

Time, Topic and Trawl: Stories About How We Reach Our Past

Joon-Suk Lee

Dept. of Computer Science, Virginia Tech
2202 Kraft Dr., Blacksburg, VA 24060
{joonlee, tatar}@vt.edu

Deborah Tatar

Elin Rønby Pedersen

Google, Inc. Mountain View,
CA 94035 USA
elinp@google.com

ABSTRACT

Legacy web tools attempt to build on information that users have when they originally conduct web research. In contrast, we examine the information that they have at the time when they attempt to recreate their past. We interviewed 11 non-expert users twice a week for eight weeks in their own physical and computational environments. We used both Google web histories and the prototype Research Trails system as prompts to probe how the participants viewed their past web experiences and how they reconstructed them. The Research Trails system lets users utilize information about both time and topic to help themselves remember and resume *everyday research* tasks. Based on these observations, a model of users' perceived past web activities informed the iterative refinement of the Research Trails system. The user may see a past action as belonging to multiple categories at the same time or as in different categories at different times.

Author Keywords

Information retrieval, information seeking, early research, revisitation, web search.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

The predominance of web use in everyday computing has long been reported [7] and the amount of time people spend on the Internet is only growing. Multitudes of systems and software have been developed to assist everyday web activities. While studies show that the majority of everyday web visits is re-visitation [7,40], only fractions of the tools developed to assist our divergent web activities are designed to accommodate re-visitation.

Re-visitation is an act of bringing ourselves back into the past we once had lived. As Hansel and Gretel left a trail of

breadcrumbs for themselves to follow back to their home at a later time, we are leaving traces of current web activities for future references. Hansel and Gretel were only concerned with leaving breadcrumbs. They were too shortsighted to see how they would trace the trail back once the birds had eaten all the breadcrumbs. Similarly, as a field, we have been significantly more focused on how to leave traces of our everyday web activities (i.e. Bookmarks and auto generated web histories), than how to follow the traces back into our pasts. The consequences are that the tools we build to assist people in relocating themselves back to their pasts, such as bookmarks and web histories, are built primarily for leaving traces every moment we are on the web, but lack the mechanisms to effectively use these traces to recreate the past. The important issue is that recreating the past is not a neutral act. Recreating the past serves a current purpose.

Bookmarks (and their variations) typically require users to identify how their current web visits will be referenced in the future by asking the web pages to be cataloged, tagged and given meaningful mnemonic names. Web histories blindly capture every possible footprint of the current moment. Therefore, (1) users are required to provide an abstraction over the raw data at the time of data collection (or creation) as opposed to the time of data retrieval and (2) data retrieval is regarded less important, as a matter of trivial data indexing and searching. However, while *remembering* is objective and procedural for machines, *remembering* is much more complex for people. People use the traces that they left in the past to actively reconstruct past experiences and build their own narratives around the footprints of the past.

During the past year, an inspiring discussion has been initiated among Lifelogging researchers [19,36]. Lifeloggers wear cameras that capture all their activity, in an attempt to put themselves in a position from which they can recreate the past. In the past, Lifelogging researchers have been primarily concerned with data capture. Now, Sellen and Wittaker [36] argue the need to shift focus from capture to use.

This is analogous to the argument that we are making about web histories and bookmarks. Like Lifelogging, web

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

DIS 2012, June 11-15, 2012, Newcastle, UK.

Copyright 2012 ACM 978-1-4503-1210-3/12/06...\$10.00.

activity logging is ultimately interesting because of how it allows people to relocate themselves in a recreated past. As Sellen and Whittaker advocated [36], we need to move away from “an obsession with capturing everything” and to start asking what it means to support human memory and how we can go about designing useful aids. People *recollect, reminisce, retrieve, reflect* on the past and *remember past intentions* [36]. We, following [19,36], also argue the need to investigate how people recall and reconstruct their past web experiences.

In this paper, we examine how people recollect and re-interpret their past. We interviewed 11 non-expert users twice a week for eight weeks in their own physical and computational environments. We used both Google web histories and the prototype Research Trails system as prompts to probe how the participants viewed their past web experiences and how they reconstructed them. We then developed models of users’ *perceived* past web activities by putting together the findings from the two months of observation. We iteratively used the findings in refining the Research Trails system. The Research Trails system differs from legacy web tools in that it does not try to create an interpretation of users’ behavior as the web activities occur, but instead utilizes information available at the moment users need to recreate their past.

INFORMATION RETRIEVAL, SEEKING AND WEB USAGE

Information management continues to be a fascinating and important topic in web usage. Numerous models and taxonomies seek to describe and explain different aspects of user behavior and information structure. Some start from received models based on library searches. These may emphasize management of different components, including the affective, cognitive and physical [23]. Some models point out that undirected browsing, directed browsing and semi-directed browsing [28] may entail different behaviors. Some similarly distinguish between (for example) passive browsing and active search [44]. Taxonomies of active search identify subcomponents, including fact-finding, information gathering, browsing, transactions and other [22] or variants on these categories [35].

Two recurring concepts in research in this area are goals and tasks. Sometimes research is contextualized by reference to the existence of goals and/or tasks. Sometimes it investigates the significance of particular goals and/or tasks.

In general, goals are seen as central to the explanation of behavior. However, recently, White and Roth [43] have noted that a person may engage in active search that does not have clearly defined goals. The person may engage in iterative search, but new information may influence the definition of success and completeness. White and Roth propose the need to study “search trails”, a series of web pages visited in a single web session, in order to better understand such various techniques such as information foraging [32], berry-picking [3] and orienteering [30].

