S. Whittaker (2008a). *Type of Search Engine and Personal Information Retrieval: Windows vs. Linux*. In *Proceeding of Chais conference for learning technologies*, pages 17–22. Ranana, Israel. http://telem-pub.openu.ac.il/users/chais/2008/after_noon/3_2.pdf.

Bergman, Ofer, Ruth Beyth-Marom, and Rafi Nachmias (2008b). *The User-Subjective Approach to Personal Information Management Systems Design: Evidence and Implementations*. *Journal of the American Society for Information Science and Technology*, 59(2), pages 235–246. doi:10.1002/asi.20738.

Bergman, Ofer, Ruth Beyth-Marom, Rafi Nachmias, Noa Gradovitch, and Steve Whittaker (2008c). *Improved Search Engines and Navigation Preference in Personal Information Management*. *ACM Trans. Inf. Syst.*, 26(4), pages 1–24. ISSN 1046-8188. doi:10.1145/1402256.1402259.

Berlowitz, Ilana and Cornelia Gautschi (2010). *WotanEye: A Faceted View of Personal Information, Usability Test*. Technical Report, University of Fribourg, Switzerland.

Bertini, Enrico and Denis Lalanne (2007). *TotalRecall Survey Report*. Technical report, University of Fribourg. http://diuf.unifr.ch/people/lalanned/Articles/TR-survey-report0807.pdf.

Blakeslee, Sandra (1976). *La surcharge informationnelle, un danger pour les nantis*. *La Liberté*, 1, pages 1–2. Edited in Fribourg, Switzerland.

Boardman, Richard (2004). *Improving Tool Support for Personal Information Management*. Electrical and electronic engineering, Imperial College, London. http://www.iis.ee.ic.ac.uk/~rick/thesis/boardman04-thesis.pdf.

Boardman, Richard and M. Angela Sasse (2004). *"Stuff Goes Into the Computer and Doesn't Come Out": A Cross-Tool Study of Personal Information Management*. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 583–590. ACM Press, New York, NY, USA. ISBN 1581137028. doi:10.1145/985692.985766.

Boardman, Richard, Robert Spence, and M. Angela Sasse (2003). *Too Many Hierarchies? The Daily Struggle for Control of the Workspace*. In *HCI International 2003 : 10th International Conference on Human-Computer Interaction*, volume 1, pages 616–620. Crete, Greece. http://www.iis.ee.ic.ac.uk/~rick/research/pubs/struggle-hcii2003.pdf.

Boutin, Paul (2004). *Keeper Finders, five new programs that let you search your hard drive without having a seizure*. http://www.slate.com/id/2111643/. (Retrieved October 12, 2010).

Brooke, J. (1996). *SUS: A Quick and Dirty Usability Scale*. *Usability Evaluation in Industry, P. W. Jordan, B. Thomas, B. A. Weerdmeester, and I. L. McClelland, Eds. London: Taylor & Francis*. http://www.itu.dk/courses/U/E2005/litteratur/sus.pdf.

Bush, Vannevar (1945). *As We May Think*. *The Atlantic Monthly*, 176(1), pages 101–108. http://www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/3881/.

Byström, Katriina and Kalervo Järvelin (1995). *Task Complexity Affects Information Seeking and Use*. *Information Processing & Management*, 31(2), pages 191–213. ISSN 0306-4573. doi:10.1016/0306-4573(95)80035-R.

Capra, Robert and Manuel A. Pérez-Quiñones (2003). *Re-Finding Found Things: An Exploratory Study of How Users Re-Find Information*. Technical Report, Computer Science Dept., Virginia Tech. http://arxiv.org/ftp/cs/papers/0310/0310011.pdf.

Capra, Robert and Manuel A. Pérez-Quiñones (2005). *Using Web Search Engines to Find and Refind Information*. *Computer*, 38(10), pages 36–42. ISSN 0018-9162. doi:10.1109/MC.2005.355.

Capra, Robert G. and Gary Marchionini (2008). *The Relation Browser Tool for Faceted Exploratory Search*. In *JCDL '08: Proceedings of the 8th ACM/IEEE-CS joint conference on Digital libraries*, pages 420–420. ACM, New York, NY, USA. ISBN 1595939982. doi:10.1145/1378889.1378967.

Card, Stuart K., Jock D. Mackinlay, and Ben Shneiderman (Editors) (1999). *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. ISBN 1558605339.

Chau, Duen Horng, Brad Myers, and Andrew Faulring (2008a). *Feldspar: A System for Finding Information by Association*. In *PIM 2008 - Proceedings of Personal Information Management, Workshop at CHI 2008*. Florence, Italy. http://pim2008.ethz.ch/papers/pim2008-chau-etal.pdf.

Chau, Duen Horng, Brad Myers, and Andrew Faulring (2008b). *What to Do When Search Fails: Finding Information by Association*. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 999–1008. ACM, New York, NY, USA. ISBN 1605580111. doi:10.1145/1357054.1357208.

