Supporting Exploratory Browsing with Visualization of Social Interaction History

A Thesis Submitted to
The College of Graduate Studies and Research
in Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy
in the Department of Computer Science
University of Saskatchewan
Saskatoon

By
Indratmo

Permission to Use

In presenting this thesis in partial fulfillment of the requirements for a Postgraduate degree from the University of Saskatchewan, I agree that the Libraries of this University may make it freely available for inspection. I further agree that permission for copying of this thesis in any manner, in whole or in part, for scholarly purposes may be granted by the professor or professors who supervised my thesis work or, in their absence, by the Head of the Department or the Dean of the College in which my thesis work was done. It is understood that any copying or publication or use of this thesis or parts thereof for financial gain shall not be allowed without my written permission. It is also understood that due recognition shall be given to me and to the University of Saskatchewan in any scholarly use which may be made of any material in my thesis.

Requests for permission to copy or to make other uses of materials in this thesis in whole or part should be addressed to:

Head of the Department of Computer Science
University of Saskatchewan
Saskatoon, Saskatchewan S7N 5C9
Canada

Abstract

This thesis is concerned with the design, development, and evaluation of information visualization tools for supporting exploratory browsing. Information retrieval (IR) systems currently do not support browsing well. Responding to user queries, IR systems typically compute relevance scores of documents and then present the document surrogates to users in order of relevance. Other systems such as email clients and discussion forums simply arrange messages in reverse chronological order. Using these systems, people cannot gain an overview of a collection easily, nor do they receive adequate support for finding potentially useful items in the collection.

This thesis explores the feasibility of using social interaction history to improve exploratory browsing. Social interaction history refers to traces of interaction among users in an information space, such as discussions that happen in the blogosphere or online newspapers through the commenting facility. The basic hypothesis of this work is that social interaction history can serve as a good indicator of the potential value of information items. Therefore, visualization of social interaction history would offer navigational cues for finding potentially valuable information items in a collection.

To test this basic hypothesis, I conducted three studies. First, I ran statistical analysis of a social media data set. The results showed that there were positive relationships between traces of social interaction and the degree of interestingness of web articles. Second, I conducted a feasibility study to collect initial feedback about the potential of social interaction history to support information exploration. Comments from the participants were in line with the research hypothesis. Finally, I conducted a summative evaluation to measure how well visualization of social interaction history can improve exploratory browsing. The results showed that visualization of social interaction history was able to help users find interesting articles, to reduce wasted effort, and to increase user satisfaction with the visualization tool.

Acknowledgements

I have received tremendous help from many people while working on this thesis. I am fortunate to have Julita Vassileva as a supervisor. Julita has given me excellent guidance in research, financial support, and opportunity to grow as a researcher. She let me explore my research interests, yet kept me focused and prevented me from going to wrong directions. I really appreciate her kindness, encouragement, and support throughout these years.

I am grateful to my committee members—Harley Dickinson, Jim Greer, Carl Gutwin, and Eric Neufeld—who have guided me through my Ph.D. journey. I often gained insights into my research through their comments and suggestions on various occasions. I thank Carl for offering his expertise in experimental designs, and Melanie Tory for attending my thesis defense and serving as the external examiner.

My time as a student would not have been enjoyable without these wonderful people. Thanks to Jan Thompson for providing excellent administrative support and being a caring 'mom' to graduate students. Thanks to Ralph Deters and members of the MADMUC Lab for creating a pleasant working environment. Weekly meetings with members of the writing support group have also helped me make progress with this thesis (especially Miguel Nacenta with his little book). Special thanks to Bunyamin's family for offering their home as my temporary shelter when I most needed it. Finally, I thank my parents and sisters, and especially my dearest wife Joyce for their love. Joyce has been present in every step of my Ph.D. journey, which included house fire and a flooded apartment! I will always cherish memories of this wonderful journey.

Ad Maiorem Dei Gloriam

Contents

List of Tables

List of Figures

Chapter 1

Introduction

To keep up with the requirements in workplaces, knowledge workers need to search for information, acquire new knowledge, and develop their skills regularly (Kidd, 1994). In the software industry, for example, software developers must learn new programming languages, paradigms, and techniques. During this learning process, they seek and integrate new information into their previous knowledge of programming. They sometimes need to change their way of thinking about how to solve problems, such as shifting from a procedural approach to an object-oriented approach. They need to learn new concepts and terminology by exploring a new information space in which they may not be familiar with the content and organization.

Exploring information spaces is necessary not only in workplaces but also in everyday life. When people start a new hobby, they need to build basic knowledge by gathering and reading relevant materials. In bookstores, they browse bookshelves and the bestseller lists to get ideas for selecting which books to buy. When they plan a trip to a tourist destination, they need to find out about good accommodation, local events, and places to see. In these examples, people essentially perform exploratory searches: they do not aim to look up a piece of specific information, but search for and gather relevant information from various sources to construct their knowledge of a subject or to help them make a decision (Baldonado & Winograd, 1997; Bates, 1989; Marchionini, 2006; O'Day & Jeffries, 1993).

Information retrieval (IR) systems traditionally assume that users have clear search goals and know how to translate these goals into search terms. A typical search engine requires a user to initiate a search session by submitting search terms.

The system returns a list of search results, which are scanned and evaluated by the user. The user continues the search process by browsing items on the list or by refining the initial search terms to obtain the desired results.

There are several problems applying this type of interaction to an exploratory search. Users may not be familiar with the terminology used in the domain of interest because an exploratory search is usually done in the context of learning about a subject (Marchionini, 2006; White, Kules, Drucker, & schraefel, 2006). This subject is often outside the users' area of expertise, so it is difficult to specify information needs as search terms using the proper keywords in the domain. For example, the general public may not recognize buzzwords in web 2.0 such as "tagging" and "folksonomies." Therefore, they would not use these keywords in their search terms when they start learning about collaborative categorization in web 2.0 applications. Failure to use proper keywords may reduce the effectiveness of search terms in retrieving highly relevant documents.

People often do not have well-defined goals while exploring information spaces (Baldonado & Winograd, 1997; Bates, 1989; O'Day & Jeffries, 1993). Using a previous example, the search goal might be to find an interesting book to read during holidays. However, while initially browsing bookshelves, people may not have a clear idea about the topic, the author, or the genre of the book they want to buy. They may have a general idea about their interests, but are not ready to specify their information needs precisely. Requiring them to compose search terms would increase their cognitive efforts and decrease the possibility of finding interesting books serendipitously.

Multiple search iterations are usually needed to satisfy the information need of a user, as typically, no single information item is sufficient to fulfill the user's goal in exploratory searches (Bates, 1989; Marchionini, 2006; O'Day & Jeffries, 1993). Users need to find, gather, and process a set of relevant documents. One of the problems relying solely on search terms is that relevant documents may not contain the submitted search terms, and hence are not ranked highly by a typical search engine (White, Bilenko, & Cucerzan, 2007). Consequently, users would miss these relevant documents, which otherwise could increase the quality of the search results.

The nature of exploratory searches illustrates the need for tools for supporting exploration of information spaces. Due to the ill-defined information needs in exploratory searches, these tools should offer interactive interfaces that help users browse and learn about the content of an information space, rather than asking them to formulate search terms in the first place.

Supporting exploratory browsing is becoming more important with the proliferation of social web applications. For example, usage analysis reveals that browsing is dominant in YouTube,[1] and a large portion of this activity falls into casual browsing (Coyle, Freyne, Brusilovsky, & Smyth, 2008). In casual browsing, users do not look for specific videos but just want to watch interesting videos. Such casual browsing likely also dominates in other applications such as Facebook[2] and blogs. Therefore, improvement in exploratory browsing will have a wide impact, as browsing has become part of daily activity for many computer users.

1.1 Scope of Research

This section introduces several terms to define the scope of this thesis.

In general, an *information item* is used interchangeably with a document. The notion of information items has various levels of granularity depending on the context of usage. When the term is used in the context of browsing a library, information items may refer to books in the collection. However, in other contexts, an information item may refer to a web article, an email message, or a specific chapter or paragraph in a book.

An *information space* is defined as a collection of information items in which users perform their searches. An information space sets the boundary of an information collection that is processed by users or systems. Depending on the context of a search, an information space may have different meanings. People looking for particular

[1] http://www.youtube.com/
[2] http://www.facebook.com/

messages in email inboxes may consider the inboxes to be the information spaces they are exploring, whereas those doing Internet searches essentially explore the World Wide Web.

The term *information exploration* refers to an exploratory search: a search task that has an ill-defined goal and that is usually performed to support learning or problem solving (Marchionini, 1995). This includes tasks such as casual browsing where users just want to explore and find interesting items in a collection (Coyle et al., 2008). For example, looking up a phone number of a person on a contact list is not an exploratory search, whereas searching for interesting places to visit in a foreign country is an exploratory task. Chapter 2.1 elaborates the characteristics of exploratory searches.

Search strategies can be categorized into two broad classes: browsing or querying. *Browsing* is a search strategy that relies on scanning and recognition of relevant information items, whereas querying relies on recall and specification of appropriate search terms to retrieve relevant documents (Cool & Belkin, 2002; Marchionini, 1995). These search strategies complement each other and usually are used together during a search session (Belkin, Marchetti, & Cool, 1993). However, due to the lack of specific goals in exploratory searches, users need better support for browsing an information space—to help them understand the structure and content of a collection, evaluate search results, and select information items that are likely to be useful or relevant. Chapter 2.2 discusses browsing in more detail.

Social interaction history is defined as traces of interaction among users in information spaces. Traces of conversations in email inboxes or discussion forums are examples of social interaction history, and so is interaction among users that happens in the blogosphere through the commenting and cross-referencing facility. The focus of social interaction history is on *user-user* interaction rather than on user-object interaction (although it does not exclude user-object interaction). This concept is an extension of computational wear (Hill, Hollan, Wroblewski, & McCandless, 1992; Wexelblat & Maes, 1999), which has been proven to be useful in many applications. Chapter 3 discusses the foundations and elements of social interaction history.

The scope of this research is how to support information exploration and browsing with visualization of social interaction history. The proposed solution developed in this thesis is not intended to solve all browsing problems or to replace content-oriented approaches to facilitating browsing. This research focuses on evaluating the effects of social interaction visualization on casual browsing tasks (Coyle et al., 2008) where users do not have specific goals or information needs but just want to find any interesting information items.

1.2 Research Question and Hypotheses

This thesis explores the following research question: *Does visualization of social interaction history improve exploratory browsing?* The basic hypothesis is that social interaction history indicates the potential value of information items. Therefore, visualization of social interaction history would provide heuristic guidelines for users to find valuable information items.

Based on the basic hypothesis, I formulate these more specific hypotheses and test them in a summative evaluation (Chapter 5):

- Visualization of social interaction history helps users find interesting articles;
- Visualization of social interaction history reduces wasted effort to find interesting articles;
- Visualization of social interaction history increases user satisfaction with the visualization tool; and
- Overall, users prefer the visualization tool to the other systems.

I derive these hypotheses from the following presumptions. First, interesting articles usually attract a lot of attention. For example, inspiring blog entries, useful posts, or well-written news articles often trigger a lot of discussion and receive many comments from the audiences. Visualization of social interaction history would afford users to identify popular articles, which are expected to be more interesting than less popular articles. Second, the visualization would increase the visibility of popular articles and subsequently attract users' attention to explore these articles first. Based on the first presumption, these articles would be interesting or valuable, which means that users do not have to deal with uninteresting articles such as spam and advertisements. Finally, receiving these benefits, users would feel more satisfied with the visualization tool and prefer to use it to browse an information space.

1.3 Overview of the Proposed Solution

My approach to supporting exploration of an information space is based on two design principles. First, an information space should provide an overview of its content and structure to facilitate browsing. Second, traces of social interaction in an information space should be visualized to help users locate potentially useful information items.

The first design principle is similar to having a map while exploring a city, or having a table of contents while reading a book. An overview enables users to learn about an information space at a glance, helping them to quickly analyze and form an impression of whether the document collection matches their interests and information needs. Without an overview, they have to spend considerable effort browsing and skimming individual items to learn about the structure and content of the information space.

The second principle—visualizing social interaction history—offers navigational cues for finding interesting items in an information space. This principle is developed in response to the emerging social web applications in recent years. The Web is no longer a medium for publishing documents only, but has become a medium for interaction and communication. Web applications such as online newspapers and blogs allow users to add content or to conduct conversations with other users through the commenting facility on the sites. Such social interaction leaves electronic trails, which can be used to help other users navigate and explore the information spaces.

When people are faced with uncertainty, they tend to use social navigation to guide their decision (Dieberger, Dourish, Höök, Resnick, & Wexelblat, 2000). Simply stated, social navigation is an act of following other people (Dourish & Chalmers, 1994). This behaviour is found in various settings. For example, lacking information to make an informed decision, people tend to follow a visible path while hiking, to watch a box-office movie instead of a less popular one, and to choose a crowded restaurant over a quiet one. In this thesis, social navigation refers to users' browsing behaviour in moving from one piece of article to another piece.

While exploring an information space, users have to deal with uncertainty too. They may not be familiar with the content and organization of the space, and their search goals are often vague and evolving (Baldonado & Winograd, 1997; Bates, 1989; O'Day & Jeffries, 1993). Rather than selecting information items arbitrarily, users would prefer to receive help in finding potentially useful items, for example, items that have received a lot of attention from other users (e.g., being discussed a lot). Visualization of social interaction history would allow users to explore popular items first without having to skim through individual items one by one, hence saving their time and efforts. Furthermore, such social cues can reduce the cognitive effort of the users, as they can rely on the notion of "the wisdom of crowds" (Surowiecki, 2004) in selecting items to read: if many people found an item interesting, the item must contain valuable information that is worth reading.

In order to validate the design principles above, I developed a visualization tool for browsing a blog archive. The tool was evaluated using a tool-specific questionnaire and the Questionnaire for User Interaction Satisfaction (QUIS) (Chin, Diehl, & Norman, 1988). Overall, responses from the study participants confirmed both the usefulness of these design principles and the usability of the tool. The user satisfaction was high, supported by a low error rate of the given tasks. Qualitative feedback from the participants also indicated the value of having an overview and visualization of social interaction history. Chapter 4 discusses this study in detail, covering the background information, the evaluation methodology, and the findings.

Following these initial findings, I conducted a summative evaluation to test further whether visualization of social interaction history can assist users in exploring an information space. Both the quantitative and qualitative results supported the research hypotheses. The results showed that visualization of social interaction history was able to improve exploratory browsing, to reduce wasted effort, and to increase user satisfaction. Chapter 5 discusses the results of this study and their implications for the design of information systems.

1.4 Preview of Contributions

This thesis takes and integrates ideas from information visualization, human-computer interaction, and social computing. It provides a case study of the synthesis and application of design principles in these fields to improve information exploration.

There are three main contributions of this research:

- First, I develop the concept of social interaction history and outline its potential to improve information management, retrieval, and exploration. In Chapter 3, I discuss four foundations of social interaction history: computational wear (Hill et al., 1992), information foraging (Pirolli & Card, 1999), social navigation (Dourish & Chalmers, 1994), and multi-dimensional notions of relevance (Barry, 1994; Saracevic, 2007).

- Second, I demonstrate the value of social interaction visualization in supporting information exploration. Evaluating a visualization tool empirically has remained challenging due to the absence of standard tasks, data sets, and performance measures (Ellis & Dix, 2006). Despite these challenges, through a series of studies with realistic tasks and data sets, I provide both quantitative and qualitative evidence of how social interaction history can improve exploratory browsing (Chapters 3, 4, and 5).

- Third, I outline practical applications of social interaction history throughout this thesis. Despite their great potential, traces of social interaction are still underutilized. I expect to see increased utilization of social interaction history in the coming years.

Chapter 2

Background

This chapter reviews and organizes literature in exploratory searches and browsing into three main sections. The first section describes exploratory searches and their characteristics. The second section discusses browsing in the context of information exploration, and argues why browsing is an appropriate strategy for conducting exploratory searches. Finally, the third section discusses relevant work in supporting browsing and exploratory searches.

2.1 Exploratory Search

An exploratory search is a search task driven by an information need that is often vague (in the beginning) and evolving, that has multidimensional aspects, and that requires analyses of sets of information items retrieved in multiple iterations (Baldonado & Winograd, 1997; Bates, 1989; Marchionini, 2006; O'Day & Jeffries, 1993; White, Kules, & Bederson, 2005). Looking up the date of Canada's Day on the Web is not an exploratory search because the goal is specific and well defined, and can be satisfied with a piece of information. In contrast, searching and gathering information for planning a trip to a foreign land is an exploratory search. This task usually starts with ill-defined goals and involves multidimensional aspects such as budget, accommodation, time constraints, and popularity of the tourist attractions. To make a plan, people need to consult information from websites, brochures, and travel books. Reviews and feedback from other travellers also provide valuable input in the decision-making processes. Furthermore, upon receiving new information, people may diverge from their initial plans and formulate new plans. This iterative and evolving process illustrates the nature of exploratory searches.

While exploring an information space, users must deal with uncertainty. They need to deal with the uncertainty of the search goals, as information exploration is usually motivated by a broad, ill-defined information need (Jansen, 2006; White et al., 2005). Typically, people start their exploration by searching and gathering general information to get an overview of the subject; then, as they read and assimilate the information at hand, they either narrow down or modify their search goals to obtain more detailed or other relevant information in the domain of interest (Baldonado & Winograd, 1997; Bates, 1989; Komlodi, Marchionini, & Soergel, 2007; O'Day & Jeffries, 1993; Qu & Furnas, 2008). For example, while planning to visit a city, people may initially consider staying in downtown and start searching for good accommodation with reasonable prices. They may begin to explore hotels in downtown area and gather information to support their decision-making processes. However, as they learn about the city and find out that it has good public transportation, they may change their initial goal to stay in downtown and instead choose another place accessible by the public transportation that offers comparable accommodation with lower prices. This scenario illustrates the ill-defined and evolving aspects of search goals in exploratory searches.

Users may not be familiar with the content and structure of the information space they are exploring (Jansen, 2006; White et al., 2006, 2005). Consider people who visit a company's website for the first time. Lacking knowledge of the website, users are uncertain whether the information they are looking for is available on the website. Even if the information is available, they do not know where or in which category the information is stored. Submitting search terms to the website does not always yield the expected results. In order to use and navigate through the website effectively, they have to spend considerable time and effort to learn about the website first. Without support, this process can be long and tedious, especially when users deal with a large, complex website, as they have to skim through the available information items to get a good sense of the content and organization of the website.

Performing an exploratory search, users may feel uncertain about the quality of their search results. There is always a possibility to overlook important documents or to interpret information incorrectly. The search engine may leave highly relevant documents out of the query results. The sources of information may not be reliable and trustworthy. Moreover, there is no exact measure of what constitutes successful exploratory search results and how to evaluate exploratory search systems (White et al., 2006). The performance measures often just rely on subjective user satisfaction (O'Day & Jeffries, 1993). This situation calls for the development of research methodology for evaluating exploratory search systems.

An exploratory search requires a high level of user engagement in the search process (Marchionini, 2006). Users need to process information and interact with the information space. Their role extends beyond submitting search terms and looking up needed information in the query results. Users need to collect, interpret, and integrate pieces of information from various sources to fulfill their information needs (Baldonado & Winograd, 1997; Bates, 1989; O'Day & Jeffries, 1993). The quality of the search results depends more on the searchers' background knowledge, analytical skills, and ability to navigate through information spaces, rather than on the performance of the underlying search engine in returning relevant documents (Bhavnani, 2001). In other words, users are in the centre of exploratory search processes.

Due to its complexity, an exploratory search requires several iterative steps to complete. In each step, users need to sort out, analyze, and assimilate all relevant materials (including those from previous steps) to decide the next direction which brings them closer to their search goals (Bates, 1989; O'Day & Jeffries, 1993). During a search session, people usually use multiple information-seeking strategies such as browsing and querying (Belkin et al., 1993). A typical search engine currently only focuses on supporting the querying strategy: that is, given a query, how a system can improve the precision and recall of the search results.[1] In this way, the effectiveness of an exploratory activity depends mainly on the users' capability in specifying

[1]Precision is the number of relevant documents in a search result compared to the number of documents in the search result. Recall is the number of relevant documents in a search result compared to the number of relevant documents that should have been included in the search result.

their information needs as search terms and in processing a set of information items returned by the search engine. Having time constraints and other responsibilities, users cannot put their best efforts to formulate all possible search terms and examine every single item in the query results. Furthermore, requiring users to come up with search terms in the first place is not the best strategy in exploratory searches, considering that they usually lack background knowledge and well-defined goals. Therefore, new user interfaces and interaction techniques are needed to support exploratory browsing and to reduce user effort to find relevant items in an information space.

Before going into approaches to supporting exploratory browsing, the next section defines the notion of browsing, discusses its characteristics, and argues that this search strategy has potential to alleviate difficulties in exploratory searches.

2.2 Browsing

The notion of browsing is multidimensional and can be analyzed from different perspectives, for example, from the perspective of library and information science to that of marketing and advertisement (Chang & Rice, 1993). This thesis, however, focuses on information exploration and retrieval in electronic environments, and hence confines the discussion of browsing to this context.

Based on the information retrieval literature, Chang and Rice (1993) define browsing as an interactive information-seeking strategy that is usually unplanned, that relies on scanning and recognition of relevant information items, and that generally aims to explore and learn about the structure and content of an information space. Compared to querying, browsing is often seen as an informal, heuristic, and opportunistic activity (Marchionini, 1995). Browsing does not require searchers to specify their information needs explicitly as search terms. To initiate browsing, people need neither specific search goals nor predefined, formal search strategies such as looking up thesauri in advance to find appropriate search terms (Chang & Rice, 1993). They may start simply with vague goals, which can be refined and narrowed down later after they gain new knowledge and figure out what they really want or need.

A browsing strategy has three main components: location, scan, and attention (Spence, 2003). *Location* concerns the area of an information space that is viewed or explored. Movement from one location to another can refer to spatial movement (e.g., viewing a different area of a map) or semantic movement (e.g., viewing a different category of a collection). *Scan* refers to how users skim through information items in a location. Attempting to find relevant information, users may scan information items sequentially or even randomly; or, they may use their previous knowledge of the domain to scan the items systematically. The last component, *attention*, concerns the selection of information items that are considered relevant and that should be examined in more detail.

Marchionini (1995) distinguishes between within-document and across-document browsing. Within-document browsing refers to exploration and navigation in a single document, and is often associated with skimming or reading, such as finding a relevant paragraph in a journal article to support an argument. Across-document browsing refers to exploration of a set of documents to select a few highly relevant documents for more detailed examination. As more and more electronic materials are becoming widely available, supporting both types of browsing is important. This thesis, however, focuses on across-document browsing.

Browsing in general refers to a search strategy that is unplanned rather than aimless (Chang & Rice, 1993). Although browsing is often used when the search goal is exploratory, people sometimes use this strategy even when the search target is well defined and known in advance. For example, some users prefer to navigate through websites to look up specific information such as phone numbers of professors: that is, rather than composing search terms, they visit the websites of the departments where the professors work, and then follow the available links to retrieve their contact information (Teevan, Alvarado, Ackerman, & Karger, 2004). Teevan et al. speculate that this "orienteering" behaviour reduces cognitive load because users do not have to specify their search targets exactly, and they already know the paths that lead to the needed information items.

As summarized by Marchionini (1995), browsing has limitations:

- Browsing is time consuming, as users have to process all potentially relevant information items. The effectiveness of browsing depends on users' capability to process information, to focus on their search goals, and to navigate through information spaces. Lacking these skills, users may be distracted and wander around the information spaces without meeting their goals.

- When information needs are well defined and can be captured as indexes, using the querying strategy is easier and faster. Looking up words in an online dictionary, for example, can be performed more efficiently by querying than browsing.

- Browsing may result in information overload when users deal with large information spaces such as the Internet, or when the structure of information items is unclear. For example, browsing a large website to find specific information can be frustrating when users do not know how information items are organized, as they cannot use the navigational menus effectively.

Browsing is suitable particularly when searchers aim to obtain an overview of an information space, when they do not have well-defined search criteria, or when they lack knowledge to express their information needs using appropriate keywords (Chang & Rice, 1993; Marchionini, 1995). Browsing an information collection allows users both to see how information items are organized and to learn about the terminology used in a domain. Users can get a sense of what is or is not available in the collection, enabling them to quickly assess whether the collection is relevant to their information needs. Furthermore, they may encounter highly relevant documents serendipitously.

The characteristics of browsing show that this strategy is suitable for conducting exploratory searches. Many difficulties in exploratory searches—composing appropriate search terms, getting an overview of both the structure and content of an information space, learning about terminology in a domain—can be reduced by browsing. However, to fully get these potential benefits, people need interactive user interfaces that facilitate browsing. Without proper support, browsing a large information space to find relevant documents can be tedious and frustrating. People often have difficulty in navigating through the Web (Cockburn & Jones, 1996; Komlodi et al., 2007). Finding specific information by browsing can be challenging when users do not know the location or the path to the information item. Understanding the purposes of the categories/navigational menus in an information space is not always straightforward. To alleviate these problems, researchers have designed and developed interactive systems to support browsing and information exploration. These approaches are reviewed in the next section.

2.3 Supporting Browsing and Exploratory Search

Tools for supporting exploratory searches should aim at reducing uncertainty and cognitive load of the users. Uncertainty of the content and organization of an information space can be reduced by providing an overview of the space, akin to providing a table of contents of a book. The complexity of analyzing information and formulating queries can be reduced by enabling users to interact with an information space through various interaction techniques. This section reviews earlier work on facilitating browsing and exploratory searches, and organizes the discussion into three sections: information representation, organization, and interaction.

2.3.1 Representing Information Items

In electronic environments, one of the basic design issues is how to represent information items. File systems use icons and filenames to represent documents. A typical search engine presents search results as a list of items displaying titles and snippets of the documents. Other systems try to improve search interfaces by using information visualization to reveal the structure, content, and patterns of an information space. This section discusses various ways to summarize and represent information items.

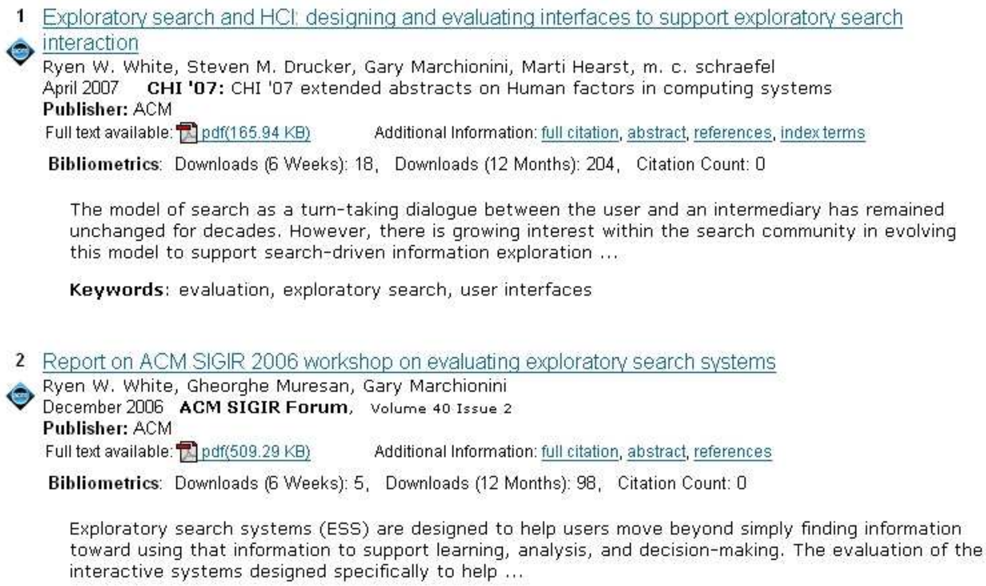
1 Exploratory search and HCI: designing and evaluating interfaces to support exploratory search interaction
Ryen W. White, Steven M. Drucker, Gary Marchionini, Marti Hearst, m. c. schraefel
April 2007 **CHI '07:** CHI '07 extended abstracts on Human factors in computing systems
Publisher: ACM
Full text available: pdf(165.94 KB) Additional Information: full citation, abstract, references, index terms
Bibliometrics: Downloads (6 Weeks): 18, Downloads (12 Months): 204, Citation Count: 0

The model of search as a turn-taking dialogue between the user and an intermediary has remained unchanged for decades. However, there is growing interest within the search community in evolving this model to support search-driven information exploration ...

Keywords: evaluation, exploratory search, user interfaces

2 Report on ACM SIGIR 2006 workshop on evaluating exploratory search systems
Ryen W. White, Gheorghe Muresan, Gary Marchionini
December 2006 **ACM SIGIR Forum**, Volume 40 Issue 2
Publisher: ACM
Full text available: pdf(509.29 KB) Additional Information: full citation, abstract, references
Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 98, Citation Count: 0

Exploratory search systems (ESS) are designed to help users move beyond simply finding information toward using that information to support learning, analysis, and decision-making. The evaluation of the interactive systems designed specifically to help ...

Figure 2.1: Examples of document surrogates in the ACM Digital Library (retrieved on April 24, 2008 from http://portal.acm.org/dl.cfm).

Document Surrogates

Browsing is an interactive process that requires searchers to continually assess and select information items that are most likely to be relevant to their information needs. While browsing, people may scan actual information items (e.g., books) or their surrogates (e.g., catalogues). On one hand, browsing actual items offers rich experience, as it enables close examination of the items and allows people to see, touch, and smell the actual objects. Think of the difference between visiting a bookstore and browsing books on the Internet. On the other hand, browsing actual objects can take longer because people have to deal with physical constraints such as the sizes and exact locations of the objects. Such constraints limit the amount of information that can be accessed within a short period. For example, trying to find a specific book in a library by browsing bookshelves directly would be difficult and time consuming; most likely, people would consult the library catalogue first to locate the book.

To allow rapid exploration of an information space, information retrieval (IR) systems usually use surrogates to present search results (see Figure 2.1). Surrogates are compact representations of information items, and may include information such as titles, keywords, snippets, abstracts, summaries, and thumbnails of documents (Ruthven et al., 2008). This information can help users preview and reason about the content of an information item quickly. The use of surrogates fits in with user practices in processing documents: to lower processing costs, people often only read titles and skim through documents while assessing the potential relevance of the documents (Petrelli, 2008; Pirolli & Card, 1999). Furthermore, presenting surrogates instead of full content of documents saves screen real estate and provides an overview of a collection.

Marchionini (1995) classifies surrogates into two categories: descriptive or semantic. Descriptive surrogates refer to the attributes of a document, such as its author, size, creation date, publisher, and type (e.g., Word, JPEG, PDF). This information in general is available as metadata and can be extracted automatically. Although descriptive surrogates do not reveal the content of a document, they provide contextual information which can be useful for sorting out irrelevant items. For example, when users look for movies written by a particular person, they can eliminate irrelevant movies quickly if the surrogates include information about the movies' writers.

Semantic surrogates describe the content of a document (Marchionini, 1995). When available, titles, abstracts, and subject categories of documents are commonly used as semantic surrogates. For textual items, document summarization techniques can also be used to extract the semantics of the items automatically, whereas the content of non-textual items such as image or video collections are usually represented using thumbnails. One of the main challenges in dealing with an image or video collection is how to choose an image or screenshot that can give a representative overview of the content of the collection. To facilitate browsing image collections, research on image retrieval has developed various algorithms to analyze and retrieve images based on some similarity features (Smeulders, Worring, Santini, Gupta, & Jain, 2000).

Issues pertinent to the use of surrogates include how to present surrogates to

Figure 2.2: Some examples of enhanced thumbnails displaying search results of the search terms "car hybrid mileage" (retrieved on June 12, 2008 from http://www.parc.com/research/projects/enhancedthumbnails/demo-hybrid.html).

users. Currently, the common way is to use text summaries, as illustrated in Figure 2.1. Woodruff, Faulring, Rosenholtz, Morrison, and Pirolli (2001), however, show that enhanced thumbnails can improve user performance in a variety of search tasks. In their design, enhanced thumbnails enrich graphical representations of documents (plain thumbnails) by increasing the readability of document headings in the thumbnails, adding an overlay displaying highlighted keywords, and reducing the contrast level in the thumbnails to make these keywords more apparent (see Figure 2.2). Although it offers additional benefits, the application of enhanced thumbnails in large-scale IR systems faces practical challenges, as generating and transferring thumbnails across networks require more computation and communication bandwidth than processing text summaries (Woodruff et al., 2001).

Another issue is concerned with the accuracy of surrogates in representing the content of documents (Lynch, 2001). In controlled environments such as company intranets and digital libraries, IR systems may safely assume that surrogates are created with intention to describe documents accurately. But, in open systems like the Web, this assumption is no longer true. Lynch provides examples of how some systems attempt to manipulate indexing and retrieval algorithms of IR systems. For example, misleading metadata (e.g., thousands of junk keywords) may be inserted to documents to affect the behaviour of search engines in processing these documents. Or, web servers may be programmed to return deceptive information to web crawlers in order to influence the results of their indexing or document summarization techniques. Thus, when IR systems operate in open, distributed environments, they must consider the trust and reputation of information providers, and use algorithms that can minimize the impact of such deception.

Information Visualization

Information visualization is a process of transforming data into a graphical representation, which aims to help users gain insights into the data (Spence, 2007). In this process, data is usually transformed into graphical attributes such as position, size, shape, and colour. Information visualization reduces users' cognitive load in processing data by exploiting human perceptual ability to recognize patterns (Ware, 2000). Looking at visualization, users can notice outliers or irregular patterns in data without spending a lot of cognitive efforts.

Data generally can be categorized into three types: nominal, ordinal, or quantitative (Stevens, 1946). *Nominal* data refers to data that is used to label or identify things. There is no specific order in nominal data. Types of fruits, phone numbers, and student numbers are examples of nominal data. *Ordinal* data refers to data that is ordered, but this order is only used to preserve the sequence of the data. The intervals between two successive values on an ordinal scale do not necessarily have the same quantity. Examples of ordinal data include military ranks (e.g., major, captain, lieutenant) and scales used in psychological experiments (e.g., disagree, neutral, agree). *Quantitative* data refers to data that allows the computation of the interval or ratio between data values. Stevens distinguishes between interval and ratio data, but, in the context of designing information visualization, these types of data are usually referred to as quantitative data. Price, weight, and length are examples of quantitative data.

The basic step to design information visualization is to choose graphical representations of data. This choice can affect the perceptual accuracy of the encoded

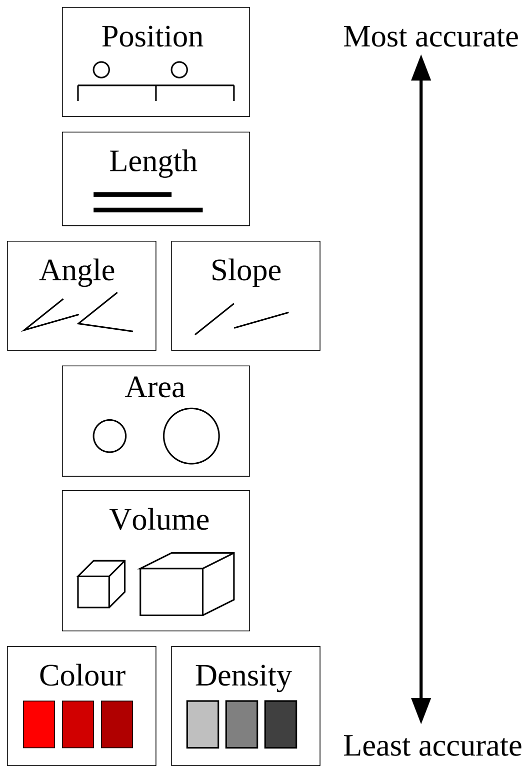

Figure 2.3: Perceptual accuracy of various graphical representations of quantitative data (Cleveland and McGill, 1984).

Table 2.1: Guidance for choosing graphical representations of quantitative, ordinal, and nominal data (Mackinlay, 1986).