Other research has focused on how people conduct particular tasks. Prior research on task-based web behaviors has focused on seeing and understanding different kinds of web browsing—whether one task across multiple sessions [21] or multiple tasks in one session [15]. However, the utility of the notion of task as an organizing structure has also been called into question. Keller *et al.* define the Information Gathering phase of web search as “a task that involves the collection of information... over multiple days...” and go on to say “It is not always clear when the task is completed and there is not always one specific answer...” (p. 7) [22].

We follow White and Roth’s suggestion to study search trails [43], and take their idea further, by investigating trails across multiple sessions over a two-month period of time in order to investigate web behaviors. We, like many other previous studies, are interested in the case in which neither the goals nor the tasks are well defined. Yet, there exists a subtle but critically significant difference between previous web behavior models and our work. We investigate how people *perceive* their past web experiences, while most previous studies aim to provide models of web behaviors based on the observations only of how people conducted the web activities originally.

PICKING UP THE TRAILS

Our main interest lies in investigating how people can pick up the trails, i.e., re-construct and develop postmortem understandings of their past web experiences to guide their future behavior, using extant and prototype tools. The current research replicates and extends an earlier ethnographic study that indicated that “ordinary” people (e.g. not information professionals) engage in what they call “research,” in ways that are distinct from traditional models of scholarly and investigative research [31]. That is, people do “research” in their everyday lives to carry out mundane tasks—planning family trips to unknown countries, or buying a new digital camera—either sporadically or in intensely focused ways. Sometimes, those activities last days, weeks or even months. Sometimes, people want to re-find some of the pages they visited previously. In previous work, we called these activities *early research* to distinguish the activity from a more scholarly and systematic approach, and to indicate a certain lack of development. Early research activities were not, as we found them, well developed. Participants seldom went back over material they had gathered, few research tasks materialized in any external representation and most insights were developed and held in people’s heads, remaining a very personal property [31].

This ethnographic study suggested a class of designs, currently embodied in a prototype Research Trails system. The Research Trails system provides an alternative kind of tool compared with bookmarks or web histories. It uses information about both topic and time to infer patterns of activities that may be useful. All of the users’ web browsing and searching activities, up to the present moment,

influence the pattern detection. In other words, Research Trails use time and topic to trawl through web pages and cluster them into putatively helpful groupings when users want to look back into their past. These clusters work as triggers that, if the system is good enough, serve to evoke a memory of a task. In an ideal world, a user would look at the grouping and say “Oh, now I remember what I was doing!” Research Trails do not require the user to identify categories in real time, nor do they purport to provide enduring and absolutely correct categories. Instead, they provide an extra layer of contextualization above and beyond simple chronology that may be useful.

BOOKMARKS AND WEB HISTORIES

Our prior work [27,32] confirmed findings widely reported in the information retrieval literature and also in pointed studies of system features such as bookmarks and web histories. In particular, we also found that context was a) crucial to the user experience and b) difficult to establish and maintain. People typically ask, “Where is all the stuff I just worked on?” or “Where was I?” [18].

Re-visiting previously viewed web pages is common. One empirical study reported that revisitation accounts for 81% [8] of web visits while an older study indicated that 58% of web browsing is revisitation [41]. These are very large percentages. Yet many studies show that people use neither bookmarks nor browser histories [1,2,17,18,43]. Arguably, bookmarks are used in only 2.7% of web navigation [42] while web histories are invoked in a miniscule 0.2% of all web page initialization requests [43].

Bookmarks are simple tools for keeping references to pages, but they require that users immediately recognize the value of a page, hindering their usefulness [8,19]. Also, bookmarks easily become overly cluttered with broken links, no longer useful sites, and dynamically generated one-time links, making maintaining current and usable bookmarks a formidable task for users [1,19].

Researchers have tried to overcome these shortcomings with different approaches. For example, different document classification techniques such as “finite mixture model” [26], “semantic treemaps” [12] and keyword matching [40] were used to automatically generate bookmark categories. Those approaches required users either to predefine document categories [26] or at least to select a category from system generated suggestions [40]. There have also been attempts to improve bookmarks by integrating bookmarks and web histories [14,19].

However, those are utilitarian approaches to fixing or improving one single aspect of current bookmarking technologies. Those approaches often neglect the structural problem that bookmarks are really nothing more than scattered fragments of past experiences and therefore that bookmarks are fundamentally inadequate in capturing the rich contextual information of everyday web activities.

Browser histories are chronological sequences of past activities and experiences. Indeed, most web browsers retain the users' browsing history. In addition to client-based history tools, users can also use server-based tools such as Google Web History. These enable users to search for entries in their web history using text queries, akin to web searches.

Though useful, existing browsing-history tools are limited by their simplicity; only chronology guides their use, and people have a poor memory for chronologies that are not tied to salient episodic events [4,10,38]. Users often elect to re-find information by issuing new web queries rather than search their history [17].

There have been attempts to improve the utility of web history through better visualization. This often takes the form of page thumbnails displayed with some meaningful structure, including path-based [14], hub-and-spoke [8], and 3D [46]. LeeTiernan *et al.* showed that clustering pages by URL similarity and temporal proximity created effective visualization [24]. Won *et al.* studied users' problems utilizing web histories to re-find pages, and used the findings to inform the design of a contextual history search tool [45]. This tool allows filtering using date ranges and gives contextual cues such as thumbnails. The Eyebrowse system records and displays users' web page visits, computes aggregate statistics, and visualizes the information for users [29].

