

# YouPivot: Improving Recall with Contextual Search

Joshua Hailpern<sup>1</sup>, Nicholas Jitkoff<sup>2</sup>, Andrew Warr<sup>2</sup>, Karrie Karahalios<sup>1</sup>, Robert Sesek<sup>3</sup>, Nikita Shkrob<sup>4</sup>

<sup>1</sup>University of Illinois  
Urbana, IL USA 61801

<sup>2</sup>Google  
Mountain View, CA USA 94043

<sup>3</sup>Boston University  
Boston, MA USA 02215  
rsesek@bu.edu

<sup>4</sup>University of Waterloo  
Waterloo, Ontario, Canada N2L 3G1  
nshkrob@uwaterloo.ca

## ABSTRACT

According to cognitive science literature, human memory is predicated on contextual cues (e.g., room, music) in the environment. During recall tasks, we associate information/activities/objects with contextual cues. However, computer systems do not leverage our natural process of using contextual cues to facilitate recall. We present a new interaction technique, **Pivoting**, that allows users to search for contextually related activities and find a target piece of information (often not semantically related). A sample motivation for contextual search would be, “what was that website I was looking at when *Yesterday* by The Beatles was last playing?” Our interaction technique is grounded in the cognitive science literature, and is demonstrated in our system **YouPivot**. In addition, we present a new personal annotation method, called **TimeMarks**, to further support contextual recall and the pivoting process. In a pilot study, participants were quicker to identify websites, and preferred using YouPivot, compared to current tools. YouPivot demonstrates how principles of human memory can be applied to enhance the search of digital information.

## Author Keywords

Search, Human Memory, Contextual Cue, Interface, Contextual Search, Pivoting, Pivot, Time Mark

## ACM Classification Keywords

H5.2. Information interfaces and presentation: User Interfaces.

## General Terms

Design, Human Factors

## INTRODUCTION

“What was that PDF I was looking at yesterday? The name escapes me, but I know that AC/DC was playing when I was reading it,” or “I know that I read it during that meeting last week.” Many people have had moments like this. They occur when trying to retrieve a document, website, or file, when the semantic information (name, url, system path, etc.) cannot be recalled, but environmental factors can. A user may remember the song that was playing, their

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

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

Copyright 2011 ACM 978-1-4503-0267-8/11/05...\$10.00.

physical location, a phone call that interrupted them, or even other files or websites that were open at the same time. These temporally related activities are referred to as **contextual cues**.

*Contextual cues* are a key component of human memory. Consider the advice given when someone loses their car keys: “retrace your steps since the last time you know you had them.” People are told to remember the context, and often the lost item is uncovered. We can define a *contextual cue* as any event/object that has a temporal relationship to a search target. For example, consider a Twitter message (a “Tweet”) that a user sent while working on a document. If she searches for that Tweet in order to locate her document, the Tweet is a *contextual cue*. We define search based on contextual cues as **Contextual Search**<sup>1</sup>. It is important to note that a *contextual cue*, and thereby the contextual search, does not have to match semantically to the search target. Consider the earlier example of searching for a song playing while a document is open. The song’s name and content is unlikely to be a semantic match to the document the user is working on.

While using contextual cues is a natural method of recall, [5] modern computers (and their extensive search applications) do not yet support this form of contextual search. Moreover, research in cognitive science has shown that leveraging context improves speed and accuracy in recall tasks [3]. At present, personal file history search is limited to file metadata (e.g., file name, last date accessed, file size, file type and sometimes file content). If a user does not know this information, she cannot search for and find a file. Without a form of contextual search, users cannot leverage the natural process of recall. We do not know how pervasive the desire for contextual search on computers is or could be, because modern systems do not support it. Working from the extensive theoretical background on memory, it stands to reason that search based on contextual cues will enhance recall and computer usability.

We present **YouPivot**, a contextual history based search tool that bridges the gap between the literature on human memory/recall and search. YouPivot leverages our natural method of recall by allowing a user to search through their digital history (e.g., files, URLs, physical location, meetings, and events) for the context they *do* remember. The user can then **Pivot**, or see everything that was going on

---

<sup>1</sup> While the term “contextual search” is used in the IR community (e.g. query context), we draw our usage from the psychology research on human memory, which relies on contextual cues.

while that context was active. Further, YouPivot displays a visualization of the user's activity, providing another method for finding context. By visually browsing their own computer usage, a user can find key moments in time by identifying "landmarks" of activity (e.g., browsing Facebook while working in Photoshop). Lastly, YouPivot introduces a new annotation method for contextual recall called **TimeMarks**, in which a user marks a moment in time as being important. This effectively leaves a temporal landmark for later contextual recall. Because YouPivot logs a user's personal activity, TimeMarks effectively bookmarks all the user's activity at that moment (open webpages, files, songs, physical location etc.) for easier recall.

We discuss how YouPivot builds on existing tools for facilitating search. We then present YouPivot's features and implementation, followed by a preliminary pilot study comparing YouPivot with current tools. We conclude by discussing future work, features, and experiments.

### **CONTEXTUAL MEMORY AND RELATED WORK**

We describe the underpinnings of human memory, how we facilitate and leverage it, and describe how our work builds upon and extends tools that examine searching/collecting a person's digital data/activities.

#### **Human Memory and Contextual Cues**

"It would be hard to overstate the importance of understanding the profound influence of environmental context on human memory." [5] Consider the difficulty in recalling a person's name when we encounter them in an unexpected situation. Yet by "placing" their face in the original context, we are able to more easily recall the individual's name and personal information about them [30]. Human memory functions by activating these contextual cues [5]. The probability that the human mind recalls a piece of information is a function of the strength of the association between the context used for recall, and the target information [5]. Therefore the richer the set of contextual cues available to an individual, the higher the probability of successful recall.