Quantitative	Ordinal	Nominal
Position	Position	Position
Length	Density	Colour hue
Angle	Colour saturation	Texture
Slope	Colour hue	Connection
Area	Texture	Containment
Volume	Connection	Density
Density	Containment	Colour saturation
Colour saturation	Length	Shape
Colour hue	Angle	Length
	Slope	Angle
	Area	Slope
	Volume	Area
		Volume

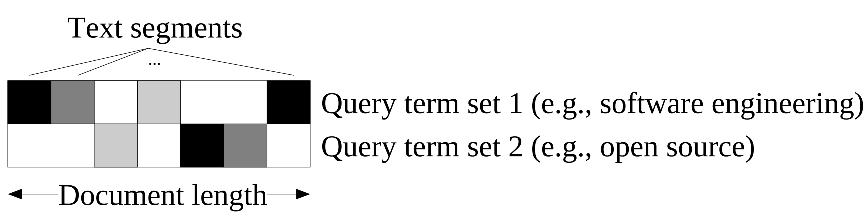

Figure 2.4: TileBar visualization (Hearst, 1995).

information. Accuracy refers to the closeness of human judgments to the actual values of encoded data. For example, using colour to represent quantitative data is a poor choice because there is no general consensus on the order of colours: is green greater or less than red? To provide guidance for encoding quantitative data, Cleveland and McGill (1984) establish a theory outlining different degrees of perceptual accuracy in extracting quantitative information from graphical attributes. This order of accuracy is illustrated in Figure 2.3. Mackinlay (1986) extends Cleveland and McGill's ranking by addressing the encoding of nominal and ordinal data as well (see Table 2.1).

Information visualization has been applied to facilitate information exploration and retrieval in various ways. TileBars, for example, enhance typical surrogates in search results by visualizing the length of documents, along with the frequency and distribution of query terms in the documents (Hearst, 1995). The basic idea is that a document is partitioned into text segments, and then query term hits in each segment are computed. A TileBar visualizes the relative length of a document as a rectangle, which contains squares representing text segments in the document (see Figure 2.4). The darker a square, the higher the frequency of query term hits in the corresponding segment is. A white square means no query term hits in the segment. If users submit two or more sets of query terms (e.g., "software engineering" and "open source"), query term hits in each set are visualized in separate rows. This visualization portrays the relation between documents and search terms in a compact and intuitive way, helping users select potentially relevant documents from search results.

One of the challenges in dealing with large information spaces is to understand the structure and content of the spaces. Responding to this need, researchers have developed techniques for visualizing large data sets. For example, Treemaps (Johnson & Shneiderman, 1991) and Cone Trees (G. Robertson, Mackinlay, & Card, 1991) are developed to visualize hierarchical data sets, whereas Perspective Wall (Mackinlay, Robertson, & Card, 1991) and LifeLines (Plaisant, Milash, Rose, Widoff, & Shneiderman, 1996) are designed to portray linear information.[2] These visualizations share a common feature: they provide overviews of information spaces to enable users to see the big picture of the spaces.

Overviews are also useful for monitoring and keeping track of dynamic data. Treemaps (Johnson & Shneiderman, 1991), for example, have been applied to visualize a stock market (see Figure 2.5). The performance of more than 500 stocks is displayed in a single view. Companies are grouped by their business sectors and visualized as rectangles. The size and the colour of a rectangle represent a company's market capitalization and its stock price performance. This visualization offers rich

[2]The next section discusses organizational structures of information, including hierarchical and linear structures.

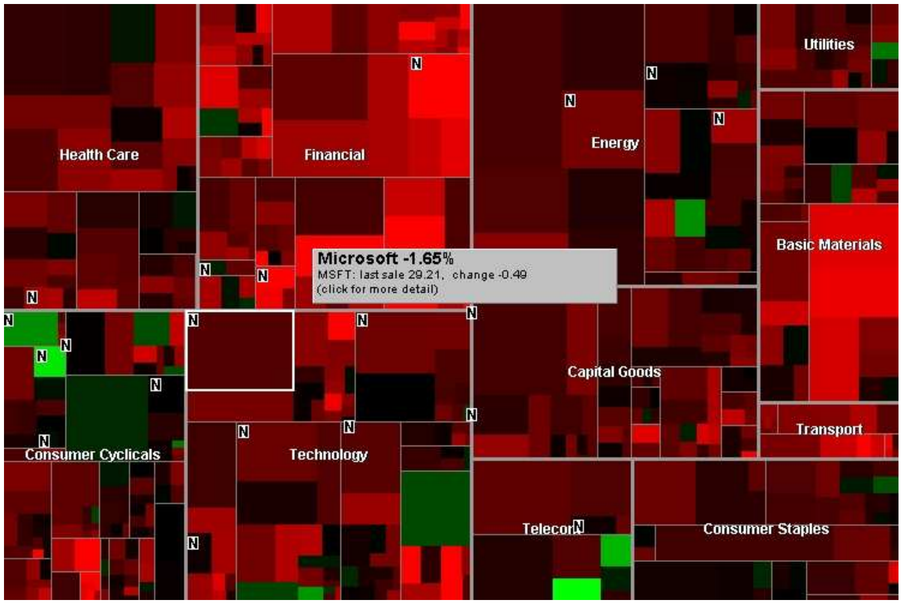

Figure 2.5: An application of treemaps in visualization of a stock market. This visualization portrays information about more than 500 stocks in a single view (retrieved on May 7, 2008 from http://www.smartmoney.com/map-of-the-market/).

information, which is useful for those wanting to explore and learn about a stock market. Moreover, the visualization reduces users' cognitive efforts by transforming complex data into graphical representations that are easy to understand.

2.3.2 Organizing Information Items

After designing how to represent information items, the next step in developing IR systems is to decide how to present and organize these representations. Logical organization of information facilitates browsing and exploration. For example, libraries organize their collections according to a standard classification scheme. Search engines typically sort query results by relevance to search terms. Blogs arrange entries in reverse chronological order. Without such organization, it would be harder to explore an information space and find a specific document within the space effectively. This section discusses the strengths and limitations of five basic organizational structures of information—hierarchical, flat, linear, spatial, and network—and provides examples of their applications in the context of information seeking and exploration.

Hierarchical

Hierarchical approaches use a tree structure to arrange information items. A node in a tree represents an information item or a collection of items. A hierarchical structure is currently the predominant way to organize information items. Most operating systems use hierarchical file systems to manage and organize files. Many websites provide hierarchical menus to support browsing and navigation. Email clients allow users to create folders and subfolders to facilitate information classification, management, and retrieval. Thus, browsing information items organized in a hierarchy is an intuitive, familiar process.

A hierarchy facilitates information seeking by enabling users to reduce a search space and to eliminate any ambiguity of a term that has different meanings. Since the classification of information items is usually done by humans, users would find relevant items categorized into the same branch of a hierarchy with high accuracy. Consider the hierarchy of subject categories used to classify books in libraries. When people look for a Java programming book in a digital library, they can navigate through the computer science branch and explore the subcategories within it. Within this branch, they do not have to deal with non-computer-science books, which are irrelevant to their information needs. Books about the island of Java (part of Indonesia), for instance, will not be found in the computer science category because those books belong to a different branch of the hierarchy (e.g., geography).

To browse a hierarchy effectively, however, users must be familiar with the structure and semantics of the hierarchy. They must have sufficient knowledge of the domain so that they can navigate through the hierarchy and choose the correct category to find a certain item. Without this background knowledge, users would not be able to take full advantages of hierarchical approaches, especially when they deal with large-scale hierarchies such as Yahoo! Directory[3] or domain-specific hierarchies such as biological classification systems.

[3] http://dir.yahoo.com/

Figure 2.6: Faceted categories to support browsing Nobel Prize winners. This screen shot shows top level categories, which provide an overview of the data set (retrieved on May 14, 2008 from http://orange.sims.berkeley.edu/cgi-bin/flamenco.cgi/nobel/Flamenco).

To improve the usability of a hierarchical structure, facet-based approaches assign hierarchical faceted categories to information items and enable users to browse the collection using these faceted categories (Allen, 1995; Capra, Marchionini, Oh, Stutzman, & Zhang, 2007; schraefel et al., 2005; Yee, Swearingen, Li, & Hearst, 2003). Faceted categorization allows multiple classifications of an information item. For example, a Nobel Prize winner can be classified into multiple categories based on his/her gender, country, affiliation, prize, and year (see Figure 2.6). These categories provide the gist of the collection. The prize category, for instance, shows that Nobel Prize is awarded in chemistry, economics, literature, medicine, peace, and physics, while the year category informs that Nobel Prize has been given since 1900s. This information supports learning and exploration of the data set.

The categories in each facet allow users to filter the data set by various attributes. Users essentially formulate and submit queries by browsing these faceted categories. Using the faceted categories in Figure 2.6, users can retrieve all female Nobel Prize winners in literature by selecting "GENDER: all > female" and "PRIZE: all > literature." If users want to restrict or broaden their current selection, they can select a subcategory or a more general category in the available facets. For example, choosing "COUNTRY: all > Austria" would narrow the current search results to female Nobel Prize winners in literature from Austria. Thus, faceted navigation reduces the complexity of formulating queries involving multiple data facets.

The main drawback of facet-based approaches is that, to work best, they require manual construction of the faceted categories, although research on automating this process is ongoing (Hearst, 2006). This shortcoming limits the scalability and maintainability of these approaches, particularly when they are applied to large collections with dynamic content. In dynamic environments, new documents need to be processed and categorized regularly, which may also result in the need to create new facets. This maintenance process requires a lot of human resources. To solve this problem, people usually apply facet-based approaches to topical data sets, such as collections of classical music,[4] Nobel Prize winners, and fine arts images (see the Fla-

[4]http://beta.mspace.fm/

menco search interface project[5] for the last two examples). In this way, the resulting faceted categories can be kept small, coherent, and stable.

Flat

While a hierarchical structure is pervasive in information management, organizing information items into a flat structure has recently gained popularity. In a flat structure, information items are organized into sets. These sets are usually labelled with tags and may overlap each other. Tagging or assignment of tags can be done both manually and automatically. Tags serve as a mechanism for grouping and retrieving information items, and provide associative access to the items (Dourish et al., 2000; Gifford, Jouvelot, Sheldon, & O'Toole, 1991; Gopal & Manber, 1999).

Tagging provides a flexible way to organize information items. Users can classify an information item into multiple categories by assigning multiple tags to the item. Grouping and re-grouping information items can be done flexibly. For example, a bioinformatics article might be tagged "paper," "bioinformatics," "protein," and "algorithms." Having multiple tags, this article can be grouped with other documents by different categories: paper, bioinformatics, protein, and algorithms. Users can also use any of these keywords to retrieve the paper.

The application of tagging in web-based systems has become popular in recent years (Golder & Huberman, 2006; Marlow, Naaman, boyd, & Davis, 2006). These applications allow users to manage and assign tags to their shared collections on the Web. Examples include tools for organizing and sharing photos,[6] web bookmarks,[7] and articles.[8] Web-based tagging systems can be considered to be social, lightweight IR systems. These systems aggregate user-defined tags and enable a user to find other people who share the same items. Using this information, people can visit and browse the collections of other users who have similar interests (e.g., those who use the same tags or keep the same information items). Browsing similarly tagged

[5]http://flamenco.berkeley.edu/demos.html
[6]http://www.flickr.com
[7]http://delicious.com
[8]http://www.citeulike.org

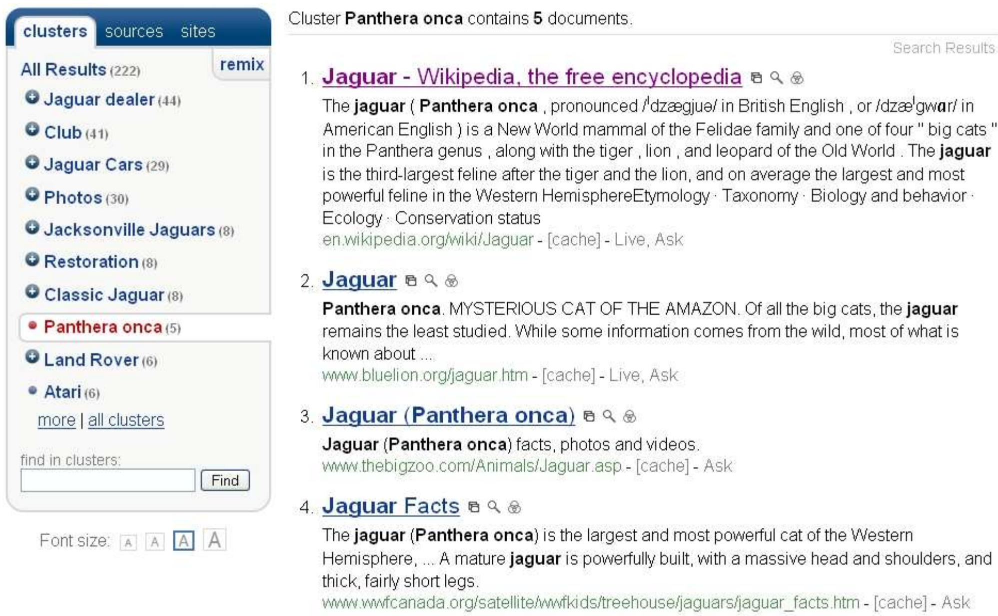

Figure 2.7: Clustering search results of a search term "jaguar" (retrieved on June 13, 2008 from http://clusty.com). Each cluster indicates the number of documents in the cluster.

collections may lead to serendipitous discovery of useful information items.

Social tagging systems can also recommend tags to users based on common keywords used to tag shared items. Such recommendations allow users to see related keywords and buzzwords in a domain. Furthermore, these recommendations result from a simple tag sharing instead of from complex computation and analysis. Thus, in social tagging systems, users indirectly collaborate with one another to organize and explore information collections.

Another application of flat structure in information retrieval is exemplified by clustering algorithms that group search results into flat categories according to content similarity (e.g., Zeng, He, Chen, Ma, & Ma, 2004). Content similarity is usually computed by analyzing the occurrence of words and phrases in documents. Based on some similarity measures, clustering algorithms aim to find topically coherent sets of documents, which are also relevant to a given query. Figure 2.7 provides a screenshot of a search engine that clusters its search results.

Clustering-based approaches offer several potential benefits. The resulting clusters can reduce users' efforts in sorting out query results. Instead of examining a long list of search results, users can focus their exploration on a smaller set of documents grouped in a topically coherent cluster. In this way, they do not have to deal with other clusters that appear irrelevant. For example, a search for a term "jaguar" returns several clusters of documents (see Figure 2.7). When users look for articles about the jaguar (animal), they can select the "Panthera onca" cluster to retrieve topically coherent documents about this big cat. Since relevant articles tend to be similar to each other (Hearst & Pedersen, 1996), users would be able to find relevant documents easily while browsing clustered search results. Without clustering, relevant documents may be scattered throughout a result list or ranked low by the underlying search engine, and thus are less visible and more difficult to find.

Clustering may reveal groups of items that may not be directly correlated with the submitted search term, but that are relevant to the overall search context. For example, when people explore classical music by Beethoven, besides browsing a cluster of his compositions, they may be interested in browsing additional clusters of classical music written by other composers. Because clustering attempts to present search results in a logical way, users can also learn about other composers or eras of classical music by exploring these clusters. Such exploration promotes serendipitous discovery and helps users obtain an overview of a domain.

Another potential benefit of clustering is that the process can be done fully automatic, so no human resource is needed to define categories and classify documents into the categories (Hearst, 2006). Linear-time algorithms have also been developed to cluster search results on-the-fly without jeopardizing system performance too much (e.g., Cutting, Karger, Pedersen, & Tukey, 1992; Zamir & Etzioni, 1999; Zeng et al., 2004).

Clustering algorithms, however, are not perfect. They may put documents in a wrong cluster or choose poor category names. The resulting clusters can be unpredictable and illogical, and thus decreasing the user acceptance of these approaches (Hearst, 2006). Furthermore, for performance reason, the algorithms often only process titles and snippets of documents, thereby reducing the accuracy of the clustering results (Zamir & Etzioni, 1999).

Linear

In a linear structure, information items are arranged on a list based on certain order. The location of an item on the list is determined according to a particular attribute used to compare the item with other items. Examples of a linear structure include the order of words in dictionaries (alphabetical), entries in weblogs (chronological), and incoming messages in email inboxes (chronological). As long as users know the attribute of the information items they are looking for, having a sorted list allows them to traverse the list in a logical way, which is helpful in the retrieval process.

Among the common attributes used to sort a collection of information items, temporal attributes have received a lot of attention from IR researchers for the following reasons:

- Time is a key retrieval cue in many situations (Dumais et al., 2003; Graham, Garcia-Molina, Paepcke, & Winograd, 2002; Rodden & Wood, 2003). People in general can roughly recall temporal attributes of information items, such as when an email message was received or when a picture was taken, and use this information to assist in finding needed information.

- If information items are organized based on their temporal attributes, then information organization can be automated. This automation means that users do not have to spend a lot of efforts in categorizing documents.

- A chronological structure may help maintain contextual information, as related information items are usually created or accessed around the same time, resulting in clusters of relevant items (Card, Robertson, & Mackinlay, 1991).

Although time is an important retrieval cue, relying solely on it may hamper retrieval performance. Particularly when the number of information items on a list is large, traversing the list to find a specific item is not easy. Moreover, there are other key retrieval cues as well, such as visual, spatial, topical, and social cues (Malone,

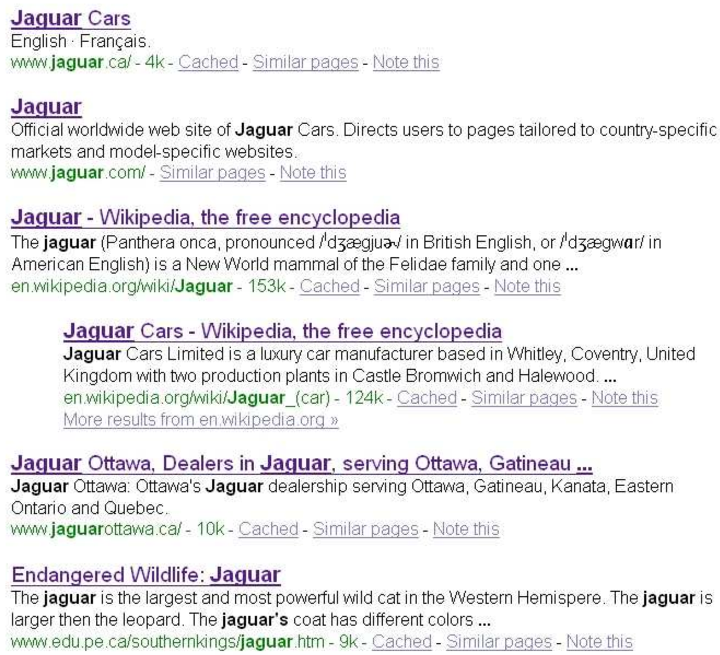

Figure 2.8: The top results of Google search for "jaguar." This list contains mixed articles about an animal, a car, and car dealers (retrieved on June 13, 2008 from http://www.google.ca).

1983; G. Robertson et al., 1998; Whittaker et al., 2004). Therefore, instead of relying on a temporal attribute only, IR systems should enable users to view a collection from multiple perspectives and to use as much contextual information as possible in the retrieval process.

Figure 2.8 illustrates another limitation of a linear structure. Given a query, a typical search engine returns a ranked list of documents sorted by the documents' relevance to the query. Constructing a ranked list is problematic when a system deals with ambiguous search terms such as "jaguar." For example, the first page of Google search for "jaguar" contains mixed web pages about an animal, a car, and car dealers. Relevant documents about Jaguar cars, for instance, are scattered throughout this result list. Moreover, if users look for information about the Jaguar operating system, which appears on page three of the search results (not shown), they have to first browse and sort out many irrelevant documents on the list.

Spatial

Spatial approaches use location as the main method of organizing information items. Spatial organization is pervasive in everyday life. People arrange things in their homes and offices so that they can find and use them easily. They separate important documents from unimportant ones on their desks. They keep frequently-used books in places that are easily accessible. In the digital world, a common example of spatial organization is computer desktop organization, where program shortcuts, file folders, and other items are arranged spatially to facilitate quick access to them (Ravasio, Schär, & Krueger, 2004).

Spatial organization on computers is supported by two-dimensional (2D) and three-dimensional (3D) user interfaces. Researchers have investigated and compared the effectiveness of 2D and 3D interfaces (e.g., Cockburn, 2004; Tavanti & Lind, 2001; Tory, Sprague, Wu, So, & Munzner, 2007). However, the study results are inconclusive and even conflicting (e.g., Cockburn vs. Tavanti and Lind). The effectiveness of 2D and 3D interfaces seems to depend on the characteristics of data and user tasks.

A spatial structure maintains good visibility of information items. Spatially arranged documents can be scanned quickly, promoting recognition over recall and supporting quick access, finding, and reminding (Malone, 1983; Whittaker et al., 2004). For example, a spatial approach to organizing web pages (Data Mountain) is shown to result in faster retrieval performance than the bookmarking facility in the Microsoft Internet Explorer (IE) (G. Robertson et al., 1998). Data Mountain is a 3D space in which users can store and place documents at arbitrary locations. In G. Robertson et al.'s study, participants organized 100 web pages using both Data Mountain and IE, and then their retrieval performance was measured. Czerwinski, Dantzich, Robertson, and Hoffman (1999) conducted a subsequent study a few months later, and found that the participants were still able to recall their spatial layouts of web pages in Data Mountain, and were able to retrieve the web pages as fast as their performance in the first study. This finding indicates that humans have a good spatial memory.

Spatial organization in computer systems, however, is limited by the size of computer monitors. With hundreds or thousands of information items, users simply cannot arrange all of the items spatially on their computer desktops without cluttering the desktops. Users want to keep their desktops tidy by filing away unneeded items, to avoid distraction and to allow them to concentrate on their work. Moreover, relying on pure spatial memory yields poor performance in information retrieval, particularly as the number of items increases (Jones & Dumais, 1986).

Network

A network structure allows information items to be linked to one another arbitrarily. It does not impose any structural constraints on how users create links between information items. An example of a network structure on the global scale is the World Wide Web.

The ability to link information items arbitrarily characterizes both the strength and the weakness of a network structure. On one hand, this flexibility allows users to create information networks regardless of the types and physical locations of information items. As long as a document has a unique identifier and is accessible through a computer network, this document can be linked to other documents. Users are free to wander around information networks by browsing and following links between information items. These features enable easy information sharing and transparent access to information.

On the other hand, since a network is less structured than other organizational structures, it is hard both to get an overview of an information network and to navigate through the network effectively (Cockburn & Jones, 1996; Komlodi et al., 2007). Furthermore, in network structures, broken links are a common error, which occurs when information items are removed from their previous locations on the servers.

2.3.3 Interacting with Information Spaces

To support exploratory browsing, a good representation and organization of information should be augmented with mechanisms that allow users to interact with information spaces. Users should be able to control their search spaces, such as filtering out irrelevant information or adjusting search criteria interactively. Besides getting an overview of a collection, they should be able to drill down and examine potentially relevant documents closely. To further engage users' interest in exploring information spaces, these features should be provided through highly interactive interfaces (Marchionini, 2006). For example, dynamic query interfaces promote interaction by providing visual gadgets such as sliders to accept user queries (Ahlberg, Williamson, & Shneiderman, 1992). By adjusting these gadgets, users can compose queries dynamically and observe changes in the query results immediately. Thus, they can focus on their main tasks in exploring and processing information, instead of switching their attention back-and-forth between specifying textual queries and assessing the query results. Below I review user interfaces and techniques for supporting interaction between users and information spaces.

Brushing

Data analysts often want to examine whether there is a correlation between two variables in a data set. When they deal with two dimensional data, they can visualize the data using a scatterplot to support their analyses. A scatterplot is a two dimensional space for plotting data values, where each axis represents a variable in the data set (see Figure 2.9). In scatterplots, data values are usually represented by points. As illustrated in Figure 2.9, scatterplots are effective for revealing correlations (if any) between two variables. In this case, the scatterplot shows that the concentrations of carbon dioxide in general increase with the temperature.

Data analysis is becoming more complex as the number of variables in data increases. People can no longer use a scatterplot effectively for analyzing data that has many variables. To reduce data complexity, they may create various views to

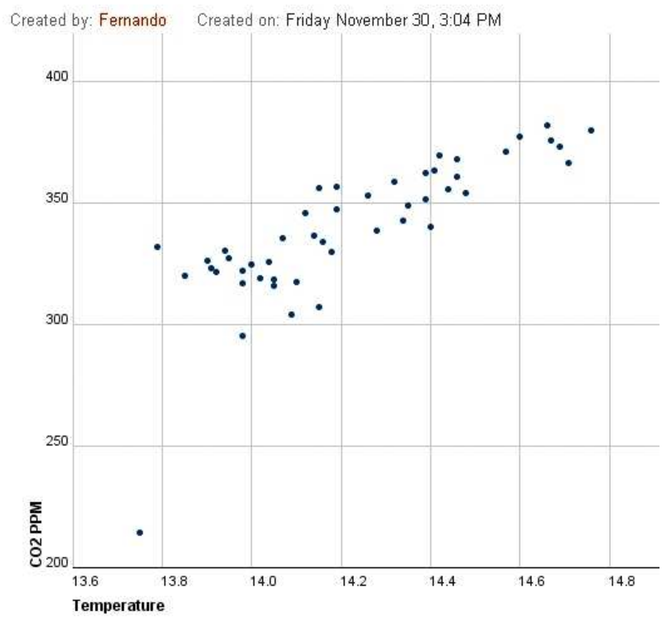

Figure 2.9: A scatterplot visualizing the concentrations of carbon dioxide (CO_2) versus temperature (retrieved on June 24, 2008 from http://services.alphaworks.ibm.com/manyeyes/view/SQAAXKsOtha6o2kApZWYK2%7E).

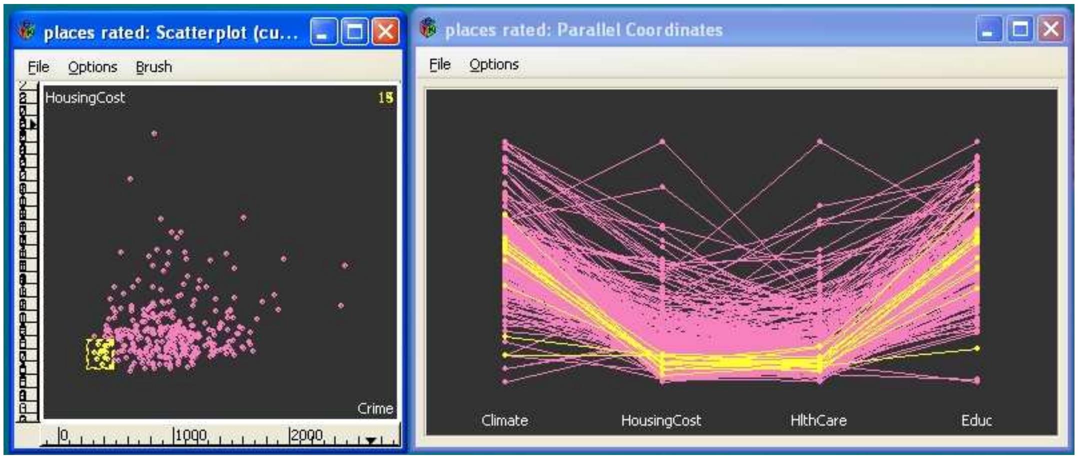

Figure 2.10: Using a brush technique to highlight selected data both in a scatterplot (left) and a parallel coordinate view (right) (GGobi data visualization system, http://www.ggobi.org/).

visualize different aspects of the same data. This approach, however, does not solve the problem completely. For example, detecting general patterns across views is difficult, since each view only provides partial information about the data.

To assist in analyzing multidimensional data, Becker and Cleveland (1987) developed a collection of brush techniques. The basic purpose of brushing is to highlight selected values in multiple representations of the same data. The brush object is used to select data values. In this example, data about housing cost, crime, climate, health care, and education in US cities are visualized in a scatterplot and a parallel coordinate view (see Figure 2.10). When the user selects certain values in the scatterplot using the brush (a rectangle), and applies the *highlight* operation, the corresponding values in the parallel coordinate view are also highlighted. The view where the user uses the brush is called the active panel (i.e., the scatterplot).

Becker and Cleveland (1987) outline several variations in this basic operation of brushing. For example, the *shadow highlight* operation works like highlight, except that it only shows highlighted values in non-active panels. So, if this operation is applied to the scatterplot in Figure 2.10, then the parallel coordinate view only displays the highlighted values; other values are hidden. The *delete* operation removes the selected values instead of highlighting them. Finally, the *label* operation enables users to display the names of the selected values.

These four brushing operations—highlight, shadow highlight, delete, and label—can be combined with the following three paint modes: transient, lasting, and undo (Becker & Cleveland, 1987). In the *transient* paint mode, the result of a brushing operation is temporary and applies only to the currently selected values. In the *lasting* paint mode, the effect of a brushing operation remains even after the brush object has been moved to cover different values. The *undo* paint mode is available to remove the effects of brushing operations performed in the lasting paint mode. This variety of brush techniques can facilitate analysis of high dimensional data, especially when the data is visualized in multiple views.

Overview+Detail

While browsing, users typically need to move between overviews and detail views. An overview provides contextual information, which allows users to understand a domain and to see the big picture, whereas a detail view enables users to examine items of interest closely. For example, consider the process of browsing hotels in a city. Before selecting a hotel, people may want to see its location relative to downtown or tourist attractions in the city (overview). After finding several options, they need to check detailed information such as the address, price, and availability (detail view). Ideally, when users pay their attention to a detail view, they still can maintain the contextual information as well. Researchers have developed techniques for supporting this interaction, and one of the techniques is called overview+detail.

Overview+detail interfaces use *spatial* separation to display both focused and contextual information simultaneously (Cockburn, Karlson, & Bederson, 2008). These interfaces usually divide a presentation space into two regions: a large space displays detailed information, while a small space provides an overview (see Figure 2.11). In these examples, an inset in Google Maps shows the surrounding area of a focused place to support spatial orientation and exploration, whereas the left panel of Microsoft PowerPoint displays thumbnails to help users browse and navigate through a document.

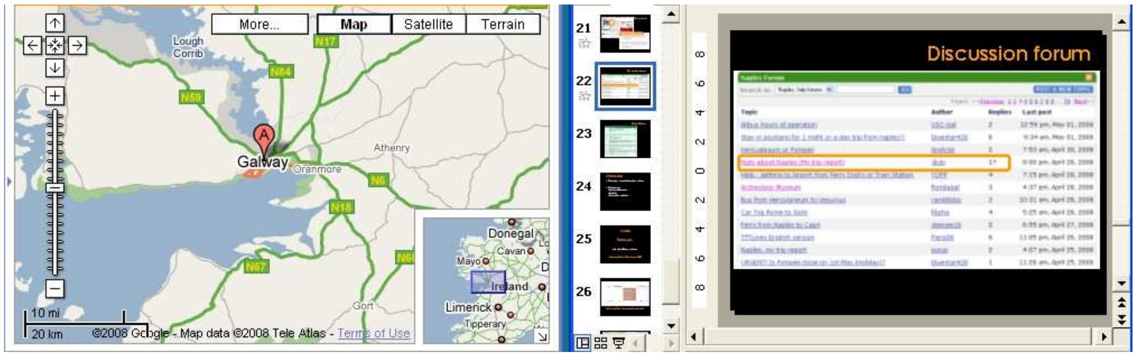

Figure 2.11: Overview+detail interfaces in Google Maps (left) and Microsoft PowerPoint (right). A large space displays detailed information while a small space provides contextual information to facilitate navigation.

There are several design issues pertinent to overview+detail interfaces (Cockburn et al., 2008):

- What is the appropriate scale ratio of an overview to a detail view, and how much space should be allocated for each view? Displaying more thumbnails in Microsoft PowerPoint, for example, reduces the need to scroll the overview panel and enables quick access to distant pages, but it also decreases the readability of the thumbnails. So, there is a trade-off here between accessibility and readability. A possible solution is to permit users to adjust the scale ratio and the space allocation of overview+detail interfaces.

- When should an overview and a detail view be synchronized? Should an application allow users to explore each view independently, or should it reflect any change in one view to the other immediately? The current common practice is to update overviews immediately when users modify detailed views. However, change in overviews is not always reflected right away in detailed views: applications such as Adobe Reader and Microsoft PowerPoint let users scroll the thumbnail panel without changing the currently displayed page.

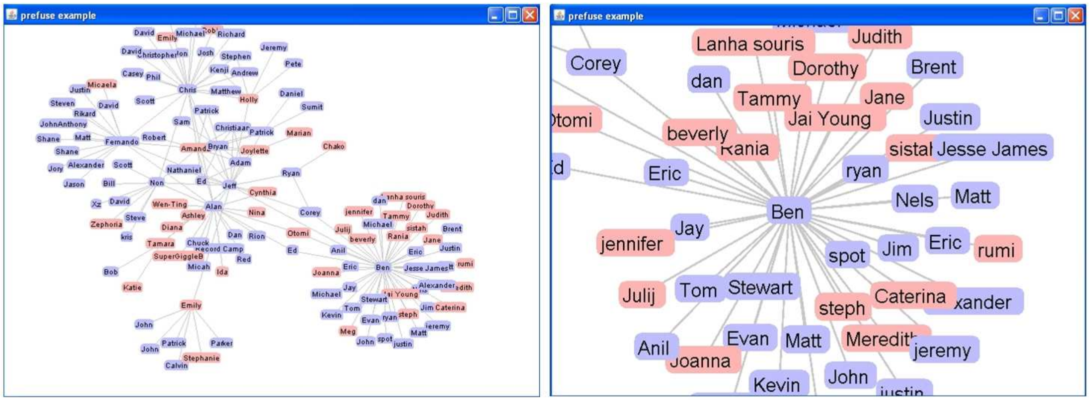

Figure 2.12: Temporal separation in a zoomable interface (Prefuse demo, http://prefuse.org/). Two different levels of magnification are shown in a single display at different times.

Zooming and Panning

Zooming is another technique for facilitating access to detailed and contextual information. Zoomable interfaces use *temporal* separation to enable users to access overviews and detail views (Cockburn et al., 2008). Users obtain an overview by reducing the sizes of documents in an information space so that they get a panoramic view of the space. When they need access to detailed information, they can select and magnify a specific area in the space. Figure 2.12 illustrates how an overview (left) and a detail view (right) can be obtained by adjusting magnification levels of a zoomable interface. In computer desktop environments, applications typically use the scroll button of a mouse to enable zoom in/out operations.

As shown in Figure 2.12 (right), when users zoom in on a certain area of an information space, they get only a partial (narrow) view of the space. If they want to focus on a different area, they have to change their focused area. Moving across an information space to select a focused view is called panning, and it is usually supported by a click-and-drag interaction technique.

Furnas and Bederson (1995) provide a framework for analyzing and understanding zoomable interfaces. They represent a zoomable interface as a space-scale diagram, where the horizontal axes capture the dimensions of a zoomable interface

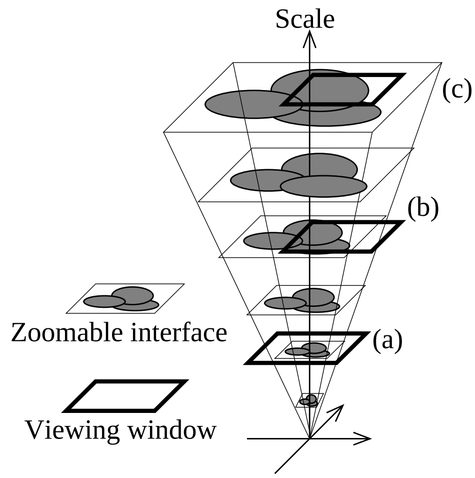

Figure 2.13: A space-scale diagram illustrating a zoomable interface at various levels of magnification (Furnas and Bederson, 1995). At a low zoom level (a), the viewing window is able to capture the whole information space, providing an overview of the space. At higher zoom levels (b) and (c), the viewing window can only display partial views of the information space, and it also takes longer distance to pan across the space.

while the vertical axis portrays scale or magnification levels of the interface (see Figure 2.13). The size of the viewing window is fixed regardless of the magnification level of the interface. Space-scale diagrams can be used to describe the properties of zooming and panning operations, and to illustrate how these operations relate to each other.

A space-scale diagram shows that there is a trade-off between zooming and panning (see Figure 2.13). In a zoomed-out view, users can move across a large information space efficiently, but the focused area may appear small and thus is not readable. A zoomed-in view, on the other hand, can improve the readability of information items, but it takes longer to move around the space. Furthermore, a zoomed-in view may cause disorientation when users are not familiar with the information space, as they lose contextual information, which facilitates navigation. Consequently, users may have to zoom in and out repeatedly to deal with this limitation of zoomable interfaces.