Personalized and task-based searches have been applied with some success in capturing users' past task-based activities. Umea [20] and TaskTracer [9] improve the process of keeping files organized according to task, but require users to predefine tasks, which can be difficult. Some personalized searches assist users in re-finding previously viewed information, first by building semantic profiles from terms appearing in pages from their web history or PC, then applying these profiles to add or rank results [5,27,33].

All of these approaches are interesting, but none have yet emerged as a dominant paradigm. Moreover, these approaches embed an unsaid implication that there is one (or a fixed number of) correct ways to view past web activities. Most of these solutions prematurely identify and create the interpretation of users' pasts at the moment of data collection, even though users' perception over the same data might fluctuate over time. In the current work, our approach is to focus on providing a layer of abstraction over what would be an overwhelming amount of accumulated data to provide opportunities for users to actively and conveniently re-construct their lived pasts when they need to look back at their past.

RESEARCH TRAILS SYSTEM

The Research Trails system is similar to a personalized search in that it uses data from users' web histories to help re-find information. It attempts to cluster the history in ways that mimic what the user perceives as somehow

belonging together. “Belonging together” is a notion that is broader than topical similarity, typically adding a pragmatic notion of activity or task. It can also be used to re-establish working contexts, rather than simply find single web pages.

The Research Trails system uses two perspectives in clustering past web histories: time and topic. These in essence propose a delineation of different web activities that correspond to the task. They may be spot on in predicting the user’s notion of the task, or they may operate by provoking and reminding the user of a prior intent, goal or strategy. In this way, users can reconstruct and reanimate their past experiences on-demand.

An *activity-based perspective* focuses on inferences based on how a user interacts with data, e.g., how long a page is viewed. A *semantic perspective* focuses on the material the user worked on, e.g., how data are related to each other by text contents. We also define three main entities: *event*, *segment* and *topic*. An *event* is a page visit from a user’s activity history. A *segment* is a temporal clustering of events. A *topic* is a semantic descriptor obtained from a suitable statistical/linguistic technique. A *trail* is a semantic clustering of temporal groups of events.

Detailed algorithms for *semantic* and *activity-based* analyses and how they are used in the construction of the trails are explored elsewhere [31].

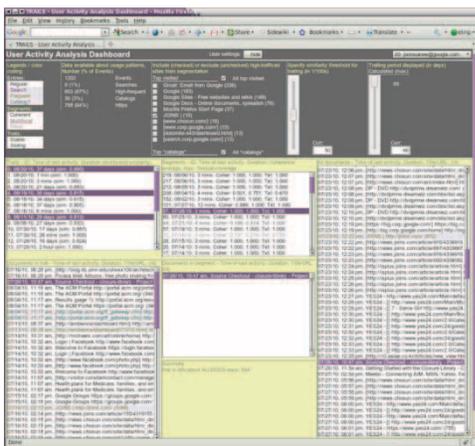


Figure 1: Research Trails Dashboard (used for reference during interviews)

THE STUDY

The work presented in this paper was conducted in the summer of 2010 with 11 participants over a period of 8 weeks.

The goals of this study included (1) testing the concept of clustering, from a phenomenologically-situated perspective [12], (2) presenting and formatively evaluating information from a prototype system to support refining and revisiting tasks, and (3) developing general design implications for tools targeted to people conducting everyday web research.

In particular, the Research Trails system proposed an operational definition of clustering for evolving tasks in a

non-expert user population, but we did not know how well it would correspond to the felt experience of participants.

The study had three parts, each involving two interviews/week and continuous gathering of web histories. Phase I (two weeks) started with a questionnaire and focused on normal participant web use; Phase II (four weeks) continued an interest in normal web activity but also used the nascent Research Trails system as a prompt for interview discussions; Phase III (two weeks) obtained participant reflection via a final questionnaire, and conducted a final interview based on the these reflections.

Participants did not use the Research Trails system on their own; it was used only during contextual interviews. Instead, in the tradition of design experimentation and paper prototyping [34], we gave them the information as manifested in Figure 1, and asked them to discuss how well the algorithm grouped their activities as compared to the simple chronologically-based web histories that we used as discussion prompts in the first phase.

Participant Recruiting

Participants were recruited from the large Usability Study Participants pool at a high technology company in Silicon Valley. Participants were pre-screened based on four criteria; proximity (to enable face-to-face interviews), and (lack of) computer expertise were necessary for the conduct of the study. Fulltime employment and at least a three-month hiatus in study participation were requirements of the company. In particular, search experts and computer engineers were excluded. Eighty-eight people responded to invitation emails. People were ruled out who intended to take long vacations during the eight weeks of the study and who did not, on consideration, think that they would be available for two interviews a week. Ultimately, 11 participants were selected for the study.

Researchers were given permission to investigate participants’ web histories using a server side history tool (Google Web History). Participants were asked to pursue their web activities as normal.

Participants were compensated at the end of each phase of the study with either message coupons or gift certificates. Among the 11 participants, one was a male and 10 were females. This imbalance was a result of the demographics of the company where the research occurred. Although it was a high tech company, the non-technical staff from which we recruited was overwhelmingly female. Participant employment included day care teacher, administrative assistant, human resource specialist, and so forth. Participants ranged in age from 25 to 35 ($M = 31$ $SD = 3.50$) One participant chose to drop out of the study during phase 2 and two other participants were not able to do the final interview due to conflicts in their schedules. Audio recordings of all interview sessions were collected and analyzed. In total, 136 interviews and 21 questionnaires plus web history logs provide the basis for analysis.