Chi, Ed (1999). *A Framework for Information Visualization Spreadsheets*. PhD Thesis, University of Minnesota. http://www-users.cs.umn.edu/~echi/phd/chi-thesis.pdf.

Clauset, Aaron, M. E. J. Newman, and Cristopher Moore (2004). *Finding Community Structure in Very Large Networks*. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, 70(6), 066111. doi:10.1103/PhysRevE.70.066111.

Cole, I. (1982). *Human Aspects of Office Filing: Implications for the Electronic Office*. In *Proceedings of the Human Factors Society - 26th Annual Meeting*. Seattle, WA. ISSN 1071-1813.

Cselle, Gabor, Keno Albrecht, and Roger Wattenhofer (2007). *BuzzTrack: Topic Detection and Tracking in Email*. In *IUI '07: Proceedings of the 12th international conference on Intelligent user interfaces*, pages 190–197. ACM Press, New York, NY, USA. ISBN 1595934812. doi:10.1145/1216295.1216331.

Cutrell, Edward, Susan T. Dumais, and Jaime Teevan (2006a). *Searching to Eliminate Personal Information Management*. *Commun. ACM*, 49(1), pages 58–64. ISSN 0001-0782. doi:10.1145/1107458.1107492.

Cutrell, Edward, Daniel Robbins, Susan Dumais, and Raman Sarin (2006b). *Fast, Flexible Filtering with Phlat*. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 261–270. ACM, New York, NY, USA. ISBN 1595933727. doi:10.1145/1124772.1124812.

Dourish, Paul, W. Keith Edwards, Anthony LaMarca, John Lamping, Karin Petersen, Michael Salisbury, Douglas B. Terry, and James Thornton (2000). *Extending Document Management Systems with User-Specific Active Properties*. *ACM Trans. Inf. Syst.*, 18(2), pages 140–170. ISSN 1046-8188. doi:10.1145/348751.348758.

Dredze, Mark, Tessa Lau, and Nicholas Kushmerick (2006). *Automatically Classifying Emails into Activities*. In *IUI '06: Proceedings of the 11th international conference on Intelligent user interfaces*, pages 70–77. ACM Press, New York, NY, USA. ISBN 1595932879. doi:10.1145/1111449.1111471.

Dumais, Susan, Edward Cutrell, JJ Cadiz, Gavin Jancke, Raman Sarin, and Daniel C. Robbins (2003). *Stuff I've seen: A System for Personal Information Retrieval and Re-use*. In *SIGIR '03: Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pages 72–79. ACM Press, New York, NY, USA. ISBN 1581136463. doi:10.1145/860435.860451.

Eick, Stephen G., Joseph L. Steffen, and Eric E. Sumner, Jr. (1992). *Seesoft-A Tool for Visualizing Line Oriented Software Statistics*. *IEEE Trans. Softw. Eng.*, 18(11), pages 957–968. ISSN 0098-5589. doi:10.1109/32.177365.

Ellis, Geoff and Alan Dix (2008). *An Explorative Analysis of User Evaluation Studies in Information Visualisation*. In *AVI Workshop on BELIV, Position Paper*. doi:10.1145/1168149.1168152.

Elsweiler, D. (2007). *Supporting Human Memory in Personal Information Management*. PhD Thesis, Department of Computer and Information Sciences, University of Strathclyde. http://www.cis.strath.ac.uk/cis/research/publications/papers/strath_cis_publication_2328.pdf.

Elsweiler, David and Ian Ruthven (2007). *Towards Task-Based Personal Information Management Evaluations*. In *SIGIR '07: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 23–30. ACM, New York, NY, USA. ISBN 1595935977. doi:10.1145/1277741.1277748.

Elsweiler, David, Ian Ruthven, and Christopher Jones (2007). *Towards Memory Supporting Personal Information Management Tools*. *J. Am. Soc. Inf. Sci. Technol.*, 58(7), pages 924–946. ISSN 1532-2882. doi:10.1002/asi.20570.

Evéquoz, Florian (2009). *An Ethnographically-Inspired Survey of PIM Strategies*. Technical Report, University of Fribourg, DIUF. http://diuf.unifr.ch/diva/publications/evequoz_techreport09.pdf.

Evéquoz, Florian and Denis Lalanne (2007a). *Indexing and Visualizing Digital Memories Through Personal Email Archive*. In *Supporting Human Memory with Interactive Systems, Workshop at HCI 2007 (HCI 2007)*, pages 21–24. Lancaster (UK). http://diuf.unifr.ch/diva/publications/evequoz_hci07.pdf.

Evéquoz, Florian and Denis Lalanne (2007b). *Personal Information Management Through Interactive Visualizations*. In *Poster Compendium of the IEEE Information Visualization Conference (Doctoral Colloquium InfoVis 2007)*, pages 158–160. Sacramento, CA (USA). http://diuf.unifr.ch/diva/publications/evequoz_infovis07.pdf.