Zooming can be classified into two categories: geometric or semantic. Both

Figure 2.14: An example of focus+context interfaces. A fisheye lens is used to show details of a focused object within context (retrieved on May 14, 2008 from http://demo.quietlyscheming.com/fisheye/TileExplorer.html).

types of zooming modify the sizes of information items. Semantic zooming extends the notion of geometric zooming by also changing the levels of details shown as information items are zoomed in/out (Perlin & Fox, 1993). For example, when a document is zoomed out, it may appear as a dot. But, as the document is zoomed in, it reveals more detailed information such as its title, abstract, headings, and eventually the whole contents. Thus, semantic zooming not only concerns with the size of an item, but also controls how much information is displayed at each magnification level.

Focus+Context

While overview+detail and zoomable interfaces use spatial and temporal separation to enable access to contextual and detailed information, focus+context interfaces eliminate this separation by integrating an overview and a detail view into a single continuous view (Cockburn et al., 2008). These interfaces typically use distortion to display a focused area in detail while also showing the surrounding context in successively less detail. This integration is intended to reduce users' cognitive load in assimilating detail and contextual views, and to facilitate navigation and comprehension of information spaces.

The concept of focus+context interfaces stems from observation that humans tend to develop "fisheye views" of the world (Furnas, 1986). Fisheye views refer to representations of local context in great detail and those of global context in less detail. Such views occur naturally in everyday settings. For example, people are familiar with landmarks in their hometowns but only recognize a few major landmarks in the neighbouring cities. In workplaces, they can name all colleagues in their divisions, but beyond their divisions, they only recognize the senior management.

Furnas (1986) formalizes fisheye views as a "degree of interest" (DOI) function. A DOI function consists of two variables: *a priori importance* and *distance*. This function computes a number representing how important an item is relative to the current focus in a fisheye view. This number is used to determine which information items should be displayed in a fisheye view: if the DOI of an item is lower than a certain threshold, then the item is considered uninteresting and hence is not displayed in the fisheye view. Formally, the DOI of item x from the current focus y is calculated using the following formula: $DOI(x|y) = API(x) - D(x, y)$, where $API(x)$ is a priori importance of x, and $D(x, y)$ is the distance between x and y. A priori importance is a predefined score assigned to each information item, while distance refers to the spatial or semantic gap between the current focus and other items. For example, major landmarks such as Niagara Falls and the Statue of Liberty may have a high score of a priori importance. However, for travellers in Asia, these landmarks have a low degree of interest because they are located far away from Asia (the current focus).

Sarkar and Brown (1992) extend Furnas's work (1986) by developing an algorithm to transform graphs into fisheye views. This algorithm computes a *visual worth* of each vertex in a graph by considering a priori importance of the vertex and its distance to the current focus. The notion of visual worth is equivalent to Furnas's degree of interest. Compared to some threshold, a visual worth determines whether a vertex should be displayed in a fisheye view. The algorithm also applies some functions to calculate the position, size, and level of detail of vertices in a fisheye view. Figure 2.14 illustrates how a graphical fisheye view typically looks: it magnifies focused objects and demagnifies those that are far away from the focus.

Based on Furnas's DOI function (1986), Van Ham and Perer (2009) propose an interaction model to explore graphs called "search, show context, expand on demand." The basic idea is that given an initial node of interest, the function computes and shows a subgraph that is most relevant to the given node. The interaction model then allows users to successively expand their current subgraph to get more contextual information surrounding a selected node. This model is suitable for users who have a semi-specific information need or interest and want to explore a collection using that interest as a starting point.

Filtering

With advances in information visualization techniques, the amount of information that can be displayed in a single view has increased significantly. However, humans have limited capacity to process information, and therefore need to focus their attention on a subset of document collections to avoid information overload. Filtering enables users to narrow a collection and select a subset of documents that match their search criteria. Although its function is simple, implementation of filtering, if not done properly, can hinder users in exploring and browsing information spaces.

IR systems typically enable filtering by providing a set of search criteria that can be modified by users. These criteria vary depending on the application domain. For real estate applications, search parameters may include the range of house prices, the number of bedrooms and bathrooms, and the size of the properties; whereas for hotel reservation systems, the parameters may comprise the accommodation type, the star rating, the price, and the location of hotels. To ensure valid input and to reduce users' cognitive load, these applications usually put constraints on the possible input values by using radio buttons, check boxes, and combo boxes (see Figure 2.15).

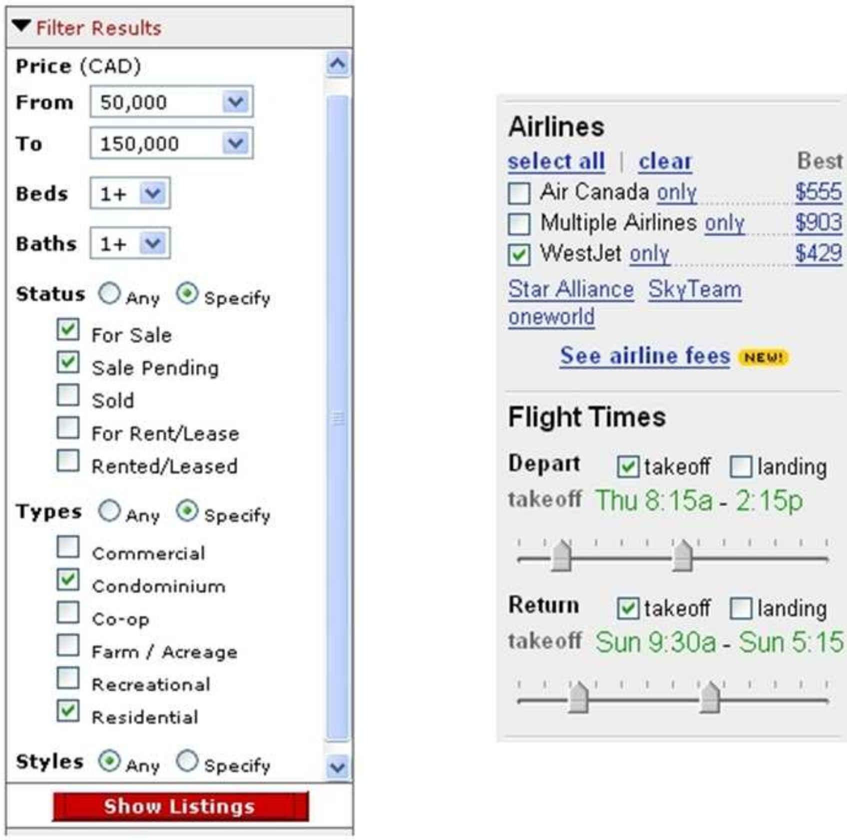

Figure 2.15: Facility to filter search results in commercial applications. Combo boxes, check boxes, and radio buttons are commonly used to limit input values. The screenshot on the right shows an example of dynamic query interfaces: the application provides sliders to filter search results by flight times (retrieved on July 2, 2008 from http://www.remaxsaskatoon.com/ (left) and http://www.kayak.com/ (right)).

To further assist users in adjusting search criteria, Ahlberg et al. (1992) developed a concept called dynamic query interfaces. Dynamic query interfaces use graphical widgets such as sliders to accept user queries (see the right screenshot in Figure 2.15). There are several underlying principles of dynamic query interfaces:

- Users formulate queries by manipulating graphical widgets (e.g., sliders) instead of typing query syntax.
- These widgets reveal the possible query ranges visually, helping users both learn about the characteristics of data sets and avoid making mistakes in adjusting search criteria (e.g., entering invalid values).
- Users receive feedback immediately after changing search criteria using these widgets, allowing them to spot trends in data sets or to reverse their actions if necessary.

Ahlberg et al. validate these design principles in an experiment and show that a dynamic query interface can improve user performance in various search tasks.

Based on the principles of dynamic query interfaces, Dörk, Carpendale, Collins, and Williamson (2008) developed interactive visualization widgets (VisGets) to support web-based information exploration. VisGets visualize publication dates, locations, and topics of a collection as a bar chart, a map, and a tag cloud (see Figure 2.16). Besides providing overviews of a collection, these widgets allow users to formulate queries visually. For example, users can adjust time sliders in a bar chart to select articles published during a certain period. They can select a region on a map to display articles that refer to that area. And when they want to find articles about a specific topic, they can use a tag cloud as a filter to remove irrelevant items. These widgets work in a coordinated manner, enabling users to filter an information collection by its temporal, spatial, and topical dimensions.

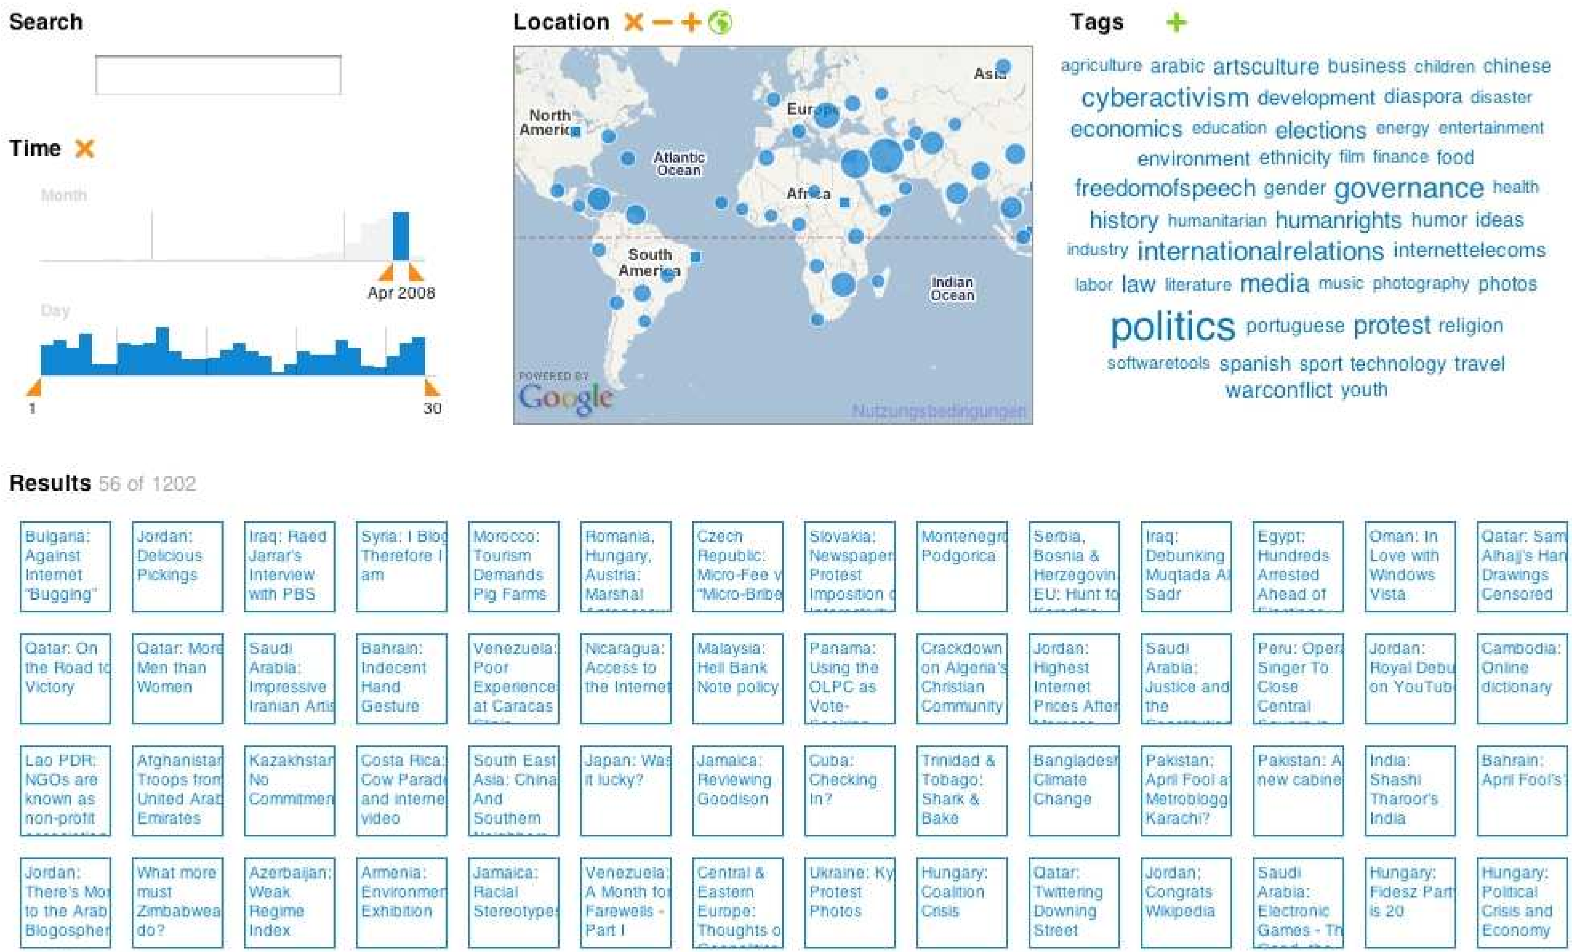

Figure 2.16: VisGets: Coordinated interactive visualization widgets (retrieved on November 21, 2009 from http://www.mariandoerk.de/md/uploads/Research/DiplomaThesis/view1.png).

2.4 Summary

This chapter reviewed the foundations of and earlier work related to this thesis. It started with a definition and characteristics of exploratory searches. The discussion illustrated the evolving and multidimensional aspects of exploratory search goals, and emphasized the active participation of users in the search process. The next section then defined the notion of browsing in the context of information retrieval, and discussed its potential to ease exploratory search processes. After that, this chapter summarized earlier work on supporting browsing and exploratory searches. The review was structured into three sections: information representation, organization, and interaction. Information representation addressed various ways to provide a compact description of information items. Information organization described basic structures for arranging and presenting information to users. The last section, information interaction, summarized information visualization techniques for supporting user interaction with information spaces.

Chapter 3

Social Interaction History

This chapter presents a conceptual framework for supporting information exploration using social interaction history. Social interaction history refers to traces of social interaction in information spaces. The basic hypothesis of this framework is that social interaction history can serve as a good indicator of the potential value of information items. To test this hypothesis, I collected a random sample from a social web application[1] and performed statistical analysis of the sample. The statistical analysis supported this hypothesis: there were significant positive relationships between traces of social interaction and the degree of interestingness of web articles. In this chapter, I elaborate on the foundations and the elements of social interaction history, describe the study methodology and the experimental results, and discuss the potential applications of social interaction history in various domains.

3.1 Foundations

This section summarizes four pillars of the social interaction history framework. First, the idea of visualizing social interaction history comes from the concept of computational wear (Hill et al., 1992). Second, visualization of social interaction history can serve as "information scent" (Pirolli & Card, 1999), which helps users assess information value of documents based on activity traces of other users. Third, enabling users to follow traces of interaction history is the core idea of social navigation (Dourish & Chalmers, 1994). Finally, this framework acknowledges the multidimensional notions of relevance (Barry, 1994; Saracevic, 2007) and argues that social

[1] http://digg.com

interaction history can serve as an indicator of cognitive, situational, and affective relevance.

3.1.1 Computational Wear

The basic principle of computational wear is to record traces of interaction between users and digital objects (e.g., documents, spreadsheets, menus), construct a graphical representation of this interaction history, and then embed this visualization in the objects—resulting in history-rich objects (Hill et al., 1992). Visualization of interaction history is analogous to use wear in physical objects. However, unlike physical wear, computational wear is traditionally hidden from users. For example, a typical file system displays documents in the same way: the graphical representations of the documents (icons) do not indicate which document the user has accessed frequently or opened recently. Computational wear is intended to reveal such information to users.

The development of computational wear is motivated by observation that wear in the physical world often represents useful information (Hill et al., 1992). For example, popular books in a library show a lot of wear and tear because they have been borrowed by many people; skid marks on the road may indicate a slippery surface; and footprints of bears in a hiking trail send a warning about the presence of bears in the area. Thus, there is a lot of useful information that can be derived from observing physical wear, and this information can help people adjust their actions (e.g., people slow their cars when they see skid marks on the winding road).

Hill et al. (1992) discuss two applications of computational wear:

- *Edit wear* records and visualizes authorship history of documents. The history contains information such as parts of a document that were changed in the last editing session, that have been edited a lot, and that have become relatively stable.

- *Read wear* works in a similar manner, except that it concerns readership history. Read wear enables users to identify popular sections in a document, various categories of readers and when they read particular sections.

Edit wear and read wear use an enhanced scrollbar to convey these interaction histories. The scrollbar thus serves two purposes: first, it allows users to extract patterns from the visualization, and second, it provides a navigational tool to locate things of interest in documents.

By recording and visualizing interaction history, computational wear can reveal users' activity patterns over time (Hill et al., 1992; Kaptelinin, 2003; Krishnan & Jones, 2005). Besides helping users understand their own activities, emerging activity patterns can be used to support collaborative work. Consider a team of programmers trying to fix a bug in an application. After working for some period, they may be able to notice parts of source code that have been read or changed a lot and by whom. This emerging pattern may indicate a possible location of the bug and subsequently help the programmers focus their attention on a specific part of the source code.

Computational wear is a by-product of user-object interaction, so it does not require any effort from users to update or maintain interaction history (Hill et al., 1992). From a user's perspective, computational wear comes for free. Information contained in interaction history is recorded and updated automatically. Recent wear is immediately reflected in the visualization. In this way, visualization of interaction history evolves following changes in user activity.

Building upon Hill et al.'s work (1992), Wexelblat and Maes (1999) developed a theory of interaction history. The theory suggests six properties that can be used to design and analyze history-rich systems:

- Property 1: Proxemic versus distemic.

  The terms *proxemic* and *distemic* refer to the degree of closeness between people and spaces. Wexelblat and Maes (1999) consider the notion of closeness or proximity in terms of both physical and cognitive distance.

  In a proxemic space, users feel that they are part of and familiar with the space, and understand its structure and semantics well. Operating systems, for example, try to create a proxemic space by using a desktop metaphor, allowing users to relate the terminology used in the systems (e.g., folders, documents) to their experience of managing physical documents.

A distemic space, in contrast, gives strange, unfamiliar feelings to users. Users have difficulties in navigating through the space and in understanding what is going on in this space. These difficulties may be due to the lack of background knowledge that is necessary to comprehend local culture and practices. Such feelings sometimes occur, for example, when people visit a foreign country, join a new company, or use a new software application.

- Property 2: Active versus passive.

  Wexelblat and Maes (1999) distinguish between active and passive interaction history. Active interaction history refers to interaction traces that are created intentionally by users, such as a list of websites bookmarked by the users. Passive interaction history, on the other hand, is recorded by a system automatically and usually results from a by-product of interaction between users and objects, such as a list of visited websites in browsing history. Thus, the distinction between active and passive interaction history is determined by the presence or absence of user intention to leave interaction traces.

- Property 3: Rate/form of change.

  Rate/form of change concerns the accumulation of interaction history over time. As people use and interact with objects regularly, more and more interaction traces are recorded. What to do with these records and how to present them to users are among the challenges faced by designers developing history-rich systems. Simply presenting all interaction records in a raw format would overwhelm users. Therefore, designers should find ways to summarize and represent the history of interaction in a format that is meaningful and easy to understand, such as by using graphs and information visualization.

- Property 4: Degree of permeation.

  Degree of permeation describes how closely interaction history is embedded in an object. Wear as a trace of interaction in the physical world is always part of an object: a coffee stain on a book, a hole in a T-shirt, skid marks on the road, and so on. Based on this observation, Hill et al. (1992) suggest embedding computational wear in the objects themselves to produce history-rich objects. In some cases, however, it is better to detach interaction history from the objects. For example, a library maintains a list of users who have signed out a book in its database application and not in the book itself. While this example illustrates a practical purpose—to keep the book clean—in the digital world, the choice is more open and flexible. It is possible to attach a list of previous borrowers to an e-book, for example, to enable discussion among the library patrons. Designers should determine how closely they want to embed visualization of interaction history in the object itself, considering the context of an application and other design constraints and requirements.

- Property 5: Personal versus shared.

  Interaction history can be kept personal or shared with other users. Personal interaction history allows individual users to recognize and learn about their own patterns of interaction with objects. Shared interaction history contains collective interaction history and enables users to see things from a more general perspective (e.g., recognizing popular books in a library). The difference between personal and shared interaction history is exemplified by bookmarking facility in web browsers versus social bookmarking applications such as Delicious[2] in which people share their bookmarks with other users. Shared interaction history opens new opportunities to facilitate collaborative work. However, individual users should also have options to control shared materials and to protect their privacy.

- Property 6: Kinds of information.

  The last property of the theory of interaction history concerns the types of information that should be stored by an application. While it is possible to

[2] http://delicious.com/

store all interaction history information in detail, from a practical point of view, not all of this information is useful and relevant to the user's task. With respect to the user's task, application developers should select and process only relevant information, which in general can be categorized into *what* (what was done), *who* (who did it), *why* (why it was done), and *how* (how it was done).

Researchers have used interaction histories to enhance browsing and searching (e.g., Freyne, Farzan, Brusilovsky, Smyth, & Coyle, 2007; Komlodi et al., 2007; Wexelblat & Maes, 1999; White et al., 2007). Interaction histories vary from a simple list of recently visited websites to a complex aggregate of interaction traces left by many users. The main principle of these approaches is to monitor user activities in browsing and searching, and then use this data to facilitate search activity, such as providing suggestions to users who have similar information needs to those of other users in the past. The suggestions can be in various forms. For example, the system may suggest relevant search terms to retrieve similar information items (Freyne et al., 2007; White et al., 2007); provide links to popular websites in which many users end up after submitting a query (White et al., 2007); or augment the user's search history with meta-information such as the last access time to particular web resources or keywords used to find particular articles, and then allow the user to re-find these items using this local, contextual information (Cutrell, Robbins, Dumais, & Sarin, 2006; Komlodi et al., 2007). Such suggestions enable users to learn from past search trails (e.g., to see common terms used in a domain) as well as to follow the trails of others if necessary.

3.1.2 Information Foraging

Information foraging theory is a framework for understanding how people adapt their strategies and environments to optimize their efforts to obtain valuable information (Pirolli & Card, 1999). The basic hypothesis of this theory is that people would always attempt to maximize the return on their efforts to find useful information. Besides modelling human behaviour in information foraging, this theory provides some models (discussed below) that can serve as a framework for analyzing the role of social interaction history visualization in information exploration.

In information foraging, strategic and environmental adaptations are necessary to enable humans to deal with abundant amounts of information, to cope with their limited capacity to process information, and to find useful information items with minimum costs (Pirolli & Card, 1999). Evidence of how people adjust their strategies and environments for these purposes appears in both physical and digital spaces. For example, studying office organization, Malone (1983) found that people arrange their offices to minimize costs of accessing task-relevant information items, such as putting frequently used documents on their desks while storing other items in bookshelves. The distance between a document and the user's work area often indicates the degree of relevance of that document to the user's current task. The practice of modifying the structure of physical environments to improve access to information is also reflected in the digital space. Computer users often create program shortcuts and keep work-in-progress documents on their desktops to ease access to these items (Ravasio et al., 2004).

Drawing analogies between information foraging and food foraging, Pirolli and Card (1999) developed some models to analyze human behaviour in information gathering and sense making:

- *Information patch models* concern the allocation of time for foraging activities. These models assume that information exists in patchy environments. An information patch may refer to various collections, such as a file folder, an email inbox, a search result, or a pile of documents. While looking for information, users need to decide whether to spend more time exploring an information patch (within-patch foraging) or to look for a better patch (between-patches foraging). They may modify their environments to improve the foraging results. Examples of environmental modifications include sorting email messages by sender to ease finding a message from a particular person, refining search criteria to filter out irrelevant items from search results, and storing piles of relevant documents close to each other to reduce costs of between-patches foraging.

- *Information scent models* concern the identification of information value and the associated costs of pursuing the information based on the available cues. To maximize the foraging results, users would try to find and follow information scent that seems to lead to relevant documents. In doing so, they observe and use existing cues, such as the titles, abstracts, authors, and publishers of documents. Based on this information, users can assess the degree of relevance and reliability of the documents. Besides the contents, the documents' representation and appearance also affect the user's perception of their potential values. When possible, information systems therefore should help users notice and follow information scent, for example, by increasing the font size of highly relevant documents in a search result (e.g., Olston & Chi, 2003).

- *Information diet models* concern the process of selecting which information item or collection users should pursue in order to maximize the possible return on their foraging efforts. Due to information overload, people must evaluate information sources and choose those that offer high values. For example, while reviewing literature, a researcher may focus on high quality journal and conference publications and ignore articles on popular magazines. During office hours, email users may give priority to work-related messages over personal messages; they may also use a filter to remove junk messages from their inboxes, reducing the costs of email processing and management.

  When people select which information item to pursue, they consider both the potential return and the costs of processing the item (Pirolli & Card, 1999). They do not just select the best quality information sources which are available. For instance, students working on a class project may prefer to use online resources that are free and easy to access, whereas those working on their theses may be willing to put more efforts to get interlibrary loan if necessary.

People will be able to select appropriate information sources if they receive sufficient information scent from the environment. Otherwise, they would wander around the information space without guidance. Therefore, information systems should aim to provide adequate information scent to assist users in optimizing their efforts to find relevant information.

3.1.3 Social Navigation

People have used the term "social navigation" broadly. Dieberger (1997) uses this term to describe common practices such as asking other people for information, sharing interesting news with friends, and publishing a list of favourite websites on the Web. Konstan and Riedl (2003) consider collaborative filtering to be another kind of social navigation. Collaborative filtering systems provide recommendation to individual users by comparing their profiles to other users' profiles, and those who share similar profiles are recommended similar information items. An example of collaborative filtering can be seen at Amazon,[3] an online store which provides a list of recommended items based on its customers' behavioural patterns (e.g., "Customers Who Bought This Item Also Bought").

In this thesis, social navigation refers to the movement of a user from one piece of information to another piece where this movement is influenced by the activity of other users in the information space (Dourish & Chalmers, 1994). For example, while browsing a discussion forum, users may select a particular message because they see that the message has received a high rating from other users. Such movement is essentially made based on observation of what others have done to the message. The rating hence serves as information scent (Pirolli & Card, 1999), indicating the value of that message from other users' perspective.

Svensson and Höök (2003) distinguish between direct and indirect social navigation. In *direct* social navigation, a user receives direction or recommendation from another person through dialogue. This dialogue may happen in a face-to-face setting or through a communication tool. In *indirect* social navigation, a user does not receive any instruction from another person, but observes the activity (or the trace

[3]http://www.amazon.com/

of activity) of other people and decides to follow it. For example, watching a movie because our friend recommends it to us is an example of direct social navigation. However, if we select that movie because we see a long queue in front of the box office, then we use indirect social navigation to make this decision.

Social navigation is enabled when people can observe traces of activity and social interaction in their environments. In physical environments, human activity leaves traces naturally, and these traces are easily observable. For instance, we can navigate through a hiking trail by following a well-worn path. A well-worn path implies that many people have used this path frequently, and based on this social information, we can infer that we must be on the right track without checking our map too often. Unfortunately, such activity traces traditionally are not available in digital systems, preventing users from using social navigation to guide their movement. For example, when we visit a website, there is no indicator which web page is most visited or how many people are currently accessing the website. In a chat room, we may see the number of online users, but we do not know exactly who is listening actively to our conversation. To facilitate interaction and collaboration in digital environments, it is necessary to reveal such social information to users.

To support social interaction in digital systems, Erickson and Kellogg (2003) developed the concept of "social translucence." Social translucence is an approach to revealing and visualizing the presence and activity of users in digital environments in order to promote shared awareness and to support social processes. This approach is developed based on observation that shared awareness facilitates and shapes the structure of social interaction in everyday life. Socially translucent systems have three basic properties:

- *Visibility:* A system should allow users to see the presence and activity of other users in the system. In face-to-face settings, we observe and use social information to guide our interaction with other people. For example, when we see that our colleagues are busy, we refrain from asking them to help us with our problems. While giving a lecture, we adjust our action based on how our audience respond: we elaborate on our talk when the audience seem to be interested, but only briefly discuss topics that seem to bore them.

- *Awareness:* When a system visualizes the activity of its users, it promotes shared awareness among the users. Users can see what others are doing, and at the same time they know that others can also see what they are doing. This awareness, however, does not mean that the activity of users becomes transparent. A system must protect user privacy and only reveal information that is necessary to support social awareness.

- *Accountability:* Shared awareness leads to user accountability. Since users know that others can notice their actions, and vice versa, they have to be accountable for their actions. This situation encourages users to follow social rules and norms, thereby supporting social interaction in the system.

In recent years, researchers have transformed visualization tools into social spaces to enable collaboration and social navigation in data analysis. For example, Many Eyes (Danis, Viégas, Wattenberg, & Kriss, 2008; Viégas, Wattenberg, van Ham, Kriss, & McKeon, 2007) and sense.us (Heer, Viégas, & Wattenberg, 2007) enable users to create, annotate, and share visualizations of data sets. These applications facilitate social data analysis by allowing users to leave comments and highlight patterns on each other's visualizations. In their study, Heer et al. observe that people use social navigation in their data exploration by checking other people's comments on visualizations to see what others have found or were looking at.

To facilitate social navigation, Willett, Heer, and Agrawala (2007) developed guidelines and a software framework for enriching graphical widgets with visualizations of interaction traces in information spaces. This approach exemplifies an application of computational wear (Hill et al., 1992) and the information foraging theory (Pirolli & Card, 1999) in supporting social navigation in information spaces. The embedded visualizations reveal the patterns of aggregate user activity and serve as information scent to help other users navigate an information space. By looking at these visual cues, users may decide whether to follow the crowds or to explore less popular information items in a collection.

3.1.4 Relevance

The main goal of an IR system is that, given a query, the system attempts to retrieve relevant documents that fulfill the information needs of the user. To determine a document's relevance to a query, the system uses some algorithm, which usually involves computing the frequency of search terms in the document compared to that in the whole collection. In this case, the notion of relevance refers to *aboutness* of a document seen from a system perspective (Saracevic, 2007), and traditionally, evaluations of IR systems only consider this type of relevance in their performance measures (i.e., precision and recall).

While a document's subject is an important factor that determines its relevance to a query, there are other factors that affect user judgments of relevance (Barry, 1994; Saracevic, 2007). Knowledge of the subject is one of these factors. For example, experts at a certain subject may consider introductory articles about that subject to be irrelevant to their needs, because they already know about it. However, the same articles may be rated highly relevant by those with limited knowledge of the subject. Time constraints can also affect relevance judgments. For instance, users who have short time to seek information may ignore long documents, although these documents discuss the topic of interest, and hence are topically relevant. Since such situations cannot be adequately captured in user queries, it is difficult for a system to know the actual needs of individual users and subsequently deliver personalized responses to them. In other words, only users can know exactly what they need and assess whether a document is relevant to their needs. Therefore, there have been proposals to involve users in the evaluations of IR systems to acknowledge the multidimensional aspects of relevance (e.g., Borlund, 2000, 2003; Dumais & Belkin, 2005).

Relevance can be considered to be the relation between two or more entities (Saracevic, 2007). These entities may refer to both abstract and concrete things such as a query, an idea, an information item, or a user's state of knowledge. The relation has some measurable property that describes how these entities are connected to each other. For example, a document can be considered relevant to a query because the subject of the document matches that of the query. In this example, the document and the query are the interacting entities, while the subject is the property that establishes that relation.

Reviewing information science literature, Saracevic (2007) summarizes the manifestations of relevance as follows:

- *System or algorithmic relevance:* relation between a query and an information item as determined by the underlying algorithm used by an IR system. System relevance is affected by the ways a system represents, organizes, and indexes its document collection. Besides processing the contents of the document collections, some algorithms also consider the structure of the collections in their relevance judgment processes. For example, the PageRank algorithm uses the link structure of the Web to infer the value or importance of a web page (Brin & Page, 1998). This assigned value subsequently affects the relevance score of the web page and its position in search results.

- *Topical or subject relevance:* relation between the subject of a query and an information item. Topical relevance assumes that each query and document represents some topic or subject. Thus, if a query and a document are about the same subject, then both items are considered topically relevant.

- *Cognitive relevance:* relation between a user's state of knowledge and an information item. Cognitive relevance is measured by the impact of an information item on a user's cognitive state. Quality and novelty of information are among the criteria used to assess cognitive relevance.

- *Situational relevance:* relation between the current contexts of a user and an information item. Situational relevance concerns the degree of usefulness of a document in helping users solve their problems, make informed decisions, reduce uncertainty, or complete their tasks.

- *Affective relevance:* relation between the goals, motivations, and emotions of a user and an information item. Affective relevance concerns the subjective user satisfaction with a document. Various factors, such as the style and clarity of writing or the relationship between the author and reader of a document, affect this notion of relevance.

In summary, there are various manifestations of relevance which should be considered in the evaluations of IR systems. While the standard performance measures of precision and recall can assess the effectiveness of a system to retrieve topically relevant documents, these measures use stated requests (i.e., queries) to measure relevance and neglect other notions of relevance. In practice, relevance is always judged based on the information needs of the user instead of the user queries (S. E. Robertson & Hancock-Beaulieu, 1992). Therefore, IR researchers need to involve users in their studies so that they can measure the performance of a system in fulfilling the information needs of the users.

3.2 Social Interaction History

3.2.1 Definition

Social interaction history refers to records of social interaction among users in information spaces. While interaction history focuses on *user-object* interaction (Hill et al., 1992), social interaction history expands this concept by including *user-user* interaction regarding an information item (see Figure 3.1). To see the difference between these concepts, consider interaction in a classroom. When students read a textbook, highlight important passages, and put annotations on the book, they leave traces of interaction history of the book. But when they discuss the book contents with their professor and classmates, this discussion is considered to be a part of social interaction history of the book.

The basic hypothesis of this framework is that social interaction history indicates the potential value of an information item. Traces of social interaction can be used to measure the degree of popularity of an information item. The more an item

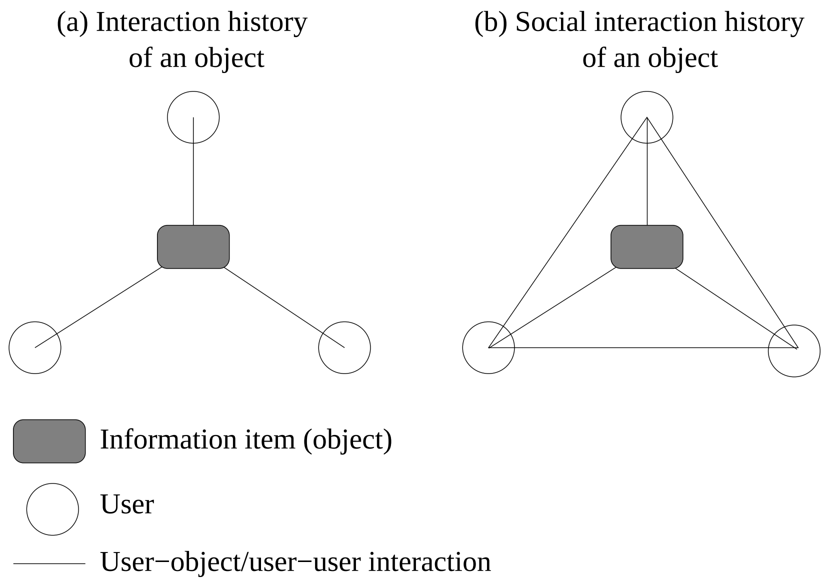

Figure 3.1: (a) Elements of interaction history: user, object, and user-object interaction; (b) Elements of social interaction history: user, object, user-object interaction, and user-user interaction.

contains traces of social interaction, the more popular the item is. However, these traces cannot tell exactly the reason why the item attracts a lot of attention from the audience.

The underlying reason why an information item becomes popular depends mainly on the context of a collection. Depending on the context, social interaction history may indicate the degree of usefulness, interestingness, timeliness, novelty, or quality of an information item. For example, in a collection of academic papers, popularity may indicate the quality and impact of research in a particular field. In a movie collection, popularity may imply that a movie has a good story line and features good actors and actresses. In another case, a recipe may become popular because it is easy to follow and produces a healthy, delicious meal.