Consider the example of two groups of students being taught material in room A. If tested on the material, those students who take the exam in the same room (A) have better recall and perform better, than students who take the exam in a different room (B) [5,30]. In 1979, Smith demonstrated that recalling even the context of the room can improve performance on a recall test [31]. Additional research has shown that as we increase the number of contexts in which information is acquired (studying for a test in 5 rooms vs. 2 rooms), performance on recall tasks also improve. This performance increase is due to the influence of room context on recall, or the ability of individuals to remember information based on their environment. The greater the variety of environmental artifacts (e.g., room setup, music playing, etc.) the greater number of "links" or associations that human memory can leverage during recall [4,15].

In short, contextual cues aid human memory; users can remember events/items in their environment, even if they cannot immediately remember the item in question. By following this preverbal "trail" of cues, we remember. And in the computer age, people more often recall non-title information (e.g. associated events) [6]. As our digital lives become increasingly complex and dispersed [3], facilitating search based on contextual cues could be one way to improve information recall in our digital lives.

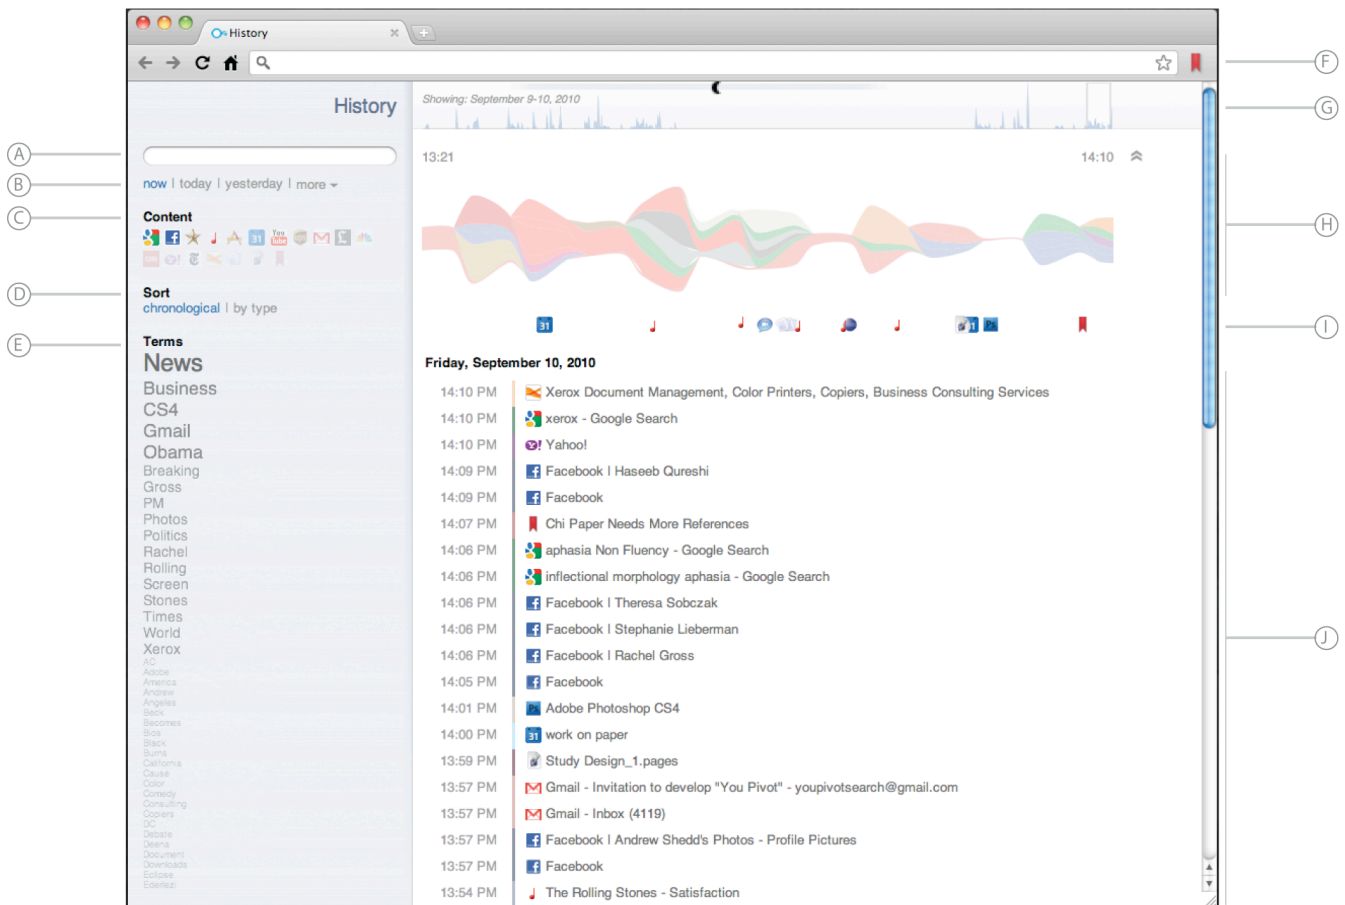
#### **Data Aggregation and File Management**

To create a rich set of cues to draw upon, computer systems must collect disparate sets of user data. One notable project in the domain of data aggregation and file management is MyLifeBits [14]. This groundbreaking work challenged researchers to explore computers that have the capability to store and archive all the data in which we encounter, perceive, and interact. MyLifeBits demonstrated that computers, simply through data collection, can provide a meaningful narrative to our digital and physical lives. A major contribution was developing a collection infrastructure to manage and store the diverse set of data MyLifeBits collected, while also providing a user interface (UI) to facilitate browsing of this data. However, this project did not explore the rich set of interactions needed for contextual search.