Data Collection & Analysis

During the initial interview, we helped participants to install and enable Google Web History on their web browsers. Google Web History records all web browsing histories on personal computers while it records only search terms performed on Google search and first level page visits from the search result if those activities were done on mobile phones. Participants were told how to manually delete any web histories they did not want to share and were reminded before the researcher looked at any information in subsequent interviews. Then we asked participants if they had any on-going non-work-related activities, if any of those activities involved web searching/browsing, and if they had any information intensive hobbies or if they were members of any online communities. During the consecutive interviews, we visited participants either at their homes or at their work places depending on participants' preferences. We went over the Google Web History entries one by one and asked participants why they visited the web site, what they were doing before, during and after visiting the site, how the visited web site was related to their daily lives, and how long they stayed on the site. Since we conducted the interviews twice a week, people had a relatively easy time remembering their past web activities and were able to give very detailed descriptions of their past visits.

Throughout the study periods, we were able to identify different topics and tasks participants were pursuing. A sampling of topics that participants reported as interests that they pursued using the web is shown in Table 1.

In the course of the study, some of those activities ended while new sets of activities emerged.

During the second phase, in addition to going over the past Google Web History, participants were introduced to the Research Trails system through the dashboard shown in Figure 1. We did not explain how the system came up with the groupings, but asked the participants if the groupings of web pages meant anything to them.

Participant Pseudonym	Identified Web Tasks and Tasks Supported by Web Use
Allison	<ul style="list-style-type: none"> - Participate in 3-year-old daughter's preschool activities - Research extra-curricular activities for 3-year-old daughter - Volunteer work for two non-profit organizations (consultation service for first time parents) - Research a potential diaper store business - Plan 3 different trips next year - Plan new car purchase - Knitting
Martha	<ul style="list-style-type: none"> - Research a rare disorder her son has

	<ul style="list-style-type: none"> - Research activities for two children - Creating baby albums - Research for different church events - Follow hockey - Research hotel information for her niece's wedding - Plan two family trips - Plan a business trip
Kelly	<ul style="list-style-type: none"> - Search for artists-in-residence program for her husband - Search for different local incidents (shootings, trials, a riot) - Research on gardening - Research on house remodeling /re-decoration - Search for activities for one-year old son
Julie	<ul style="list-style-type: none"> - Plan four trips - Explore urban (San Francisco) activities - Plan camping and hiking - Yoga
Nicole	<ul style="list-style-type: none"> - Volunteer work for Humane Society - Volunteer work for Make-A-Wish foundation - Participate in Yahoo dog owners group - Plan Machu Picchu trip - Plan winery visits - Read celebrity gossip - Search info & spoilers for TV shows and movies - Emails from a special mailing list (She gets 5 points for clicking a link from her email. Points can then be redeemed for gift items.)
Michelle	<ul style="list-style-type: none"> - Games on Facebook - Surf Craigslist for free stuff - Search price information for stuff to sell

Table 1: Web Tasks Identified by Selected Participants

During the four weeks, we tested the system with different parameter values for activity-based temporal clustering algorithms and semantic analysis algorithms. In some cases, the system drew surprised reactions from participants, capturing seemingly unrelated but indeed very related web activities under one task successfully.

FINDINGS

Research vs. Non-Research (Project vs. Non-Project)

During the interviews, participants referred to some of their web activities as “projects” or “research work.” When asked what they meant by these terms, participants included everyday web activities with the following characteristics: (1) short bursts of web searching and browsing activities that (2) could potentially spread over multiple sessions; and (3) complex and information-intensive tasks that (4) might have an array of sub-goals. The set of behaviors participants referred to as “projects” or “research” was analogous to the behaviors we characterized as “everyday web research” in our previous study [31]. A common example of “research,” as the users used the term, was vacation planning.

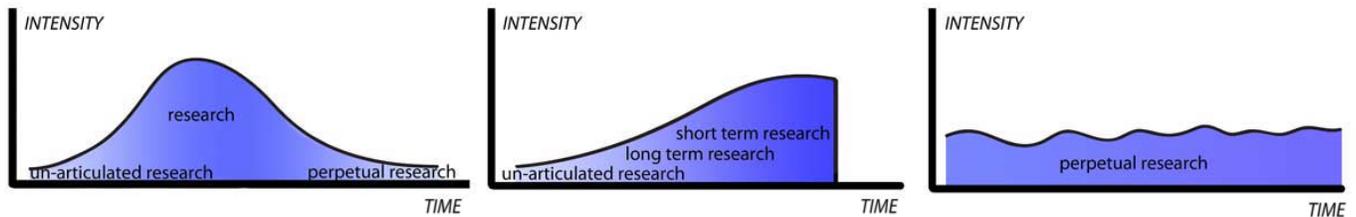


Figure 2: Idealized Temporal Intensity of Kinds of Web Activities

Participants also viewed some of their web activities as non-research or non-projects. Those activities typically either did not have clearly set goals, or filled a spontaneous need for a specific piece of information. Browsing through newspapers or blog sites was generally considered as non-project; so too were one-time information seeking activities such as searching for the phone number of a restaurant.