Besides the context of a collection, the characteristics of users also determine what social interaction traces imply. A community of serious scholars may appreciate thoughtful posts and dislike jokes or silly comments, whereas in another community, the sillier a comment is, the better the community will like it. Thus, people must be familiar with the characteristics of a collection and its users to understand the implied meaning of social interaction history.

An information item that has rich social interaction history is likely more useful or interesting than those that have little or nothing. Traces of social interaction history therefore can serve as information scent (Pirolli & Card, 1999) and can be used as criteria to infer cognitive, situational, and affective relevance (Saracevic, 2007). Furthermore, social information can be incorporated into IR algorithms to affect the order of retrieved documents in search results. For example, if two documents, A and B, discuss the same subject, but A receives more comments than B, then A should be ranked higher than B in search results.

Social interaction history is present in social media such as wikis, online newspapers, and blogs. Web applications have shifted from static publishing media to interactive media that allow users to communicate and discuss the contents of an article with other users by leaving comments on the article. There is a lot of information that can be derived from these traces of interaction. The number of comments on an article may indicate the degree of usefulness, interestingness, or controversy of the article. Data analysis may reveal the characteristics of user interaction and answers the following questions: How many users are involved in a discussion? Is there a dominant user? Who has been active for a long time? What are the sentiments of the discussion? Does the discussion last for a short or long period? Which articles have triggered a lot of discussion? This information can help users reason about the social dynamics within a community, identify core members of the community, and understand the social norms in the community. Furthermore, such information is also useful for establishing figures of authority on certain topics and for assessing the reliability of posted information.

Analyzing social interaction history is the centre of research on social networks (Garton, Haythornthwaite, & Wellman, 1999). However, the objective of this thesis is different from that of social network research. Research into social networks focuses on the relations among social entities and the overall structure of a network, whereas this thesis focuses on the role and application of social interaction history in supporting information exploration and browsing.

3.2.2 Elements of Social Interaction History

The first step in planning to use social interaction history is to analyze and extract potentially useful information from the history. Figure 3.1(b) illustrates that there are four elements of social interaction history: user, object, user-object interaction, and user-user interaction. This framework uses these elements to analyze typical information that is available in a history of social interaction. In each element, developers can use interrogative words such as *who*, *what*, *when*, *where*, and *how* to extract specific information of interest. Table 3.1 provides examples of information that is commonly preserved in social interaction history, and that is potentially useful for supporting information exploration.

A key component of the social interaction history framework is user-user interaction. This aspect has not been considered in the manifestations of relevance (Saracevic, 2007), which only focuses on the relation among queries, documents, and information needs of individual users. To acknowledge the influence of social interaction on relevance judgments, the social interaction history framework introduces and defines an additional manifestation of relevance called social relevance.

Social relevance concerns the relation between a user and an information item in a social context. Social relevance considers the user's social network to be the determining factor in relevance judgments. For example, a user decides to read a book because her friends have been excited about the book recently. Initially, she may not have personal interest in the book. Her primary motivation for reading the book is driven by her friendships: she wants to be able to relate to her friends when they are talking about the book. This example illustrates the influence of social interaction on relevance judgments.

Having collected and analyzed social interaction history, an application may use this information for various purposes, for example:

- Promoting articles that have been read by many users and that have triggered

Table 3.1: Typical information in social interaction history.

Elements	Information
User	Who are the users? Who are the active users? What are the user demographics? What are their interests or expertise? When did they start using the application? Where are their locations of residence? How many users do use the application? How often do they use the application?
Object (information item)	Who wrote or posted the article? What is the topic of the article? What are the categories of the article? When was the article posted? Where was the article posted? How long is the article?
User-object interaction	Who has read the article? Who likes or dislikes the article? What is the general user sentiment towards the article? When was the last time a user read or commented on the article? Where did a user save or annotate the article (e.g., in a public/private space)? How many comments have been received by the article? How many times has the article been read?
User-user interaction	Who interacts with whom frequently (e.g., leaving comments on each other's posts)? What are the hot topics of the discussion? When was the last reply posted? Where did users interact with each other (e.g., in a public/private space)? How many users are involved in the discussion? How quickly do they reply to one another? How long has the discussion been going on?

a lot of discussion;

- Promoting articles that are read by friends of the user;
- Recommending articles to users who show similar taste and characteristics;
- Integrating the popularity of documents into its algorithm to compute the relevance score of the documents in search results;
- Visualizing social interaction history to help users select potentially useful information items (e.g., Indratmo, Vassileva, & Gutwin, 2008; Willett et al., 2007);
- Granting special privileges to users who have made good contribution to the community (e.g., Slashdot[4]).

In summary, social interaction history has a lot of potential and opens many opportunities to improve information management, retrieval, and exploration in various domains. By mining social interaction history, an application can identify the characteristics of users and the patterns of their interactions, and subsequently use this information to provide personalized services and recommendation (e.g., Worio[5]). While social interaction history has many potential applications, this thesis provides a case study of visualizing social interaction history to improve exploratory browsing and focuses the discussion on this context.

3.3 Study Methodology

3.3.1 Objective

The social interaction history framework is based on a conjecture that traces of social interaction can serve as an indicator of the potential value of information items. I aimed to test the validity of this hypothesis by performing statistical analysis of

[4]http://slashdot.org/
[5]http://www.worio.com/

Digg's data set.[6] Digg is a social web application used for sharing web resources such as news, videos, and images. Digg users may submit new stories to Digg or vote for existing stories. A "digg" number is the number of user votes for a story and represents the degree of interestingness of the story (Lerman & Galstyan, 2008). Digg's data set was appropriate for this study because it contains social interaction history such as the number of user votes for each story ("diggs"), comments from readers, and users who were involved in a discussion.

While the scope of social interaction history is very broad, this experiment focused on two factors: the numbers of comments on stories and the numbers of users who commented on the stories. Following earlier work (Lerman & Galstyan, 2008), this study used the number of diggs to represent the degree of interestingness of a story. The research hypotheses were as follows:

1. There would be a positive relationship between the numbers of comments and diggs;

2. There would be a positive relationship between the numbers of commenters and diggs; and

3. The correlation between the numbers of commenters and diggs would be stronger than the correlation between the numbers of comments and diggs.

I derived these hypotheses from the following presumptions. First, an article that receives many comments must contain valuable or interesting information; otherwise, readers would not spend any effort to discuss it. Second, besides the number of comments, the number of commenters is also an important, if not more important, indicator of the degree of interestingness of an article. An article that receives many comments from few readers may be less interesting than an article that triggers a lively discussion involving many people.

[6] http://digg.com

3.3.2 Data Set

The data set consisted of 266 random news stories which were promoted to popular stories in Digg's Business & Finance section. The status of a story—whether it is popular or upcoming—is determined by Digg. To collect this sample, I developed software using a Java toolkit[7] to interact with the Digg Application Programming Interface. The tool generated random numbers representing the Unix time between January 1 and December 31, 2008, and used these numbers to retrieve stories that were promoted around particular time. Besides collecting news stories, the tool also retrieved all comments on each story. For data analysis, I extracted the numbers of diggs, comments, and commenters from the data set. Since there were deleted comments which could not be retrieved, I used the actual numbers of retrieved comments in the analysis. The statistical tests were done using SPSS 16.0 for Windows.

I focused on popular stories because I wanted to examine if there are correlations between traces of social interaction and the degree of interestingness of these stories. Popular stories tend to have rich social interaction history. They usually receive many diggs and trigger a lot of discussion, and hence were suitable for this study. On the contrary, upcoming stories do not fit the purpose of this study because they typically receive very few diggs and no comments.

Instead of using recent stories, I chose stories which were promoted in 2008 to give enough time for each story to reach a stable state: that is, the numbers of diggs on stories remain relatively constant after a few days, and the final counts represent the degree of interestingness of the stories from the general community's perspective (Lerman, 2007). Based on my observation, the numbers of comments on stories also stabilize after a short period, although this period may vary from several days to a few months. Overall, the numbers of diggs, comments, and commenters in the data set had become stable when the data were retrieved in March 2009.

[7] http://bleu.west.spy.net/~dustin/projects/digg/

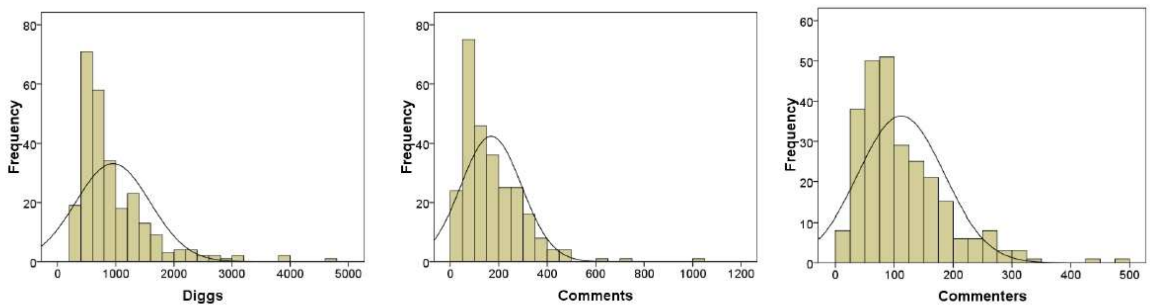

Figure 3.2: The distributions of the numbers of diggs, comments, and commenters.

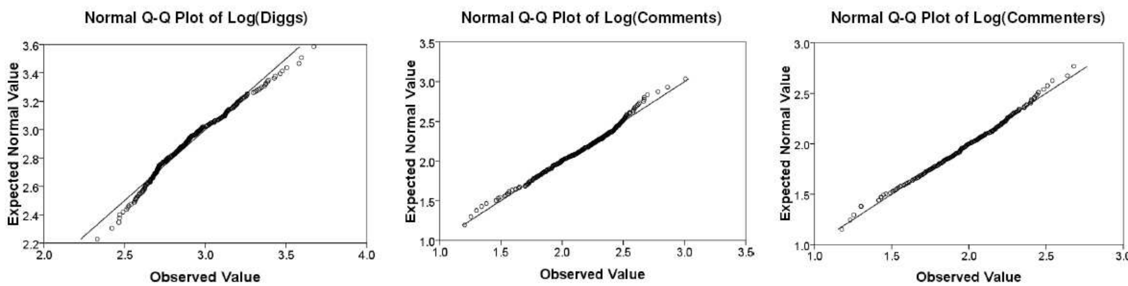

Figure 3.3: The normal Q-Q plots of log(diggs), log(comments), and log(commenters).

3.4 Results

The distribution of diggs in the sample ranged between 214 and 4701 (M = 952, SD = 641). The numbers of comments varied from 16 to 1019 (M = 169, SD = 125). The numbers of users involved in a discussion were between 15 and 478 (M = 112, SD = 73).

Figure 3.2 shows that the distributions of diggs, comments, and commenters were positively skewed. There were few stories that received large numbers of diggs and comments. These stories might have been promoted as top stories in the past and were easily accessible from Digg's front page, hence having a better chance to receive more diggs and comments than other stories. Figure 3.3 illustrates that the results of a logarithmic transformation ($\log_{10}$) of these data approximately followed a normal

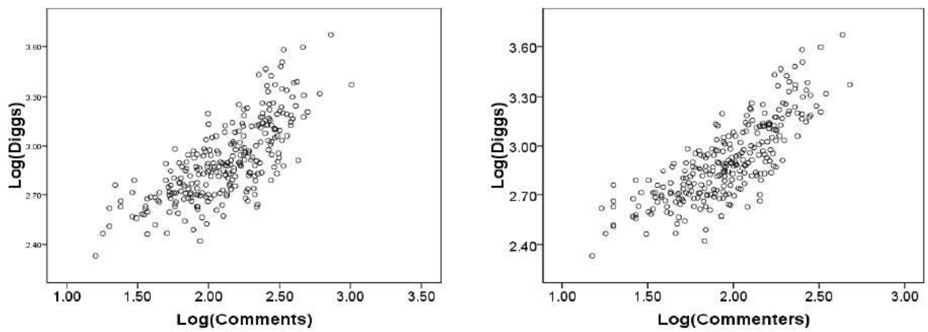

Figure 3.4: The scatter plots of diggs-comments and diggs-commenters (log data).

Table 3.2: Pearson correlation matrix (log data).

	Diggs	Comments	Commenters
Diggs	1.00	0.75	0.80
Comments	0.75	1.00	0.98
Commenters	0.80	0.98	1.00

distribution, thereby allowing parametric tests on the transformed data.

The Pearson correlation coefficient (r) was used to compute the correlations between the numbers of diggs, comments, and commenters. Since this statistical test assumes that a sample is taken from a bivariate normal distribution, the test was run on the transformed data. Figure 3.4 plots the relationships between the numbers of diggs, comments, and commenters (log data), whereas Table 3.2 summarizes the correlations between these variables. These results showed that the numbers of comments and diggs were highly correlated, $r = 0.75$, $p < 0.001$ (one-tailed), and so were the numbers of commenters and diggs, $r = 0.80$, $p < 0.001$ (one-tailed). To test whether the diggs-commenters correlation was significantly higher than the diggs-comments correlation, I ran a t test for comparing two nonindependent correlation coefficients (Howell, 2002, p. 281). The result suggested that the difference between these correlation coefficients was significant, $t(263) = 5.88$, $p < 0.001$ (one-tailed).

3.5 Discussion

The results of the statistical tests supported all of the research hypotheses:

1. There was a significant positive relationship between the numbers of comments and diggs;

2. There was a significant positive relationship between the numbers of commenters and diggs; and

3. The correlation between the numbers of commenters and diggs was significantly higher than the correlation between the numbers of comments and diggs.

If the number of diggs represents the degree of interestingness of a story (Lerman & Galstyan, 2008), then these results also supported the basic hypothesis of the social interaction history framework: that is, records of social interaction history can serve as a good indicator of the potential value of an information item. As a measure of story interestingness, the number of commenters was a better indicator than the number of comments. A few talkative users can leave a lot of comments on a story, but this does not imply that the story is interesting for the general audience. However, if these comments are written by many people, then it indicates that the story attracts attention from a wider audience.

Factors that make a story interesting are varied and may be related to the notion of cognitive, situational, and affective relevance (Saracevic, 2007). Two stories discussing the same topic may receive different responses from the audiences due to differences in the clarity of the writings, the reputation of the authors, the timeliness of the publication, or other non-topical aspects. Social interaction history can assist in the assessment of the quality of such non-topical factors.

3.5.1 Implications for Design

Social interaction history has a lot of potential to improve IR systems. Search engines can take social information into account while processing queries and returning search results. Articles receiving many comments can be ranked higher than similar articles receiving few comments. Collaborative filtering systems can analyze discussion threads and recommend articles to users with similar interests, such as those who often participate in the same discussion threads. Traces of social interaction can be visualized to enable social navigation. Users will be able to select appropriate information sources if they receive sufficient information scent from the environment (Pirolli & Card, 1999). Thus, information systems can use social interaction history to provide additional clues to assist users in finding relevant information.

In an application where users can rate or vote for information items, social interaction history can complement user votes to improve the underlying algorithm that measures the degree of interestingness of information items. Besides relying on the user ratings, an algorithm can incorporate other factors such as the number of comments on an article and the number of users who are involved in a discussion to compute the popularity of the article. In fact, engaging in a discussion by leaving comments or replying to other comments requires more effort than simply submitting a vote. Therefore, social interaction traces can serve as a valuable source of information that indicates the potential value of information items. To help users find potentially useful items, an application can promote articles that have received many comments and that have triggered lively discussions involving many users. Visualization of social interaction traces can also assist users in exploring and learning about the characteristics of social information spaces.

The role of social interaction history is even more important when users cannot rate information items. Consider blogs as an example. Bloggers generally do not allow their audiences to rate their posts. As a result, it is hard to recognize potentially useful entries without reading or skimming over the individual entries. When a blog archive contains many entries, browsing these entries sequentially requires a lot of time and effort. Providing an overview of a blog archive and visualizing the traces of social interaction enable users to reason about the archive at a glance and provide heuristic guidelines for selecting entries to read (Indratmo et al., 2008). Such visualization is particularly useful for supporting exploratory browsing and serendipitous discovery of useful information items.

The basic principles of social interaction history can be applied to other domains which traditionally do not consider social interaction that happens in the information spaces, such as digital libraries. Research on digital libraries often focuses on how to improve retrieval performance. In line with the emerging social media, developers can take a different approach to enhancing user experience of using digital libraries. For example, digital libraries can be transformed into a social space where users can interact and conduct discussion about a book with other patrons. Books that have been borrowed or discussed a lot subsequently can be made more visible so that people new to an area can find prominent books easily. Without compromising the user privacy, a system can provide a communication channel that allows a user to contact other users who have read a particular book or who appear to be an expert in a particular field. These features may improve the user experience of using digital libraries significantly.

3.5.2 Basic Requirements of Social Information Spaces

Social applications must maintain a low noise-to-signal ratio. Social interaction in a public space always attract "noise" such as spam messages. Noise does not represent meaningful social interaction, and hence should be removed from the historical records. Digg maintains this requirement by requiring users to login if they want to participate in a discussion. Thus, we can assume that comments on Digg stories are generated by real users. Other applications use a simple test or a "captcha" to ensure that a request or information is submitted by a human. If this requirement is not fulfilled, the number of comments on a story may not correlate to the degree of interestingness of the story.

Social applications must provide some mechanisms for maintaining social order, preventing damage, and resolving disagreement while giving flexibility to users to shape the information spaces. To maintain the quality of its contents, Digg enables users to "bury" submitted stories and comments that are deemed to be poor. If a story or a comment has been buried by many users, then the content of this item is hidden. On the contrary, Digg promotes stories that have received many votes from users. This feature allows social control over the contents of Digg.

Wikipedia[8] is another good example of systems that enable social control (Viégas, Wattenberg, & Dave, 2004). Wikipedia is an online encyclopaedia in which anyone can write and edit its contents. The quality of Wikipedia contents is controlled by communities of users. A user's contribution is continually reviewed and edited by other users. Wikipedia keeps track of and stores revision history of each article. The revision history in Wikipedia exemplifies social interaction history, as authors of an article often communicate with one another through comments on their revisions. To support exploration of social dynamics in Wikipedia, Viégas et al. developed a tool to visualize revision history in Wikipedia. Several patterns of interaction and collaboration emerge from this visualization. For example, the "vandalism and repair" pattern portrays a user's attempt to vandalize an article by deleting its contents, and how the Wikipedia community fixes this vandalism by putting back the previous version of the article. The "negotiation" pattern reveals how authors repeatedly change the contents of an article to their favoured versions. Without visualization, such social interaction patterns are hidden in the massive, textual records of revision history.

Another approach to maintaining social order is to provide an incentive to improve the quality of a community. For example, to promote high quality comments on posted articles, Slashdot[9] keeps track of the ratings of comments from its users. Users whose comments have received good ratings may be given a special privilege to be a moderator. On the other hand, Slashdot reduces the visibility of posted articles or comments from those who have posted spam messages or low quality comments in the past. Besides granting special privileges to good users, an application may complement such incentive mechanisms with visualization of the levels of user participation in a community (Vassileva & Sun, 2007). These mechanisms encourage users to contribute high quality information to their communities.

---

[8] http://wikipedia.org/
[9] http://slashdot.org/

3.6 Summary

This chapter outlined a framework for using social interaction history to support information exploration and navigation. The hypothesis of this framework is that social interaction history can serve as a good indicator of the potential value of information items. To test this hypothesis, I collected a random sample from Digg and performed statistical analysis on the data set. The study focused on two factors of social interaction history: the number of comments on a story and the number of users who are involved in a discussion. The results supported the research hypothesis: there were significant positive relationships between these factors and the popularity of stories.

With the emerging popularity of social media and social networking applications, social interaction history is becoming more prevalent. Blogs and online newspapers, for example, allow users to leave comments on posted articles, promoting conversations among the users. Social networking applications maintain users' friend lists and keep track of their interactions. Analyzing and visualizing these kinds of information may yield interesting results and patterns. The next chapter reports on a case study of visualizing a blog archive to explore the potential of using social interaction history to facilitate browsing.

Chapter 4

Exploratory Study

To evaluate the potential of using social interaction history for supporting information exploration, I developed and evaluated an interactive visualization tool for browsing a blog archive (iBlogVis). Blogs were selected because they are social media, and hence contain social interaction history. Furthermore, browsing a blog is an exploratory activity in which people wander around the information space to find the entries of interest. Thus, the nature of this activity fits the objective of this research: to find ways to support exploratory searches.

To give contextual information about this study, I start with an introduction to blogs, the lack of support for browsing blog archives, and related work in blog visualization. Then I discuss the design and rationale behind iBlogVis. After that, I describe the study methods for evaluating the usability of iBlogVis and report on the findings from this study.

4.1 Blogs

Blogs have emerged as a new medium for communication. According to Lenhart and Fox (2006), in the US, about 12 million people maintain blogs, and about 57 million Internet users read blogs. Blogs promote conversation between the bloggers and their audiences by allowing users to comment on published entries, or to cite the entries of interest in their own blogs and then send notification to the source using the TrackBack protocol.[1]

A blog archive contains a collection of entries arranged in reverse chronological

[1] http://www.sixapart.com/pronet/docs/trackback_spec

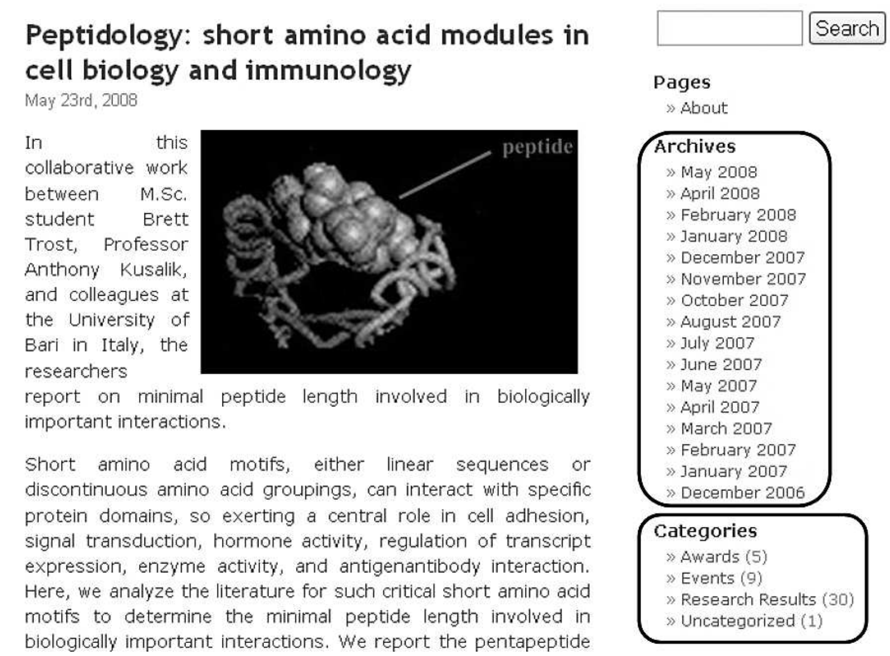

Figure 4.1: Typical navigational support in blogs: links to monthly archives and a list of tags or categories (retrieved on June 3, 2008 from http://www.cs.usask.ca/csweblog/).

order. Typically, only a few of most recent entries are displayed on the front page of a blog. Useful content, however, is not limited to the most recent entries; there may be many old entries in a blog that are worth reading and that offer valuable information. Therefore, users may find useful entries by browsing a blog archive.

Browsing a blog archive, however, is not well supported. Typically, blogs only provide links to monthly archives and a list of tags or categories for selecting a subset of entries assigned with a specific keyword (see Figure 4.1). This navigational support does not offer the users any cue where to find potentially useful entries—entries that have sparked lively discussion or that have been read by many people. Should there be a valuable entry posted two months ago, users would likely miss it because the entry is no longer (less) visible in the blog. Consequently, users have to rely on their own capability to explore blog archives by filtering entries by tag or by browsing the monthly archives sequentially. This task becomes tedious when users browse a large collection of entries.

This limitation results from the nature of reverse chronological order of blog entries, and the presumption that the value of an entry decays with time. While this presumption may be true for diary-kinds of blogs, the value of entries in some blogs—especially topic-oriented blogs—may last for a long time. The value of an article about designing good visualization, for example, does not necessarily decrease with time, just like a seminal paper that is worth reading over and over again. Since many people now have used blogs as a medium to express and elaborate their thoughts, which are often enriched by comments from the audiences, there are many valuable discussions preserved in blog archives. To assist users in locating such discussions, blogs need to provide additional support for browsing their archives.

To support exploration of blog archives, I developed iBlogVis—an interactive visualization tool that offers a new way to browse a blog archive (see Figure 4.2). There are two key features of iBlogVis. First, it provides a rich overview of a blog to enable users to reason about the blog at a glance. Second, it visualizes the history of social interaction in a blog to help users identify potentially useful entries in the blog. This prototype exemplifies the synthesis and application of existing visualization and interaction techniques to a new domain (blogs). The evaluation showed that the prototype was successful in realistic tasks.

There are two generalizable principles embodied in iBlogVis that can be applied to other domains. First, findings from this study indicated that visualizing social interaction history has potential for supporting exploratory tasks—initial findings which were examined more rigorously in a summative evaluation (Chapter 5 discusses the study method and results). Second, the findings added further evidence that providing an overview of an information space is valuable. These principles can have immediate impact on real-world sites such as Blogger[2] and LiveJournal[3] to complement their current user interface designs.

[2]http://www.blogger.com
[3]http://www.livejournal.com

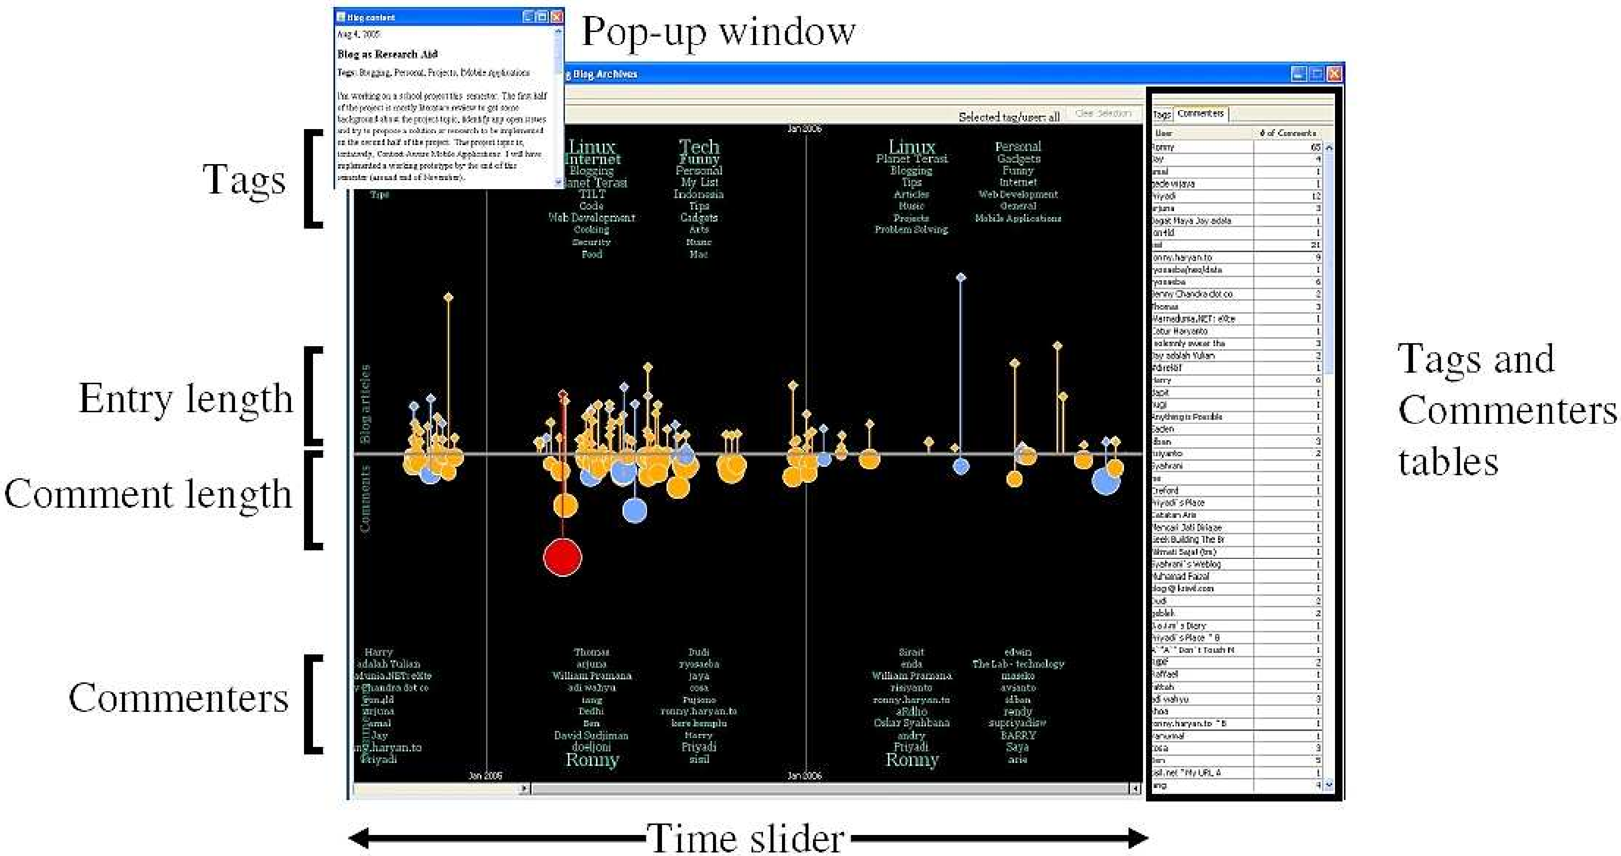

Figure 4.2: A screenshot of iBlogVis visualizing blog entries along a timeline.

4.2 Related Work in Blog Visualization

To provide context and background information about this study, this section reviews related work in blog visualization and shows how this study differs from other work.

Research on blog visualization currently focuses on analyzing and visualizing the link structure and the content of blogspace—a large-scale collection of blogs. Link analysis, for example, has been used to track information flow through blogspace (Adar & Adamic, 2005; Gruhl, Liben-Nowell, Guha, & Tomkins, 2004). Gruhl et al. developed a model to track the routes of topic propagation through individual blogs. Their model is similar to disease epidemic models—a blogger gets "infected" with a topic and then spreads the topic further to his/her contacts. In this context, contacts refer to the audience of a blog. Similarly, Adar and Adamic developed a technique to infer the source of information spread in blogspace based on the timestamps of entries and the link structure of blogs. They visualize the inferred routes as "infection trees" where the nodes and the edges represent blogs and propagation paths.

Herring et al. (2005) analyze link structure of blogspace to examine the interconnectedness among blogs. They use social network visualization to plot information such as inbound and outbound links and whether the links are one-way or reciprocal. The visualization reveals clusters of topic-oriented blogs that are more interconnected and reciprocally linked than "A-list" (most popular) blogs. A-list blogs, however, have more one-way inbound links, and hence are more reachable than other blogs.

Using visualization, social network analysis, and a 'sense of community survey,' Chin and Chignell (2006) developed a model to discover communities in blogspace. Unlike other link analysis methods, their method does not use links mentioned in blog entries. Instead, it uses links provided by bloggers while leaving comments on other blogs (the links usually point to the bloggers' personal blogs). Chin and Chignell consider such links more explicit in suggesting a social relationship among the bloggers.

As more and more people use blogs to express their opinions, mining blog content can yield useful information. Companies are interested in knowing what people say about their products or what the current hot topics are in blogspace. Analytical tools and search engines have been developed to meet such demands, including BlogScope,[4] Digg,[5] and Technorati.[6]

Tirapat, Espiritu, and Stroulia (2006) developed an interactive tool to examine whether there is a correlation between the success of a movie and the "buzz" in blogspace. They use document-clustering techniques to analyze blog entries and construct a topic map, capturing associations between movies and blog entries. Based on the topic map, the tool generates multiple views to allow users to explore different aspects of the data set.

Harris and Kamvar[7] developed a tool that looks for the phrases "I feel" and "I am feeling" in blog entries and extracts human feelings from the entries along with information about the bloggers (age, gender, and location) and the local weather conditions while the entries were posted. This information is saved; the feeling is identified (e.g., happy, sad) and then visualized as a particle. The attributes of a particle (e.g., colour) represent some encoded information. A particle can be clicked to display a full sentence that describes a human feeling. The visualization tool allows users to search and sort data by feeling, gender, age, weather, geographic location, and date.

Other related projects, not specific to blog visualization, include LifeLines (Plaisant et al., 1996), TagLines (Dubinko et al., 2006), and Timeline.[8] These systems provide interactive environments or an application programming interface (Timeline) to visualize data sets along a timeline. Specifically, LifeLines visualizes personal histories such as medical records, whereas TagLines identifies most representative tags at Flickr[9] during a certain time period and visualizes the evolution of these tags.

Compared to the existing work on blog visualization, this project differs in the following way. Instead of analyzing and visualizing blogspace, this study focuses on

[4]http://www.blogscope.net
[5]http://digg.com
[6]http://www.technorati.com
[7]http://www.wefeelfine.org
[8]http://simile.mit.edu/timeline/
[9]http://www.flickr.com

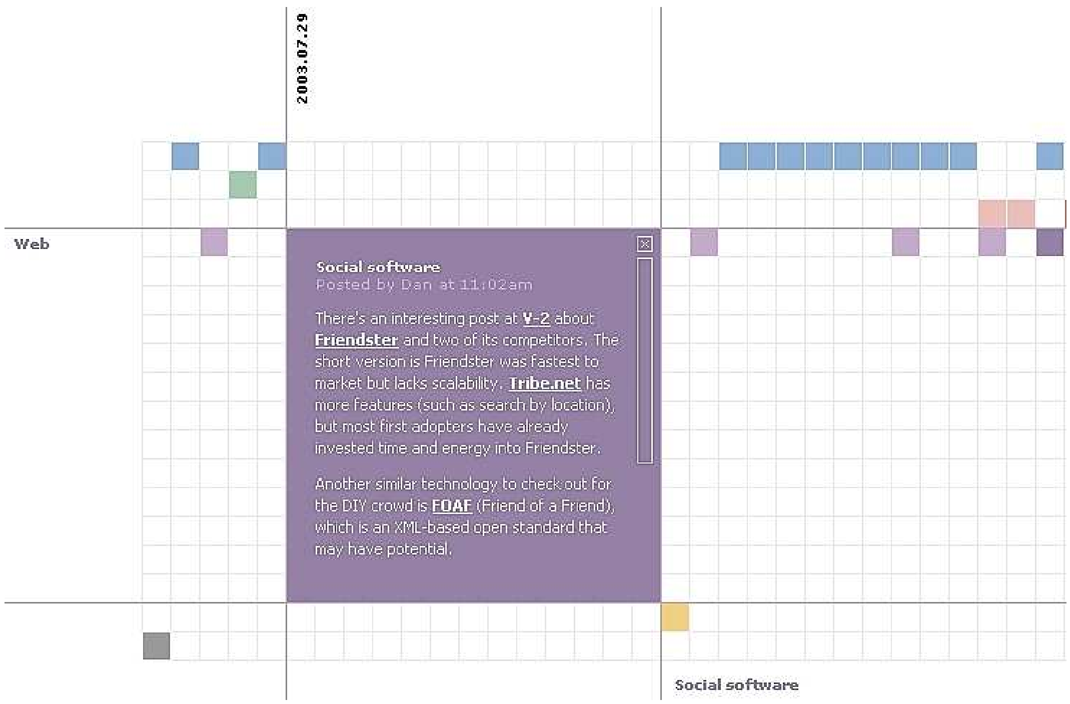

Figure 4.3: Box Grid visualization (retrieved on February 1, 2008 from http://phiffer.org/projects/box-grid/).

supporting exploration of individual blogs. To some extent, the research objective is similar to that of Box Grid[10] (see Figure 4.3). Box Grid visualizes blog entries in a grid, where the position of an entry in the grid is determined by the entry's category on the vertical axis against its posting date on the horizontal axis. Unlike iBlogVis, however, Box Grid neither visualizes the social interaction history preserved in a blog archive nor allows the users to filter the visual items. In the following section, I discuss the design and rationale behind iBlogVis.