MyLifeBits was not the only foray into new techniques to store and represent personal data activity. Poyozo [26] sought to aggregate all of a user's web-activity into one central location. Lifestreams [13], TimeScope [28] and others [25,27,29] have explored displaying and organizing personal activity (both digital and real-world) chronologically, rather than the traditional directory-based organization method. DejaView takes [24] another temporal-based approach by treating desktop activity as a video that can be replayed. While all of these approaches advance personal file organization, they do not, nor is it the purpose of these systems to, address a user's need to search for items based on semantically unrelated contextual cues.

#### **Other Search Techniques**

Other research has examined the creation of robust techniques for searching through personal data. Haystack [23] implements a "fuzzy text" search allowing users to locate semantically related files/events. Balakrishnan [1] and others [2,22] examined search and organizing based on semantically identified tasks/activities. Similarly, Connections [32] allows a user to find semantically related files that were open at the same time. Horvitz identifies key landmarks in a user's history based on calendar activity [20]. Jensen has explored document provenance as a technique for tracking file relationships [21]. Most notably, a number of approaches provide additional related (e.g. semantic) meta-data to files or websites, that is not in the file/website itself (Feldspar [9], Stuff I've Seen [12], CWH[34], Faceted Metadata [35], Phlat [11]). Thus, a user has the ability to search based on the visual structure of a



**Figure 1. YouPivot Interface and Its Elements**

*A) search box B) date selector C) source list represented by icons whose opacity is based on the frequency of their occurrence in the currently viewed time period D) sort history elements by type or chronologically E) key term list whose opacity and size is based on the frequency of their occurrence in history elements in the currently viewed time period F) create TimeMark button G) 24 hour visualization H) detailed Modeled Activity Visualization I) detailed NonModeled Activity Visualization J) History List*

web page, or execute queries such as, “who emailed me this file?”

These techniques provide new and innovative approaches to personal history but do not facilitate searches based on contextual cues to find semantically unrelated items in a user’s activity history. We have developed YouPivot to support this need.

### YOUPIVOT

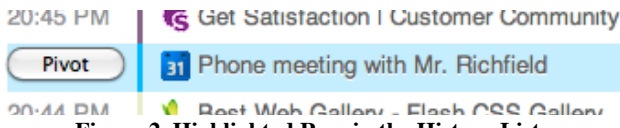
Given the strong evidence from researchers in the cognitive sciences, we envision computer systems (desktop, portable and mobile) that augment and leverage the natural recall methods of the human mind. While personal data collection (e.g., MyLifeBits) and storage is not a modern problem, the creation of efficient and timely retrieval techniques and interfaces is central to leveraging this near flawless memory stores of computers [3].

To this end, we have developed **YouPivot**, a novel system that directly addresses and demonstrates how computers can implement contextual search. At the heart of our system is the ability for a user to shift their point of view to

that of a file, website, or some other activity in their own digital history. From that perspective, they can see everything else that was active during that period of time. Because access to websites, songs, files, other digital activity and physical location is temporal (not just a split second moment, but with start and end times), a user can now think of search in terms of the context of a file rather than just a meta-data title, or keyword.

We illustrate the features and functionality of YouPivot through a scenario of Sarah, a graphic designer. Through her interactions with YouPivot, we will highlight the interface, interactions, and use cases therein. To illustrate the breadth of YouPivot’s features, we present the following overly extensive scenario to cover multiple approaches to YouPivot search. In practice, users can apply any subset of the following search techniques at any time or in any order depending on their contextual cue. Providing users with flexibility in leveraging contextual cues is a key strength of YouPivot.

Sarah is developing a new website layout presentation for the WeSaySo Cooperation. As Sarah is working on her



**Figure 2. Highlighted Row in the History List**

On hover, rows highlight and users are presented with a Pivot button that allows them to switch their context to see everything that was active while the selected item is open

presentation for WeSaySo, she realizes that it would benefit with screen shots of websites on which she is basing her design. She first needs to find the websites.


### Search and Pivoting

At first, Sarah attempts to recall the target website's name. She cannot remember the title, domain or even when she last saw the webpage. A traditional web history cannot aid her - there are far too many items and the titles in the history do not help her find the one she wants. The only relevant information Sarah recalls is that she last viewed the website around the same time of her last meeting with the CEO of WeSaySo, Mr. Richfield. Sarah launches her YouPivot interface (Figure 1), and searches for Mr. Richfield (Figure 1A). YouPivot returns a list of history/activity (replacing the content in Figure 1 section J) that match her semantic query. Finding the meeting she is thinking of, she hovers over the entry (Figure 2), which highlights the row and brings up a Pivot button. When Sarah presses the Pivot button, a new history list loads (in Figure 1J) and shows approximately 38 minutes of data. When Sarah entered her meeting with Mr. Richfield into her calendar, she listed it as being 30-minutes long. YouPivot used the calendar entry's lifespan as the pivot time period (plus a small buffer on each side for temporal context).

The updated history list now displays all of Sarah's activity that occurred during her meeting with Mr. Richfield. This includes any websites that were open on Sarah's computer, local files or applications open, Tweets sent, etc. If a website or file was opened *before* the visible time period, it would still appear in the history list so long as it was not closed prior to the start of the meeting. With a revised and shorter history list at hand, Sarah can now check to see if her target website is listed, and if found, open it.

### Temporal Context: Before, During and After

A review of the history of activities during her meeting with Mr. Richfield reveals that the website in question was not actually open during the meeting, so Sarah thinks it must have been open a short while before or after the meeting. Running along the top of the YouPivot interface is a 24-hour visualization of Sarah's activity (Figure 1G). We used area graphs for this purpose as they have been proven to allow users to easily identify overall trends through visual analysis [10].