Evéquoz, Florian and Denis Lalanne (2009). *"I Thought You Would Show Me How to Do It" - A Study on Self-Analysis of PIM Practices*. In *ASIS&T PIM Workshop, Vancouver, BC, Canada.* http://pimworkshop.org/2009/papers/evequoz-pim2009.pdf.

Evéquoz, Florian, Julien Thomet, and Denis Lalanne (2010). *La navigation par facettes appliquée à la gestion de l'information personnelle*. In *IHM'10: Proceedings of the 22nd International Conference on Association Francophone d'Interaction Homme-Machine*, pages 41–48. ACM Press, Luxembourg. ISBN 1450304109. http://diuf.unifr.ch/diva/publications/evequoz_ihm10.pdf.

Fertig, Scott, Eric Freeman, and David Gelernter (1996). *Lifestreams: An Alternative to the Desktop Metaphor*. In *CHI '96: Conference companion on Human factors in computing systems*, pages 410–411. ACM, New York, NY, USA. ISBN 0897918320. doi:10.1145/257089.257404.

Few, Stephen (2007). *Infovis as Seen By the World Out There*. http://www.perceptualedge.com/articles/11-13-07.pdf. Infovis 2007 Capstone presentation.

Fitzpatrick, Geraldine (2009). *Applying an Ethnographic Approach in Design*. Talk given at 3eme cycle romand in Computer Science (CUSO Winter School).

Forsythe, Diana E. (1999). *"It's Just a Matter of Common Sense": Ethnography as Invisible Work*. *Comput. Supported Coop. Work*, 8(1-2), pages 127–145. ISSN 0925-9724. doi:10.1023/A:1008692231284.

Freeman, Eric and David Gelernter (1996). *Lifestreams: A Storage Model for Personal Data*. *SIGMOD Rec.*, 25(1), pages 80–86. ISSN 0163-5808. doi:10.1145/381854.381893.

Fruchterman, Thomas M. J. and Edward M. Reingold (1991). *Graph Drawing by Force-directed Placement*. *Software: Practice and Experience*, 21(11), pages 1129–1164. doi:10.1002/spe.4380211102.

Gantz, John F. and David Reinsel (2010). *The Digital Universe Decade – Are you Ready ?* IDC, EMC Corporation, http://www.emc.com/digital_universe.

Gemmell, Jim, Gordon Bell, and Roger Lueder (2006). *MyLifeBits: A Personal Database for Everything*. *Commun. ACM*, 49(1), pages 88–95. doi:10.1145/1107458.1107460.

Gemmell, Jim, Gordon Bell, Roger Lueder, Steven Drucker, and Curtis Wong (2002). *MyLifeBits: Fulfilling the Memex Vision*. In *MULTIMEDIA '02: Proceedings of the tenth ACM international conference on Multimedia*, pages 235–238. ACM Press, New York, NY, USA. ISBN 158113620X. doi:10.1145/641007.641053.

Golemati, Maria, Akrivi Katifori, Eugenia G. Giannopoulou, Ilias Daradimos, and Costas Vassilakis (2007). *Evaluating the Significance of the Windows Explorer Visualization in Personal Information Management Browsing tasks*. In *11th International Conference on Information Visualisation, IV 2007*, pages 93–100. Zürich, Switzerland. doi:10.1109/IV.2007.46.

Hartel, Jenna (2006). *Information Activities and Resources in an Episode of Gourmet Cooking*. *Information Research*, 12(1). http://informationr.net/ir/12-1/paper282.html.

Hearst, Marti A. (1995). *TileBars: Visualization of Term Distribution Information in Full Text Information Access*. In *CHI '95: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 59–66. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA. ISBN 0201847051. doi:10.1145/223904.223912.

Heer, Jeffrey and Danah Boyd (2005). *Vizster: Visualizing Online Social Networks*. *Information Visualization, IEEE Symposium on*, 0, page 5. ISSN 1522-404x. doi:10.1109/INFOVIS.2005.39.

Henry, Nathalie, Jean-Daniel Fekete, and Michael J. McGuffin (2007). *NodeTrix: a Hybrid Visualization of Social Networks*. volume 13, pages 1302–1309. IEEE Computer Society, Los Alamitos, CA, USA. ISSN 1077-2626. doi:10.1109/TVCG.2007.70582.

Hurst, Mark (2007). *Bit Literacy*. Good Experience Press. ISBN 0979368103.

Jones, William (2007). *Keeping Found Things Found: The Study and Practice of Personal Information Management*. Morgan Kaufmann. ISBN 0123708664.