4.3 Design and Rationale

The design of iBlogVis follows design principles in information visualization. First, the graphical representations of data adhere to the guidelines on how to visualize quantitative, ordinal, and nominal data (Cleveland & McGill, 1984; Mackinlay, 1986). For example, iBlogVis uses graphical attributes such as length and area to represent quantitative data while using colours to encode nominal data (see Table 4.1). Second, iBlogVis uses a timeline metaphor to visualize blog entries because a timeline is easy to understand and has been deployed successfully in visualization tools (e.g., Plaisant et al., 1996; Viégas & Smith, 2004). Third, the interactive components of iBlogVis apply Shneiderman's mantra (1996): "overview first, zoom and filter, then details-on-demand." I elaborate the rationale behind this design decision below.

4.3.1 Design Objective

The main design objective of iBlogVis is to facilitate exploration of individual blog archives. To achieve this objective, iBlogVis provides an overview of a blog and visualizes not only the content of the blog, but also the history of user interaction, thereby providing cues for social navigation (Dieberger et al., 2000). By using iBlogVis, users are expected to be able to answer the following questions:

[10] http://phiffer.org/projects/box-grid/

Table 4.1: The encoding of blog attributes in iBlogVis.

Blog Attributes	Visual Attributes
Posting time	Position
Tag popularity	Font size
Entry length	Line length
Comment length	Line length
Number of comments	Circle area
Commenter activity	Font size
Read wear status	Colour

- *Content*: What is the blog about? Does the blogger post entries regularly? Do the blogger's interests change over time?

- *Social interaction history*: Did the blog receive many comments from the audience? Which entries received many comments? When did the blog start getting popular? Who are the regular commenters?

Without a tool, gathering such information normally requires that users follow a blog: reading entries and comments regularly, observing interests of the blogger, and identifying active users in the blog. Even so, some of this information is an aggregate of values over a long time, which cannot be estimated easily.

4.3.2 Visual Design

Blogs display entries in reverse chronological order. To reflect this organizational structure, iBlogVis arranges visual items along a timeline (see Figure 4.2). The visualization panel consists of two main parts. The upper part (above the timeline) visualizes the content of a blog: entries and the associated tags. The font size of a tag represents the tag's popularity during a certain year. The larger a tag, the more frequently the tag is used during the corresponding year. A blog entry is represented by a diamond shape and a line. The diamond shape provides an interface to view the content of an entry, while the length of a line represents the number of characters in an entry.

The lower part (below the timeline) displays the social interaction history contained in a blog archive. The line length and the font size have similar meanings as above, except that these visual items represent comments and commenters, respectively. The length of a line represents the total number of characters in all comments received on a particular entry, while the area of a circle represents the number of comments in that entry.

The right panel contains two tables. One displays a list of all tags and the number of entries labelled with each tag. Another table displays a list of all commenters and the number of comments they have written on the blog. These tables can be sorted alphabetically (by tag or commenter) or numerically (by the number of entries or comments). These sorting functions are to ease retrieval and reveal the popularity of tags and commenters. Users may select an item from these tables to filter items displayed in the visualization panel.

Visualization of blog content allows users to get an overview of the blog's subjects and the length of the entries. Tag visualization is supported because tagging is currently the main method for classifying entries. Tag visualization is useful for assessing whether a blog matches a user's interests, while the length of entries can serve as a retrieval cue. Furthermore, the decision to select which messages to read (in Usenet newsgroups) is influenced strongly by the subject and the size of a thread (Fiore, LeeTiernan, & Smith, 2002). Tags and the length of entries and comments in blogs are similar to the subject and the size of a thread in Usenet, and hence may be influential as well in the exploration of a blog.

Visualization of social interaction history can serve as social navigational cues. As shown in Figure 4.4, entries that receive many comments can be spotted easily by looking at the size of the circles. Since other users have left many comments on these entries, the entries may offer useful information. With hundreds of entries available in a blog, such social navigational cues can help users select potentially interesting entries without spending too much effort on skimming over every entry in the blog.

To further facilitate the browsing activity, iBlogVis uses the idea of computational wear (Hill et al., 1992) to help users keep track of entries that have been read

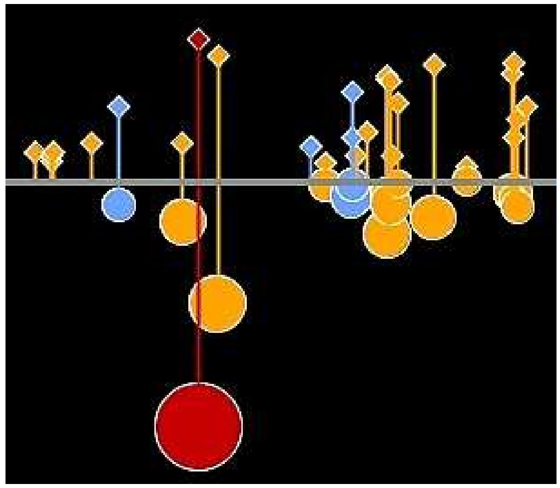

Figure 4.4: Computational wear and social navigational cues.

(blue), have not been read (orange), or the one that is currently being read (red) (see Figure 4.4). Computational wear and its variants have been shown to improve navigation through information spaces (e.g., Skopik & Gutwin, 2005; Wexelblat & Maes, 1999). Table 4.1 summarizes the encoding of blog attributes in iBlogVis.

4.3.3 Interactive Components

The interactive components in iBlogVis are designed based on heuristic guidelines on information visualization: "overview first, zoom and filter, then details-on-demand" (Shneiderman, 1996). These components also serve as a dynamic query interface (Ahlberg et al., 1992), which allows users to formulate queries dynamically and get feedback immediately by clicking on the items in the visualization panel (e.g., tags, commenters), selecting items from the tables in the right panel, or adjusting the time slider (see Figure 4.2). Highly interactive interfaces are engaging and support exploratory tasks (Marchionini, 2006), and hence fit the characteristics of browsing a blog archive.

The visualization tool starts by providing an overview of a blog archive (see Figure 4.2). The overview displays all entries and comments, most popular tags, and most frequent commenters along a timeline. The popularity of tags and commenters is aggregated on a yearly basis. This overview reveals temporal posting patterns of the blogger and enables viewers to scan the content of a blog quickly and to reason about its general structure and community dynamics.

Using iBlogVis, users can filter entries by tag, commenter, and posting time. Displaying a subset of entries labelled with a particular tag (filtering by tag) or those commented on by a particular person (filtering by commenter) can be done by clicking on the tag or the commenter. To perform this task, users can click on a tag or a commenter either in the visualization panel or in the right panel (listed in tables). These filters exist to allow users to limit their search space by removing irrelevant or uninteresting items. Furthermore, since current blogging tools use tagging as the main mechanism for labelling, organizing, and retrieving blog entries, users would expect to be able to explore an archive by tag.

Filtering by posting time is supported by a time slider, located at the bottom of the visualization panel (see Figure 4.2). The slider is positioned across the visualization panel to give users a good sense of the possible query range, as the range corresponds directly to the width of the visualization. This design allows users to derive the query range visually, simply by looking at the timeline in the visualization. Users, therefore, can pay their attention to the visualization while adjusting the time slider. The time slider is also used to zoom in on the area of interest in the visualization. As the visible time range decreases, the level of details in the visualization increases: more tags and commenters are displayed whenever there is more space on screen.

Finally, users can view the content of an entry through a pop-up window by clicking on a diamond shape (representing an entry) or a circle (representing the number of comments on the entry) in the visualization (see Figure 4.2). When users view an entry, iBlogVis changes the colour of the entry to indicate its read wear status.

4.4 Evaluation Methodology

iBlogVis offers a new way to browse individual blog archives, and to the best of my knowledge, there was no other blog browser that I could use for comparison. Therefore, this study focused on soliciting feedback on the visualization design and assessing the utility of the underlying principles in iBlogVis to support browsing tasks. Due to the similar purpose of the tool (i.e., facilitating information exploration) and the characteristics of user tasks (i.e., exploratory search), the study methods were based on earlier studies (Marchionini, 2000; Viégas & Smith, 2004).

The usability of iBlogVis was evaluated using both objective and subjective performance measures. The objective measure was the error rate of the completed tasks, while the subjective measure was the user satisfaction with the tool. The user satisfaction was measured using a set of tool-specific questions adapted from earlier work (Marchionini, 2000; Viégas & Smith, 2004), and a short version of the Questionnaire for User Interaction Satisfaction (QUIS 7.0) (Chin et al., 1988). I included only relevant QUIS items from the following categories (the number indicates the number of questions used in that category): overall user reactions (6), screen (3), terminology and system information (5), learning (4), and system capabilities (4). Besides using Likert-scale items to measure user satisfaction, I also asked the participants to explain their ratings. These exploratory questions were intended to shed some light on user practices in exploring a blog and to get constructive feedback about the tool.

Since iBlogVis offers an alternative way to browse a blog, the decision whether to use the tool depends much on the user satisfaction with the tool, which makes this criterion important. Quantitative results from this measurement, combined with observation and comments from the participants, can give valuable feedback for improving the visualization tool.

The tasks in the study were designed to evaluate three aspects of the visualization tool:

1. *Overview design*: How effective the tool was in giving the users an overview of the content and community dynamics (i.e., social interaction history) of a

blog;

2. *Basic functionality*: How well the tool worked in helping the users do things that they could do with typical blogs; and

3. *Visualization of social interaction history*: How effective the tool was in giving the users information about community dynamics in a blog (e.g., revealing regular commenters and popular entries).

4.4.1 Participants

Nineteen students (13 males, 6 females) from the University of Saskatchewan participated in the study. The ages of participants ranged from 23 to 37 years old. Each participant received a $10 honorarium. As a prerequisite, the participants had to be familiar with browsing the web and have some experience reading blogs. On a scale of one (beginner) to five (expert), the participants self rated their computer skills as an end-user at level three or above (two at level three, six at level four, and eleven at level five). On average, twelve participants browsed the web between one to five hours daily, while the rest spent more than six hours daily.

4.4.2 Apparatus

I implemented iBlogVis as a desktop application using Java and the prefuse toolkit (Heer, Card, & Landay, 2005). iBlogVis has a pre-processing module that both transforms the data structures of a blog into tables and the GraphML file format[11] and computes the aggregate values required by the visualization, such as the popularity of tags and commenters, the lengths of entries and comments, and the number of comments on each entry. The data set used in the study contained approximately 100 entries and 300 comments posted from October 2004 to December 2006. The blog attracted a regular audience and received comments regularly. All information presented in the visualization was publicly available.

[11]http://graphml.graphdrawing.org/

4.4.3 Procedure

Participants were introduced to the purpose of the study and asked to sign an informed consent form. After that, they were introduced to the features and the meaning of the visualization using an example data set. This demonstration took approximately five minutes. Then participants were given an opportunity to familiarize themselves with the visualization tool.

After a short period of practice, participants were given a blog data set and a set of tasks. There was no time limit for completing these tasks. There were 16 tasks divided into three categories:

- *Overview tasks*: understanding the timeline visualization, identifying the main topics of the blog, recognizing regular commenters, and getting a sense of popularity of the blog.

- *Typical browsing tasks*: finding the most recent entry, filtering entries by tag (by browsing and selecting a tag from the tag cloud or the tag table), and browsing monthly archives.

- *Social navigational tasks*: identifying popular entries and finding entries commented on by specific persons.

During the data collection session, I observed how the participants used the visualization tool, and took notes of comments and difficulties faced by them.

After performing each task, the participants rated their satisfaction with the tool. Some of the questionnaire items were open-ended questions asking the participants to explain their ratings. The participants were also asked about their favourite and least favourite features of the visualization tool, desired functions that did not exist at the time, whether they would be interested in using the tool if it were integrated into blogs, and whether they have privacy concerns with the information presented in the visualization. At the end of the data collection session, the participants completed a short version of QUIS 7.0 (Chin et al., 1988).

4.5 Results

Each data collection session took up to one hour. In general, the participants did not have difficulties in learning to use iBlogVis and completing the tasks. After listening to a brief introduction about iBlogVis, all participants needed less than five minutes to feel comfortable in using the tool and ready for the tasks.

Overall, iBlogVis received positive reviews from the participants. Subjective user satisfaction was high, supported by a low error rate in the completed tasks. Quotations included in this section were from written comments from the participants.

4.5.1 Error Rate

Out of 16 given tasks, the number of errors made by the participants ranged from zero to three. Seven participants performed all tasks correctly; six made one mistake; five made two mistakes; and one participant made three mistakes. The average error rate was 6.25% (1 out of 16 tasks).

The most common mistake occurred when the participants were asked to find the most recent entry in the blog. Among 19 participants, five performed this task incorrectly. As some of the posting dates of the entries were close to one another, there were some overlapping visual items on the overview of the blog. Mistakes occurred when the participants simply selected a visual item on the overview that seemed to be at the right most position of the timeline. The selected item happened to be the second most recent entry in the blog. The more accurate way to perform this task was to zoom in on the area of interest to separate the overlapping items before selecting the most recent entry.

4.5.2 User Satisfaction

The tool-specific questionnaire for measuring user satisfaction used forced-choice fixed-scale items with four points on the scale, for example, not at all easy, not easy, easy, and extremely easy. Table 4.2 provides a summary of the results. Overall, the

Table 4.2: Results of the tool-specific questionnaire ($N = 19$).

No	Questionnaire items	Not at all easy	Not easy	Easy	Extremely easy
1.	Understanding the timeline visualization was ...	0	3	9	7
2.	Identifying the main topic of the blog was ...	0	0	8	11
3.	Identifying regular commenters in the blog was ...	0	0	5	14
4.	Getting a sense of popularity of the blog was ...	0	2	12	5
5.	Finding the most recent entry in the blog was ...	0	0	5	14
6.	Finding blog entries tagged by a specific keyword was ...	0	0	4	15
7.	Finding blog entries posted in a specific month was ...	0	4	9	6
8.	Finding popular entries in the blog was ...	0	0	8	11
9.	Finding blog entries commented on by a specific user was ...	0	1	7	11

participants expressed high satisfaction with the tool. Most items were rated easy or extremely easy. The results of this questionnaire, however, also indicate room for improvement (e.g., problems related to the timeline—see statements 1 and 7 in Table 4.2).

Overview Design

Besides evaluating task-specific functions above, the participants gave their overall reactions to the visualization tool. They rated how effective the tool was in giving overviews of the content and community dynamics of the blog. Most participants thought that the tool presented an overview of the blog content effectively. Four participants rated it highly effective; thirteen rated it effective; and two rated it not effective. In terms of providing an overview of the community dynamics, iBlogVis was rated highly effective by four users and effective by fifteen users.

All participants thought that, while exploring a blog, having access to an overview of the blog was useful (eleven rated it extremely useful, eight rated it useful). A common reason was that an overview can help users learn about the content and the characteristics of a blog and the blogger quickly:

> Because then I don't waste my precious time. I can get a quick overview about the blog.
>
> It gives a good overall sense of the blog: the topics covered, the level of participation, the frequent commenters.
>
> It lets me know what the blog is about overall, the author's evolving interests over time, and how popular it is based on how many readers it got.

Basic Functionality

The participants rated how well iBlogVis helped them do things that they could and could not do with typical blogs. Compared to what the participants could do with typical blogs, iBlogVis was rated extremely well by seven users and well by twelve users. Additional functions of iBlogVis (things that the participants could not do with typical blogs) were rated extremely well by nine users and well by ten users.

Visualization of Social Interaction History

Most participants thought that having access to the community dynamics (i.e., social interaction history) in a blog was useful (nine rated it extremely useful, eight rated it useful, and two rated it not useful). By observing social cues, users can follow the crowds to find most interesting or contentious entries in a blog archive. In other words, visualization of social interaction history offers indirect recommendation about potentially interesting entries to users:

> It always can recommend the most popular article to me.
>
> [The visualization] helps identify which threads contain interesting/useful information.
>
> To know how popular certain posts are, and to know what sort of posts were appreciated by those other than the writer.

Two participants, however, did not find the visualization of social interaction history to be useful. One was only interested in reading the latest entry: "Just want to see the latest news." Another one was "not overly concerned with comments a blog gets." The latter, however, acknowledged that "it's sometimes useful to know what the most popular/contentious entry was."

Favourite Features of iBlogVis

The favourite features of iBlogVis included (1) the visualization of the length of entries and comments; (2) the lists of tags and commenters; and (3) the visualization of the number of comments on blog entries. In particular, the visualization of the length of entries and comments was mentioned most frequently by participants as a favourite feature of iBlogVis:

> I like that the length of a line signifies how long the blog entry/comment was. I haven't seen that before and [the visualization] displays the "tempo" of the blogger (lots of short posts, etc.).
>
> Displaying the length of posts and comments. That would be an incredibly good feature.

Some participants mentioned that the decision to read an entry is influenced by the length of the entry and its comments. One participant mentioned that he preferred to read relatively short entries with many comments, because sometimes he only had a short time to check a blog entry. Thus, visualization of the length of entries and comments could help him make a decision whether to check an entry without having to actually open the entry. Another participant wrote, "I enjoy reading the comments but prefer many short comments to a few long comments." Having access to entry and comment length visualization would help him make a qualitative judgment regarding the comments on blog entries in order to select entries that match his preference.

Least Favourite Features of iBlogVis

The least favourite features of iBlogVis mostly related to the timeline visualization and the use of a time slider to filter visual items by date. Some participants expected to see a clearer boundary between periods (e.g., a monthly boundary). They also wanted to be able to quickly select exact dates or periods by having predefined filters (e.g., by year, month, or week). Other desired features mentioned by the participants included filtering entries by multiple criteria, searching by keywords, and better highlighting the currently selected tag or commenter.

Privacy

iBlogVis computes and displays aggregates of values over a long time, which otherwise cannot be seen clearly in a typical blog. Although all information processed by iBlogVis is available in the public domain, the tool reveals some hidden patterns such as regular commenters and allows users to filter entries by commenter, enabling them to learn about a particular person and his or her thoughts. To get a sense of attitude toward privacy, I asked the participants whether they had privacy concerns regarding the information displayed in the visualization.

Most of the participants did not have privacy concerns regarding information presented by the visualization tool, as the information was already in the public domain. They believed that protecting privacy in the blogosphere is the responsibility of the individuals. None of the participants mentioned or demanded technical support for protecting privacy (e.g., access control). They thought that people should develop a good practice to protect their privacy, such as using pseudonyms while leaving comments on blogs:

> When you place information on the Internet you have already made a decision on privacy. Anonymity can be used to hide identity when placing comments.
>
> I recognize that when I attach my name to an entry it is public and now in the public domain. If I was concerned usually I can post with a pseudonym or anonymously.

Two participants, however, indicated privacy concerns, and one was not sure about her attitude to this issue. One wrote:

> Information presented in a typical blog is not very easy to analyze. But with a visualization tool, participant identifiers are readily available with their comments and demonstrated behaviors, and therefore easy to stereotype.

Overall User Satisfaction

All participants agreed that having access to the visualization tool would affect their choices of which entries to read (ten strongly agreed, nine agreed). The most common reason was that, by looking at the visualization, they would be able to see the popularity and the length of entries:

> I would like to know both how long an entry is and how popular it is before reading, since I would prefer to read short entries that are thought provoking/noteworthy.

When asked whether they would be interested in using the tool if it were integrated into blogs, all participants showed interest in the tool (ten were extremely interested, nine were interested).

After completing a tool-specific questionnaire, the participants filled out a short version of QUIS 7.0 using a nine-point scale. Table 4.3 presents the mean score of each QUIS category used in the study and the lower and upper limits of the mean at 95% confidence level. Based on QuantQUIS,[12] the midpoint scale (five) can be used

[12] http://lap.umd.edu/quis/QuantQUIS.htm

Table 4.3: QUIS results (1 = poor, 9 = excellent, $N = 19$).

QUIS Category	Mean	95% Confidence Interval	
		Lower	Upper
Overall	7.67	7.22	8.12
Screen	7.86	7.30	8.43
Terminology	8.32	7.99	8.64
Learning	8.43	8.09	8.78
Capabilities	8.08	7.69	8.48

to represent mediocre user satisfaction. Compared to this value, the results showed that, in all categories, the user satisfaction with iBlogVis was significantly higher than mediocre.

4.6 Discussion

iBlogVis was designed based on the hypothesis that providing an overview of a blog and revealing social interaction history would help users explore a blog archive. As presented in the previous section, both the quantitative and the qualitative responses from the participants supported this hypothesis. A common reason was that an overview and visualization of social interaction history enable users to learn about a blog at a glance and to identify popular entries in the blog: "I can easily know when the blogger posted entries frequently and find which entries are more popular." This section discusses the roles of an overview and visualization of social interaction history in more detail, along with design implications and limitations of the study.

4.6.1 Overview

An overview serves like a map for spatial navigation or a table of contents in a book. An overview of a blog enables users to quickly assess whether the blog matches their taste or information needs. An overview is especially useful for people who encounter and explore new blogs. Having an overview, people visiting a blog for the first time do not need to spend a lot of effort such as skimming through individual entries to get general information about the content of the blog and the characteristics of the blogger. Thus, an overview helps users form the first impression of a blog, assess its quality and relevance, and decide whether to continue exploring the blog:

> If, really, I am now deciding whether to blogroll it.[13]
>
> Most of the times you can form an impression of a blog in two seconds, but sometimes you may want to look at it more closely to see if it's really worth your time.
>
> Because in blogs, many posts are very useless. It's nice to see the good ones.

4.6.2 Social Interaction History

While exploring a new information space, users need heuristic guidelines to select potentially useful documents in the space. iBlogVis provides such guidelines by visualizing not only the content of a blog but also social interaction history such as the popularity of entries and regular commenters. Visualization of social interaction history can serve as a collaborative filter to reduce the number of entries that potentially contain useful information: that is, entries that have attracted a lot of attention from the audience. After reading a few most popular entries, users will get some idea about the writing style of the blogger and be able to select other entries that may be worth reading:

> I'd probably just read the most popular stories first to get a sense of the blogger's style and focus. Then I'd read a few of the least popular entries for comparison.
>
> It helps me to choose the entry that might interest me most. There might be hundreds of entries with the tag I want, and the community dynamics [visualization of social interaction history] can help me to filter them.
>
> It gives me an idea if the entry is worth reading. My experience is that most of the entries I am interested in reading are usually heavily commented.

Due to the versatility of blogs as a communication medium, not all blogs offer useful information to their visitors. The majority of blogs are used to express

[13] A blogroll is a list of favourite blogs.

personal feelings and serve as personal diaries (Herring, Scheidt, Bonus, & Wright, 2004). Such blogs are probably relevant only to the bloggers and their close friends and families. In contrast, blogs attracting a large audience most likely contain information that is relevant and useful for the general public. Therefore, the presence of social cues can be used to assess the relevance, quality, and influence of a blog:

> [Visualization of social interaction history] shows the support of a blog and its influence.
>
> It helps find the most popular/comprehensive blog or popular blogger.

Visualization of social interaction history can also help people reduce uncertainty of the quality and influence of a blog. With support from easy-to-use blogging tools, setting up a blog and posting entries to it involve simple steps, requiring no or little technical knowledge of web publishing. Almost everyone can have a blog, each serving various purposes. Consequently, to find good quality blogs, people have to choose from a large number of blogs in which the majority is not relevant to their information needs. To ease this task, social cues offer heuristic guidelines for selecting good blogs: popular blogs with large audiences are most likely more interesting and useful than those with no audience. Thus, based on such social cues, people can quickly assess whether they should explore a blog in more detail or just try to find another one to fulfill their information needs.

The usefulness of visualizing social interaction history, however, depends on the kinds of blogs and what users look for in the blogs. Visualization of social interaction history is particularly useful when users want to *explore* an information space: that is, they do not have specific information to retrieve, but want to learn about an information space—its content and social dynamics within it—while hoping to find useful or interesting information within the space. In such cases, social interaction history can provide navigational cues for the users so that they can follow the crowds to find entries that have attracted a lot of attention in the community:

> Usually the topics [entries] receiving many comments are 'hot.' Thus most likely they will be interesting for me too.

In everyday life, people use social navigation to guide their decision when they are faced with uncertainties (e.g., choosing a crowded restaurant for dinner rather than going to a quiet one). Visualization of social interaction history uses the same principle to reduce the cognitive load of users. When people are faced with many blog entries, they have difficulties to identify the good ones without examining the entries individually. Social cues reduce this cognitive load by enabling users to follow the crowds. This approach is based on the assumption that popular entries are better or more useful than non-popular entries. Of course this assumption is not always true. However, having social cues is better than nothing. Moreover, such social recommendation requires no explicit effort from other people. The recommendation is simply resulted from visualizing the interaction traces preserved in a blog archive.

Visitors to blogs containing information that is easily outdated may receive less benefit from the visualization of social interaction history compared to those visiting topic-oriented blogs. For example, consider a diary blog used for sharing news with friends. Information contained in these kinds of blogs may become out-of-date or irrelevant quickly. Knowing that a friend was visiting our city last week is no longer useful, as we could not meet up with him or her. Furthermore, there is little need to revisit old entries in such blogs. The readers may just want to follow the most recent entry in the blog. From their perspective, comments from other people on old entries are not important.

The content of topic-oriented blogs (e.g., programming tips, education, and research) does not easily become outdated. Old entries may still contain relevant information that is worth reading. Comments from the audience can enrich the discussion, as they may offer different perspectives or add new content to the entry. In these kinds of blogs, visualization of social interaction history can provide guidance for visitors to select which entries to read and to explore the blogs effectively.

4.6.3 Design Implications

Responses from the study participants revealed that the decision to select which entries to read was affected by factors such as the posting time, the topic, the length of entries and comments, and the number of comments on entries. To facilitate exploration of blog archives, blogs should provide navigational support that allows users to search for particular entries using these criteria. Relying only on time- and topic-oriented navigation is not enough.

The favourite feature of iBlogVis was the visualization of the length of entries and comments. Some participants mentioned that these aspects influenced their decision whether to read an entry—an initial finding that was similar to the results of a Usenet study (Fiore et al., 2002). Some users preferred to read a short, popular entry to a long one, while others might have different preferences.

A challenge for designers is how to present various attributes of blog entries effectively to users. Users should be able to see, compare, and analyze entries from different aspects simultaneously. Simply providing additional sorting functions is insufficient, as there are multiple factors involved, and it is hard to maintain all contextual information that is important to the users while the entries are rearranged based on different criteria. Moreover, users do not always want the most popular entry. What they want may be recently published entries that are not too long and receive many short comments. Formulating such queries is complicated because the criteria are vague: recently published, not too long, many short comments. Requiring users to come up with exact criteria, however, will increase their cognitive load, and hence is not a desirable solution.

Information visualization is a viable solution to this problem. Designed properly, a visualization tool can present multiple attributes of blog entries simultaneously while allowing users to compare, analyze, and select entries that match their search criteria intuitively without having to formulate complex query statements. Contextual information can be maintained by providing an overview of a blog and enabling users to interact with the information space through a dynamic query interface (Ahlberg et al., 1992). The power of information visualization relies on the fact that human vision is excellent at comparing, extracting, and recognizing patterns (Ware, 2000).

Despite its potential, information visualization is not a panacea for all navigational problems in blogs. For lookup tasks such as fact retrieval, using search engines or keyword-based retrieval is more appropriate than using visualization tools because queries can be formulated easily in these tasks, and there is no or little need to compare the query results (Marchionini, 2006). Visualization, therefore, should be seen as a *supplement* to the current navigational support for blogs, especially to ease exploratory tasks.

From my observations, several participants tried to click on items displayed in the visualization panel, such as months, when the given tasks were relevant to the items. For example, when participants were asked to find entries posted in a specific month, some of them tried to highlight the entries by clicking on the corresponding month in the timeline. Repeated attempts to click on items in a visualization panel were also observed in a Usenet study (Viégas & Smith, 2004). This observation implies that users expect that each item in a visualization panel, whenever relevant to their task, can be clicked to help them perform the task at hand. This user expectation may develop into a design principle of information visualization.

4.6.4 Critical Reflection

Scalability in general is a challenging issue for visualization, and applies to iBlogVis as well. In practice, however, our experiences suggest that users want to deal with a manageable data set at a time. Using iBlogVis, users can narrow down a larger data set by filtering entries by tag, posting time, or commenter, and then zoom in on the area of interest to reduce visual occlusion. Another approach (not available in iBlogVis used in this study) is to provide another level of overview such as a miniature view—a global overview on a small scale (Gutwin, Roseman, & Greenberg, 1996). When the main visualization panel cannot show all entries, a miniature view would allow users to maintain contextual information of their position in a blog archive and then select a smaller set for exploration.

As elaborated by Ellis and Dix (2006), empirical evaluation of visualization tools poses several problems, such as the absence of standard data sets and user tasks. Furthermore, people usually use a visualization tool to perform exploratory tasks, making it even more difficult to come up with a set of standardized tasks and reliable, objective performance measures. While this study also inherited these limitations, I have tried to incorporate Ellis and Dix's suggestions into the study methods as follows.

This study used both quantitative and qualitative methods. The questionnaires gave quantitative results of the study: that is, a set of numbers indicating how satisfied or dissatisfied the participants were with the visualization tool. The observation and open-ended questions, however, produced insightful information beyond these numbers. This qualitative data gave some explanation for *why* users rated certain features of the visualization tool as useful or not useful. This explanation shed some knowledge of user practices in exploring blogs and contributed to understanding in which context revealing social interaction history contained in blog archives is perceived to be useful.

This study used a single data set. Although taken from a real blog, the data set might not be representative. There is no guarantee that the participants would give similar ratings if the visualization tool was used to visualize different data sets, particularly those having different characteristics (e.g., photo blogs). Therefore, the utility of iBlogVis is currently limited to a certain class of blogs: that is, topic-oriented blogs that contain mostly textual entries, have medium posting frequencies (a few entries per week), and receive comments from the audience regularly. Visualization of social interaction history is one of the main features of iBlogVis. Without using a data set that contains social interaction history, iBlogVis would not be able to deliver its full functionality, which consequently could affect the ratings of its utility.

Finally, most participants considered themselves to be advanced computer users, and most of them had a background in computer science. Their ability to learn and use iBlogVis might not represent the average user's ability. On the positive side, their expertise was valuable in providing constructive feedback about the system.

4.7 Summary

This chapter discussed the design, implementation, and evaluation of iBlogVis—an interactive visualization tool for facilitating exploration of blogs. The tool was evaluated using various methods. First, the design rationale was explained and justified based on existing research on human-computer interaction and information visualization. Second, the usability of the tool was evaluated using both subjective and objective performance measures. The results of these measures showed that the user satisfaction was high, and the average error rate of the given tasks was low. Third, the study explored the reasons behind the user satisfaction ratings qualitatively, using observation and comments from the participants. These qualitative responses have added to the understanding of blog reading behaviour and how to apply visualization techniques to ease exploratory tasks in blogs.

Comments from the participants indicated that the decision to select which entries to read was affected by multiple factors. Besides the topic and the posting time of an entry, the length and the number of comments on the entry also influenced the decision. The important role of these factors was reflected in the participants' responses about their favourite feature of iBlogVis: the visualization of the length of entries and comments. Thus, to facilitate exploratory tasks, blogs should provide additional support beyond the current time- and topic-oriented navigation.

Chapter 5

Summative Evaluation

The main objective of this thesis is to examine the role of social interaction history in information exploration. To achieve this objective, I explored the feasibility of social interaction history to support exploratory tasks by developing and evaluating a prototype for browsing a blog archive. As discussed in Chapter 4, findings from the study were promising. To follow up these initial findings, I carried out a summative evaluation to assess how well visualization of social interaction history can improve exploratory browsing. This chapter describes the objective, method, and results of the summative evaluation.

5.1 Objective and Hypotheses

The objective of this study was to compare the outcomes of exploratory browsing and subjective user satisfaction when people use three interfaces. In exploratory browsing, people explore an information collection with intention to find interesting articles to read. The interfaces used in the experiment differ in the ways they present information items and in the absence or presence of visualization of social interaction history. More specifically, the experiment was intended to test these research hypotheses:

1. Visualization of social interaction history helps users find interesting articles.

2. Visualization of social interaction history reduces wasted effort to find interesting articles.

3. Visualization of social interaction history increases user satisfaction with the visualization tool.

4. Overall, users prefer the visualization tool to the other systems.

5.2 Method

5.2.1 Experimental Design

This study used a within-subjects design with one independent variable and five dependent variables. System/interface served as the independent variable and had three levels: List, TimeVis, and SocialVis. The dependent variables were interest score, perception of support, wasted effort, user satisfaction, and overall preference. The next sections discuss these independent and dependent variables in detail.

Since this experiment was designed to measure the effects of social interaction history visualization on exploratory browsing, other factors affecting the results of exploratory tasks must be controlled. These factors include the effects of systems, searchers, and topics (Dumais & Belkin, 2005; Lagergren & Over, 1998). For example, a system may perform well only with a certain data set. Due to their interests in a topic, searchers who have good knowledge of the topic may be able to explore an information space better than other users. Furthermore, in information retrieval (IR) experiments, participants are usually asked to save their search results, which then are judged by independent assessors. Differences between searchers' and assessors' relevance judgments also affect the performance evaluations of IR systems.

I took several approaches to controlling searcher, topic, and system effects. First, this study used a within-subjects design to control searcher effects. A within-subjects design allows performance comparisons of individual participants when they use different systems, thereby isolating differences in background knowledge and search experiences among the participants. Second, instead of independent assessors, participants judged the outcomes of their browsing sessions themselves (Borlund, 2000). This approach fits the nature of exploratory tasks: the criteria of success are per-

Figure 5.1: The List user interface displays document surrogates in reverse chronological order. A document surrogate consists of a title, a short description, a website address, a digg count, and a comment count.

sonal and subjective, and searchers are free to follow their interests because they do not have specific goals. Third, participants browsed three different data sets (topics) to control learning effects that are associated with within-subjects designs. By exploring different data sets, participants would not be able to take advantage of their familiarity with a data set while using a different system. Finally, topic assignments and system trials were counterbalanced using a Latin square to reduce bias towards a particular system (Lagergren & Over, 1998).

5.2.2 Experimental Conditions

To study the effects of social interaction history on exploratory browsing, I developed three systems—List, TimeVis, and SocialVis—for browsing collections of Digg stories.[1]

[1]http://digg.com

Figure 5.2: Summary of a story on Digg (retrieved on October 26, 2009 from http://digg.com/food_drink/Teens_in_Fatal_Crash_Have_Fast_Food_Delivered_to_The_Scene).

List

The List user interface is similar to that of a typical search engine. It displays a list of document surrogates in reverse chronological order (see Figure 5.1). Each page contains 10 stories, and users can move from page to page by clicking on a page number, "Previous," or "Next" links that are available on the bottom of each page.

When users click on a story's title, List opens a new tab and takes the users to the Digg website where a summary of the story is stored (see Figure 5.2). On this page, users can read a short description about the story, the person who submitted the story, and comments on the story if any (comments are not shown in the figure). Users can view the content of the story by clicking on its title on this summary page. Because List opens a new tab for each story, users can open multiple stories simultaneously. Similar to the standard web interface, List changes the colour of visited links to help users keep track of stories that they have opened.

