Eyebrowse: Selective and Public Web Activity Sharing

Amy X. Zhang
MIT CSAIL
Cambridge, MA
axz@mit.edu

Joshua Blum
MIT CSAIL
Cambridge, MA
joshblum@mit.edu

David Karger
MIT CSAIL
Cambridge, MA
karger@mit.edu

CSCW '16 Companion, February 27 - March 02, 2016, San Francisco, CA, USA
ACM 978-1-4503-3950-6/16/02.
http://dx.doi.org/10.1145/2818052.2874341

Abstract
Browsing the Internet today is for the most part a private experience, with the exception of a few websites. We explore the possibilities and design considerations around a system to make browsing the web a more social activity. In order to address this, we present a system called *Eyebrowse* that allows users to selectively share their web browsing activity publicly and with friends, using a whitelist at the domain level. This opens up the capabilities to conduct discussions both in real-time and asynchronously around webpages that are designated by the user as a "public space", much like public spaces in real life. Eyebrowse also allows users to find interesting content recommendations, collect personal, web-scale, and friend browsing analytics, and maintain a public persona around their browsing data.

Author Keywords
web browsing; web tracking; activity traces; self-presentation; web analytics; social media

ACM Classification Keywords
H.5.3. [Group and Organization Interfaces]: Web-based interaction; Computer-supported cooperative work

Introduction
Today, users do not have much ability to interact with others or see where others have been while traversing across

the web. The only exceptions to this are on certain websites that have built in social capabilities such as a real-time presence on Google Hangouts or asynchronous interaction via comments and posts on Twitter. At the same time, many corporations are collecting large amounts of browsing data often without users' awareness or consent in order to provide targeted advertising, personalized products, or simply to sell to other corporations. Users get limited benefit from or control over the tracking of their own visits, while corporations often benefit greatly.

Taking inspiration from existing social media, if users could also publish chosen parts of their web activity traces like they do social media posts, then instead of only benefiting organizations, shared data could directly benefit its sharers. Because much of our web browsing is a reflection of our interests, following the web activity of people who are friends or have similar interests could be a useful way to find new and interesting content. Previous research has also shown how the web history of other users can usefully annotate news articles [4] and search results [1]. Second, current social interactions on the web are very limited and dependent on the desires and design of each website. By giving users the ability to interact in the context of a webpage, regardless of the website, this greatly broadens the capabilities for in-place discussion. Third, users interested in self-presentation or social accountability may find personal benefit in sharing what webpages they visit with others, as is often the case with social media accounts. Finally, capturing browsing information in real-time adds many capabilities for real-time discovery of popular or trending content, as well as real-time chat.

Beyond the benefits to users, collecting a public repository of web activity has many benefits for research and developers. Many companies use various means such as installing cookies or using ads to collect large datasets of web-wide browsing activity. These data sets of web activity are valuable because browsing history can provide deep insights into individual users. However, unlike companies that can collect web-wide data, researchers and open-source developers have few means to access aggregate web activity across people and across websites. While large repositories of web activity are certainly valuable, they also present privacy problems. Publicly released datasets of anonymized web activity have generated significant controversy, such as in the case of the AOL Query Log Dataset, where many people were easily de-anonymized. People are also often unhappy about the myriad of companies that track their web activity indiscriminantly. In the past several years, many tools have cropped up to help users block data collection of their web activity. However, these tools fail to address situations where people may not mind and may even like sharing their browsing data, but want to do so in a informed and controllable way [3, 2].

Our system, *Eyebrowse*, is an attempt to both tackle both the problematic incentives in modern-day collection of web-wide browsing activity and the opportunities for social interaction that arise, in a way that provides value to users yet is still privacy-sensitive. The system consists of an extension to the Chrome web browser along with a companion website. Eyebrowse allows users to publicly share select parts of their web browsing activity by whitelisting specific domains or checking in to certain pages that they are comfortable sharing their activity on. By providing social capabilities such as following other users and allowing discussions and chat within pages, we aim to provide the necessary incentive to contribute to Eyebrowse. Early interviews and field studies of the Eyebrowse system presented at this CSCW conference have suggested that the social features

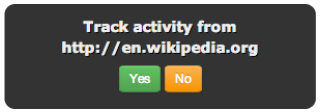

Figure 1: Prompt asking whether to whitelist a domain.