Jones, William, Predrag Klasnja, Andrea Civan, and Michael L. Adcock (2008). *The Personal Project Planner: Planning to Organize Personal Information*. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 681–684. ACM, New York, NY, USA. ISBN 1605580111. doi:10.1145/1357054.1357162.

Jones, William, Ammy Jiranida Phuwanartnurak, Rajdeep Gill, and Harry Bruce (2005). *Don't Take my Folders Away!: Organizing Personal Information to Get Things Done*. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*, pages 1505–1508. ACM, New York, NY, USA. ISBN 1595930027. doi:10.1145/1056808.1056952.

Jones, William and Jaime Teevan (Editors) (2007). *Personal Information Management*. University of Washington Press, Seattle. ISBN 0295987375.

Kaptelinin, Victor (2003). *UMEA: Translating Interaction Histories into Project Contexts*. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 353–360. ACM, New York, NY, USA. ISBN 1581136307. doi:10.1145/642611.642673.

Kaptelinin, Victor and Mary Czerwinski (Editors) (2007). *Beyond the Desktop Metaphor*. The MIT Press, Cambride, Massachusetts. ISBN 026211304X.

Karger, David R., Karun Bakshi, David Huynh, Dennis Quan, and Vineet Sinha (2005). *Haystack: A General-Purpose Information Management Tool for End Users Based on Semistructured Data*. In *CIDR*, pages 13–26. `http://www-db.cs.wisc.edu/cidr/cidr2005/papers/P02.pdf`.

Kelly, Diane (2006). *Evaluating Personal Information Management Behaviors and Tools*. *Commun. ACM*, 49(1), pages 84–86. ISSN 0001-0782. doi:10.1145/1107458.1107497.

Kim, Pilho, Mark Podlaseck, and Gopal Pingali (2004). *Personal Chronicling Tools for Enhancing Information Archival and Collaboration in Enterprises*. In *CARPE'04: Proceedings of the the 1st ACM workshop on Continuous archival and retrieval of personal experiences*, pages 56–65. ACM, New York, NY, USA. ISBN 1581139322. doi:10.1145/1026653.1026662.

Kluev, Vitaliy (2002). *OASIS at NTCIR-3: Monolingual IR Task*. In *Proceedings of the Third NTCIR Workshop Meeting*. ISBN 4860490169. http://research.nii.ac.jp/ntcir/workshop/OnlineProceedings3/NTCIR3-CLIR-KluevV.pdf.

Krishnan, Aparna and Steve Jones (2005). *TimeSpace: Activity-based Temporal Visualisation of Personal Information Spaces*. *Personal Ubiquitous Comput.*, 9(1), pages 46–65. ISSN 1617-4909. doi:10.1007/s00779-004-0291-x.

Kwasnik, Barbara (1989a). *How a Personal Document's Intended Use or Purpose Affects its Classification in an Office*. *SIGIR Forum*, 23(SI), pages 207–210. ISSN 0163-5840. doi:10.1145/75335.75356.

Kwasnik, Barbara (1989b). *The Influence of Context on Classificatory Behavior*. PhD Thesis, Rutgers, the State University of New Jersey. http://www.worldcat.org/oclc/22939889.

Kwasnik, Barbara (1991). *The Importance of Factors that are Not Document Attributes in the Organization of Personal Documents*. *Journal of Documentation*, 47, pages 389–398. doi:10.1145/75335.75356.

Lamming, Mik and Mike Flynn (1994). *Forget-me-not: Intimate Computing in Support of Human Memory*. In *Proceedings FRIEND21 Symposium on Next Generation Human Interfaces*. http://www.cc.gatech.edu/classes/AY2001/cs7470_spring/readings/lamming-forget-me-not-intimate.pdf.

Lamping, John and Ramana Rao (1996). *Visualizing Large Trees Using the Hyperbolic Browser*. In *CHI '96: Conference companion on Human factors in computing systems*, pages 388–389. ACM, New York, NY, USA. ISBN 0897918320. doi:10.1145/257089.257389.

Landauer, Thomas K. (1986). *How Much do People Remember? Some Estimates of the Quantity of Learned Information in Long-Term Memory*. *Cognitive Science*, 10(4), pages 477–493. doi:10.1207/s15516709cog1004_4.

Lansdale, Mark (1988). *The Psychology of Personal Information Management*. *Applied Ergonomics*, 19(1), pages 55–66. doi:10.1016/0003-6870(88)90199-8.

Lee, Bongshin, Greg Smith, George G. Robertson, Mary Czerwinski, and Desney S. Tan (2009). *FacetLens: Exposing Trends and Relationships to Support Sensemaking within Faceted Datasets*. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 1293–1302. ACM, New York, NY, USA. ISBN 1605582467. doi:10.1145/1518701.1518896.

Lesk, Michael (1997). *How Much Information is There in the World?* http://www.lesk.com/mlesk/ksg97/ksg.html.