TimeVis

TimeVis uses the Timeline visualization tool[2] to display Digg stories along a timeline (see Figure 5.3). The visualization panel consists of three parts. The main part displays the titles of stories in a data set. These stories are arranged based on their submitted time. When users click on a title or an icon, a pop-up window appears and displays information about the story: its title, its numbers of diggs, comments, and

[2]http://www.simile-widgets.org/timeline/

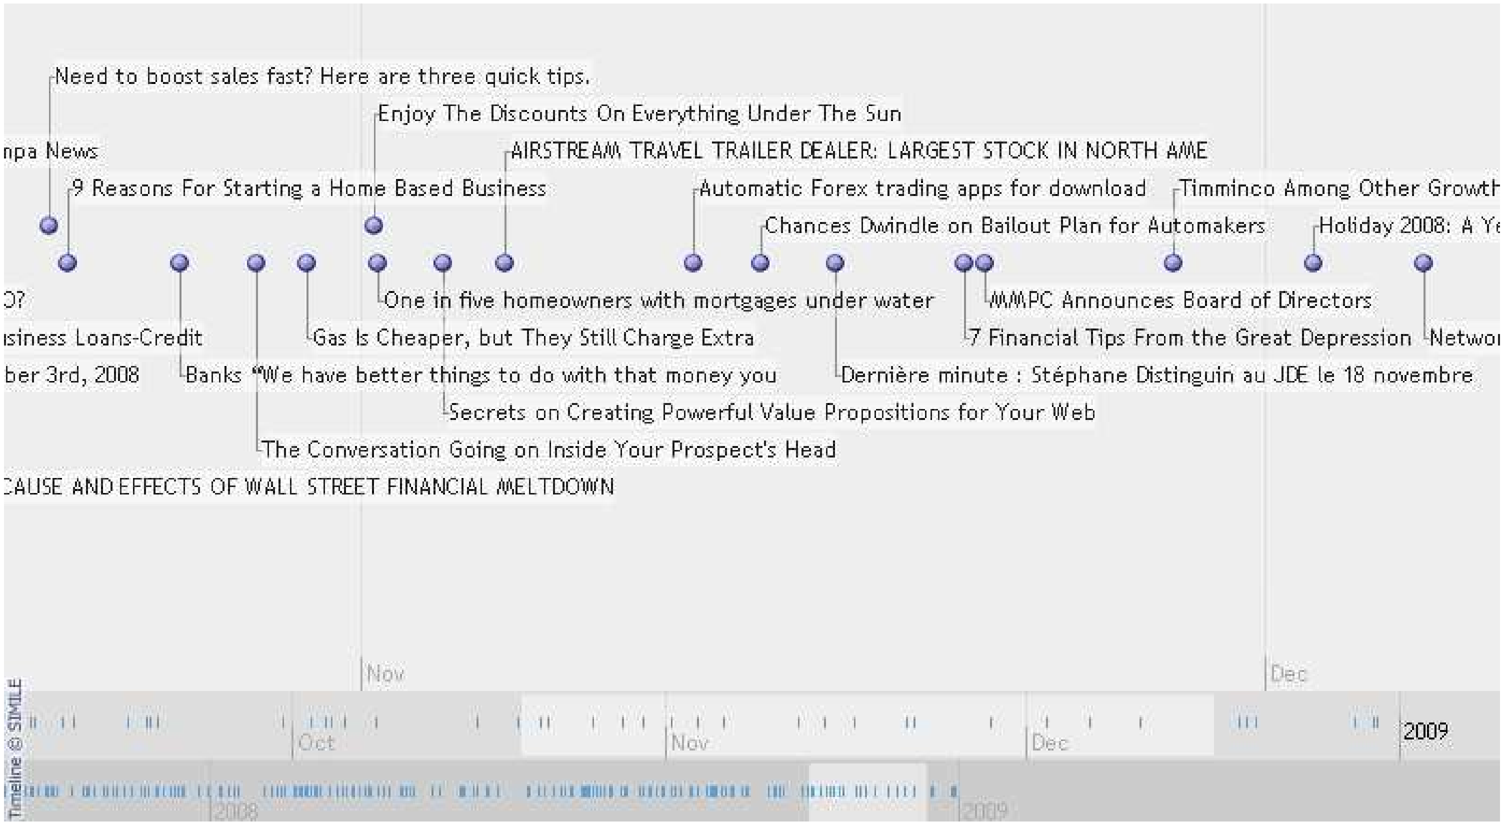

Figure 5.3: The TimeVis user interface displays titles of Digg stories along a timeline. The light colour background in the overview panels indicates the currently visible time range, while ticks in the foreground represent stories. For example, the bottom panel in this screenshot indicates that there are no stories beyond the year 2009 in the data set.

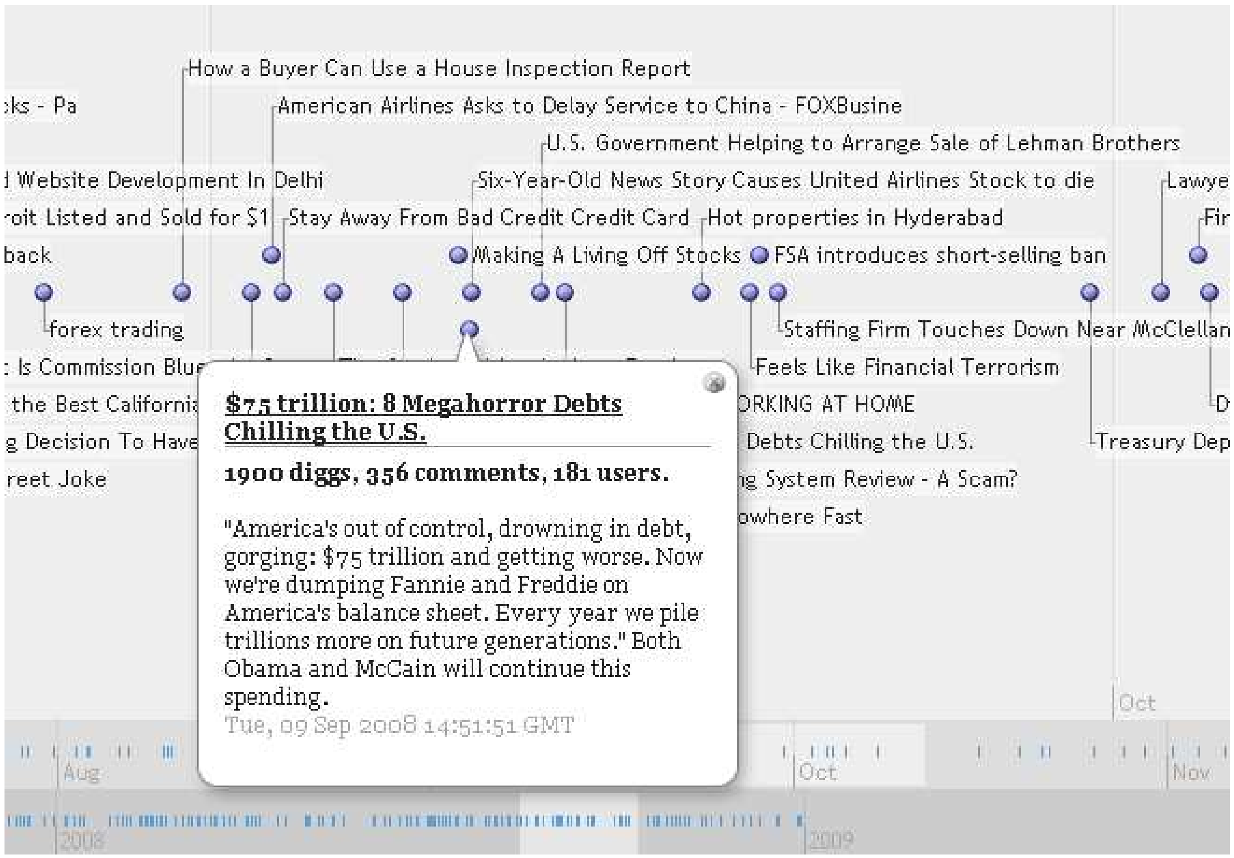

Figure 5.4: When users click on a title or an icon, a pop-up window appears and displays information about the story. Users can access the story on Digg by clicking on its title in this pop-up window.

commenters, its short description, and its submitted time (see Figure 5.4). Users can access the story on Digg (Figure 5.2) by clicking on the title in this pop-up window. Users may view multiple stories by opening a new tab for each story. Similar to List, TimeVis changes the background colour of visited stories (in the beginning, all stories are highlighted in white to indicate that they have not been opened).

The lower parts of the visualization panel provide overviews of a data set (see Figure 5.3). The light colour background indicates the visible time range, while ticks in the foreground represent stories. For example, Figure 5.3 shows that there are no stories beyond the year 2009 in the data set.

TimeVis offers some interactive features. Panning is supported by a click-and-drag interaction style. Zooming in/out is achieved by scrolling a mouse wheel forward/backward. When users zoom in on a particular area, the visible time range decreases, and a more detail timeline is shown (see Figure 5.5).

SocialVis

SocialVis and TimeVis are basically identical. All features of TimeVis work in the same way in SocialVis. The only differences between these systems are that SocialVis visualizes the numbers of diggs (blue bars) and commenters (pink bars) on each story and allows users to filter stories by these numbers (see Figure 5.6). SocialVis visualizes this social interaction history because a digg count indicates the degree of interestingness of a story from a social perspective (Lerman & Galstyan, 2008), and so does the number of commenters on a story (Indratmo & Vassileva, 2009).

Filtering in SocialVis is implemented using jQuery UI range sliders[3] that are placed below the visualization panel (see Figure 5.7). Users can adjust these range sliders to filter out unwanted stories. Based on the values of these sliders, SocialVis updates the visible stories in the visualization panel.

List and TimeVis serve as baseline interfaces in this experiment for the following reasons. First, the List interface represents typical information presentation on

[3]http://jqueryui.com/home

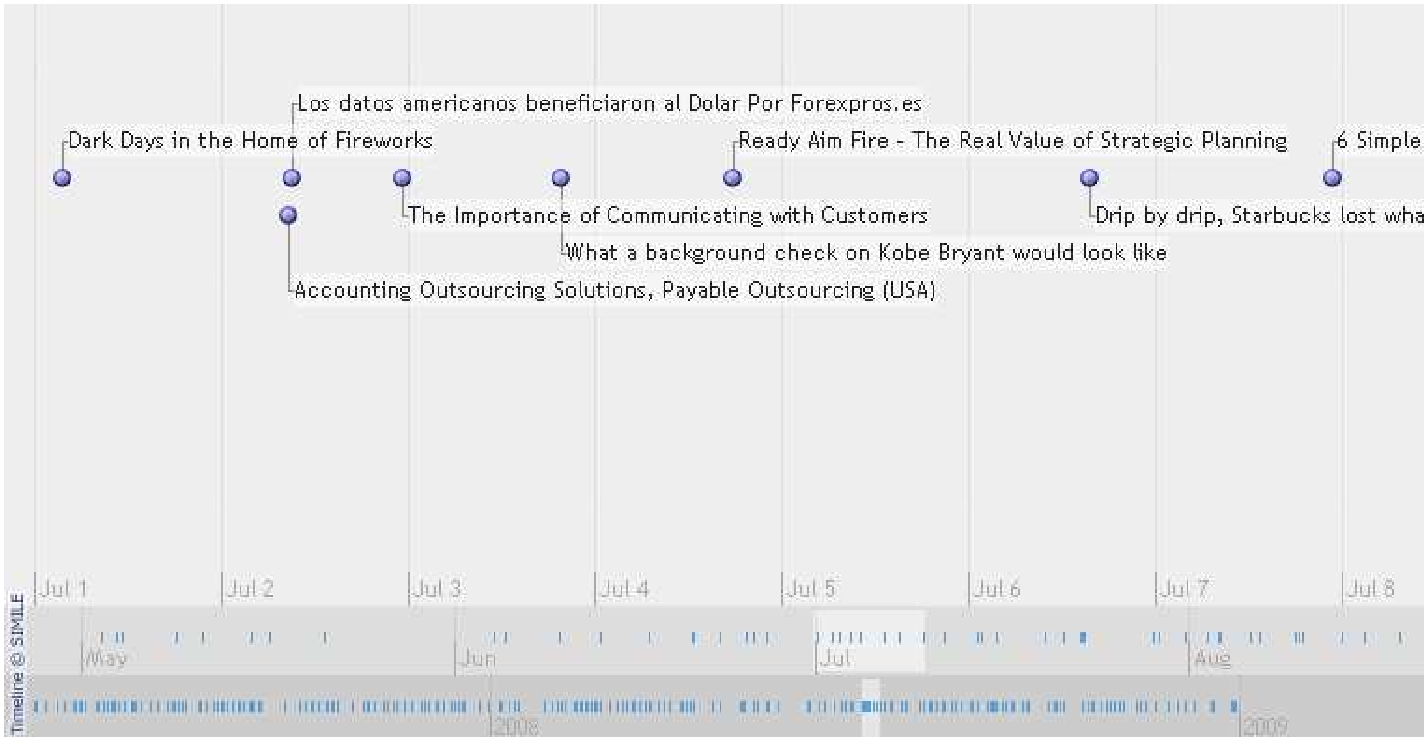

Figure 5.5: Zooming in/out is supported by scrolling a mouse wheel forward/backward. Stories spread across a more detail timeline in a zoomed-in view.

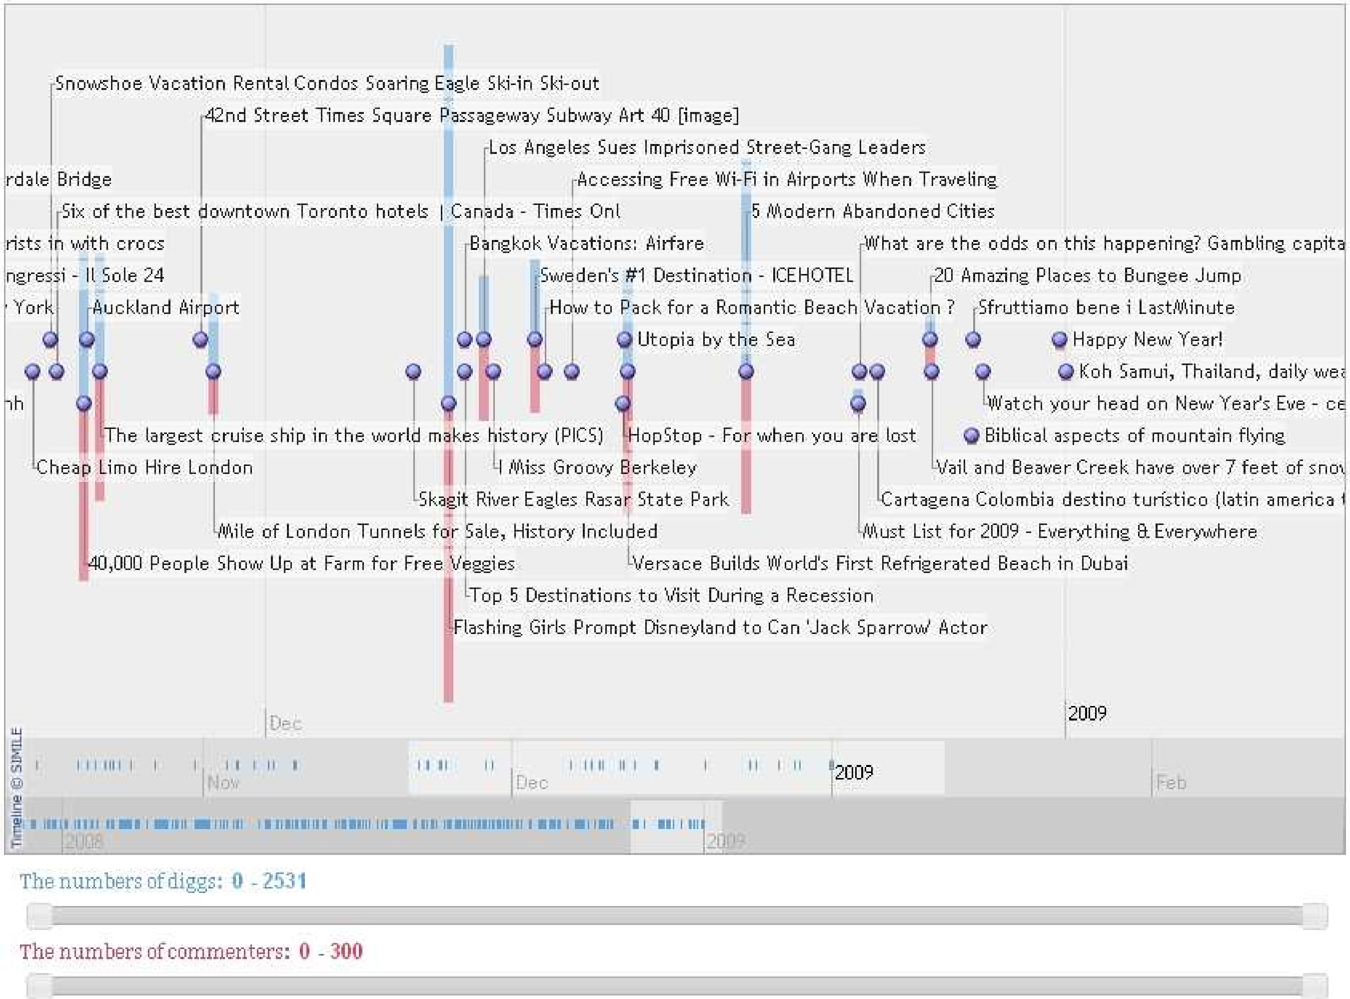

Figure 5.6: The SocialVis user interface is identical to TimeVis, except that SocialVis visualizes the numbers of diggs (blue bars) and commenters (pink bars) on each story and allows users to filter stories by these numbers.

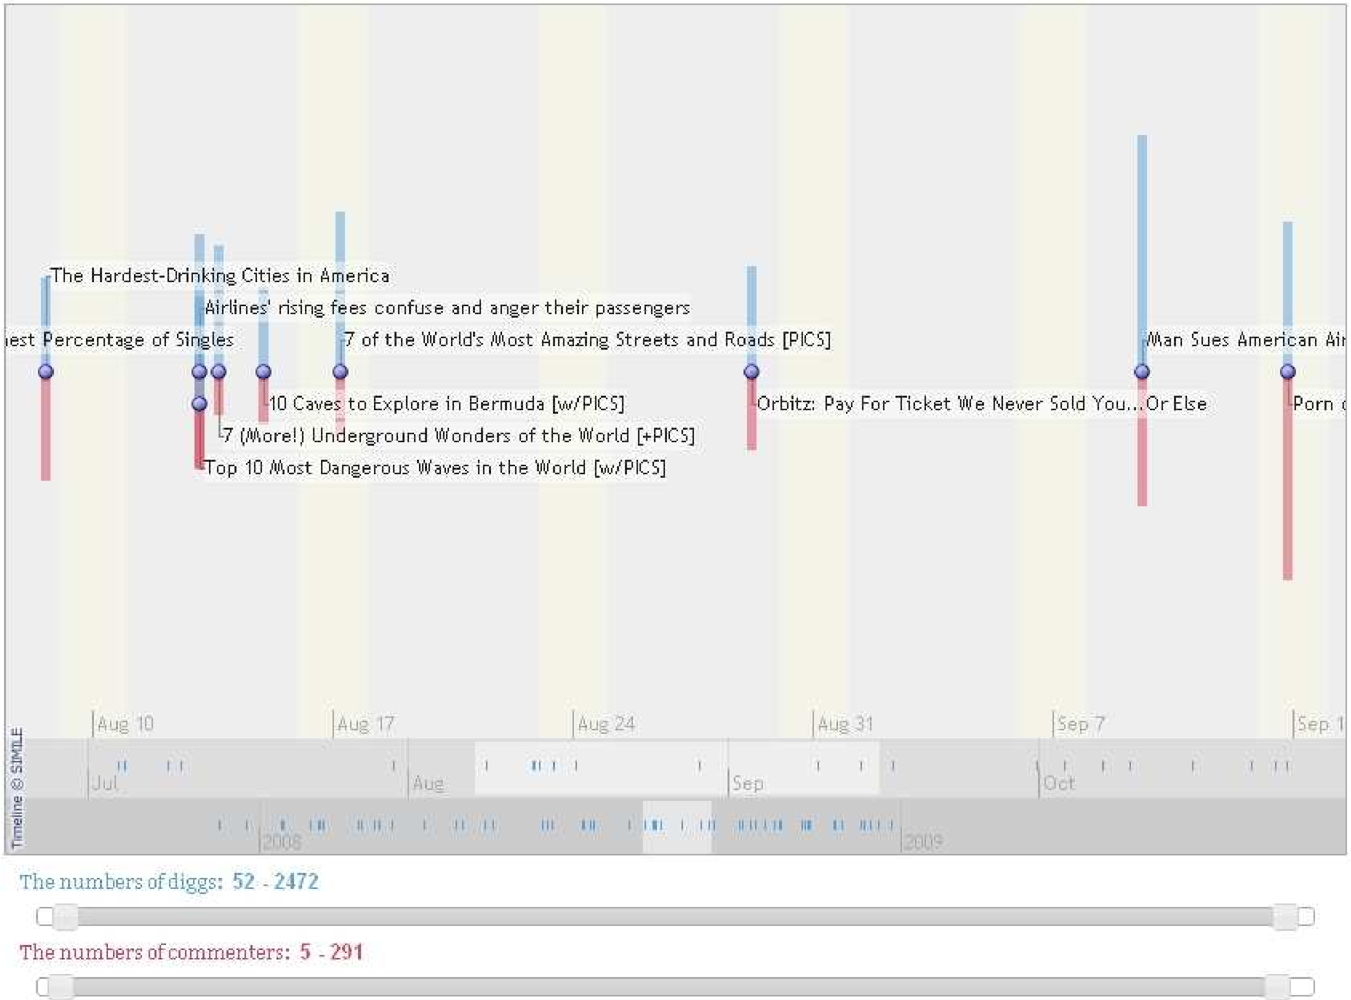

Figure 5.7: To filter stories by the numbers of diggs and commenters, users can adjust the range sliders below the visualization panel. In this screenshot, the user displays stories that have 52 – 2472 diggs and that receive comments from 5 – 291 users.

the Web. Many applications—such as search engines, discussion forums, and email clients—use this kind of standard interface to arrange information items. Second, TimeVis and SocialVis are identical except for the presence of the social interaction visualization. By using these baseline interfaces, we would be able to know whether differences in the experimental results are due to the different information presentation (List vs. TimeVis) or due to the absence or presence of the visualization (TimeVis vs. SocialVis).

Since the purpose of this study was to evaluate the systems' performance in supporting exploratory browsing—not the indexing or retrieval algorithms of the systems—all systems do not enable users to search stories by keyword (He et al., 2008). The systems also do not visualize the contents of stories (e.g., the length of a story), as the study focused on social aspects of information spaces.

5.2.3 Performance Measures

There were five dependent variables in this experiment: interest score, perception of support, wasted effort, user satisfaction, and overall preference.

1. *Interest score* is a measure of the outcome of exploratory browsing and indicates how well a system helps users find interesting articles. It is collected by asking participants to rate each article that they open while browsing (1: not interesting, 9: interesting). It is expected that users will find articles that are more interesting when using a better system.

2. *Perception of support* is a subjective measure of how well a system supports the ability of users to find interesting articles. This satisfaction measure is independent from the actual outcome of browsing (i.e., interest score). Ideally, a system should provide a good sense of support and improve the outcome of browsing. Perception of support is measured using this questionnaire item (White et al., 2007)[4]:

[4]White et al. (2007) use a 5-point instead of a 9-point Likert-scale item and focus on finding relevant instead of interesting information.

Using this system enhances my effectiveness in finding interesting information.
Strongly disagree 1 2 3 4 5 6 7 8 9 Strongly agree

3. *Wasted effort* is computed based on interest scores: it is defined as *the number of articles rated three or below divided by the total number of articles opened during a browsing session*. The notion of wasted effort is adapted from He et al. (2008). Opening and reading an article takes more effort than just browsing document surrogates. If a viewed article turns out to be interesting, then the additional effort to open the article is worthwhile; otherwise, users have wasted some effort. This measure assumes that articles with interest scores between one and three to be uninteresting. It is expected that a good system should be able to minimize wasted effort.

4. *User satisfaction* is a subjective measure of how users like a system. It is expected that a better system will have higher user satisfaction scores. User satisfaction is computed based on the mean of the following semantic differentials (Kules & Shneiderman, 2008):

Terrible	1	2	3	4	5	6	7	8	9	Wonderful
Difficult to use	1	2	3	4	5	6	7	8	9	Easy to use
Dull	1	2	3	4	5	6	7	8	9	Stimulating
Frustrating	1	2	3	4	5	6	7	8	9	Satisfying
Overwhelming	1	2	3	4	5	6	7	8	9	Manageable
Complex	1	2	3	4	5	6	7	8	9	Simple
Too slow	1	2	3	4	5	6	7	8	9	Fast enough

5. *Overall preference* is a subjective measure of the overall functionality of a system. This measure assumes that users will choose a system that provides best support for exploratory browsing. It is collected after users have used all systems so that they can compare one system with the others.

5.2.4 Participants

Thirty people (16 males, 14 females) participated in the study. Participants were recruited primarily through email invitation. Subject ages varied in the following ranges (in years): $20-29$ (14 people), $30-39$ (15 people), and 40 or above (1 person). Each participant but one received a \$10 honorarium.[5] As a prerequisite, participants had to be familiar with browsing the Internet. On a scale of one (beginner) to nine (expert), participants self rated their computer skills as an end-user at level five or above ($M = 7.63$, $SD = 1.25$). On a daily average, most participants browsed the Internet between one and five hours (21 people), while the rest spent between six and ten hours (5 people) or more than ten hours (4 people). Only half of the participants were familiar with Digg.

5.2.5 Setting and Apparatus

The study was conducted in a computer laboratory using the same equipment for all participants. The computer had dual processors (Xeon 3.20 GHz) and 1 GB of RAM, running on Windows XP. The display was a 23-inch flat panel monitor with screen resolution set at 1920 by 1200 pixels.

Participants used the Firefox web browser to explore three different data sets. These data sets consisted of a total of 900 random stories that were submitted to Digg between January 1, 2007 and December 31, 2008. Each data set consisted of 300 random news stories about a certain topic—Business & Finance, Travel & Places, or Food & Drink. Out of these stories, 20% were promoted to popular stories on Digg.

5.2.6 Procedure

Table 5.1 illustrates a 3×3 Latin square that was used to construct an experimental matrix for the study. This matrix was replicated as necessary to accommodate 30 participants. During the experiment, the systems were called Arjuna (List), Tambora

[5]One participant declined the payment.

Table 5.1: A 3-system-by-3-topic experimental matrix to run the study.

Participants	Data Sets/Topics		
	Business & Finance	Travel & Places	Food & Drink
1	List	TimeVis	SocialVis
2	SocialVis	List	TimeVis
3	TimeVis	SocialVis	List

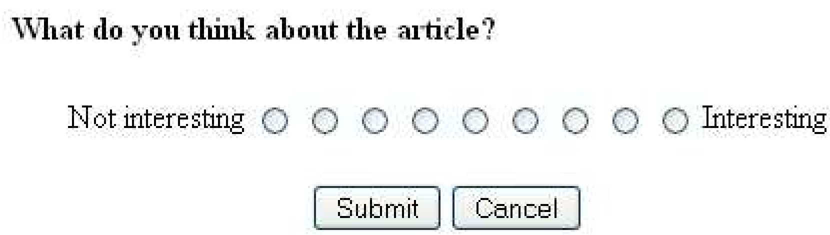

Figure 5.8: A pop-up window for collecting an interest score of an article.

(TimeVis), and Rinjani (SocialVis).

Upon arrival, participants were introduced to the objective of the study and the consent process. After signing an informed consent form, participants filled out a pre-test questionnaire that collects background information about them. If participants were not familiar with Digg, they received a brief introduction about it. Following the experimental matrix in Table 5.1, participants were introduced to the features of a system and were given a short time to practice using the system.

When participants felt ready, they received this scenario:

> You have a 15-minute free time period and want to spend the time for browsing articles about Business & Finance/Travel & Places/Food & Drink. You are not looking for specific information but just want to find interesting articles to read. Please browse the given collection using the system, and try to browse the collection in a natural way. You may skim through or read an article in detail. If you open an article, please submit your rating about the article.

Participants submitted their ratings about articles through a pop-up window

Table 5.2: Means and standard deviations of performance measures for all interfaces ($N = 30$).

Measures	List		TimeVis		SocialVis	
	M	SD	M	SD	M	SD
Interest score	5.58	1.00	5.86	1.13	6.48	1.09
Perception of support	5.00	2.49	6.07	1.91	7.47	1.20
Wasted effort	0.24	0.17	0.19	0.17	0.11	0.14
User satisfaction	6.45	1.43	6.38	1.42	7.55	0.77

(Figure 5.8). The system saved and converted a user rating into a 9-point scale (1: not interesting, 9: interesting). I observed all browsing sessions to make sure that participants submitted an interest score for every article that they opened.

After 15 minutes browsing a data set, participants answered questions about perception of support and user satisfaction. After that, they were introduced to the next system and repeated the same task using a different data set.

At the end of a data collection session, participants indicated their overall preferences along with the reasons. They also had an opportunity to ask me any questions about the study.

5.3 Experimental Results

A data collection session took about one hour. All participants were able to understand the features of all systems easily and did not spend much time to practice using the systems.

5.3.1 Quantitative Results

All statistical tests used an alpha level of 0.05. Mean differences were tested using one-way repeated-measures analysis of variance (RM ANOVA). When the assumption of sphericity was violated, degrees of freedom were corrected using the appropriate estimates of sphericity. Post hoc analyses used Tukey's honestly significant difference (HSD). However, when the assumption of sphericity was not met, the post

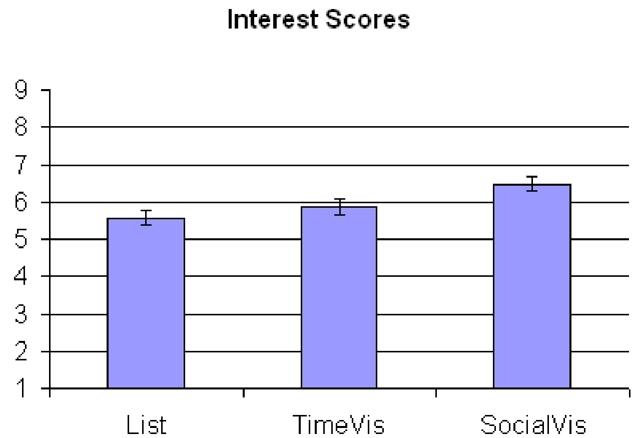

Figure 5.9: Effect of interface on interest scores (mean ± standard error). Higher scores indicate that participants were able to find articles that were more interesting.

Table 5.3: Post hoc pairwise comparisons of means of interest scores (Tukey's HSD = 0.54 at the 0.05 alpha level).

$M1$ vs. $M2$	Mean Difference ($M1 - M2$)	p
SocialVis vs. List	0.90	< 0.01
SocialVis vs. TimeVis	0.62	< 0.05
List vs. TimeVis	-0.28	> 0.05

hoc analyses used the Bonferroni method instead of Tukey's HSD (Maxwell, 1980). Table 5.2 provides a summary of the quantitative results. The effect of interface on each measure is reported in the following sections.

Effect of Interface on Interest Scores

Visualization of social interaction history was expected to help users find interesting articles. This hypothesis was tested by looking at the outcomes of exploratory browsing (interest score) and the participants' subjective assessment of system support (perception of support). The experimental results showed that social interaction visualization increased both interest scores and perception of support.

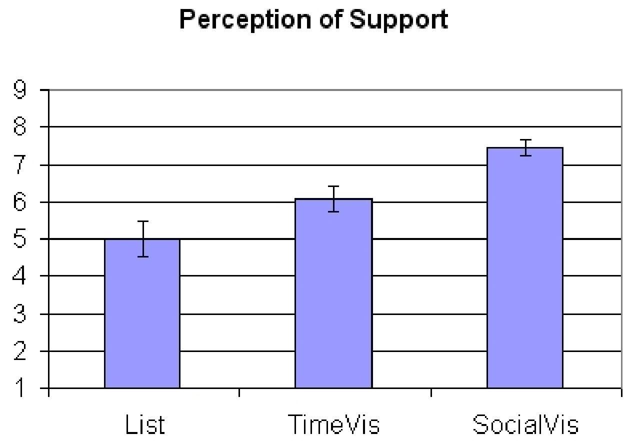

Figure 5.10: Effect of interface on perception of support (mean ± standard error). Higher scores indicate that participants felt a system enhances their effectiveness better in finding interesting articles.

As illustrated in Figure 5.9, the mean of interest scores for SocialVis was the highest among the experimental systems. On the 9-point scale, the mean of interest scores for SocialVis was 0.90 point higher than the List system (16% improvement), and 0.62 point higher than the TimeVis system (11% improvement). RM ANOVA suggested that there was a significant effect of interface on interest scores, $F(2, 58) = 8.52$, $p = 0.001$. Tukey's HSD revealed that participants were able to find articles that were more interesting when they used SocialVis compared to when they used List ($p < 0.01$) or TimeVis ($p < 0.05$). No other significant differences were found (see Table 5.3).

Effect of Interface on Perception of Support

As seen in Figure 5.10, participants perceived SocialVis to best enhance their ability to find interesting information. The mean of perception of support for SocialVis was 2.47 points higher than the List system (49% improvement), and 1.40 points higher than the TimeVis system (23% improvement). Mauchly's test suggested that the data did not meet the assumption of sphericity, so degrees of freedom were corrected using Huynh-Feldt estimate of sphericity ($\epsilon = 0.80$). RM ANOVA indicated that

Table 5.4: Post hoc pairwise comparisons of means of perception of support (p values have been adjusted using the Bonferroni correction).

$M1$ vs. $M2$	Mean Difference ($M1 - M2$)	p
SocialVis vs. List	2.47	< 0.001
SocialVis vs. TimeVis	1.40	< 0.001
List vs. TimeVis	-1.07	$= 0.08$

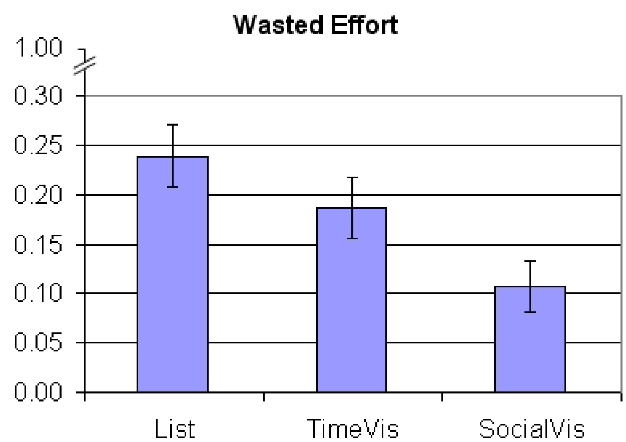

Figure 5.11: Effect of interface on wasted effort (mean ± standard error). Lower scores indicate that participants opened fewer uninteresting articles compared to the number of articles they opened during their browsing sessions.

there was a significant effect of interface on perception of support, $F(1.61, 46.64) = 17.42$, $p < 0.001$. The Bonferroni test revealed that participants perceived SocialVis to increase their effectiveness in finding interesting information compared to List ($p < 0.001$) or TimeVis ($p < 0.001$). No other significant differences were found (see Table 5.4).

Effect of Interface on Wasted Effort

Visualization of social interaction history was expected to reduce wasted effort. As illustrated in Figure 5.11, participants wasted the least effort when using SocialVis. The mean of wasted effort for SocialVis was 0.13 point lower than the List system

Table 5.5: Post hoc pairwise comparisons of means of wasted effort (Tukey's HSD = 0.08 at the 0.05 alpha level).

$M1$ vs. $M2$	Mean Difference ($M1 - M2$)	p
SocialVis vs. List	-0.13	< 0.01
SocialVis vs. TimeVis	-0.08	< 0.05
List vs. TimeVis	0.05	> 0.05

(54% reduction), and 0.08 point lower than the TimeVis system (42% reduction). RM ANOVA indicated that there was a significant effect of interface on wasted effort, $F(2, 58) = 7.40$, $p = 0.001$. Tukey's HSD revealed that SocialVis reduced wasted effort significantly compared to List ($p < 0.01$) or TimeVis ($p < 0.05$). No other significant differences were found (see Table 5.5).

Effect of Interface on User Satisfaction

Visualization of social interaction history was expected to increase user satisfaction with the visualization tool (SocialVis). As illustrated in Figure 5.12, participants were most satisfied with SocialVis. The mean of user satisfaction scores for SocialVis was 1.10 points higher than the List system (17% improvement), and 1.17 points higher than the TimeVis system (18% improvement). Due to violation of the sphericity assumption, degrees of freedom were corrected using Huynh-Feldt estimate of sphericity ($\epsilon = 0.85$). RM ANOVA indicated that there was a significant effect of interface on user satisfaction, $F(1.71, 49.48) = 16.33$, $p < 0.001$. The Bonferroni test revealed that user satisfaction with SocialVis was significantly higher compared to List ($p < 0.001$) or TimeVis ($p < 0.001$). No other significant differences were found (see Table 5.6).