**Figure 3. Selectable Region in 24 Hour View**

A) The selectable, resizable, draggable "in focus" region in the 24 hour view is denoted by a white box with a gray border B) red sympathetic highlighting in the 24 hour view indicating all the times when a particular website was open and how important

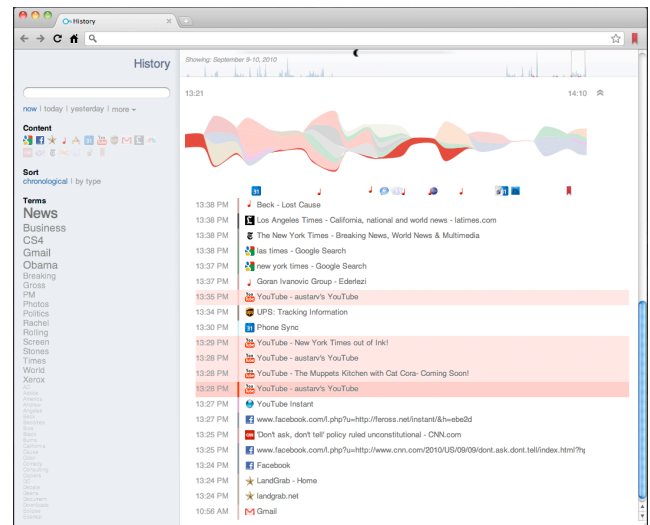
Much like visualizations of stock data, this 24-hour view has a draggable and resizable region (Figure 3A) representing the time range over which Sarah is seeing data. She can easily resize, drag and reposition the selected area to adjust the amount of data she sees in her history list.

Sarah notices that to the left and right of the selected region (representing the meeting with Mr. Richfield), there are approximately three hours of high activity. She resizes the selection area to encompass the active time before and after her meeting with Mr. Richfield.

### Visualization of Activity

After quickly scanning the list of activity she has selected, the website Sarah is looking for still has not jumped out at her. The six hour time range selected has too much data. As she racks her brain, she remembers that at the time she saw this website, she was accessing many Gmail and Facebook pages. Sarah needs a way to identify a context based on surrounding sources of activity.

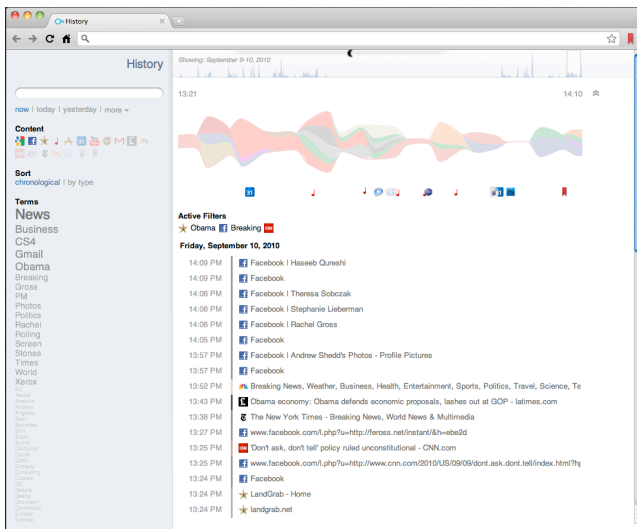
Directly under the 24-hour visualization, is a more detailed representation<sup>2</sup> (Figure 1H) of the selected region in the



**Figure 4. Sympathetic Highlighting**

YouTube 13:28 is highlighted with other websites from YouTube having a lighter colored background. The visualization highlights the change in activity of the selected YouTube history element.


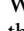
<sup>2</sup> The visualization represents change in website's utility to the users based on the modeling work of [16]. Discussion of website utility is covered in more detail later in the paper.



**Figure 5. Live Filter List**

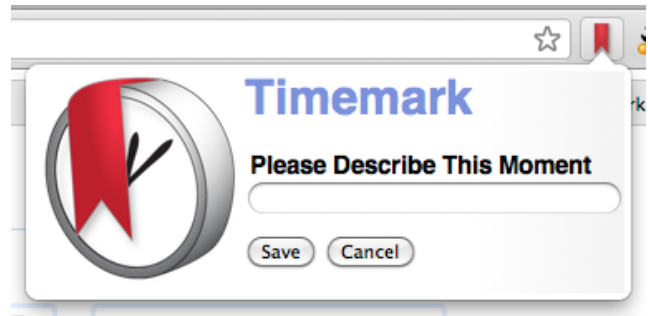
When users filter on terms of sources, those filters appear above the history list. Clicking on filters them removes them

24-hour visualization. This visualization (Figure 1H) shows each history item's activity over time within a stream graph. We used a stream graph for this data, because users can more easily identify and quantify individual layers, offering a more enhanced and legible version of a stacked area graph [8,19,33]. [17] illustrated that visualizations of user activity over time can serve as recall cues, allowing a user to identify their own activity from only the visualizations. Each history item is a different color based on the dominant color in the icon associated with its source (the source icon). YouPivot extracts color, using K-means, from the favicon of websites, file icons of local files, or representative icon from other sources (e.g., music, calendar, etc.). This allows connections to be made between colors in the visualizations and history items.

Sarah can easily see a large swatch of blue (for Facebook's ) icon), and a large swatch of red (for Gmail's ) icon). When Sarah hovers over the blue area in the visualization, the individual layer highlights (Figure 4), while also highlighting the corresponding row in the history list (Figure 2). Further context is provided by highlighting that history element in the 24-hour visualization (Figure 3B). For a website like Facebook, the 24-hour highlight is a small blip. If Sarah highlights Gmail (whose webpage is open all day long), it highlights during multiple periods of time.

At this point, Sarah can resize and refine her viewable area using the 24-hour visualization, or can Pivot on one of the Facebook webpages she visited (discussed at the beginning of this chapter). Either way, Sarah refines her time frame, and examines her activity in the context of the time period when she was browsing Facebook and answering a lot of emails on Gmail. If her target website is listed, she can open it.