Vacation planning and one-time restaurant phone number lookup are both goal-driven activities, yet most participants described the former as project and the latter as non-project. They seemed to be influenced by 1) the extent of the information searching activities, 2) the relevance between the collected information and the decision-making process, and 3) the intensity of goal finding activities. Projects usually had multiple goal-oriented sub-tasks that focused on decision-making rather than finding a single correct answer. For instance, to plan a Hawaii trip, a participant needed to book a room at a hotel, reserve plane tickets, obtain rental car information and find different local events or activities to attend. Information was collected in order to make interrelated decisions and involved collecting information on alternatives.

Six categories of web activities

We identified six kinds of everyday web activities, which, following the participants, we divide into projects and non-projects. Non-projects include *Fabric of Life* and *one-time info seeking/browsing* activities. Projects include *unarticulated research (early research)*, *short-term research*, *long-term research*, and *perpetual research*.

Fabric of Life: Most participants visited the same set of web sites regularly, often daily. They referred to these visits as their “daily routine”, “a fabric of their life” or “habit.” Participants invariably said that they did not have any fixed goal nor a need for any particular information in visiting these sites. They browsed thorough the web sites to get an update on things they are interested in. The information found on the sites is consumed right away rather than being filed or categorized.

For instance, a participant who has an avid interest in Las Vegas, visits a couple of Las Vegas related blog sites on a regular basis. During the interview, he referred to those blog sites as “my Vegas blogs.” Clearly those several different blog sites were clustered under a “Vegas” theme from the participant’s perspective. While the time he spent on each of those sites was relatively small compared to

most other research-like web activities, the cumulative time spent on these regular sites was very large.

One-time info seeking/browsing: User explanations of one-time searches featured the relationship between these searches and social events. For instance, Lindsay, one of our study participants, was watching a baseball game on TV and having a conversation with her husband. When her husband asked if Jose Canseco had been a good baseball player, she searched for his career records on the web. She did not have an enduring interest in either baseball players’ career records or Jose Canseco. When she found what she was looking for, she mentioned it to her husband. The information was consumed in the moment and ended when the answer was found. If participants could not find the information, they often ended search activities after a few trials. Very little time was spent on one-time info seeking.

Projects: In contrast to *Fabric of Life* or one-time info seeking/browsing, other tasks were described as projects or research. The consensus was that these were tasks, assignments, or investigations that could be either short-term or a more significant undertaking. They usually consisted of information-gathering activities, including collecting information from various places on the web, from traditional media including books and magazines, and from social interactions with other people. Projects often, but not always, require organizing or reproducing gathered information into a form that can be referred back to at a later time.

When we asked participants to list some of their on-going projects, they differentiated between them, leading to a sub-categorization:

Long-term and short-term research: A transparent difference among the various research projects was the lifespan of the activity. While some projects tended to last months or years, others only lasted days or weeks. Hence we categorized them as long-term and short-term research. In addition to duration, participants identified qualitative differences between long-term and short-term projects.

Short-term projects had more pre-defined goals. For instance, Martha was planning a family trip to Michigan for her friend’s wedding. The destination, duration, start and end dates for the trip were already predefined, and Martha only needed to research clearly defined sub-tasks: finding a hotel, booking airplane tickets, reserving a rental car, and planning extra activities for herself and her family.

In contrast, Candace had a family vacation trip project she had been planning for a while. She was planning the trip around Thanksgiving or Christmas breaks, but where to go was not decided, and what to do at the destination location remained unthought-through. She had a vague idea that she would prefer to go somewhere in the South Pacific but she also needed to coordinate the trip planning with her friend's family.

Long-term projects were correspondingly more susceptible to *topic sliding* (slight change of the theme during the research process) [31] than short-term projects. Candace started by researching different islands in the South Pacific but shifted her focus to the continental US once she found out that her friend did not want to travel so far.

Perpetual research: Perpetual research resembles the Fabric of Life activities in that it may be very important to the user and also has no end, but it differs in that it has clear goals and does not constitute a regular activity. Instead, actions taken for this type of research were interspersed among other web-activities in infrequent and irregular intervals. For instance, Martha has a two-year-old son who has been diagnosed with a rare medical disorder. She is an active member of a nonprofit organization dedicated to supporting patients affected by the disorder. She sometimes spends time answering some of the questions posted on the forum or reads new articles related to the disorder. Topic sliding is less evident in perpetual research since the research activities are more or less stable over a long time.

Unarticulated research (early-research): In distinguishing short-term, long-term and perpetual research in our participants, we also found a revised interpretation of early research, as *unarticulated*. It comprises activities still in the incipient phase, in which people either do not perceive their actions as a part of the research or only sense them as research without any associated trail. However, since any one-time web search or serendipitous information encounter has a potential to grow into a project, unarticulated research may be retrospectively identifiable.

Lindsay is a passionate cook. She visits different cooking blogs and watches food shows daily. She said she has not repeated a recipe in the past two years. During the study period, we saw one web search entry on the mercury levels of tuna fish that seemed to be part of the cooking cluster. However, when asked about why she was interested in this, she replied that she might get pregnant some day and she wanted to know if tuna was safe for expectant mothers. In the final interview, she said she had decided to have a baby, and was now conducting intensive research on pregnancy. This example shows how un-articulated research grows into a fully formed research work and also how it may appear in a context that is difficult for others to distinguish as a beginning.

We investigate how people *perceive* their past web experiences, while most previous studies aim to provide models of web behaviors based on the observations only of

how people conducted the web activities originally. For instance, Choo et al. showed how information gathering activities progress from undirected viewing to conditioned viewing [6]. It is important to note that what we are presenting in this section is not a *web behavior model*. Instead, we establish that, viewed retrospectively, a single browsing/searching behavior can be interpreted differently by the person who conducted the original search, and that their interpretations fluctuate overtime. In the next section we discuss one way to use this fluctuating multiplicity of interpretations in the design of a system to help that user.