Figure 2: Small popup of recent notes posted and visitors to a page.

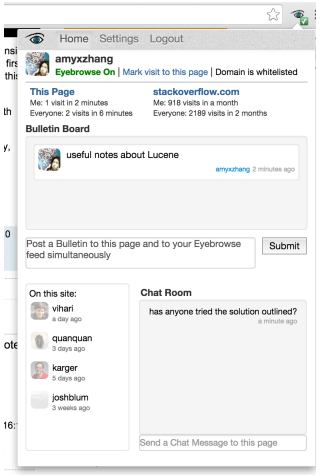

Figure 3: Extension popup shows statistics about the page, all notes posted, chat messages, and recent visitors to the site.

of the system are interesting and that domain-level whitelisting is a plausible control mechanism [5].

Eyebrowse Interface

Eyebrowse[1] is made up of a Chrome browser extension and a companion website. After the user initially installs the extension, a small Eyebrowse icon appears on the top right area of the Chrome browser, where they can then log in. After logging in, Eyebrowse will periodically ask the user using a small popup on the top right of the screen while they are browsing the web whether they wish to whitelist a particular domain (Figure 1). This is done only occasionally so as to not disturb the user. If the user clicks no, then the domain is added to an internal blacklist, so that the user is not asked about this domain again. We also automatically blacklist certain websites that we felt should not be whitelisted due to sensitive information in the url, such as popular search engines, email, and social media such as Facebook.

While browsing, Eyebrowse shows via a small popup also in the top right corner the visitors and comments most recently left behind by other Eyebrowse users on that page (Figure 2). If a user visits a webpage that is in their whitelist, the Eyebrowse logo shows a small green icon with a check mark on it, as shown in the corner of Figure 3. Because we found that users often switch back and forth between tabs, multiple visits to the same page within a few minutes of each other are grouped together as one event. Also, since users will often open a page and quickly leave it, Eyebrowse does not track a page on a whitelisted domain until the user has spent at least 5 seconds on the page.

Upon clicking the Eyebrowse icon, a larger popup window appears (Figure 3). This window contains some statistics

[1] http://eyebrowse.csail.mit.edu

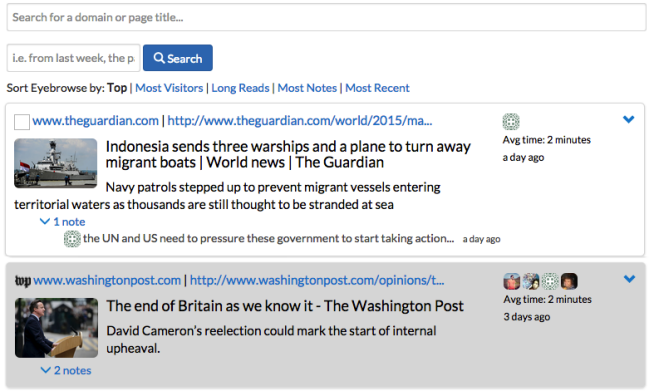

Figure 4: Screenshot of the live stream of data presented on one of the the Eyebrowse feeds with attached notes expanded and already read pieces grayed out.

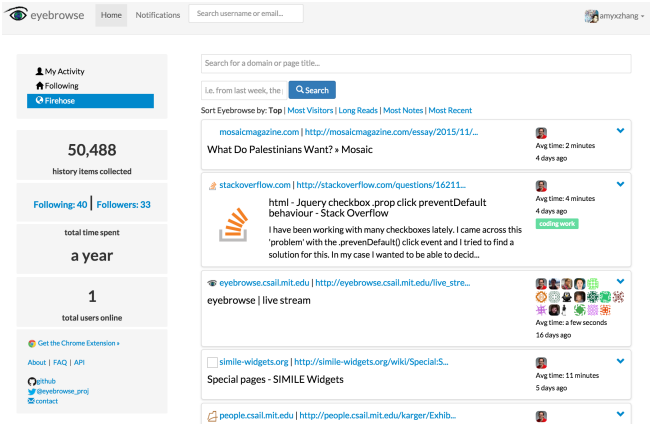

Figure 5: Screenshot of the Eyebrowse website page showing the firehose feed containing all public web activity from Eyebrowse users sorted by a measure of popularity.