<sup>3</sup> Websites icons originate from their domain, local files from the application that open them (e.g. word documents), other sources (e.g. calendar or music) are grouped by source.



**Figure 6. TimeMark**

When the TimeMark icon next to the omni box is clicked, users can enter their TimeMark descriptions

### Ambient Context Display & Filtering

When Sarah refocuses to her new context, she sees almost a “sea of blue,” with the amount of activity at Facebook. For websites which display new content with pageloads, or when users open many files at the same time (e.g., picture browsing), they generate a lot of activity that can muddle and cloud a history list.

Running along the left side of the YouPivot interface, Sarah browses a list of the source icons<sup>3</sup> (Figure 1C) and key terms extracted from the pages (Figure 1E) that represent all the activity during the displayed time period. Order and opacity of each source icon and key-word vary based on the occurrence count within the visible time-range. For the key-term list, size also changes.

By looking at her list of terms and domains, Sarah can quickly see that terms on *gardening* and *finance* are large, not opaque, and high on the term list. Sarah remembers, the websites she was looking for were from banking and gardening websites.

Sarah clicks on the terms in the list, dynamically applying them as filters on her data (source icons can also be used as filters). Now only activity that had the words gardening or finance are displayed (Figure 5). Sarah quickly opens all the websites she found in new tabs, and locates her desired website.

As she is opening her links, Sarah notices a Microsoft Word document that matched the filters in her list. She forgot that while doing research on website layouts, she had started a Word document with notes to herself. She proceeds to open the document by clicking the element in the YouPivot history list.

### TimeMarks

After she has the Word document and all her websites open, Sarah wants to make she she can find all these resources again. Bookmarking 20+ websites is a time consuming task, and she would need to organize them in a manner she could easily find again. In addition, there is no way to bookmark local files, or non-web activity.

Sarah leverages the YouPivot pivot architecture to resolve this problem and creates a TimeMark. **TimeMarks**, are a new mechanism to allow a user to decide that “this moment” is worth remembering. By clicking the TimeMark button (Figure 1F), the user is presented with a window asking them for a description about that moment (Figure 6). YouPivot then stores that TimeMark in their history, creating a custom contextual cue.

Sarah creates a TimeMark, and labels it “Website Layout Inspiration for WeSaySo.” Now Sarah can easily locate all these files and websites again by browsing or search for this TimeMark (example in Figure 1 at time 14:07). She can then Pivot, and reclaim this movement from her past, and the full context that comes with them.

## YOUPIVOT ARCHITECTURE

YouPivot natively supports a handful of data types (Twitter, LastFM, ICS calendar data, and web activity), in addition to logging user physical location, local file and application activity. We now discuss the more technical aspects of the YouPivot system, its client-server architecture, and implementation.

### Interface

The YouPivot interface is constructed using web technology (CSS, HTML5, Javascript, Python) and the Protivis visualization toolkit [7]. It can therefore be accessed on any standards-compliant web browser, regardless of OS or make of computer.

The visualizations of user activity, within the YouPivot interface, represent the change in “importance” of individual tabs/websites in a web browser over time. YouPivot leveraged and adapted the work of Hailpern, et. al. [16] who created models that can quantitatively predict the utility of a given application on a scale from 1 (high utility) to 0 (low utility). According to [16], at a given point in time, users can have multiple applications, websites, and documents open, but only a small subset of them are considered “important.” They termed this subset of important applications “high utility applications.” Further, visualizations of changes in utility can provide a meaningful interface for recalling daily activity [17].

We likewise use visualizations of website utility to allow a user (e.g., Sarah) to easily identify trends in activity. As each website has more and less utility to the user, its thickness grows and shrinks. A website like Gmail (that is open all day) will have thick moments, when the user is actively responding to email, and thin moments when Gmail has not been recently used. The 24-hour visualization (Figure 1G) represents the sum of all importance values, easily highlighting lulls in activity (e.g., sleep, or lunch), while the detailed visualization (Figure 1H) allows users to find source specific trends in activity.

In a realistic representation of real world data, not all data types have, or will have, models of utility level. Therefore YouPivot appropriately handles activity that does not have

a model. While activity that does not have utility values cannot be visualized in a stream graph, these events can act both as contextual cues and temporal landmarks in the user’s daily activity (e.g., seeing activity relative to a phone call or a meeting). We therefore locate them under the visualization (Figure 1I) by placing their icon at the correct time. In this manner, the visualization shows the user when these contextual cues are accessed in relation to each other and in relation to the modeled user activity (Figure 1H).

### Backend

Users’ web activity is logged via the Chrome Extension and sent to the YouPivot server. The YouPivot server also pulls LastFM and Twitter data from the services’ open APIs; and retrieves calendar data via public ICS files.

YouPivot’s default storage system uses Google’s AppEngine. This provides secure authentication through a user’s Google account, as well as secure cloud-based storage, and a dynamic and fast architecture. This infrastructure does not store the user’s activity (the actual files), rather it retains a record of the occurrence (and its associated meta-data). This further improves performance speed, security, and privacy.

### *Multiple Device History*

By using cloud-based storage, a user with multiple devices can send all her data to a central repository. Thus the user’s contextual history can be far richer than one based on the history of a single device. YouPivot shows each user their history from all submitting devices.

### *Ubiquitous Access to Contextual Data*

Because YouPivot is based in the Cloud, it can be accessed on any standards-compliant web browser (desktop or mobile). With the YouPivot storage system living in the cloud, a user’s context effectively travels with them. Even when on a mobile device, a user can search their own history based on contextual cues. In addition, The YouPivot interface can be accessed via an extension in Chrome browser/OS, that replaces the default Chrome history page with YouPivot.

### Communication Protocol

Whenever a user would like to look at their history, they simply log into their Google account. This provides a secure method for viewing and interacting with their data. All data is transferred over secure HTTP (HTTPS).