Research Trails Clustering

All participants noticed that some Trails reflected different tasks they had been engaged in. Other Trails were less successful. Research Trails was highly successful in the many cases that did not involve difficult grouping (e.g. where topic slide was minimal and/or the research maintained temporal continuity). There were also many successes in inferring non-obvious but correct links, as exemplified in Case 1 and Case 2.

Case 1 – Diaper store.

Allison, a 31-year-old female participant, has a lifetime goal of opening a diaper store in the Mountain View area. She constantly visits various online diaper and baby product stores to gather information about different merchandise. She also actively searches for information on the process of incorporating a business as well as on real-estate prices and locations in the area. Often, she looks up real-estate information right before or after browsing merchandise information on the web. The Research Trails system was able to group those two seemingly unrelated topics under one trail.

Case 2 – High school alumni & basketball player.

One day during the second phase, Lindsay, a 27-year-old female participant, went to her college and high school alumni sites to update her status information. When she was done updating her profile on both sites, she spent time browsing through status information on different people with whom she went to school. She found out that one of her friends is now living somewhere in Europe and has a boyfriend who is a professional basketball player. Curious about the player, she looked up his name on Google. The system was able to capture browsing alumni site activities and the one-time search of the basketball player into one group.

However, the Research Trails system failed to capture participants' past activities in two ways. It sometimes failed to group all the related histories under one trail, resulting in multiple trails for a single task (false negative). In other times, the system put unrelated web pages together under one trail, resulting in trails that reflected multiple activities (false positive). In general, most participants were more forgiving for false negative results than false positive ones. When they saw a group that, as far as they were concerned,

consisted of meaninglessly clustered pages, they became impatient. This often happened when the user had engaged in parallel browsing behaviors [15] and can be reduced by putting more weight on topical clustering algorithms as compared to the temporal ones.

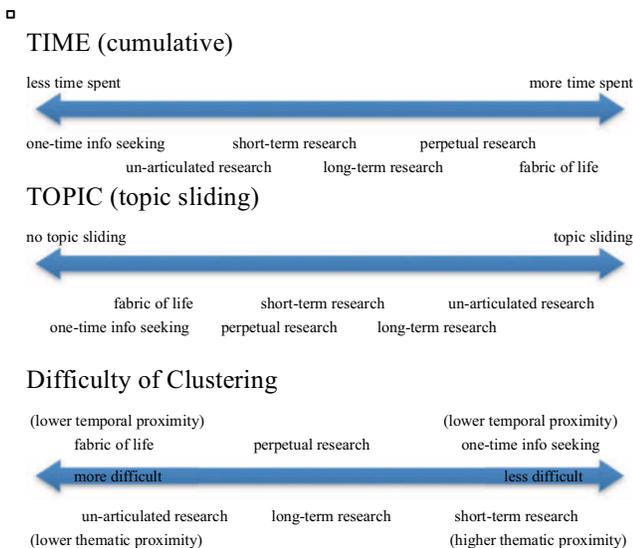


Figure 3: Web Activity Continua

DISCUSSION

Transforming the Notion of Early Research

In prior work, we had the notion of supporting *early research* in a certain class of users—those who are not information handling experts. We had previously found that this class of user engaged in activities for their personal consumption that were fragmented in time, subject to topic sliding, and often suffered from premature attempts to categorize or determine goals [31].

Through deeper interaction with this class of users, we observed that (1) *early research* per se does not stand up as a unique and distinct class of behavior; however, (2) the activities we originally folded into the concept of early research suggest potentially powerful avenues of support for a wide range of research activities over different temporal durations. The phenomena associated with early research exist on continua that we now conceptualize as including *unarticulated research*, *short-term* and *long-term research*, and *perpetual research*. These are distinct from *fabric of life* and one-time web based activities because they are pursued with respect to future action rather than being consumed entirely in the moment, but tied to research activities in other respects.

Each of these six kinds of projects and non-projects should be viewed as nuanced characteristics of different overlapping behaviors. Figure 3 shows how each of the six behavioral categories varies along the continua of *time* and *topic*. It also shows how they are projected to vary along the difficulty of attempting to infer the grouping of web pages

that will most help the user’s tasks. Of course, habitual web browsing does not usually have clearly defined goals, so clusters that infer user groupings of such web activity will almost certainly be less useful. In contrast, it is relatively easy to infer short-term research clusters where the sub-goals and temporal extent are directed and confined, harder to infer long-term research clusters, and yet harder to infer perpetual research clusters. Unarticulated research clusters are by definition nascent and can only be identified depending on how they emerge after the initial foray.

These categories also differ via the *intensity* or temporal contiguity with which they are pursued, as shown in idealization in Figure 2. This change in intensity means that changing the algorithmic focus between time and topic will vary the likelihood with which different types of activity are correctly detected.