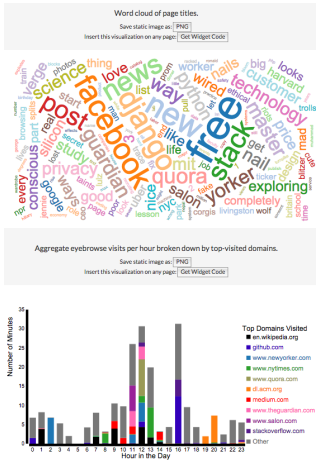

Figure 6: Personal informatics visualizations available about a user.

about the page and domain, all previous notes and chat messages left by Eyebrowse users, as well as shows the users that were recently on the page and domain. It also allows the user to enter a comment or participate in chat on the page. The user can easily toggle off Eyebrowse tracking at any point by clicking the “Eyebrowse On" text, causing the Eyebrowse logo to show a closed eye instead of an open one. The user can also mark a one-time visit to the page or choose to whitelist the current domain from this window.

On the Eyebrowse website, a user can edit their own profile, view their own feed, and manage their settings, such as their whitelist. They can also find and follow other users. There are two feed views set up for the user, one that shows all web activity of their followees (Figure 4) and a firehose feed containing all public web activity (Figure 5). From these feeds users can see the top webpages recommended to them to visit, calculated by each page’s recency, number of visitors, time spent on the page, and number of notes left on the page. They can also see notes other people have left behind, see how many people have been to a page, sort the visits by recency and other ways, and search for specific URLs or keywords. In the small dropdown menu by each visit, users can choose to mute a domain or a key term from their feeds, tag a domain to give it a label, or on their own profile, delete any of their prior visits.

On the website, there is an API and API documentation[2] for developers and researchers to access the public web activity data. For each user, we also provide visualizations around the top keywords from pages visited and top domains visited split up by day of the week and time of the day (Figure 6). These can be aggregated over any time period, such as “last week” or “last year”. The visualizations can be added as widgets to any page or downloaded as a static PNG file.

[2] http://eyebrowse.csail.mit.edu/api_docs

Demo Goals and Conclusion

Our goal in demonstrating Eyebrowse at CSCW is to receive feedback, especially what new features users would like to have given the data and how users might use the system. We also would like to invite CSCW attendees to use the tool over the course of the conference and hopefully beyond. We will create a special page for CSCW that attendees can join to have their whitelisted web visits aggregated. The data collected from Eyebrowse attendees will provide an interesting real-time feed of content viewed by members of the CSCW community. Attendees will be able to use this during the conference to get interesting webpage recommendations and visualizations of the aggregate data, as well as conduct discussions within webpages.

References

[1] Jennifer Fernquist and Ed H. Chi. 2013. Perception and Understanding of Social Annotations in Web Search. In *Proc. WWW*.

[2] Batya Friedman, Daniel C Howe, and Edward Felten. 2002. Informed consent in the Mozilla browser: implementing Value-Sensitive Design. In *Proc. HICSS*.

[3] Jason I Hong and James A Landay. 2004. An architecture for privacy-sensitive ubiquitous computing. In *Proc. MobiSys*.

[4] Chinmay Kulkarni and Ed Chi. 2013. All the News That’s Fit to Read: A Study of Social Annotations for News Reading. In *Proc. CHI*.

[5] Amy X. Zhang, Joshua Blum, and David Karger. 2016. Opportunities and Challenges Around a Tool for Social and Public Web Activity Tracking. In *Proc. CSCW*.