To log data *into* YouPivot, applications send the data via secure web GET requests. When a user first signs up for YouPivot, they are assigned a unique alpha-numeric string. This string is used by any application wishing to log data into YouPivot, by passing the alpha-numeric sequence with a submission request.

### Extensibility & 3rd party data support

While YouPivot natively can log multiple types of data, many users may wish to log additional data for contextual search or to act as additional contextual cues. While some of these data types may be more universal (e.g., local file access), others may be unique to the individual or their profession (e.g., actions in CAD for architects, or statistical tests run in STATA for epidemiologists).

Custom third-party logging software can easily be written to extend the data contained within YouPivot. Developers/users simply need to write the appropriate logging software that sends a secure GET request to the server with the data to be logged and the user's identifying alpha-numeric string. Thus, YouPivot is neither limited by the initial data feeds, nor access to the YouPivot source code.

We developed two such third-party desktop logging applications. The first logs a user's geographic locations using Mac OS X's CoreLocation; and the other logs a user's local file and application access.

### PILOT STUDY

To show the benefit of YouPivot we conducted a preliminary pilot study to test if YouPivot could support and improve recall given some contextual cues, compared to existing tools. This experiment should be viewed a first piece of a larger set of possible studies that could be conducted to evaluate the effectiveness and value of the contextual history. It was not the purpose of this experiment to prove the use of contextual cues to support recall. Rather, the purpose was to demonstrate that YouPivot could support and improve recall given some contextual cues, compared to existing methods.

### Dependent and Independent Measures

This experiment followed a within subject design that manipulated one independent variable, the browser's history, which consisted of two levels:

1. The traditional browsing history user interface condition (referred to as *Traditional Condition*)
2. The contextual browsing history user interface condition (referred to as *Contextual Condition*)

The dependent variable measured was the time taken to identify a website given a contextual cue. It was hypothesized that participants would be able to identify website(s) quicker with YouPivot, than using the traditional browsing history user interface in Chrome. In addition to this measure, we also administered a questionnaire to understand user satisfaction and preference.

### Participants

Seven participants (four female and three male, aged 18 - 44) took part in this study. Participants were volunteers from a variety of backgrounds (e.g., research and development, consumer operations, finance and sales) at a large software corporation. Furthermore, all participants recruited were frequent users of the Internet, who used

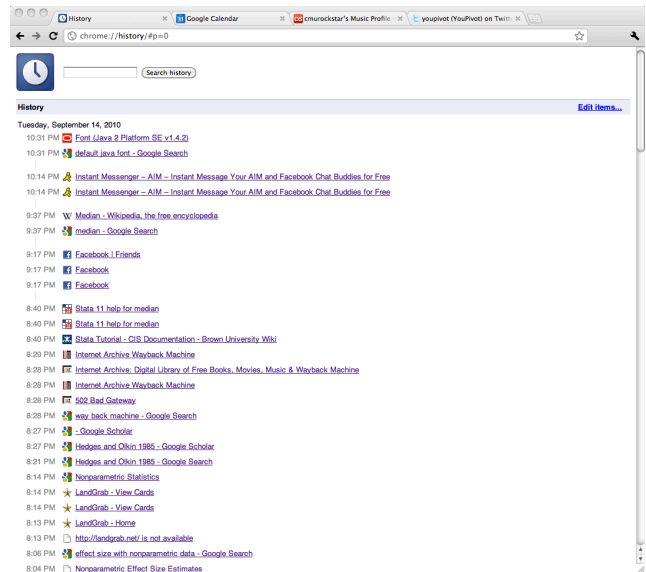


Figure 7. Traditional browsing history user interface set-up

Chrome to browse the Internet. None of the participants had prior knowledge or experience with YouPivot. Participants received a \$40 gift card for participating in the study.

### Set-up

Participants used a 15" MacBook Pro laptop running the Chrome web browser. In the *Traditional Condition*, a Chrome web browser window was open with 4 open tabs (see Figure 7):

1. The Chrome browsing history - allowing the participants to see their browsing history.
2. Google Calendar - allowing the participant to see their appointments/meetings
3. LastFM - allowing users to see the music they listened to.
4. Twitter - allowing users to see the tweets they posted.

In the *Contextual Condition*, a Chrome web browser was open with a single tab displaying the YouPivot interface (see the YouPivot section for more information).

The data used in this experiment (i.e., browsing history, appointments/meetings, music and tweets) was not the participant's real data, but was rather constructed by the experimenters for the purpose of this experiment. While a fictitious dataset would not be appropriate for recall tasks, the goal of this study was to test if users, when given a specific cue, can find an experimental target. Therefore using a consistent dataset across participants is essential for an objective analysis of interface performance. We do envision a future study test the effectiveness of YouPivot in recall tasks (see future work).

## Methods

Each participant was assigned to one of the two conditions. The assignment of conditions was counterbalanced across participants to mitigate ordering effects. The procedure in each condition is detailed below.

Participants were demonstrated either the traditional browsing history user interface, or the contextual browsing history interface, depending on the condition assigned to them. Then they were given time to explore it on their own. Following their self exploration of the interface, participants were asked to perform a 15 tasks. For example, some tasks used in the experiment were:

- What website were you viewing before you watched the New Kindle Pool Ad on YouTube?
- What websites were you viewing when listening to the song Jack's Mannequin - Dark Blue?
- What website did you view after writing your last tweet?

Tasks required participants to identify a website(s) viewed (if any) based on different types of contextual cues: websites viewed, music listened to, tweets posted, and appointments/meetings. In addition, different contexts of use were included: before; during; after. In addition to identifying the website(s) in the browsing history, each participant was given an additional three tasks that asked them to identify a website (or websites) that were not in the browsing history - for all tasks, the participant did not know if the website(s) existed or not. This was to address if the user interfaces could support a user identifying when something did not exist, and thereby reducing wasted time. Tasks were administered randomly between conditions and participants to mitigate ordering and learning effects.