Mackinlay, Jock D. and Polle T. Zellweger (1995). *Browsing vs. Search: Can we Find a Synergy? (panel session)*. In *CHI '95: Conference companion on Human factors in computing systems*, pages 179–180. ACM, New York, NY, USA. ISBN 0897917553. doi:10.1145/223355.223490.

Malone, Thomas W. (1983). *How Do People Organize their Desks?: Implications for the Design of Office Information Systems*. *ACM Trans. Inf. Syst.*, 1(1), pages 99–112. ISSN 1046-8188. doi:10.1145/357423.357430.

Marchionini, Gary (1997). *Information Seeking in Electronic Environments*. Cambridge University Press, New York, NY, USA. ISBN 0521586747.

Marchionini, Gary (2006). *Exploratory Search: from Finding to Understanding*. *Commun. ACM*, 49(4), pages 41–46. ISSN 0001-0782. doi:10.1145/1121949.1121979.

Mugellini, Elena, Denis Lalanne, Bruno Dumas, Florian Evéquoz, Sandro Gerardi, Anne Le Calvé, Alexandre Boder, Rolf Ingold, and Omar Abou Khaled (2009). *MEMODULES as Tangible Shortcuts to Multimedia Information*, chapter 2, pages 103–132. Springer Berlin / Heidelberg. doi:10.1007/978-3-642-00437-7_5.

Munzner, Tamara (2009). *A Nested Process Model for Visualization Design and Validation*. *IEEE Transactions on Visualization and Computer Graphics*, 15(6), pages 921–928. ISSN 1077-2626. doi:10.1109/TVCG.2009.111.

Naumer, Charles M. and Karen E. Fisher (2007). *Naturalistic Approaches for Understanding PIM*, in *Personal Information Management*, chapter 5, pages 76–89. University of Washington Press. ISBN 0295987375. doi:10.1103/PhysRevE.69.026113.

Newman, M. E. J. and M. Girvan (2004). *Finding and Evaluating Community Structure in Networks*. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, 69(2), 026113. `http://link.aps.org/abstract/PRE/v69/e026113`.

North, Chris (2006). *Toward Measuring Visualization Insight*. *IEEE Comput. Graph. Appl.*, 26(3), pages 6–9. ISSN 0272-1716. doi:10.1109/MCG.2006.70.

Palla, Gergely, Imre Derenyi, Illes Farkas, and Tamas Vicsek (2005). *Uncovering the Overlapping Community Structure of Complex Networks in Nature and Society*. *Nature*, 435(7043), pages 814–818. ISSN 0028-0836. doi:10.1038/nature03607.

Papa, Steve (2006). *The Faceted Navigation and Search Revolution*. KM World Saturday. `http://www.allbusiness.com/technology/computer-software/4072803-1.html`. (Retrieved October 12, 2010).

Perer, Adam and Ben Shneiderman (2006). *Balancing Systematic and Flexible Exploration of Social Networks*. *IEEE Transactions on Visualization and Computer Graphics*, 12(5), pages 693–700. ISSN 1077-2626. doi:10.1109/TVCG.2006.122.

Perer, Adam, Ben Shneiderman, and Douglas W. Oard (2006). *Using Rhythms of Relationships to Understand E-mail Archives*. *J. Am. Soc. Inf. Sci. Technol.*, 57(14), pages 1936–1948. ISSN 1532-2882. doi:10.1002/asi.v57:14.

Perer, Adam and Marc A. Smith (2006). *Contrasting Portraits of Email Practices: Visual Approaches to Reflection and Analysis*. In *AVI '06: Proceedings of the working conference on Advanced visual interfaces*, pages 389–395. ACM Press, New York, NY, USA. ISBN 1595933530. doi:10.1145/1133265.1133346.

Pratchett, Terry (1983–2010). *Discworld (series of novels)*. Corgi Books.

Pujol, Josep M., Javier Béjar, and Jordi Delgado (2006). *Clustering Algorithm for Determining Community Structure in Large Networks*. *Phys. Rev. E*, 74(1), page 016107. doi:10.1103/PhysRevE.74.016107.

Ranganathan, S. R. (2006). *Colon Classification, 6th Edition*. Ess ess Publications, Dehli, India. ISBN 8170004608.

Rekimoto, Jun (1999). *Time-machine Computing: A Time-centric Approach for the Information Environment*. In *UIST '99: Proceedings of the 12th annual ACM symposium on User interface software and technology*, pages 45–54. ACM, New York, NY, USA. ISBN 1581130759. doi:10.1145/320719.322582.

Ringel, Meredith, Edward Cutrell, Susan T. Dumais, and Eric Horvitz (2003). *Milestones in Time: The Value of Landmarks in Retrieving Information from Personal Stores*. In Rauterberg, Matthias, Marino Menozzi, and Janet Wesson (Editors), *INTERACT*, pages 184–191. IOS Press, Zurich (Switzerland). ISBN 1586033638. `http://research.microsoft.com/~sdumais/SISLandmarks-Interact2003-final.pdf`.