Phenomenologically-situated take on web-behaviors

Although we presented taxonomical descriptions of various web-activities and transitional models of web behaviors, finding out the meanings of what people do on the web is always subject to different interpretations and sensemaking. (1) The same set of past web pages on two individuals web-histories does not necessitate the same interpretation. For instance, planning a trip to a Comic-Con event was just one small part of one participant’s lifelong trip-planning project while most participants considered planning a trip to a particular place as a standalone project. (2) The same set of past web pages may entail more than a single interpretation from the same user. For Allison, who visits diaper related web sites quite often, one visit entailed three projects: buying diapers for her 3-year old daughter; purchasing a gift for a baby shower; and browsing merchant related information for a project to start a diaper store. Moreover (3) the same set of past web pages could hold different interpretations at different times for the same user. For instance, browsing an article on pregnancy and air travel holds characteristics of an un-articulated project. However, the un-articulated pregnancy project could possibly evolve into long-term research or short-term research as the participant decides to seriously consider having a baby or as she carries out more intense research on pregnancy-related information.

The four components of early research that we identified in previous work [31]—fragmented process, personal consumption, topic sliding and premature structure—are not traits unique to early research, but rather properties common to a wider variety of web activity.

Under this conception, classes of web-based activities are not descriptors for a fixed set of behaviors, but rather momentary characteristics of constellations of web behaviors. The meanings of web activities are constantly created, altered and ended by the social contexts in which they are embedded. A project pertaining characteristics of a certain category could easily evolve into a different kind.

Good Enough Clustering?

This view of user activities as constituted by transformation, evolution, and continua makes it very clear that no automatic system will ever be able to cluster web pages unerringly. Indeed, no person can cluster his/her own pages unerringly.

The question raised is whether combining topic and time information can represent a significant improvement over using just one or just the other—is it “good enough” to be useful? On one hand, given the low rates of bookmark and web history usage, it would be hard for this approach to be worse. On the other, even if it was developed sufficiently to be released as a user system, it might fail in actual adoption.

We now live in a world in which we work with many “good enough” systems. For example, Google searches do not always produce the desired information. They just produce it often enough to have changed the world of computing. People are able to handle their disappointment when a search does not work, and decide whether to pursue it or not.

The current study has produced some successes, and some evidence of noise in the system. As it was illustrated in Case 1 and Case 2, the trails produced by the system were often “Good Enough” interpretations of users’ past web experiences. On multiple occasions, the system produced trails of web histories corresponding to users’ perception of the past experiences, to the amazement of the users. We are currently tuning the system to minimize the noise that users found least pleasant. As we move forward, one possibility is to hand the tuning over to the user, so that the user can move between different clustering possibilities depending on why s/he is looking at history as they try to gaze back their past experiences. At one end of the spectrum, the user could minimize the number of irrelevant pages that are clustered, while dividing many conceptually single groupings into multiple clusters. At the other end, the user could maximize the clustering of pages that belong to a single group. We could also attempt to provide “détentes” for different kinds of web projects. An easily browsable user interface (explored in a separate project [26]) could alleviate the need for perfection in the groupings.

CONCLUSION

Most everyday web activities do not exist in isolation but are constantly caused, altered and ended by the social contexts in which they are embedded. While web browser histories are trails of our past activities and experiences, current browser histories invariably plot our past web events only against time in chronological order. Even though chronological ordering of past web events is an easily implementable way to represent the past from a system’s point of view, people tend to perceive their past in a more nuanced way. Chronological ordering represents too many events at the same level of abstraction, without differentiation. The result is an overwhelming and counter-intuitive prompt to reconstruct past experiences in

unfamiliar terms. Bookmarks identify topic areas, but also require considerable management and lose context.

We found six different ways that people categorized their past activities overall, but at any given time a particular person could change his/her characterization. People viewed their past activities through multiple, flexible lenses. They may see many trails in past activities. In particular, their perception of the meaning of their past is dependent upon what they will be doing in the future.

Both time and topic are concepts that people used to describe how past activities. We can record information that allows us to infer clusters of web pages that may remind users of the tasks that they are trying to resume. The Research Trails system is an example of a system that “do(es) not ‘capture experiences’ but instead provide(s) cues that might trigger different kinds of memories (p. 77) [36].” It is an example of a system that is not only “open to multiple interpretations” [37], but also designed to *encourage* multiple interpretations from the onset.

The Research Trails prototype system shows promising results in identifying and clustering related events into trails and segments. Even though the system also reveals some problems, system development and longitudinal works will enable us to further test and develop this promising approach.

ACKNOWLEDGMENTS

We thank all the participants for their availability and engagement in this project.

REFERENCES

1. Abrams, D. and Baecker, R., How people use WWW bookmarks, *Ext. Abstracts CHI 1997*
2. Aula, A, Jhaveri, N, and Käki, M., Information search and re-access strategies of experienced web users, *In Proc. of WWW 2005*
3. Bates, M.J., The design of browsing and berrypicking techniques for the online search interface. *Online Review*, 1989. **13**(5): p. 407-424.
4. Brown, N.R., Rips, L.J., and Shevell, S.K., The subjective dates of natural events in very-long-term memory. *Cognitive Psychology*, 1985. **17**: p. 139-177.
5. Chirita, P.A., Firan, C.S., and Nejdl, W., Personalized query expansion for the web, *In Proc. of Research and development in information retrieval*. 2007
6. Choo, C.W., Detlor, B., and Turnbull, D., Information Seeking on the Web: An Integrated Model of Browsing and Searching. American Society for Information Science. 2000.
7. Cockburn, A. and McKenzie, B., What do web users do? An empirical analysis of web use. *International Journal of Human-Computer Studies*, 2001. **54**(6)
8. Cockburn, A., et al. WebView: A Graphical Aid for Revisiting Web Pages. *In Proc. of OZCHI*. 1999.