It should be noted, that while the YouPivot interface could have been used to identify items other than websites (e.g., music playing when viewing a website, or files open while at a meeting), these questions were not asked. These types of queries are not answerable with traditional browsing history user interfaces, and would not be meaningful to compare YouPivot against something traditional history could not accomplish.

After completing all tasks for a condition, each participant was asked to complete a questionnaire about their experience with the browsing history user interface. Upon completing the questionnaire, the participant repeated the experiment in the other condition. As noted above, conditions were counterbalanced across participants to mitigate ordering effects. Once the participant had completed both questionnaires, the participant was asked to complete a third questionnaire about their experience with both browsing history user interfaces.

## RESULTS & DISCUSSION

While the primary goal of this experiment was intended to elicit qualitative feedback, some quantitative measures

were taken (e.g. time taken to identify website). Due to the scale of the study, these findings are preliminary.

For each participant, the total time to identify a website(s) for all tasks was calculated, thereby control for any potential systematic differences between tasks. Tasks with the traditional history had a mean time to identify a website(s) of 569.00 seconds ( $\sigma = 61.07$ ) while contextual search had a mean time of 451.86 seconds ( $\sigma = 96.77$ ). Statistical significant differences were found (Student's t-test) between the traditional history user interface condition and the contextual history user interface condition ( $p = 0.01$ ). Suggesting that YouPivot performs faster than traditional history.

Participants were asked to rate their satisfaction with each of the browsing history user interfaces on a scale of 1 (very dissatisfied) to 5 (very satisfied). The mean satisfaction rating for the traditional browsing history user interface was 2.29 ( $\sigma = 0.95$ ); and the mean satisfaction rating for the contextual browsing history user interface was 3.43 ( $\sigma = 1.33$ ). The data was analyzed using a Wilcoxon signed rank test. Statistical significant differences were found between participants satisfaction ratings for the traditional history user interface condition and the contextual history user interface condition ( $p = 0.03$ ). Therefore, participant were more satisfied with the contextual history user interface condition, compared to the traditional history user interface condition. Further, five out of the seven participants reported that they preferred the contextual browsing history user interface.

Qualitative feedback tended to focus on having all the information in one place. Participants also liked being able to search and sort their information in different ways;

*It's really helpful to be able to sort based on website. It's also nice to be able to see what you were looking at when you were doing another activity. Searching by time period is also super useful. This one seemed more all-encompassing than traditional history.*

*I like the different search functions with the contextual history which made it easier to search through the data.*

The two participants that preferred the traditional history user interface focused on speed and simplicity;

*I prefer the simplified UI and lack of clutter, despite my hatred for the 'Older' link and inability to pivot on websites*

*fast, simple, though the visualization [in the contextual history user interface] was cool!*

This suggests that future tools must allow users to have a simplified view, and run as fast as traditional history systems built into web-browsers. Yet overall, our initial study suggests that participants were quicker identifying web-



sites, more satisfied and preferred using the contextual browsing history user interface.

### **FUTURE WORK**

YouPivot is not intended to be the perfect incarnation of contextual search. Rather, it is the first demonstration of how computer systems can leverage human memory to enhance recall and recovery of personal information, such as files access and other computer-related activity. As a result, there are many avenues of further exploration into contextual search.

### **Applications Beyond Personal Search**

While YouPivot focuses on personal data/computer usage, contextual search may also be applicable to many other domains, augmenting the capacity of human memory. Consider a doctor who sees a patient who reminds her of another case she had. A contextual search-based computer system would allow the doctor to search on the context (symptoms) of other patients, thus discovering diagnoses, treatments and other relevant information. This could lead to faster diagnoses and treatment. Similarly, lawyers could leverage a context search tool based on cases, allowing them to search on client name, arguments made, resources cited, and then seeing the contexts in which they were applied.

### **Long Term Use**

As the literature in cognitive psychology suggests, the more contextual cues users can leverage, the greater the ability for recall. While a controlled lab study that forces users to leverage their own contextual cues is difficult to design. Diary studies could provide a unique insight into the real world use of a contextual search tool like YouPivot. We propose testing YouPivot over a long period of time to see how contextual search can integrate into the daily functioning of users.

### **CONCLUSION**

We leveraged theory in cognitive psychology to motivate how contextual cues can greatly improve the quality and speed of recall. We developed a novel system called YouPivot that addresses the need for personal file and activity search through digital history based on the context users remember. A core part of our work is the ability of users to Pivot, and change their point of view to that of a file, website, or some other activity in their own digital history. From that perspective, they can see everything else that was active during that period of time. As a result, a user can now think of search in terms of the context of the activity surrounding a file or website rather than just the title, keywords, or other meta-data.

We also introduce a new method of personal annotation called TimeMarks, which bookmarks a moment in time, allowing a user to access all activity that was ongoing at a particular moment rather than manually keep track of individual files, websites, and conventional bookmarks.

Results from a pilot study suggest that participants were quicker to identify websites, and preferred using YouPivot, compared to current history tools. This research demonstrates how principles of human memory can be applied to enhance the search of digital information.

### **ACKNOWLEDGMENTS**

We would like to thank all our reviewers, participants, family and friends for their support. In addition, we want to thank Jason Hong, Vicki Hanson, Shari Trewin, Craig Tashman and Tessa Lau who have been supportive of this research, and the project that led to it [18].