Rodgers, Peter and Paul Mutton (2003). *Visualizing Weighted Edges in Graphs*. *iv*, 00, page 258. ISSN 1093-9547. doi:10.1109/IV.2003.1217988.

Rosenberg, Scott (1999). *Let's Get this Straight: Personal Information Mismanagement*. `http://www.salon.com/21st/rose/1999/03/05straight2.html`, April 8 2010.

Schraefel, M.C., Daniel Alexander Smith, Alisdair Owens, Alistair Russell, Craig Harris, and Max L. Wilson (2005). *The Evolving mSpace Platform: Leveraging the Semantic Web on the Trail* *of the Memex*. In *Hypertext, 2005*, pages 174–183. ACM Press. `http://eprints.ecs.soton.ac.uk/10710/`.

Sellen, Abigail J. and Richard H.R. Harper (2003). *The Myth of the Paperless Office*. MIT Press, Cambridge, MA, USA. ISBN 026269283X.

Shannon, C. E. (1948). *A Mathematical Theory of Communication*. *The Bell System Technical Journal*, 27, pages –. `http://cm.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf`.

Shneiderman, Ben (1994). *Dynamic Queries for Visual Information Seeking*. *IEEE Softw.*, 11(6), pages 70–77. ISSN 0740-7459. doi:10.1109/52.329404.

Smith, Greg, Mary Czerwinski, Brian Meyers, Daniel Robbins, George Robertson, and Desney S. Tan (2006). *FacetMap: A Scalable Search and Browse Visualization*. *IEEE Transactions on Visualization and Computer Graphics*, 12(5), pages 797–804. ISSN 1077-2626. doi:10.1109/TVCG.2006.142.

Soergel, Dagobert. (1985). *Organizing Information: Principles of Data Base and Retrieval Systems*. Academic Press, Orlando, Fla. :. ISBN 0126542600.

Spence, Robert (2006). *Information Visualization: Design for Interaction*. Second Edition. Prentice Hall. ISBN 0132065509.

Taylor, A. G. (1992). *Introduction to Cataloging and Classification. 10th ed.* Libraries Unlimited. ISBN 1591582350.

Teevan, Jaime, Christine Alvarado, Mark S. Ackerman, and David R. Karger (2004). *The Perfect Search Engine is Not Enough: A Study of Orienteering Behavior in Directed Search*. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 415–422. ACM, New York, NY, USA. ISBN 1581137028. doi:10.1145/985692.985745.

Thomet, Julien (2010). *Weena, un système de navigation par facettes de documents personnels*. Master's Thesis, University of Fribourg. Unpublished.

Tufte, Edward R. (2001). *The Visual Display of Quantitative Information*. Second Edition. Graphics Press. ISBN 0961392142.

Tunkelang, Daniel (2009). *Faceted Search*. Morgan and Claypool. ISBN 1598299999. doi:10.2200/S00190ED1V01Y200904ICR005.

Tvarozek, Michal and Maria Bielikova (2007). *Personalized Faceted Navigation for Multimedia Collections*. *Semantic Media Adaptation and Personalization, International Workshop on*, 0, pages 104–109. doi:10.1109/SMAP.2007.28.

Varian, Peter and Hal R. Lyman (2000). *Reprint: How Much Information?* Technical Report, School of Information Management and Systems at the University of California, Berkeley. http://www2.sims.berkeley.edu/research/projects/how-much-info/summary.html.

Viegas, Fernanda B., Scott Golder, and Judith Donath (2006). *Visualizing Email Content: Portraying Relationships from Conversational Histories*. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 979–988. ACM Press, New York, NY, USA. ISBN 1595933727. doi:10.1145/1124772.1124919.

Ware, Colin (2004). *Information Visualization: Perception for Design*. 2nd Edition. Morgan Kaufmann. ISBN 1558608192.

Weiland, Markus and Raimund Dachselt (2008). *Facet Folders: Flexible Filter Hierarchies with Faceted Metadata*. In *CHI '08: CHI '08 extended abstracts on Human factors in computing systems*, pages 3735–3740. ACM, New York, NY, USA. ISBN 160558012X. doi:10.1145/1358628.1358922.

Whittaker, Steve, Quentin Jones, Bonnie Nardi, Mike Creech, Loren Terveen, Ellen Isaacs, and John Hainsworth (2004). *ContactMap: Organizing Communication in a Social Desktop*. *ACM Trans. Comput.-Hum. Interact.*, 11(4), pages 445–471. ISSN 1073-0516. doi:10.1145/1035575.1035580.

Whittaker, Steve and Candace Sidner (1996). *Email Overload: Exploring Personal Information Management of Email*. In *CHI '96: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 276–283. ACM Press, New York, NY, USA. ISBN 0897917774. doi:10.1145/238386.238530.

