SocialBrowsing: Augmenting Web Browsing to Include Social Context

Michael M. Wasser – mwasser@umd.edu, Advisor – Jennifer Goldbeck

Abstract

In this paper we discuss SocialBrowsing, a Firefox extension that adds a social context to the web browsing experience. The extension works with external sources of social information to enrich a user's experience. It searches every document browsed for keywords related to specific social data, queries these aforementioned external sources, and produces a visual overlay that presents this relevant social information. We present an overview of the tool and its implementation, a discussion of the merits and drawbacks of such a system, and an outline of future work and analysis.

Introduction

Social networking on the web is a huge phenomenon, with nearly half a billion user accounts among hundreds of websites [4]. These numbers do not just reflect idle user accounts; several of the top ten most visited sites on the web are social networking sites [1], showing that people are actively participating. In many cases, users participated for entertainment, either directly through social interactions with other users, or via the features offered by the site such as blogging. However, users find that these social experiences, and the information generated as result, generally stay within the social networking websites. Put differently, social data is largely only presented on its respective social networking site.

We believe that the next step in social networking is for this data to move beyond the websites it was created on. This data will benefit applications outside of the original social networks. We have previously investigated the use of web-based social networks for spam filtering [5]. [5][7] and [10] have all studied the use of trust from social networks in recommender systems. These are good first steps, but there is much more information that social networks can make available. For example, the recommendations produced by social applications may be applicable in other contexts.

In this paper, we introduce SocialBrowing, a Firefox extension that integrates social networking information into a user's browsing experience. As a user browses the web, the extension searches for items related to specific social data sources. When these items are located, the respective data source is queried for additional information. The extension then overlays this information in a seamless manner that may provide a richer browsing experience by displaying more detailed user specific social data.

The first section of this paper introduces the extension and discusses the user experience provided. The second section presents the extensions internals. The third section touches on the motivation for creating this system of social and web data integration. The fourth notes related work and the fifth presents our conclusions. We end the paper with a discussion of possible future work.

Extension Overview

Firefox Extensions are add-ons for the Firefox Webbrowser. They alter and add to the ways a user interacts with the browsers. Existing examples include Adblock, which blocks ads on a webpage using regular expressions, Forecastfox, which adds a weather forecast to the feature bar, and Chatzilla which allows a user to connect to an IRC channel within Firefox. One extreme example is the Songbird music player which uses firefox as a base to create a full featured audio player. Extensions can use external resources and provide a rich user interface while utilizing a set of services provided by the browser itself. These services allow a developer to augment the browser's user interface and functionality.

Our extension, SocialBrowsing, performs an analysis of a pages' contents and highlights words or phrases which have contextual social information. It then presents this information in a manner that enriches the users' experience. In this section we will describe the exterior of our extension and the user experience it creates in detail.

Configuration

When users install SocialBrowsing, they can configure it to work with any services that provide an interface. These services are made up of three basic components. One component Identifies a word or phrase that is relevant to its given social context. A second fetches the social information and a third displays the information using a set of built in display tools. We have developed two services that provide these basic interfaces available for the extension.

FilmTrust (http://trust.mindswap.org/FilmTrust/) is a website with a social network and movie rating and review system. Users create a profile, rate movies, and write reviews. They also maintain a list of friends with ratings that indicate how much the user trusts each friend with respect to movies. The social network and movie ratings are used to produce recommended ratings for movies [5]. When a reference to a movie is found on a webpage, FilmTrust provides its recommended rating information to display to the user. It also allows users to submit their own rating if they have previously seen the movie.

FOAF spider - FOAF (Friend Of A Friend) is a Semantic Web vocabulary for representing information about people and their relationships. Users can create their own FOAF files, but the vast majority are automatically generated by social networking websites. There are at least twenty different social networks with a combined total of over 15 million users that produce FOAF. We have a crawled part of this network and can use the centralized model to show users how they are connected to the author of or someone mentioned on the page being viewed.