### **REFERENCES**

1. Balakrishnan, A. D., Matthews, T., and Moran, T. P. Fitting an activity-centric system into an ecology of workplace tools. In CHI '10 (2010). ACM, Atlanta, Georgia, USA New York, NY, USA.
2. Bellotti, V. et al. Taskmaster: recasting email as task management. In CSCW Workshop: Redesigning Email for the 21st Century (2002). ACM, New Orleans, Louisiana.
3. Benjamin, A. S. Memory is more than just remembering: Strategic control of encoding, accessing memory, and making decisions. *Psychology of Learning and Motivation*, 48, 2007, 175-223.
4. Bilodeau, I. M. and Schlosberg, H. Similarity in stimulating conditions as a variable in retroactive inhibition. *Journal of Experimental Psychology*, 41(3), 1951, 199-204.
5. Bjork, R. A. and Richardson-Klavehn, A. On the puzzling relationship between environmental context and human memory. In *Current issues in cognitive processes: The Tulane Flowerre Symposium on cognition*. Hillsdale, NJ: Lawrence Erlbaum Associates (1989).
6. Blanc-Brude, T. and Scapin, D. L. What do people recall about their documents?: implications for desktop search tools. In *Proc. of the international conference on Intelligent user interfaces (2007)*. ACM,
7. Bostock, M. and Heer, J. Protovis: A graphical toolkit for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 2009, 1121-1128.
8. Byron, L. and Wattenberg, M. Stacked graphs-geometry & aesthetics. *IEEE Transactions on Visualization and Computer Graphics*, 2008, 1245-1252.
9. Chau, D. H., Myers, B., and Faulring, A. What to do when search fails: finding information by association. In *Proc. CHI '08 Conference (2008)*. ACM, New York, NY.
10. Cleveland, W. S. and McGill, R. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association*, 79(387), 1984, 531-554.

11. Cutrell, E. et al. Fast, flexible filtering with phlat. In Proc. CHI'06 (2006). ACM,
12. Dumais, S. et al. Stuff I've seen: a system for personal information retrieval and re-use. In Proc. International ACM SIGIR Conference (2003). ACM,
13. Freeman, E. and Gelernter, D. Lifestreams: a storage model for personal data. ACM SIGMOD Record, 25(1), 1996, 80-86.
14. Gemmell, J., Bell, G., and Lueder, R. MyLifeBits: a personal database for everything. Commun. ACM, 49(1), 2006, 88-95.
15. Greenspoon, J. and Ranyard, R. Stimulus conditions and retroactive inhibition. Journal of Experimental Psychology, 53(1), 1957, 55-59.
16. Hailpern, J. et al. The CLOTHO Project: Predicting Application Utility. In Designing Interactive Systems (2010). ACM, Aarhus Denmark.
17. Hailpern, J., Karahalios, K., and Jitkoff, N. Clotho and Lachesis: Building, Optimizing and Validating Models to Predict Important Applications in a Real-World Context. Universit of Illinois Computer Science Research and Tech Reports, <http://hdl.handle.net/2142/17091>, September, 2010.
18. Hailpern, J. WISE: a wizard interface supporting enhanced usability. In Proceedings of ASSETS 2006 (2006). ACM, Portland, OR.
19. Havre, S., Hetzler, B., and Nowell, L. ThemeRiver: Visualizing theme changes over time. In InfoViz 2000 (2000). IEEE Computer Society,
20. Horvitz, E., Dumais, S., and Koch, P. Learning predictive models of memory landmarks. In Proceedings of the CogSci 2004: 26th Annual Meeting of the Cognitive Science Society (2004). Citeseer,
21. Jensen, C. et al. The life and times of files and information: a study of desktop provenance. In CHI '10 (2010). ACM, Atlanta, Georgia, USA New York, NY, USA.
22. Kaptelinin, V. UMEA: translating interaction histories into project contexts. In Proc. CHI '03 Conference (2003). ACM, New York, NY.
23. Karger, D. R. et al. Haystack: A customizable general-purpose information management tool for end users of semistructured data. In Proc. 2005 CIDR Conference (2005). Asilomar, CA.
24. Laadan, O. et al. Dejaview: a personal virtual computer recorder. ACM SIGOPS Operating Systems Review, 41(6), 2007, 292.
25. Lamming, M. and Flynn, M. Forget-me-not: Intimate computing in support of human memory. In Proc of the FRIEND 21: International Symposium on Next Generation Human Interfaces (1994). Tokyo, Japan.
26. Moore, B., Van Kleek, M., and Xu, C. Poyozo make life make sense. <http://mypoyozo.com/>, 2010.
27. Prante, T., Sauer, J., and Lotfy, S. Personal Experience Trace: Orienting Oneself in One's Activities and Experiences. In Extended Abstracts of SIGCHI CHI 2010 Conference (2010). ACM, Atlanta, GA.
28. Rekimoto, J. Time-machine computing: a time-centric approach for the information environment. In Proceedings of the 12th annual ACM symposium on User interface software and technology (1999). ACM,
29. Ringel, M. et al. Milestones in time: The value of landmarks in retrieving information from personal stores. In NTERACT'03 (2003). Ios Pr Inc, Zurich, Switzerland.
30. Smith, S. M. 1988 Environmental context-dependent memory. In Memory in context: Context in memory, G. M. Davies and D. M. Thomson, Eds New York: Wiley.
31. Smith, S. M. Remembering in and out of context. Journal of Experimental Psychology, 5(5), 1979, 460-471.
32. Soules, C. A. N. and Ganger, G. R. Connections: using context to enhance file search. In SOSP '05 (2005). ACM, Brighton, United Kingdom New York, NY, USA.
33. Viegas, F. B. et al. Manyeyes: a site for visualization at internet scale. IEEE Transactions on Visualization and Computer Graphics, 13(6), 2007, 1121.
34. Won, S. S., Jin, J., and Hong, J. I. Contextual web history: using visual and contextual cues to improve web browser history. In Proc. SIGCHI CHI '09 Conference (2009). ACM, New York, NY.
35. Yee, K. P. et al. Faceted metadata for image search and browsing. In Proc. of the CHI'03 (2003). ACM, Fort Lauderdale, FL.