Whittaker, Steve, Loren Terveen, and Bonnie A. Nardi (2000). *Let's Stop Pushing the Envelope and Start Addressing it: A Reference Task Agenda for HCI*. *Hum.-Comput. Interact.*, 15(2), pages 75–106. ISSN 0737-0024. doi:10.1207/S15327051HCI1523_2.

Wolte, Josef (1998). *Information Pyramids*. Master's Thesis, Technische Universitaet Graz. http://www.iicm.tu-graz.ac.at/thesis_106/jwolte.pdf.

Xiao, Huiyong and Isabel F. Cruz (2005). *A Multi-ontology Approach for Personal Information Management*. In *ISWC 2005: Proceedings of the Semantic Desktop Workshop*. http://www.cs.uic.edu/~advis/publications/dataint/sdk05.pdf.

Xobni Corporation (2010). *Xobni homepage*. http://www.xobni.com.

Yee, Ka-Ping, Kirsten Swearingen, Kevin Li, and Marti Hearst (2003). *Faceted Metadata for Image Search and Browsing*. In *CHI '03: Proceedings of the SIGCHI conference on Human fac-

tors in computing systems, pages 401–408. ACM, New York, NY, USA. ISBN 1581136307. doi:10.1145/642611.642681.

Zeldes, Nathan (2009). *Infoglut*. *IEEE Spectrum*, 46(10), pages 30–55. doi:10.1109/MSPEC.2009.5267994.

List of Figures

List of Tables

Curriculum Vitae

Florian Evéquoz
Route de St-Légier 15B
CH-1800 Vevey

+41 78 661 48 39

florian.evequoz@unifr.ch

PERSONAL DATA

Nationality:	Swiss
Date of birth:	January 3, 1981

EDUCATION

2006-2010	Fribourg, Switzerland University of Fribourg *PhD in Computer Science*
2008-2009	Fribourg, Switzerland University of Fribourg *Diploma in Didactics for Higher Education*
2006	Fribourg, Switzerland IVE Institute *Workshop « What a young entrepreneur should know »*
2000-2005	Fribourg, Switzerland University of Fribourg *Bachelor and Master of Science in Computer Science (minor: Philosophy)*
1998-2000	St-Maurice, Switzerland Collège de l'Abbaye
1995-1998	Sion, Switzerland Collège des Creusets *Maturité cantonale, type B* (latin-sciences)*
1988-1995	Conthey, Switzerland Primary and secondary school

LANGUAGES

French	Mother tongue
English	Fluent
German	Basic
Italian	Notions

TEACHING EXPERIENCE

University of Applied Sciences of Western Switzerland (Valais)

since 2010 Several classes in Business Informatics
Professor

University of Fribourg and University of Applied Sciences of Western Switzerland (Fribourg)

2006-2010 Several classes in Computer Science
Teaching assistant

2003-2005 Several classes in Computer Science
Tutor

ACADEMIC PROJECTS INVOLVED IN

2006-2010 Memodules (Hasler Foundation)
Scientific collaborator, PhD student

2006-2010 NCCR IM2 (Swiss National Science Foundation)
Scientific collaborator, PhD student

2006 CTI Netsecurity
Scientific collaborator

2005-2006 BioSecure (European project)
Collaborator

RESEARCH INTERESTS

- Personal information management
- Human-computer interaction
- Information visualization
- Data mining
- Text retrieval

LIST OF SCIENTIFIC PUBLICATIONS

Book Chapters with Peer Review

[1] Mugellini, Elena, Denis Lalanne, Bruno Dumas, Florian Evéquoz, Sandro Gerardi, Anne Le Calvé, Alexandre Boder, Rolf Ingold, and Omar Abou Khaled (2009). MEMODULES as Tangible Shortcuts to Multimedia Information, chapter 2, pages 103–132. Springer Berlin / Heidelberg. doi:10.1007/978-3642-00437-7 5.

[2] Lalanne, Denis, Maurizio Rigamonti, Florian Evéquoz, Bruno Dumas, and Rolf Ingold (2007). An Ego-centric and Tangible Approach to Meeting Indexing and Browsing, Lecture Notes in Computer Science, volume 4892/2008, chapter 8, pages 84–95. Springer Berlin / Heidelberg. doi:10.1007/978-3-54078155-4 8.

[3] Rigamonti, Maurizio, Denis Lalanne, Florian Evéquoz, and Rolf Ingold (2005). Browsing Multimedia Archives Through Intra-and Multimodal Cross-Documents Links, chapter 3, pages 114–125. Springer Verlag, Heidelberg, Germany. ISBN 3540325492. doi:10.1007/11677482.