Overall Preference

Participants were expected to prefer SocialVis to the other systems. To test this hypothesis, participants indicated their overall system preferences at the end of a data

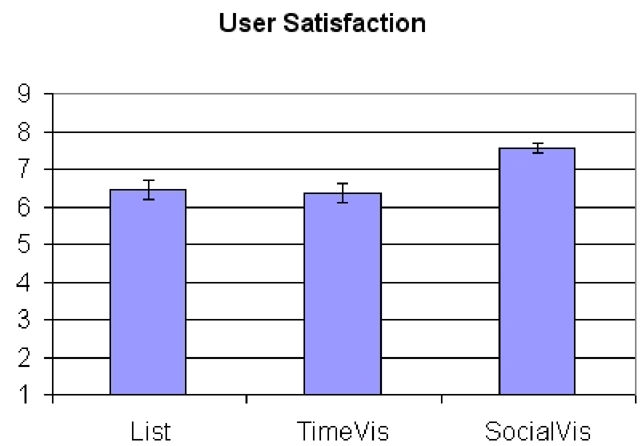

Figure 5.12: Effect of interface on user satisfaction (mean ± standard error). Higher scores indicate that participants were more satisfied with a system.

Table 5.6: Post hoc pairwise comparisons of means of user satisfaction scores (p values have been adjusted using the Bonferroni correction).

$M1$ vs. $M2$	Mean Difference ($M1 - M2$)	p
SocialVis vs. List	1.10	< 0.001
SocialVis vs. TimeVis	1.17	< 0.001
List vs. TimeVis	0.07	$= 1.00$

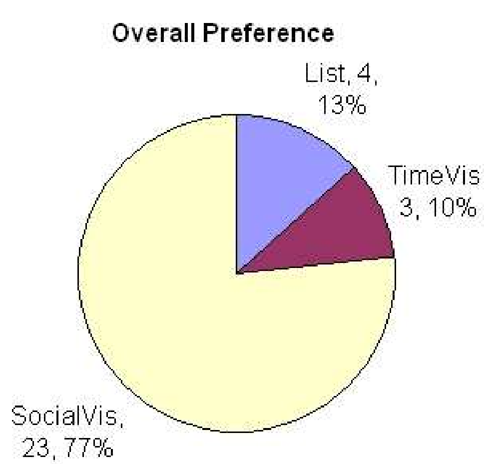

Figure 5.13: Overall preference of participants.

collection session. Out of 30 participants, four preferred List, three chose TimeVis, and the rest (23) preferred SocialVis (see Figure 5.13). A chi-square test indicated that the number of participants preferring each system differed significantly from the expected number, $\chi^2(2,\ N = 30) = 25.40$, $p < 0.001$. This result implied that participants significantly preferred SocialVis to the other systems.

5.3.2 Typical Use of Each Interface

During data collection sessions, I observed how participants used each system. This section reports on this qualitative observation to provide background information for a better understanding of the analysis of results in Section 5.4.

List

The typical use of the List interface was to browse a collection of stories sequentially. By default, List displayed the most recent stories on the first page, so participants explored a collection in reverse chronological order. Given a page, participants scanned the titles of stories from top to bottom. When they found an interesting title, they clicked on it to access the story on the Digg website. After reading or skimming over the story, they continued their exploration by moving to the next page and rarely went back to visited pages. During a 15-minute browsing period, participants usually only visited the first several pages of the collection.

TimeVis

TimeVis arranged stories along a timeline. When participants opened TimeVis, they viewed the most recent stories in a collection. Since TimeVis displayed stories in a compact way, participants might scan the titles of the stories randomly or systematically (e.g., left to right, top to bottom). When they found an interesting title, they clicked on it to open the short description of the story and immediately followed the link to go to the Digg website. After reading the content of the story, they continued browsing the collection by moving along the timeline (i.e., panning). Since panning was continuous, and a view in TimeVis had a high information density, participants were able to explore the whole collection and often revisited the same views multiple times (by moving left and right repeatedly). Some participants adjusted the zoom level to reduce the number of stories in a view (by having a zoomed-in view).

SocialVis

SocialVis and TimeVis were almost identical. Not surprisingly, participants used SocialVis in a similar way as they did with TimeVis. The main distinction was that after participants browsed a collection for a short while, they adjusted the range sliders in SocialVis to filter out stories having a few diggs and commenters. A few participants used these filters to control their views systematically, e.g., starting from the most to less popular stories. However, the majority of the participants simply used the filters to remove less popular stories. Overall, filtering played a more important role than zooming, and filtering by diggs was used more often than filtering by commenters.

In all systems, many participants seemed to pay attention only to the titles of stories. They often ignored the short descriptions of stories. The decision whether to open a story relied heavily on whether they were interested in its title. SocialVis, however, allowed participants to filter out less popular stories so that they were able to explore popular stories easily.

5.4 Discussion

The experimental results supported all research hypotheses. Visualization of social interaction history was able to help users find interesting articles, to reduce wasted effort, and to increase user satisfaction with the visualization tool. Overall, most participants preferred the visualization tool to the other systems. To help understand these results, this section offers a descriptive theory (Dix, 2008) of why visualization of social interaction history was able to improve exploratory browsing.

Table 5.7: Feature comparisons of the experimental systems. Participants favoured the highlighted features of SocialVis: continuous navigation, filtering, and visualization.

No.	Features	List	TimeVis	SocialVis
1.	Organization	Linear	Linear	Linear
2.	Document surrogate	Title + desc.	Title	Title
3.	Overview	Yes	Yes	Yes
4.	Details-on-demand	Yes	Yes	Yes
5.	Information density	Constant	Variable	Variable
6.	Zoom	No	Yes	Yes
7.	Navigation	Discrete	Continuous	Continuous
8.	Filter	No	No	Yes
9.	Social information	Textual	Textual	Visual

5.4.1 Feature Comparisons of the Systems

To examine which specific features of SocialVis had significant effects on the performance measures, I make a comparison of features of the experimental systems in Table 5.7:

1. *Organization*: All systems organize information items in a linear structure (chronological).

2. *Document surrogate*: All systems include the titles of stories in document surrogates. Besides displaying titles, List shows short descriptions, URLs, and the numbers of diggs and comments on stories. This information is also available in TimeVis and SocialVis when users click on a story.

3. *Overview*: All systems provide an overview of a collection that helps users explore the information space. In List, users can get an overview through the navigational links on the bottom of each page (e.g., page numbers indicate the current location and the available options for the next movement). In TimeVis and SocialVis, such an overview is available in the lower parts of the

visualization panel (e.g., the presence of ticks indicates that there are stories in that area).

4. *Details-on-demand*: Using all systems, people can access the contents of stories by following the provided links.

5. *Information density*: List has a constant information density (10 stories per page). In TimeVis and SocialVis, users can adjust the information density in a view using the zoom and filter functions (e.g., zooming in reduces the number of stories in a view).

6. *Zoom*: Zoom is available both in TimeVis and SocialVis, but not in List.

7. *Navigation*: Movement in List is discrete (page by page), whereas TimeVis and SocialVis allow continuous movement along a timeline.

8. *Filter*: Filtering by diggs and commenters are supported only by SocialVis.

9. *Social information*: List and TimeVis provide social information about a story (e.g., a digg count) in a textual form, that is, as part of a document surrogate. SocialVis augments this textual description with visualization.

Recall that this study used two baseline interfaces—List and TimeVis—to examine whether differences in the experimental results were caused by the differences in presenting information items (List vs. TimeVis) or by the presence of social interaction visualization (TimeVis vs. SocialVis). No significant differences between List and TimeVis suggest that improvements in browsing were not due to the feature differences between these two systems (rows 5 – 7 in Table 5.7). However, when the filters and visualization were introduced (rows 8 – 9 in Table 5.7), participants were able to explore the information space better. In their comments, participants repeatedly mentioned these features (and easy navigation) to be the reasons why they preferred SocialVis to the other systems. The following sections explain why these features enabled participants to do better in the experiment.

Filtering by Diggs and Commenters

In this study, the outcome of browsing was measured by the mean of interest scores. To get a high mean value, participants must find highly interesting stories and avoid those having low interest values. Filtering by diggs and commenters helped with both aspects as discussed below.

Participants typically filtered out stories with low digg counts after using SocialVis for a short while. This removal reduced the number of stories in their views and allowed them to focus on popular stories. While participants did not always find popular stories interesting, such stories were usually well written and contained some valuable information. Therefore, participants more likely considered popular stories interesting and gave them high interest scores. Even when participants were not interested in the topics of these stories, they rarely rated the stories very lowly. Participants reserved very low interest scores (i.e., one or two) mostly for spam and advertisements. The important role of filtering in enhancing participants' ability to find interesting stories was reflected in their comments and was one of the common reasons why most participants preferred SocialVis:

> It is easy to locate most popular and interesting stories.
>
> I could play around with the filtering criteria to get a set of articles that had been Digged enough.
>
> The ability to filter results based on the number of Diggs was useful for me. I was able to start at a high number of diggs and effectively slide the window of Digg counts from a high number downwards, to find articles others found interesting on a decreasing scale.

Besides facilitating participants to find interesting stories, filtering helped them avoid spam and advertisements. As an open social media application, Digg receives spam and advertisements from time to time. In most cases, these submissions only receive a single digg and no comments. The filtering functions in SocialVis allowed participants to remove spam and uninteresting stories from their views:

> Filtering allowed removal of articles with very small digg counts. As an article with a very small digg count quite often indicates it is an

advertisement, or uninteresting, this can generally be safely eliminated without losing any interesting articles.

The ability to filter the articles based on the number of diggs is helpful to get the approximate relevancy of the articles quickly and filter out some of the spam with only one digg as it has only been supported by the person submitting it.

Most of the articles that only have 1–10 "diggs" aren't very well-written or interesting, and the Rinjani [SocialVis] system allowed me to filter them out easily.

Unlike SocialVis, List and TimeVis lacked the functionality to remove spam and advertisements. When using List and TimeVis, participants sometimes ended up with advertisements and rated them very lowly. Such ratings lowered the means of interest scores for these systems. This effect could be big especially when participants opened only a few stories during their browsing sessions. In contrast, SocialVis did not suffer this problem, as participants were able to remove spam from their views by using the filters. Thus, the mean of interest scores for SocialVis was affected the least by very low ratings that were reserved by participants to rate spam and low quality stories.

The effect of filtering on browsing can also be explained using the information foraging theory (Pirolli & Card, 1999). This theory assumes that people always try to optimize their efforts to gain valuable information and that information exists in a patchy environment. If we consider a view in SocialVis to be an information patch, then filtering can reduce the cost of within-patch exploration. Filtering by diggs and commenters enabled participants to remove stories that were assumed to have low interest values, such as those that received only a few diggs and no comments. Having fewer distractions, participants could allocate their attention better to potentially interesting stories. They had more control of selecting stories to display in the visualization, and, as a result, they did not feel overwhelmed with a large number of stories in a collection:

I pick the first one [SocialVis] because I feel I have more control on it.

It is more manageable, and has some features that can filter some articles that we like.

> I was able to set the filters to ignore the highly commented and high digg scores and focus on articles that had SOME commentary and diggs, but not lots.

In summary, filtering by diggs and commenters helped participants browse information collections for the following reasons. First, it helped participants allocate their attention to potentially interesting stories. Second, it allowed participants to remove spam and low quality stories from their views. Third, it enabled participants to specify and refine their exploratory criteria, giving them more control of their information spaces. And, fourth, it decreased the cognitive load of participants by reducing the number of stories to process at a time.

Visualization of Social Interaction History

Visualization of social interaction history enables social navigation in an information space. It allows users to identify popular articles without skimming through individual items in a collection. While there is no guarantee that a popular article will be interesting to everyone, an article attracting a lot of attention likely has good characteristics, such as informative, well written, useful, or novel. Social interaction visualization essentially provides information scent to help users assess the potential value of articles and select those that fit their interests (Pirolli & Card, 1999).

The role of social interaction visualization in providing information scent for participants was reflected in their comments. Participants preferred SocialVis because they were able to get a sense of the popularity of stories easily (due to the visualization):

> I like the system [SocialVis] because it shows the number of diggs and commentators.
>
> I like the idea [of] showing number of diggs and comment[er]s as colored bars, which give a brief idea how popular the topic [story] is.
>
> As I gained experience looking at the Digg articles themselves and what the indicators were for an article I would be interested in, I realized seeing how many Diggs the item has was just as important as or more important than the title of the article. It was nice being able to see at a glance the most popular articles based on the bar size.

Participants actually could use social interaction history to guide their exploration in all systems, as these systems provided this information in one way or another. List displayed the numbers of diggs and comments as part of document surrogates. In TimeVis and SocialVis, this information was available in a pop-up window that appeared when users clicked on a story. The main difference was that SocialVis augmented this textual description with visualization of the numbers of diggs and commenters on stories.

The experimental results showed that visualization was more effective than a textual description for conveying social cues to users. These results were probably due to the ability of human vision to process and extract information from visualization (Ware, 2000). In SocialVis, participants were able to recognize popular articles by simply looking at the bars representing the numbers of diggs and commenters on stories. Rather than dealing with numbers, they could see the lengths of these bars to compare the popularity of different stories. In the context of exploratory browsing, participants did not need to know the exact number of diggs or commenters on a story. They just needed to get a sense of the relative popularity of stories in a collection. Visualization worked well because it provided subtle cues that were easily recognized by human vision. To take advantage of the visualization, participants needed to allocate only minimum effort. Thus, they could still pay full attention to the main task of exploring and finding interesting stories.

The effectiveness of information visualization was supported by comments from participants. For example, although the numbers of diggs and comments were also available in List, some participants seemed to ignore or did not pay enough attention to getting this information from short descriptions of stories:

> [List] couldn't tell which stories are more popular.
>
> I don't like Arjuna [List], because it does not provide any help for finding a popular or a specific story. I often ignore the short descriptions of the stories in Arjuna.

Besides guiding users to find popular articles, visualization of social interaction history also served as a collaborative filter to avoid spam. Spam and low quality articles do not attract a lot of attention. They are usually ignored and consequently do not have a rich interaction history. While it is possible non-popular articles are useful and interesting, focusing on popular articles allowed participants to reduce wasted effort, as it decreased the possibility of selecting spam.

In summary, visualization of social interaction history worked more effectively than textual description of the same information because the visualization did not distract participants from their main task, and participants could use it with minimum effort. While browsing, many participants focused on the titles of stories and paid little attention to short descriptions of the stories. Consequently, they missed social cues such as the numbers of diggs and comments on stories when this information was available only in a textual form. In contrast, when this information was visualized, participants received this social information automatically and were able to use it to guide their browsing. Using terminology in information foraging (Pirolli & Card, 1999), visualization of social interaction history provided rich information scent that helped participants assess and select stories with high interest values.

Easy Navigation

The feature comparisons in Table 5.7 suggest that continuous navigation did not have main effects on the performance measures (List vs. TimeVis). However, comments from participants indicated that easy, continuous navigation might have played an important role in supporting exploratory browsing when combined with social filters and visualization. The click-and-drag interaction style that enabled panning in SocialVis (and TimeVis) allowed participants to explore a data set easily:

> It is easy to navigate through the stories.
>
> [...] panning was easy and putting the articles on a timeline was nice, so I could see what time frame the articles were in (I could get a feel for their freshness).
>
> I liked how the flow of the titles was continuous and I didn't have to click "next" all of the time to get to a new page of entries like I did with the Arjuna [List] system.

As discussed in Section 3.1.2, information patch models in information foraging theory concern time allocation for within-patch and between-patches foraging activities (Pirolli & Card, 1999). Implementation of zooming and panning in SocialVis facilitated both within-patch and between-patches foraging. Instead of page-by-page navigation, participants could browse a whole data set in a continuous manner. By zooming out the visualization, they could move across a timeline quickly. And when they wanted to focus on a particular area, they could adjust the zoom level to get the desired view. Furthermore, due to a compact visualization design, SocialVis had higher information density than the List interface, so that participants could scan the whole collection quickly.

The nature of continuous navigation in SocialVis enabled participants to encounter interesting titles multiple times. Using this system, participants could move back and forth smoothly. Thus, they might reencounter interesting articles that they missed in the first place. In contrast, when using List, once participants moved away from a page, they could no longer see the contents of that page. If they missed interesting stories on that page, they would not encounter the stories again, unless they revisited that page. In practice, however, participants tended to browse new pages rather than going back to visited pages.

5.4.2 Implications of Findings

This experiment provided a case study where social interaction history was visualized to allow users to use social navigation to explore an information space. The results supported the basic hypothesis that social interaction history indicates the potential value of information items. This section outlines how these results generalize to other kinds of data sets, tasks, and applications.

Data Sets

The data set used in this study was taken from Digg, a real-world social web application. Digg organizes a collection of stories in chronological order based on their submitted time. To vote for a story or to leave comments on it, people must register to Digg and login using their usernames. Digg keeps track of user activities and social interaction that happen in its information space.

There are many applications that share similarities with Digg, such as blogs, online newspapers, and discussion forums. These applications arrange their information items—blog entries, news articles, and messages—in chronological order. They facilitate social interaction by allowing users to leave comments on these items. As a result, data sets in these applications contain rich social interaction history. People exploring these kinds of data sets may benefit from the approaches used in SocialVis. For example, filtering by the number of commenters may help identify popular blog entries, or separate interesting discussion from meaningless argument between two users. Social interaction visualization may help new users learn about the characteristics of information spaces (e.g., Xiong & Donath, 1999; Viégas & Smith, 2004).

Visualization of social interaction history in general has greater value when it is applied to data sets that do not become outdated easily. Social interaction history is an accumulation of traces of user activities over a long time. These traces become a meaningful indicator of the value of an information item only if the value does not decrease with time very quickly. If recency of contents is very important (e.g., articles about business and finance), people will focus on most recent items in a collection and consider older items irrelevant. In this case, social interaction history becomes less useful for supporting information exploration.

Developers must be aware of the "Matthew effect" (Merton, 1968) where "the rich get richer." This is an effect common to complex, self-organizing systems, and was first discussed in the context of social sciences. Merton discusses how famous researchers tend to get more rewards and recognition than less famous researchers who make comparable achievements. Highly cited articles tend to get cited more, as they are more visible. This tendency of allocating resources disproportionately to those that already have a lot of resources has been demonstrated in self-organizing networks such as the Web and also in physical systems. Barabási (2003) explains this phenomenon by a mechanism where each individual player (e.g., a human, a research article, a node on the Web) uses a simple rule—preferential attachment—in its interaction with others. For example, a new node on the Web tends to connect to an existing node that has many links. People prefer to buy books that have been announced as bestsellers, and prefer to read highly rated articles or articles from well-known resources. In essence, nodes that by chance or timing have managed to create more connections (or accumulate more ratings, or enter the bestseller list) have an advantage over the other nodes in attracting more connections, ratings, votes, etc. Timing is an important factor, as earlier nodes have more time to accumulate links. However, innate characteristics of a node can also play an important role in making it attractive. This innate attractiveness can be the high-quality, controversy, or another qualitative feature, related to the "value" of the node for the other nodes (as discussed in Chapter 3.2.1, the value of an information item has certainly a social aspect and depends on the community).

With all these considerations, developers must choose carefully which information to visualize. Visualization of social interaction history may amplify the rich-get-rich phenomenon where the popular items get more popular (because they become more visible) and the less popular items are left unexplored. Therefore, the choice of what features to visualize (e.g., the number of clicks, ratings, diggs, or comments) is important. For example, the number of clicks on a web article may not be a good measure of the article's quality because this number increases over time regardless of whether users like the article or not, and whether they actually read it or not. Thus, such read wear may end up being a misleading indicator of the value of an information item, particularly if the item happens to be clicked by many people initially. Other measures such as the number of comments or commenters on an article may serve as a better indicator. While it is possible that users leave comments on an article because the article is so poorly written, people in general are economic beings and will spend their effort to discuss an article only if they find it valuable.

The main point of using social interaction traces in a data set is not to lead users to popular items, but to give them more control to navigate an information space. Filtering stories by the number of comments, for example, allows users to explore popular items (by focusing on highly commented stories) or to look for hidden gems (by browsing stories with no comments). In this way, users can control their movement more flexibly. They can choose whether to follow or to avoid the crowds.

Exploratory Browsing Tasks

There are different kinds of exploratory tasks that may benefit from visualization of social interaction history. The experimental results provided evidence that social interaction visualization could improve the outcome of *casual browsing tasks*. The task given to participants involved browsing a data set to find interesting but not specific information items. Although the experimental setting was artificial, the test task simulated realistic, casual browsing tasks in the real world. For example, people use social networking applications such as Facebook[6] and Friendster[7] not to find specific information, but to keep in touch with friends. Their main activities involve browsing status updates, photos, and other shared items without having specific search goals or information needs. Casual browsing activities have also been observed at YouTube[8] where users explore a video collection to watch any interesting videos (Coyle et al., 2008). Visualization of social interaction history can help users perform casual browsing tasks because it reveals information items receiving a lot of attention from the audience, which are likely more interesting than less popular items.

Social interaction visualization can facilitate search tasks that involve *finding the most popular items*, such as finding bestsellers in a bookstore, popular books in a library, or the most active users in an online community. Bestseller items tend to be discussed or reviewed by a lot of people. Popular books in a library are borrowed by many patrons. Active users leave a lot of traces in a community. Visualization makes this information apparent to users.

'Browsing to learn' tasks may also benefit from visualization of social interaction

[6]http://www.facebook.com
[7]http://www.friendster.com
[8]http://www.youtube.com

history. Consider people who visit a personal blog for the first time. They only have a vague idea of the blog and do not know exactly what the blogger has written in the blog or how big the influence of the blog is. When faced with this situation, they would likely spend considerable time to browse the blog to learn about its contents, style, and audience. Visualization can support this learning process by depicting the characteristics and social dynamics in an information space (e.g., Indratmo et al., 2008; Xiong & Donath, 1999; Viégas & Smith, 2004). Without visualization, it is hard to recognize meaningful patterns in an unfamiliar space. In general, visualization offers greater value to people who are not familiar with an information space.

Potential Applications

From an applied perspective, there are many potential applications of social interaction history. For example, a discussion forum may visualize the number of users who are involved in a discussion thread to help its audience separate meaningful threads from "flame wars." Besides using a simple measure, the system may apply social network analysis to identify prominent members who contribute quality messages and use this information to promote their contributions. Furthermore, the system may enable users to specify and refine their search criteria based on social information, such as filtering messages by the number of replies or by the number of users involved in a discussion thread. I believe that augmenting content-oriented information exploration with social information will increase user satisfaction and the quality of search results.

Applied in a digital library, social interaction history may help library patrons find important literature based on borrowing patterns and other traces of activity left by users. Analysis and visualization of references in books and articles may help users identify relevant literature and prominent authors in a subject. For research papers, how often a paper has been downloaded and cited and by whom can be an indicator of its impact in a field.[9] If presented appropriately, this information will

[9] For example, http://portal.acm.org/dl.cfm

enhance user experience in exploring and browsing information collections.

Besides being visualized, traces of social interaction can be analyzed to provide personalized services and recommendations. For example, an e-commerce system may detect that there is a set of users who often discuss the same stuff and share the same sentiment towards the stuff. If a user from that group likes a particular item, then the system can recommend that item to other users in the group.

Consider also the area of personal information management (PIM), that is, approaches to managing, organizing, and retrieving personal collections of documents (Teevan, Jones, & Bederson, 2006). Although considered personal, PIM is often used in social settings (Erickson, 2006). People collaborate with their colleagues to prepare proposals or to write reports. They often exchange documents with other people via email. When they work on a project, they maintain and share relevant information items with other team members.

Traces of social interaction in PIM have potential to facilitate information management and to help users recall contextual information of their collections. For example, visualization of email exchanges may help users recall stakeholders of a project. Based on analysis of email exchanges, an email client may suggest important contacts to users and highlight new messages from these persons (Whittaker et al., 2004). In a study, Fisher and Dourish (2004) show that when users were presented with visualization of email archives portraying burst conversations, they were able to recognize and relate these communication patterns to meaningful events such as "the arrival of the summer interns" or "the fall patent negotiation."

5.4.3 Limitations of the Study

As discussed in section 5.2.1, I have taken several approaches to ensuring internal validity of this study. All participants used the same equipment, followed the same procedure, and were encouraged to give their honest opinions about the systems. Topic assignments and system trials were counterbalanced to avoid bias towards a particular system. Furthermore, questionnaire items were taken or adapted from previous studies.

Despite this effort, this study had some limitations. Influenced by Shneiderman's mantra (1996), the experimental design considered the filters in SocialVis to be an integral part of the social interaction visualization. As discussed in Chapter 5.4.1, comments from the participants suggested that both the visualization and the filters enhanced their browsing experiences. The filters allowed participants to focus on popular stories, to remove spam and advertisements, to refine their exploratory criteria, and to reduce the number of stories to process at a time. The visualization provided rich yet subtle information scent that helped participants assess the potential value of stories in a collection and make visual comparison of the stories' popularity. However, the experimental results were unable to tell which factor had greater effects on the performance measures. I speculate that the filters had greater contribution to the improved performance than the visualization, as they offered more control to users to explore an information space. This conjecture, however, does not mean that the visualization was less valuable. I believe that the visualization also played an important role in supporting the browsing processes and improving the user satisfaction with SocialVis. Future work may address this limitation by measuring the effect of each factor separately. For example, in one experimental condition we show only social interaction visualization, while in another condition we provide only a social filter.

Each system trial lasted only 15 minutes, and participants evaluated the system based on this short period of trial. Although I do not suspect that the results will be significantly different, when people use the systems for a longer period, there may be factors that affect their satisfaction with the systems. For example, one participant commented that "Having mild carpal tunnel syndrome, I found having to hold the mouse button while scrolling [panning] a bit painful." Fortunately, the Timeline visualization tool,[10] which was used to implement TimeVis and SocialVis, also enables users to pan across a timeline using the left and right arrow keys on a computer keyboard (in addition to click-and-drag).

The effectiveness of social interaction history in facilitating information explo-

[10]http://www.simile-widgets.org/timeline/

ration depends on the nature of a data set. For example, one participant seemed to completely ignore the visualization of the numbers of diggs and commenters while browsing articles about business and finance. Rather than browsing the whole collection, she only focused on most recent articles: "I would like to read the articles most recently listed, particularly for the articles related to business and finance." To her, old news about business and finance were not interesting because they were already outdated. Thus, social interaction history may be considered irrelevant in a collection where time of publication is important.

The types of search tasks also affect the role of social interaction history in information retrieval. If users just need to look up specific information, they need not pay much attention to social interaction history. If they know exactly what they are looking for and how to retrieve the information, they can access the information source directly. Recall that the idea of utilizing social interaction history comes from the observation in everyday life where people tend to observe activities of other people to help them make decision, especially when they are faced with uncertainty. Therefore, this approach is particularly suitable for situations where users are not familiar with a collection, have exploratory or vague search goals, or cannot formulate search terms properly due to lack of necessary background knowledge.

Finally, most participants considered themselves to be expert end-users, and all had experience browsing the Internet. While all participants had no difficulties in learning to use the systems, novice end-users may have different opinions. The experimental results currently can only be generalized to experienced computer users.

5.5 Summary

This study measured and compared the performance of three systems in supporting exploratory browsing. The experimental results showed that visualization of social interaction history was able to help users find interesting articles, to reduce wasted effort to find these articles, and to increase user satisfaction with the visualization tool. Overall, participants preferred the visualization tool to the other systems.

These results supported the basic hypothesis that social interaction history can serve as an indicator of the potential value of information items. Implications of these findings were discussed, along with their potential applications in various domains.

Chapter 6

Conclusion

This chapter discusses a summary of this thesis, research contributions, and directions for future work.

6.1 Summary

The main objective of this research was to investigate whether visualization of social interaction history can improve exploratory browsing. The research started with a basic hypothesis that social interaction history can serve as an indicator of the potential value of information items. Therefore, visualization of social interaction history would be able to improve browsing, as it provides information scent (Pirolli & Card, 1999) and affords users to use social navigation in exploring information spaces.

To test this basic hypothesis, I conducted three studies. First, I ran statistical analysis of a data set from a social web application (Chapter 3). The results showed that there were positive relationships between traces of social interaction and the degree of interestingness of web articles. In particular, the number of commenters on an article was a better indicator of the story interestingness than the number of comments. This result was then applied to the design of a visualization tool (SocialVis), which served as one of the experimental conditions in the summative evaluation (Chapter 5).

Second, I conducted an exploratory study to evaluate the feasibility of using social interaction history to improve browsing (Chapter 4). To achieve this goal, I developed a visualization tool for browsing a blog archive (iBlogVis). The tool provides an overview of a blog archive and visualizes social interaction history preserved in the archive. The tool offers a novel way to browse a blog archive and demonstrates the synthesis and application of existing visualization and interaction techniques in blogs. Responses from the study participants supported the usefulness of this approach: participants indicated that the visualization of social interaction history would allow them to identify popular blog entries quickly and effortlessly. The evaluation results also added further evidence that providing an overview of an information space is useful for supporting information exploration.

Finally, to follow up these findings, I conducted a summative evaluation to test quantitatively whether users can perform better when they receive social navigational cues (Chapter 5). The results showed that visualization of social interaction history helped users find interesting information, reduced wasted effort to find this information, and increased user satisfaction with the visualization tool. Most participants preferred the visualization tool to the other systems.

Overall, the study results provided evidence of the value of social interaction history in supporting exploration of information spaces. Throughout this thesis, I have discussed implications of these findings and how these results generalize to different kinds of applications.

Besides showing the value of social interaction history, this thesis also reaffirmed the value of design principles in information visualization. To some degree the features of the visualization tools used in the studies—iBlogVis, TimeVis, and SocialVis—exemplified the following principles:

- The mapping between data and its graphical representations in the visualizations followed the guidelines suggested by Cleveland and McGill (1984) and Mackinlay (1986).

- The design and interactive features of the tools resembled Shneiderman's mantra (1996): "Overview first, zoom and filter, then details-on-demand."

- The widgets used to support interactive features of the tools implemented the principles of dynamic query interfaces (Ahlberg et al., 1992).

In the exploratory and summative studies, comments from participants suggested that these design principles have contributed to the usability of the visualization tools, the outcome of browsing, and the user satisfaction with the tools. For example, in the exploratory study, participants appreciated an overview of a blog archive, as it allowed them to learn about the blog and its characteristics. In the summative evaluation, participants repeatedly mentioned that the ability to filter stories by the numbers of diggs and commenters was the main reason why they liked SocialVis. Overall, participants did not have any difficulties in learning and understanding the meanings and features of the visualizations, indicating that the mapping between data and their graphical representations was logical and appropriate.

To some degree the visualization tools used in my studies implement Shneiderman's interaction model (1996): the tools provide an overview of a collection then allow users to drill down. While this interaction model worked well for the exploratory browsing tasks in the studies, this model may not be the best choice to deal with different kinds of search tasks. For example, if users have a semi-specific information need, they may find an overview to be less useful; instead, they may want to start with a specific item of interest and explore relevant information surrounding this item (van Ham & Perer, 2009).

6.2 Contributions

This thesis covers multi-disciplinary areas such as information exploration and visualization, human-computer interaction, and social computing, and contributes to these fields in the following ways:

- First, this thesis developed the concept of social interaction history, which has potential to improve information management, retrieval, and exploration. In Chapter 3, I discussed how social interaction history relates to computational wear (Hill et al., 1992), information foraging (Pirolli & Card, 1999), social navigation (Dourish & Chalmers, 1994), and multi-dimensional notions of relevance (Barry, 1994; Saracevic, 2007).

- Second, on the basis of three studies, this thesis demonstrated the usefulness of social interaction history in supporting information exploration. Specifically, in Chapter 5, I showed that visualization of social interaction history was able to improve browsing, to reduce wasted effort, and to increase user satisfaction. Besides providing quantitative evidence, this thesis also offered a theory of why social interaction history works in facilitating information exploration.

- Third, this thesis outlined practical applications of social interaction history in various domains. Despite having great potential, traces of social interaction are still underutilized. With the proliferation of social web applications, I expect to see increased utilization of social interaction history. I believe that augmenting content-oriented navigation with social information will improve information management, retrieval, and exploration.

6.3 Future Work

This thesis has shown the value of social interaction history in facilitating exploratory browsing and has identified its potential applications in various domains. The following sections suggest several directions for future work.

6.3.1 Qualitative Aspects of Social Interaction History

A natural extension of this work is to consider qualitative aspects of social interaction history.[1] For example, an article receiving a few comments from prominent users can be more interesting and valuable than those receiving many comments from unknown users. Being cited by a few famous researchers can be a better indicator of the quality of a research paper than simply having a high number of citations. A user may consider an email message from a family member more important than a message from a colleague, although the user may exchange messages more frequently with that colleague. Thus, besides relying on quantitative measures such as the number

[1] Marc Smith (2009), personal communication.

of comments or the frequency of interaction, a system should give different weight to different people and should allow users to adjust this weight if necessary. Future work may benefit from methods for analyzing social networks (e.g., Wasserman & Faust, 1994) and the notion of trust and reputation in distributed systems (e.g., Wang & Vassileva, 2003).

A system should also consider recency of information items. Social interaction history is an accumulation of user activities in an information space. Consequently, older items will contain more traces of social interaction compared to newer items (assuming that they are of comparable quality). Without considering recency of items, a system will mistakenly assume the older items to have better quality than the newer ones. Furthermore, old information at some point may become irrelevant or less valuable. Therefore, developers should take such factors into account when utilizing social interaction history, for example, by assigning less weight to older traces of interaction and more weight to newer ones.

6.3.2 Privacy Issues

Another direction of research is to explore privacy issues tied to usage of social interaction history. Recording, analyzing, and visualizing traces of social interaction have potential to facilitate information management, retrieval, and exploration. However, this approach also brings the risk of privacy violation. Traces and patterns of social interaction, which are traditionally hidden, now become visible. People may become uncomfortable and suspicious to use a system that records every detail of their activities. They need to know what kinds of information are recorded, who can access this information, and how a system uses this information.

There must be a balance between exploitation of social interaction history and protection of user privacy. Further studies need to develop a set of design principles to achieve this balance. Recommendations from other domains can serve as a starting point to delve into these issues (e.g., Anwar, Greer, & Brooks, 2006). Open issues include the appropriate level of details of social interaction records and the proper usage, accessibility, and presentation of this information.

6.3.3 Integrated Design of Exploratory Search Systems

Finally, back to the big picture, when users perform exploratory searches, they likely use multiple information-seeking strategies (Belkin et al., 1993). They enter search terms, browse a collection, save relevant information, filter out irrelevant items, compare a set of documents, refine search criteria, and so on. Ideally, an exploratory search system should support all of these strategies. While earlier work has developed solutions to specific problems, integrating these solutions into an effective exploratory search system warrants further research. User studies observing how people search and explore information collections can help identify useful features that should be supported by exploratory search systems (e.g., Diaz, Hu, & Tory, 2009).

After a full exploratory search system has been developed, the system should be evaluated in a longitudinal, field study. Controlled experiments, as done in this thesis, are valuable for testing specific hypotheses and for identifying factors that affect performance measures of a system. However, people may develop different search strategies after they become familiar with an information space. For example, do they still consider social interaction history useful when they know the information space well? A longitudinal study will be able to answer such questions and reveal the actual usage patterns of a system in the real world.

6.4 Concluding Remarks

This thesis discusses the results of three studies that evaluate the value of social interaction history in supporting exploratory browsing. The approach proposed in this thesis is not intended to solve all browsing problems or to replace existing approaches to facilitating browsing. Instead, this research shows that visualizing social interaction history of an information space can help users explore an information collection. The transformation of the Web into social information spaces opens good opportunity to use social interaction history to improve information exploration in various domains.

References

Adar, E., & Adamic, L. A. (2005). Tracking information epidemics in blogspace. In *Proceedings of the 2005 IEEE/WIC/ACM international conference on Web intelligence* (pp. 207–214).

Ahlberg, C., Williamson, C., & Shneiderman, B. (1992). Dynamic queries for information exploration: An implementation and evaluation. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 619–626).

Allen, R. B. (1995). Retrieval from facet spaces. *Electronic Publishing*, *8*(2 & 3), 247–257.

Anwar, M. M., Greer, J., & Brooks, C. A. (2006). Privacy enhanced personalization in e-learning. In *Proceedings of the 2006 international conference on Privacy, security and trust* (pp. 1–4).

Baldonado, M. Q. W., & Winograd, T. (1997). SenseMaker: An information-exploration interface supporting the contextual evolution of a user's interests. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 11–18).