9. Dragunov, A.N., et al., TaskTracer: a desktop environment to support multi-tasking knowledge workers, *In Proc. of IUI '05*
10. Eysenck, M.W. and Keane, M.T., *Cognitive Psychology*. Psychology Press. 2010
11. Feng, Y. and Börner, K., Using Semantic Treemaps to categorize and visualize bookmark files. *In Proc. of SPIE - Visualization and Data Analysis*, vol. 4665. 2002
12. Harrison, S., Tatar, D. and Sengers, P. The Three Paradigms of HCI. *alt.chi Presentation, at CHI 2007*
13. Hascoët, M., A user interface combining navigation aids, *In Proc. of Hypertext and hypermedia*. 2000
14. Hightower, R.R., et al., Graphical multiscale Web histories: a study of padprints, *In Proc. of Hypertext and hypermedia*. 1998
15. Huang, J. and White, R.W. Parallel browsing behavior on the web, *In Proc. of Hypertext and hypermedia*. 2010
16. Jones, W., Bruce, H. and Dumais, S., Keeping found things found on the web, *In Proc. of Information and knowledge management*. 2001
17. Jones, W., et al., H., Once found, what then? A study of “keeping” behaviors in the personal use of Web information. *In Proc. of the American Society for Information Science and Technology*. 2002.
18. Kaasten, S. and Greenberg, S., Integrating back, history and bookmarks in web browsers, *Ext. Abstract CHI '01*
19. Kalnikaite, V., Sellen, A. and Whittaker, S., Now let me see where I was: understanding how lifelogs mediate memory, *In Proc. of CHI 2010*
20. Kaptelinin, V. UMEA: translating interaction histories into project contexts. *In Proc. of CHI 2003*
21. Kay, B., and Watters, C. Exploring multi-session web tasks. *In Proc. of CHI 2008*
22. Kellar, M., Watters, C., and Shepherd, M. A Goal-based Classification of Web Information Tasks. *In Proc. of the American Society for Information Science and Technology*, 2006. **43**(1): p. 1–22.
23. Kuhlthau, C., Inside the search process: Information seeking from the user's perspective. *Journal of the American Society for Information Science*, 1991. **42**(5)
24. LeeTiernan, S., et al., Two methods for auto-organizing personal web history, *Ext. Abstract CHI '03*
25. Li, H. and Yamanishi, K., Document classification using a finite mixture model, *In Proc. of Annual Meeting on Association for Computational Linguistics*. 1997
26. Liu, J., Hong, P.J. and Pedersen, E.R., Research trails: getting back where you left off. *In Proc. of WWW '10*
27. Luxenburger, J., et al., Matching task profiles and user needs in personalized web search, in *Proc. of Information and knowledge management*. 2008
28. Marchionini, G.M., *Information Seeking in Electronic Environments*. 1995: Cambridge University Press.
29. Moore, B., Kleek, M.V. and Karger, D., *Eyebrowse*. Available from: <http://eyebrowse.csail.mit.edu/>.
30. O'Day, V. and Jeffries, R. Orienteering in an information landscape: how information seekers get from here to there. *In Proc. of CHI 1993*
31. Pedersen, E.R., Gyllstrom, K. Gu S., and Hong, P.J. Automatic generation of research trails in web history. *In Proc. of IUI '10*
32. Pirolli, P. and Card, S.K., Information foraging. *Psychological Review*, 1999. **106**(4): p. 643-675.
33. Qiu, F. and Cho, J., Automatic identification of user interest for personalized search. *In Proc. of WWW 2006*
34. Sefelin, R., et al., V. Paper prototyping - what is it good for?: a comparison of paper- and computer-based low-fidelity prototyping, *Ext. abstracts CHI '03*
35. Sellen, A., Murphy, R., and Shaw, K.L. How knowledge workers use the web. *In Proc. of CHI 2002*
36. Sellen, A. and Whittaker, S., Beyond total capture: a constructive critique of lifelogging, *Communications of the ACM*, 2010.
37. Sengers, P. and Gaver, B., Staying open to interpretation: engaging multiple meanings in design and evaluation. *In Proc. of DIS 2006*
38. Shum, M. S., & Rips, L. J. (1999). The respondent's confession: Autobiographical memory in the context of surveys. In Sirken, M., et al. (Eds.), *Cognition and survey research* (pp. 95-109). Wiley.
39. Staff, C. and Bugeja, I., Automatic classification of web pages into bookmark categories, *In Proc. of Research and development in information retrieval*. 2007
40. Tauscher, L. and Greenberg, S., How people revisit web pages: empirical findings and implications for the design of history systems. *Int. J. Hum.-Comput. Stud.*, 1997. **47**(1): p. 97-137.
41. Tauscher, L. and Greenberg, S., Revisitation patterns in World Wide Web navigation, *In Proc. of CHI 1997*
42. Weinreich, H., et al., Off the beaten tracks: exploring three aspects of web navigation, *In Proc. of WWW 2006*
43. White, R. and Roth, R.A., *Exploratory Search: Beyond the Query-Response*. Morgan and Claypool Publishers.
44. Wilson, T.D., *Information behavior: an interdisciplinary perspective*. *Information Processing & Management*, 1997. **33**(4): p. 551-572.
45. Won, S., Jin, J., and Hong, J. I., Contextual web history: using visual and contextual cues to improve web browser history. *In Proc. of CHI 2009*
46. Yamaguchi, T., et al., On a web browsing support system with 3d visualization, *In Proc. of WWW 2004*