International Conferences with Peer Review

[4] Evéquoz, Florian, Julien Thomet, and Denis Lalanne (2010). La navigation par facettes appliquée à la gestion de l'information personnelle. In IHM'10: Proceedings of the 22nd International Conference on Association Francophone d'Interaction Homme-Machine, pages 41–48. ACM Press, Luxembourg. ISBN 1450304109. http://diuf.unifr.ch/diva/publications/evequoz_ihm10.pdf.

[5] Evéquoz, Florian, Maurizio Rigamonti, Denis Lalanne, and Rolf Ingold (2006). Document inquisitor : un système de validation des structures et d'élicitation de modèles de documents. In Proceedings of 9ème Colloque International sur le Document Electronique (CIDE' 9), pages 79–95. Europia, Paris, France. ISBN 2909285367. http://www.evequoz.name/publications/evequoz_cide06.pdf.

[6] Bouguila, Laroussi, Florian Evéquoz, Michele Courant, and Beat Hirsbrunner (2004). Walking-pad: A Step-in-Place Locomotion Interface for Virtual Environments. In Proceedings of the 6th international conference on Multimodal interfaces, pages 77–81. ICMI '04, ACM, New York, NY, USA. ISBN 1581139950. doi:10.1145/1027933.1027948.

Workshop, poster and demo papers

[7] Evéquoz, Florian and Denis Lalanne (2009). "I Thought You Would Show Me How to Do It" -A Study on Self-Analysis of PIM Practices. In ASIS&T PIM Workshop, Vancouver, BC, Canada. http://pimworkshop.org/2009/ papers/evequoz-pim2009.pdf.

[8] Evéquoz, Florian and Denis Lalanne (2007). Indexing and Visualizing Digital Memories Through Personal Email Archive. In Supporting Human Memory with Interactive Systems, Workshop at HCI 2007 (HCI 2007), pages 21– 24. Lancaster (UK). http://diuf.unifr.ch/diva/publications/evequoz_ hci07.pdf.

[9] Evéquoz, Florian and Denis Lalanne (2007). Personal Information Management Through Interactive Visualizations. In Poster Compendium of the IEEE Information Visualization Conference (Doctoral Colloquium InfoVis 2007), pages 158–160. Sacramento, CA (USA). http://diuf.unifr.ch/diva/ publications/evequoz_infovis07.pdf.

[10] Lalanne, Denis, Florian Evéquoz, Herve Chiquet, Matthias Mueller, Mehdi Radgohar, and Rolf Ingold (2007). Going Through Digital Versus Physical Augmented Gaming. In Tangible Play: Research and Design for Tangible and Tabletop Games, workshop at the 2007 Intelligent User Interfaces Conference (IUI'07), pages 41–44. Hawaii (USA). http://www.evequoz.name/ publications/lalanne_iui07.pdf.

[11] Chiquet, Herve, Florian Evéquoz, and Denis Lalanne (2006). Elcano, a Tangible Personal Multimedia Browser. In Adjunct proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology (UIST 2006), pages 51–52. ACM Press. http://www.acm.org/uist/archive/adjunct/ 2006/pdf/demos/p51-chiquet.pdf.

[12] Mueller, Matthias, Florian Evéquoz, and Denis Lalanne (2006). TJASS, a Smart Board for Augmenting Card Game Playing and Learning. In Adjunct proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology (UIST 2006), pages 67–68. ACM Press. http://www.acm.org/ uist/archive/adjunct/2006/pdf/demos/p67-muller.pdf.

[13] Radgohar, Mehdi, Florian Evéquoz, and Denis Lalanne (2006). Phong, Augmenting Virtual and Real Gaming Experience. In Adjunct proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology (UIST 2006), pages 71–72. ACM Press. http://www.acm.org/uist/ archive/adjunct/2006/pdf/demos/p71-radgohar.pdf.

[14] Bouguila, Laroussi, Evequoz Florian, Michèle Courant, and Beat Hirsbrunner (2005). Abstract Active Walking Interface for Human-Scale Virtual Environment (poster). In Proceedings of 11th International Conference on Human-Computer Interaction (HCII'05). Lawrence Erlbaum Associates, Inc, Las Vegas, Nevada. ISBN 0805858075. http://www.evequoz.name/publications/ bouguila_hcii05.pdf.

[15] Dumas, Bruno, Catherine Pugin, Jean Hennebert, Dijana Petrovska-Delacretaz, Andreas Humm, Florian Evéquoz, Rolf Ingold, and Didier Von Rotz (2005). MyIdea -Multimodal Biometrics Database, Description of Acquisition Protocols. In Third COST 275 Workshop (COST 275), pages 59–62. Hatfield (UK). http://www.evequoz.name/publications/dumas_cost05.pdf.

Technical reports

[16] Evéquoz, Florian (2009). An Ethnographically-Inspired Survey of PIM Strategies. Technical Report, University of Fribourg, DIUF. http://diuf.unifr. ch/diva/publications/evequoz_techreport09.pdf.

Fribourg, 16.11.2010