Barabási, A.-L. (2003). *Linked: How everything is connected to everything else and what it means for business, science, and everyday life*. Plume.

Barry, C. L. (1994). User-defined relevance criteria: An exploratory study. *Journal of the American Society for Information Science*, *45*(3), 149–159.

Bates, M. J. (1989). The design of browsing and berrypicking techniques for the online search interface. *Online Review*, *13*(5), 407–424.

Becker, R. A., & Cleveland, W. S. (1987). Brushing scatterplots. *Technometrics*, *29*(2), 127–142.

Belkin, N. J., Marchetti, P. G., & Cool, C. (1993). BRAQUE: Design of an interface to support user interaction in information retrieval. *Information Processing & Management*, *29*(3), 325–344.

Bhavnani, S. K. (2001). Important cognitive components of domain-specific search knowledge. In *Proceedings of the 10th Text REtrieval conference* (pp. 571–578).

Borlund, P. (2000). Experimental components for the evaluation of interactive information retrieval systems. *Journal of Documentation*, *56*(1), 71–90.

Borlund, P. (2003). The IIR evaluation model: A framework for evaluation of interactive information retrieval systems. *Information Research*, *8*(3). Paper no. 152, available at http://informationr.net/ir/8-3/paper152.html.

Brin, S., & Page, L. (1998). The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, *30*(1–7), 107–117.

Capra, R., Marchionini, G., Oh, J. S., Stutzman, F., & Zhang, Y. (2007). Effects of structure and interaction style on distinct search tasks. In *Proceedings of the 2007 conference on Digital libraries* (pp. 442–451).

Card, S. K., Robertson, G. G., & Mackinlay, J. D. (1991). The information visualizer, an information workspace. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 181–186).

Chang, S. J., & Rice, R. E. (1993). Browsing: A multidimensional framework. *Annual Review of Information Science and Technology*, *28*, 231–276.

Chin, A., & Chignell, M. (2006). A social hypertext model for finding community in blogs. In *Proceedings of the 17th conference on Hypertext and hypermedia* (pp. 11–22).

Chin, J. P., Diehl, V. A., & Norman, K. L. (1988). Development of an instrument measuring user satisfaction of the human-computer interface. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 213–218).

Cleveland, W. S., & McGill, R. (1984). Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association*, *79*(387), 531–554.

Cockburn, A. (2004). Revisiting 2D vs 3D implications on spatial memory. In *Proceedings of the fifth conference on Australasian user interface* (pp. 25–31).

Cockburn, A., & Jones, S. (1996). Which way now? Analysing and easing inadequacies in WWW navigation. *International Journal of Human-Computer Studies*, *45*(1), 105–129.

Cockburn, A., Karlson, A., & Bederson, B. B. (2008). A review of overview+detail, zooming, and focus+context interfaces. *ACM Computing Surveys*, *41*(1), 1–31.

Cool, C., & Belkin, N. J. (2002). A classification of interactions with information. In *Proceedings of the 4th international conference on Conceptions of library and information science* (pp. 1–15).

Coyle, M., Freyne, J., Brusilovsky, P., & Smyth, B. (2008). Social information access for the rest of us: An exploration of social YouTube. In *Proceedings of the 5th international conference on Adaptive hypermedia and adaptive web-based systems* (pp. 93–102).

Cutrell, E., Robbins, D., Dumais, S., & Sarin, R. (2006). Fast, flexible filtering with phlat. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 261–270).

Cutting, D. R., Karger, D. R., Pedersen, J. O., & Tukey, J. W. (1992). Scatter/gather: A cluster-based approach to browsing large document collections. In *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval* (pp. 318–329).

Czerwinski, M., Dantzich, M. van, Robertson, G., & Hoffman, H. (1999). The contribution of thumbnail image, mouse-over text and spatial location memory to web page retrieval in 3D. In *Proceedings of INTERACT '99* (pp. 163–170).

Danis, C. M., Viégas, F. B., Wattenberg, M., & Kriss, J. (2008). Your place or mine? Visualization as a community component. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 275–284).

Diaz, J., Hu, K., & Tory, M. (2009). An exploratory study of tag-based visual interfaces for searching folksonomies. In *Proceedings of the BCS conference on Human computer interaction.*

Dieberger, A. (1997). Supporting social navigation on the world wide web. *International Journal of Human-Computer Studies*, *46*(6), 805–825.

Dieberger, A., Dourish, P., Höök, K., Resnick, P., & Wexelblat, A. (2000). Social navigation: Techniques for building more usable systems. *interactions*, *7*(6), 36–45.

Dix, A. (2008). Theoretical analysis and theory creation. In P. Cairns & A. L. Cox (Eds.), *Research methods for human-computer interaction* (pp. 175–195). Cambridge: Cambridge University Press.

Dörk, M., Carpendale, S., Collins, C., & Williamson, C. (2008). Visgets: Coordinated visualizations for web-based information exploration and discovery. *IEEE Transactions on Visualization and Computer Graphics*, *14*(6), 1205–1212.

Dourish, P., & Chalmers, M. (1994). *Running out of space: Models of information navigation.* (Short paper presented at HCI'94. Retrieved February 11, 2009, from http://www.dourish.com/publications/1994/hci94-navigation.pdf)

Dourish, P., Edwards, W. K., LaMarca, A., Lamping, J., Petersen, K., Salisbury, M., et al. (2000). Extending document management systems with user-specific active properties. *ACM Transactions on Information Systems*, *18*(2), 140–170.

Dubinko, M., Kumar, R., Magnani, J., Novak, J., Raghavan, P., & Tomkins, A. (2006). Visualizing tags over time. In *Proceedings of the 15th international conference on World wide web* (pp. 193–202).

Dumais, S., & Belkin, N. J. (2005). The TREC interactive tracks: Putting the user into search. In E. M. Voorhees & D. K. Harman (Eds.), *TREC: Experiment and evaluation in information retrieval* (pp. 123–152).

Dumais, S., Cutrell, E., Cadiz, J., Jancke, G., Sarin, R., & Robbins, D. C. (2003). Stuff I've seen: A system for personal information retrieval and re-use. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval* (pp. 72–79).

Ellis, G., & Dix, A. (2006). An explorative analysis of user evaluation studies in information visualisation. In *Proceedings of the 2006 AVI workshop on Beyond time and errors* (pp. 1–7).

Erickson, T. (2006). From PIM to GIM: Personal information management in group contexts. *Communications of the ACM*, *49*(1), 74–75.

Erickson, T., & Kellogg, W. A. (2003). Social translucence: Using minimalist visualisations of social activity to support collective interaction. In *Designing information spaces: The social navigation approach* (pp. 17–41). London, UK: Springer-Verlag.

Fiore, A. T., LeeTiernan, S., & Smith, M. A. (2002). Observed behavior and perceived value of authors in usenet newsgroups: Bridging the gap. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 323–330).

Fisher, D., & Dourish, P. (2004). Social and temporal structures in everyday collaboration. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 551–558).

Freyne, J., Farzan, R., Brusilovsky, P., Smyth, B., & Coyle, M. (2007). Collecting community wisdom: Integrating social search & social navigation. In *Proceedings of the 12th international conference on Intelligent user interfaces* (pp. 52–61).

Furnas, G. W. (1986). Generalized fisheye views. *SIGCHI Bulletin*, *17*(4), 16–23.

Furnas, G. W., & Bederson, B. B. (1995). Space-scale diagrams: Understanding multiscale interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 234–241).

Garton, L., Haythornthwaite, C., & Wellman, B. (1999). Studying online social networks. In S. Jones (Ed.), *Doing internet research: Critical issues and methods for examining the net* (pp. 75–105).

Gifford, D. K., Jouvelot, P., Sheldon, M. A., & O'Toole, J. W. (1991). Semantic file systems. In *Proceedings of the 13th ACM symposium on Operating systems principles* (pp. 16–25).

Golder, S. A., & Huberman, B. A. (2006). Usage patterns of collaborative tagging systems. *Journal of Information Science*, *32*(2), 198–208.

Gopal, B., & Manber, U. (1999). Integrating content-based access mechanisms with hierarchical file systems. In *Proceedings of the 3rd symposium on Operating systems design and implementation* (pp. 265–278).

Graham, A., Garcia-Molina, H., Paepcke, A., & Winograd, T. (2002). Time as essence for photo browsing through personal digital libraries. In *Proceedings of the 2nd ACM/IEEE-CS joint conference on Digital libraries* (pp. 326–335).

Gruhl, D., Liben-Nowell, D., Guha, R., & Tomkins, A. (2004). Information diffusion through blogspace. *ACM SIGKDD Explorations Newsletter*, *6*(2), 43–52.

Gutwin, C., Roseman, M., & Greenberg, S. (1996). A usability study of awareness widgets in a shared workspace groupware system. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work* (pp. 258–267).

He, D., Brusilovsky, P., Ahn, J., Grady, J., Farzan, R., Peng, Y., et al. (2008). An evaluation of adaptive filtering in the context of realistic task-based information exploration. *Information Processing & Management*, *44*(2), 511–533.

Hearst, M. A. (1995). TileBars: Visualization of term distribution information in full text information access. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 59–66).

Hearst, M. A. (2006). Clustering versus faceted categories for information exploration. *Communications of the ACM*, *49*(4), 59–61.

Hearst, M. A., & Pedersen, J. O. (1996). Reexamining the cluster hypothesis: Scatter/gather on retrieval results. In *Proceedings of the 19th annual international ACM SIGIR conference on Research and development in information retrieval* (pp. 76–84).

Heer, J., Card, S. K., & Landay, J. A. (2005). prefuse: A toolkit for interactive information visualization. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 421–430).

Heer, J., Viégas, F. B., & Wattenberg, M. (2007). Voyagers and voyeurs: Supporting asynchronous collaborative information visualization. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 1029–1038).

Herring, S. C., Kouper, I., Paolillo, J. C., Scheidt, L. A., Tyworth, M., Welsch, P., et al. (2005). Conversations in the blogosphere: An analysis "from the bottom up". In *Proceedings of the Hawaii international conference on System sciences.*

Herring, S. C., Scheidt, L. A., Bonus, S., & Wright, E. (2004). Bridging the gap: A genre analysis of weblogs. In *Proceedings of the Hawaii international conference on System sciences* (pp. 101–111).

Hill, W. C., Hollan, J. D., Wroblewski, D., & McCandless, T. (1992). Edit wear and read wear. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 3–9).

Howell, D. C. (2002). *Statistical methods for psychology* (5th ed.). Pacific Grove, CA: Duxbury.

Indratmo, & Vassileva, J. (2009). Social interaction history: A framework for supporting exploration of social information spaces. In *Proceedings of the IEEE international symposium on Social computing applications* (pp. 538–545).

Indratmo, Vassileva, J., & Gutwin, C. (2008). Exploring blog archives with interactive visualization. In *Proceedings of the working conference on Advanced visual interfaces* (pp. 39–46).

Jansen, B. J. (2006). Using temporal patterns of interactions to design effective automated searching assistance. *Communications of the ACM*, *49*(4), 72–74.

Johnson, B., & Shneiderman, B. (1991). Tree-maps: A space-filling approach to the visualization of hierarchical information structures. In *Proceedings of the 2nd conference on Visualization* (pp. 284–291).

Jones, W., & Dumais, S. (1986). The spatial metaphor for user interfaces: Experimental tests of reference by location versus name. *ACM Transactions on Information Systems*, *4*(1), 42–63.

Kaptelinin, V. (2003). UMEA: Translating interaction histories into project contexts. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 353–360).

Kidd, A. (1994). The marks are on the knowledge worker. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 186–191).

Komlodi, A., Marchionini, G., & Soergel, D. (2007). Search history support for finding and using information: User interface design recommendations from a user study. *Information Processing & Management*, *43*(1), 10–29.

Konstan, J. A., & Riedl, J. (2003). Collaborative filtering: Supporting social navigation in large, crowded infospaces. In *Designing information spaces: The social navigation approach* (pp. 43–82). London, UK: Springer-Verlag.

Krishnan, A., & Jones, S. (2005). TimeSpace: Activity-based temporal visualisation of personal information spaces. *Personal and Ubiquitous Computing*, *9*(1), 46–65.

Kules, B., & Shneiderman, B. (2008). Users can change their web search tactics: Design guidelines for categorized overviews. *Information Processing & Management*, *44*(2), 463–484.

Lagergren, E., & Over, P. (1998). Comparing interactive information retrieval systems across sites: The TREC-6 interactive track matrix experiment. In *Pro-*

ceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval (pp. 164–172).

Lenhart, A., & Fox, S. (2006). *Bloggers: a portrait of the internet's new storytellers.* Pew Internet Report, http://www.pewinternet.org/PPF/r/186/ report_display.asp.

Lerman, K. (2007). Social information processing in news aggregation. *IEEE Internet Computing*, *11*(6), 16–28.

Lerman, K., & Galstyan, A. (2008). Analysis of social voting patterns on digg. In *Proceedings of the ACM SIGCOMM workshop on Online social networks* (pp. 7–12).

Lynch, C. A. (2001). When documents deceive: Trust and provenance as new factors for information retrieval in a tangled web. *Journal of the American Society for Information Science and Technology*, *52*(1), 12–17.

Mackinlay, J. (1986). Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, *5*(2), 110–141.

Mackinlay, J., Robertson, G. G., & Card, S. K. (1991). The perspective wall: Detail and context smoothly integrated. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 173–176).

Malone, T. W. (1983). How do people organize their desks? Implications for the design of office information systems. *ACM Transactions on Information Systems*, *1*(1), 99–112.

Marchionini, G. (1995). *Information Seeking in Electronic Environments*. Cambridge University Press.

Marchionini, G. (2000). *From overviews to previews to answers: Integrated interfaces for federal statistics.* http://ils.unc.edu/~march/bls_final_report_99-00.pdf.

Marchionini, G. (2006). Exploratory search: From finding to understanding. *Communications of the ACM*, *49*(4), 41–46.

Marlow, C., Naaman, M., boyd, d., & Davis, M. (2006). Ht06, tagging paper, taxonomy, flickr, academic article, to read. In *Proceedings of the 17th conference on Hypertext and hypermedia* (pp. 31–40).

Maxwell, S. E. (1980). Pairwise multiple comparisons in repeated measures designs. *Journal of Educational Statistics*, *5*(3), 269–287.

Merton, R. K. (1968). The Matthew effect in science: The reward and communication systems of science are considered. *Science*, *159*(3810), 56–63.

O'Day, V. L., & Jeffries, R. (1993). Orienteering in an information landscape: How information seekers get from here to there. In *Proceedings of the INTERACT '93 and CHI '93 conference on Human factors in computing systems* (pp. 438–445).

Olston, C., & Chi, E. H. (2003). ScentTrails: Integrating browsing and searching on the web. *ACM Transactions on Computer-Human Interaction*, *10*(3), 177–197.

Perlin, K., & Fox, D. (1993). Pad: An alternative approach to the computer interface. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques* (pp. 57–64).

Petrelli, D. (2008). On the role of user-centred evaluation in the advancement of interactive information retrieval. *Information Processing & Management*, *44*(1), 22–38.

Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review*, *106*(4), 643–675.

Plaisant, C., Milash, B., Rose, A., Widoff, S., & Shneiderman, B. (1996). LifeLines: Visualizing personal histories. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 221–227).

Qu, Y., & Furnas, G. W. (2008). Model-driven formative evaluation of exploratory search: A study under a sensemaking framework. *Information Processing & Management*, *44*(2), 534–555.

Ravasio, P., Schär, S. G., & Krueger, H. (2004). In pursuit of desktop evolution: User problems and practices with modern desktop systems. *ACM Transactions on Computer-Human Interaction*, *11*(2), 156–180.

Robertson, G., Czerwinski, M., Larson, K., Robbins, D. C., Thiel, D., & Dantzich, M. van. (1998). Data Mountain: Using spatial memory for document management. In *Proceedings of the 11th annual ACM symposium on User interface software and technology* (pp. 153–162).

Robertson, G., Mackinlay, J. D., & Card, S. K. (1991). Cone Trees: Animated 3D visualizations of hierarchical information. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 189–194).

Robertson, S. E., & Hancock-Beaulieu, M. M. (1992). On the evaluation of IR systems. *Information Processing & Management*, *28*(4), 457–466.

Rodden, K., & Wood, K. R. (2003). How do people manage their digital photographs? In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 409–416).

Ruthven, I., Baillie, M., Azzopardi, L., Bierig, R., Nicol, E., Sweeney, S., et al. (2008). Contextual factors affecting the utility of surrogates within exploratory search. *Information Processing & Management*, *44*(2), 437–462.

Saracevic, T. (2007). Relevance: A review of the literature and a framework for thinking on the notion in information science. Part II: Nature and manifestations of relevance. *Journal of the American Society for Information Science and Technology*, *58*(13), 1915–1933.

Sarkar, M., & Brown, M. H. (1992). Graphical fisheye views of graphs. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 83–91).

schraefel m. c., Smith, D. A., Owens, A., Russell, A., Harris, C., & Wilson, M. (2005). The evolving mSpace platform: Leveraging the semantic web on the trail of the memex. In *Proceedings of the 16th ACM conference on Hypertext and hypermedia* (pp. 174–183).

Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the IEEE symposium on Visual languages* (pp. 336–343).

Skopik, A., & Gutwin, C. (2005). Improving revisitation in fisheye views with visit wear. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 771–780).

Smeulders, A. W. M., Worring, M., Santini, S., Gupta, A., & Jain, R. (2000). Content-based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *22*(12), 1349–1380.

Spence, R. (2003). Information space navigation: A framework. In *Designing information spaces: The social navigation approach* (pp. 405–426). London, UK: Springer-Verlag.

Spence, R. (2007). *Information Visualization: Design for Interaction (Second Edition)*. Pearson Education Limited.

Stevens, S. S. (1946). On the theory of scales of measurement. *Science*, *103*(2684), 677–680.

Surowiecki, J. (2004). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations*. Doubleday.

Svensson, M., & Höök, K. (2003). Social navigation of food recipes: Designing kalas. In *Designing information spaces: The social navigation approach* (pp. 201–222). London, UK: Springer-Verlag.

Tavanti, M., & Lind, M. (2001). 2D vs 3D, implications on spatial memory. In *Proceedings of the IEEE Symposium on Information Visualization* (pp. 139–145).

Teevan, J., Alvarado, C., Ackerman, M. S., & Karger, D. R. (2004). The perfect search engine is not enough: A study of orienteering behavior in directed search. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 415–422).

Teevan, J., Jones, W., & Bederson, B. B. (2006). Personal information management. *Communications of the ACM*, *49*(1), 40–43.

Tirapat, T., Espiritu, C., & Stroulia, E. (2006). Taking the community's pulse: One blog at a time. In *Proceedings of the 6th international conference on Web engineering* (pp. 169–176).

Tory, M., Sprague, D., Wu, F., So, W. Y., & Munzner, T. (2007). Spatialization design: Comparing points and landscapes. *IEEE Transactions on Visualization and Computer Graphics*, *13*(6), 1262–1269.

van Ham, F., & Perer, A. (2009). "Search, show context, expand on demand": Supporting large graph exploration with degree-of-interest. *IEEE Transactions on Visualization and Computer Graphics*, *15*(6), 953–960.

Vassileva, J., & Sun, L. (2007). An improved design and a case study of a social visualization encouraging participation in online communities. In *Proceedings of the 13th collaboration researchers' international workshop on Groupware* (pp. 72–86).

Viégas, F. B., & Smith, M. (2004). Newsgroup Crowds and AuthorLines: Visualizing the activity of individuals in conversational cyberspaces. In *Proceedings of the Hawaii international conference on System sciences.*

Viégas, F. B., Wattenberg, M., & Dave, K. (2004). Studying cooperation and conflict between authors with history flow visualizations. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 575–582).

Viégas, F. B., Wattenberg, M., van Ham, F., Kriss, J., & McKeon, M. (2007).

Many eyes: A site for visualization at internet scale. *IEEE Transactions on Visualization and Computer Graphics*, *13*(6), 1121–1128.

Wang, Y., & Vassileva, J. (2003). Trust and reputation model in peer-to-peer networks. In *Proceedings of the 3rd international conference on Peer-to-peer computing* (pp. 150–157).

Ware, C. (2000). *Information visualization: Perception for design*. Morgan Kaufmann Publishers.

Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications*. Cambridge University Press.

Wexelblat, A., & Maes, P. (1999). Footprints: History-rich tools for information foraging. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 270–277).

White, R. W., Bilenko, M., & Cucerzan, S. (2007). Studying the use of popular destinations to enhance web search interaction. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval* (pp. 159–166).

White, R. W., Kules, B., & Bederson, B. (2005). Exploratory search interfaces: Categorization, clustering and beyond. *SIGIR Forum*, *39*(2), 52–56.

White, R. W., Kules, B., Drucker, S. M., & schraefel m.c. (2006). Supporting exploratory search. *Communications of the ACM*, *49*(4), 36–39.

Whittaker, S., Jones, Q., Nardi, B., Creech, M., Terveen, L., Isaacs, E., et al. (2004). ContactMap: Organizing communication in a social desktop. *ACM Transactions on Computer-Human Interaction*, *11*(4), 445–471.

Willett, W., Heer, J., & Agrawala, M. (2007). Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics*, *13*(6), 1129–1136.

Woodruff, A., Faulring, A., Rosenholtz, R., Morrison, J., & Pirolli, P. (2001). Using thumbnails to search the web. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 198–205).

Xiong, R., & Donath, J. (1999). PeopleGarden: creating data portraits for users. In *Proceedings of the 12th annual ACM symposium on User interface software and technology* (pp. 37–44).

Yee, K. P., Swearingen, K., Li, K., & Hearst, M. (2003). Faceted metadata for image search and browsing. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 401–408).

Zamir, O., & Etzioni, O. (1999). Grouper: A dynamic clustering interface to web search results. *Computer Networks*, *31*(11-16), 1361–1374.

Zeng, H. J., He, Q. C., Chen, Z., Ma, W. Y., & Ma, J. (2004). Learning to cluster web search results. In *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval* (pp. 210–217).

Appendix A

Materials for the Exploratory Study

A.1 Informed Consent Form

You are invited to participate in a study entitled "Visualization of Blog Archives." Please read this form carefully, and feel free to ask questions you might have.

Researchers: Julita Vassileva, Department of Computer Science (966-2073)
Indratmo, Department of Computer Science (966-4744)

The purpose of the study is to evaluate the usability of a visualization tool for exploring blog archives. In the study, you will try the tool and answer some questions about your satisfaction with the tool. The estimate of the total time to participate in this study is 45 minutes.

There are no known risks in this study.

Findings from the study will be used to refine a tool to facilitate exploration of blog archives. You will be able to try and use the tool for your own purposes if you like it. You will receive a $10 honorarium.

The research data will be stored on a password-protected computer system and will be available only to the investigators. Identifying information will be destroyed upon completion of data collection, and then pseudonyms will be used to refer to the participants. All data will be kept for a minimum of five years upon the completion of this study.

Aggregate results will be used in a thesis and articles published in conferences and journals. However, any information that can be linked to a specific participant will be removed or altered.

Your participation is voluntary, and you may withdraw from the study for any reason, at any time, without penalty of any sort. You may refuse to answer individual questions. If you withdraw from the study at any time, any data that you have contributed will be destroyed at your request.

If you have any questions concerning the study, please feel free to ask at any point; you are also free to contact the researchers at the numbers provided above if you have questions at a later time. This study has been approved on ethical grounds by the University of Saskatchewan Behavioural Research Ethics Board on May 17, 2007. Any questions regarding your rights as a participant may be addressed to that committee through the Ethics Office (966-2084). Out of town participants may call collect. You may find out about the results of the study through MADMUC website (http://bistrica.usask.ca/madmuc/) or by contacting the researchers.

I have read and understood the description provided above; I have been provided with an opportunity to ask questions and my questions have been answered satisfactorily. I consent to participate in the study described above, understanding that I may withdraw this consent at any time. A copy of this consent form has been given to me for my records.

(Name of Participant)	(Date)
(Signature of Participant)	(Signature of Researcher)

A.2 Tasks and Questionnaires

A.2.1 List of Tasks

1. In what year did the blogger post his first entry?
2. In what year did the blogger post his last entry?
3. In what year was the most active posting period in the blog?
4. What was the most frequently used tag (overall) in the blog?
5. What was the most frequently used tag in 2004?
6. Who was the most frequent commenter (overall)?
7. Who was the most frequent commenter in 2005?
8. Did the blog receive comments regularly from the audience? Yes / No
9. What is the title of the most recent entry in the blog?
10. Please list the title of an entry tagged "Linux."
11. Please list the title of an entry tagged "Problem solving."
12. Please list the title of an entry posted in May 2005.
13. Please list the title of an entry posted in September 2006.
14. Please list the title of the most popular entry (popular entries refer to those received many comments).
15. Please list the title of an entry commented by "Ronny."
16. Please list the title of an entry commented by "Dan Williams."

A.2.2 Tool-Specific Questionnaire

Please answer the following questions using the visualization tool.

Time

- In what year did the blogger post his first entry?
- In what year did the blogger post his last entry?
- In what year was the most active posting period in the blog?

I found that getting a sense of time of the year was ...
☐ Not at all easy ☐ Not easy ☐ Easy ☐ Extremely easy

Tags

- What was the most frequently used tag (overall) in the blog?
- What was the most frequently used tag in 2004?

I found that identifying the main topic of the blog was ...
☐ Not at all easy ☐ Not easy ☐ Easy ☐ Extremely easy

Commenters

- Who was the most frequent commenter (overall)?
- Who was the most frequent commenter in 2005?

I found that identifying regular commenters in the blog was ...
☐ Not at all easy ☐ Not easy ☐ Easy ☐ Extremely easy

Popularity

- Did the blog receive comments regularly from the audience? Yes / No
- Which visualization features did you use to make such conclusion?

I found that getting a sense of popularity of the blog was ...
☐ Not at all easy ☐ Not easy ☐ Easy ☐ Extremely easy

Overall Reactions to the System

How effective was the visualization in giving you an overview of the content of the blog?
☐ Not at all effective ☐ Not effective ☐ Effective ☐ Highly effective

How effective was the visualization in giving you an overview of the community dynamics of the blog? (e.g., which entries do receive many comments, who are

the regular commenters)
☐ Not at all effective ☐ Not effective ☐ Effective ☐ Highly effective

Finding the most recent entry

- What is the title of the most recent entry in the blog?

I found that finding the most recent entry in the blog was ...
☐ Not at all easy ☐ Not easy ☐ Easy ☐ Extremely easy

Filtering by tag

- Please list the title of an entry tagged "Linux":
- Please list the title of an entry tagged "Problem solving":

I found that finding blog entries tagged by a specific keyword was ...
☐ Not at all easy ☐ Not easy ☐ Easy ☐ Extremely easy

Browsing monthly archive

- Please list the title of an entry posted in May 2005:
- Please list the title of an entry posted in September 2006:

I found that finding blog entries posted in a specific month was ...
☐ Not at all easy ☐ Not easy ☐ Easy ☐ Extremely easy

Overall Reaction to the System

How well did the visualization tool help you do things that you could do with typical blogs?
☐ Not at all well ☐ Not well ☐ Well ☐ Extremely well

Finding popular entries

- Please list the title of the most popular entry: (popular entries refer to those received many comments)

I found that finding popular entries in the blog was ...
☐ Not at all easy ☐ Not easy ☐ Easy ☐ Extremely easy

Having access to this visualization would affect my choices of which entries to read.
☐ Strongly disagree ☐ Disagree ☐ Agree ☐ Strongly agree

Why?

Finding entries commented by a specific person

- Please list the title of an entry commented by "Ronny":

- Please list the title of an entry commented by "Dan Williams":

I found that finding blog entries commented by a specific user was ...
☐ Not at all easy ☐ Not easy ☐ Easy ☐ Extremely easy

Overall Reaction to the System

How effective was the visualization in giving you information about community dynamics in the blog? (e.g., which entries do receive many comments, who are the regular commenters)
☐ Not at all effective ☐ Not effective ☐ Effective ☐ Highly effective

OVERALL

While exploring a blog, having access to an overview of the blog is ...
☐ Not at all useful ☐ Not useful ☐ Useful ☐ Extremely useful

Why?

While exploring a blog, having access to the community dynamics (e.g., which entries receive many comments) is ...
☐ Not at all useful ☐ Not useful ☐ Useful ☐ Extremely useful

Why?

What is your favourite feature of the visualization tool?

What is your least favourite feature of the visualization tool?

How well did the visualization tool help you do things that you could NOT do with typical blogs?
☐ Not at all well ☐ Not well ☐ Well ☐ Extremely well

If the visualization tool were integrated into blogs, would you be interested in using the tool?
☐ Not at all interested ☐ Not interested ☐ Interested ☐ Extremely interested

Do you have privacy concerns about the information presented in the visualization? Yes / No (e.g., if your name or comment appears in the visualization)

Why?

Please describe any other comments you may have about the visualization tool (e.g., desired functions that currently do not exist in the visualization tool).

A.2.3 Questionnaire for User Interaction Satisfaction 7.0

Identification number:

Personal information

1. Age
2. Gender   a. male   b. female
3. On a scale of 1 (beginner) to 5 (expert), how would you rate your computer skills as an end-user?
4. How many hours do you browse the web daily (on average)?

   a. Less than 1 hour   b. 1 - 5 hours   c. 6 - 10 hours   d. More than 10 hours
5. Please check all that apply:

   [ ] I read blog(s)

   [ ] I leave comments on blog(s)

   [ ] I subscribe to blog syndication(s)

   [ ] I know how to host blog(s) on my machine

Overall User Reactions

Please circle the numbers which most appropriately reflect your impressions about using this computer system.
Not Applicable = NA.

Overall reactions to the system:

Terrible	1	2	3	4	5	6	7	8	9	Wonderful	NA
Frustrating	1	2	3	4	5	6	7	8	9	Satisfying	NA
Dull	1	2	3	4	5	6	7	8	9	Stimulating	NA
Difficult	1	2	3	4	5	6	7	8	9	Easy	NA
Inadequate power	1	2	3	4	5	6	7	8	9	Adequate power	NA
Rigid	1	2	3	4	5	6	7	8	9	Flexible	NA

Screen

Characters on the computer screen

hard to read 1 2 3 4 5 6 7 8 9 easy to read NA

Highlighting on the screen

unhelpful 1 2 3 4 5 6 7 8 9 helpful NA

Screen layouts were helpful

never 1 2 3 4 5 6 7 8 9 always NA

Please write your comments about the screen here:

Terminology and System Information

Use of terminology throughout system

inconsistent 1 2 3 4 5 6 7 8 9 consistent NA

Terminology relates well to the work you are doing?

never 1 2 3 4 5 6 7 8 9 always NA

Messages which appear on screen

inconsistent 1 2 3 4 5 6 7 8 9 consistent NA

Messages which appear on screen

confusing 1 2 3 4 5 6 7 8 9 clear NA

Computer keeps you informed about what it is doing

never 1 2 3 4 5 6 7 8 9 always NA

Please write your comments about terminology and system information here:

Learning

Learning to operate the system

difficult 1 2 3 4 5 6 7 8 9 easy NA

Exploration of features by trial and error

discouraging 1 2 3 4 5 6 7 8 9 encouraging NA

Remembering names and use of commands

difficult 1 2 3 4 5 6 7 8 9 easy NA

Tasks can be performed in a straight-forward manner

never 1 2 3 4 5 6 7 8 9 always NA

Please write your comments about learning here:

System Capabilities

System speed

too slow 1 2 3 4 5 6 7 8 9 fast enough NA

The system is reliable

never 1 2 3 4 5 6 7 8 9 always NA

Correcting your mistakes

difficult 1 2 3 4 5 6 7 8 9 easy NA

Ease of operation depends on your level of experience

never 1 2 3 4 5 6 7 8 9 always NA

Please write your comments about system capabilities here:

A.3 Email Invitation

Hi,

I'm recruiting participants for a user study. The study aims to evaluate the effectiveness of an interactive tool for exploring a blog and to gain some insight into the practice of browsing blogs.

Potential participants should be familiar with browsing the Internet and have experience reading blogs. The study session will take up to 45 minutes. Participants will receive a $10 honorarium.

If you're interested in this study, please contact me at: j.indratmo@usask.ca or 966-4744.

Thank you,
Indratmo

Appendix B

Materials for the Summative Evaluation

B.1 Informed Consent Form

You are invited to participate in a study entitled "Supporting Exploratory Browsing." Please read this form carefully, and feel free to ask questions you might have.

Researchers: Julita Vassileva, Department of Computer Science (966-2073)
Indratmo, Department of Computer Science (966-4744)

The purpose of the study is to evaluate the effectiveness of tools for browsing web articles. In the study, you will use the tools and answer some questions about your experience with the tools. The estimate of the total time to participate in this study is 60 minutes.

There are no known risks in this study.

Findings from the study will be used to develop a framework for designing tools for browsing social information spaces. You will receive a $10 honorarium.

The research data will be stored on a password-protected computer system and will be available only to the investigators. Identifying information will be destroyed upon completion of data collection, and then pseudonyms will be used to refer to the participants. All data will be kept by Dr. Vassileva for a minimum of five years upon the completion of this study. If Dr. Vassileva chooses to destroy the data after the five years, the data will be destroyed beyond recovery.

Aggregate results will be used in a thesis and articles published in conferences and journals. Any information that can be linked to a specific participant will be removed or altered.

Your participation is voluntary, and you may withdraw from the study for any reason, at any time, without penalty of any sort. You may refuse to answer individual questions. If you withdraw from the study, data that you have contributed will be destroyed at your request.

If you have any questions concerning the study, please feel free to ask at any point; you are also free to contact the researchers at the numbers provided above if you have questions at a later time. This study has been approved on ethical grounds by the University of Saskatchewan Behavioural Research Ethics Board in February 2009. Any questions regarding your rights as a participant may be addressed to that committee through the Ethics Office (966-2084). Out of town participants may call collect. You may find out about the results of the study at http://madmuc.usask.ca/ or by contacting the researchers.

I have read and understood the description provided above; I have been provided with an opportunity to ask questions and my questions have been answered satisfactorily. I consent to participate in the study described above, understanding that I may withdraw this consent at any time. A copy of this consent form has been given to me for my records.

(Name of Participant)	(Date)
(Signature of Participant)	(Signature of Researcher)

B.2 Pre-Test Questionnaire

Personal information (This information will be treated confidentially)

1. Age a. below 20 b. 20 – 29 c. 30 – 39 d. 40 above
2. Sex a. Male b. Female
3. On a scale of 1 (beginner) to 9 (expert), how would you rate your computer skills as an end-user?
4. How many hours do you browse the Web daily (on average)?

   a. Less than 1 hour b. 1 – 5 hours c. 6 – 10 hours d. More than 10 hours
5. On a scale of 1 (beginner) to 9 (expert), how would you rate your general knowledge of the following topics?
   - Business and finance:
   - Travel and places:
   - Food and drink:
6. On a scale of 1 (dislike) to 9 (like), how would you rate your general interest in the following topics?
   - Business and finance:
   - Travel and places:
   - Food and drink:
7. Are you familiar with Digg (http://digg.com)? Yes No

B.3 Post-Test Questionnaire

Please select the numbers which most appropriately reflect your impressions of using this computer system.

Using this system enhances my effectiveness in finding interesting information.

Strongly disagree	1	2	3	4	5	6	7	8	9	Strongly agree

Overall reactions to the system:

Terrible	1	2	3	4	5	6	7	8	9	Wonderful
Difficult to use	1	2	3	4	5	6	7	8	9	Easy to use
Dull	1	2	3	4	5	6	7	8	9	Stimulating
Frustrating	1	2	3	4	5	6	7	8	9	Satisfying
Overwhelming	1	2	3	4	5	6	7	8	9	Manageable
Complex	1	2	3	4	5	6	7	8	9	Simple
Too slow	1	2	3	4	5	6	7	8	9	Fast enough

Overall, which system do you prefer to use? (please circle)

Arjuna Tambora Rinjani

Why do you like that system?

B.4 Email Invitation

Hi,

I'm recruiting participants for a user study. The study aims to evaluate the effectiveness of tools for browsing web articles.

Potential participants should be familiar with browsing the Internet. The study session will take up to 60 minutes, and each participant will receive a $10 honorarium.

If you're interested in participating in this study, please contact me and more details will be provided.

Thank you,

Indratmo
PhD student, Computer Science