The interface is made up of an icon in the status-bar of the browser, an options panel, and web page modifications that include highlighted terms with tooltips. The icon tells a user if the extension is currently parsing pages by displaying a colored or grayed icon for enabled and disabled parsing respectively. The options panel gives the user the ability to customize the visualization of terms within webpages as well as provide login information for social networks. The login information is necessary for receiving data specific to the user rather than just more general information from a social network.
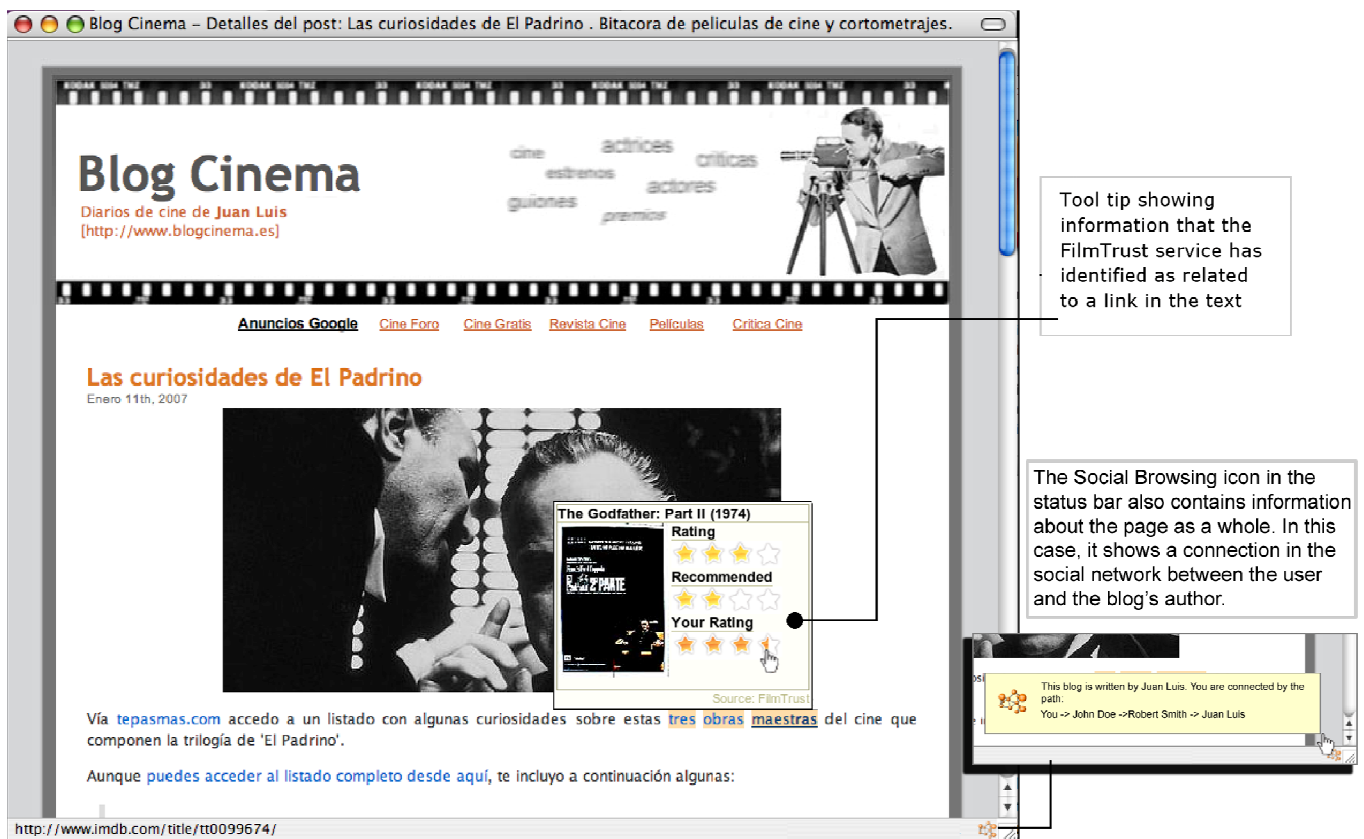
Tool tip showing information that the FilmTrust service has identified as related to a link in the text

The Social Browsing icon in the status bar also contains information about the page as a whole. In this case, it shows a connection in the social network between the user and the blog's author.

Figure 1: Illustration of the SocialBrowsing extension's functionality.
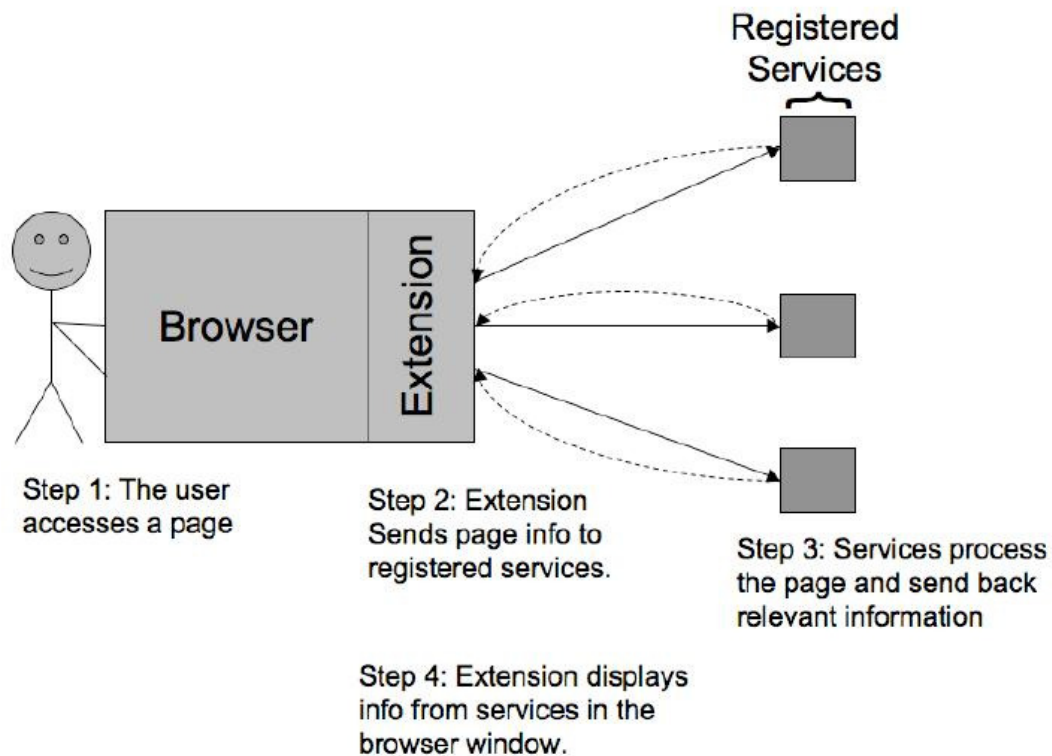
Interface

With the SocialBrowsing extension installed and configured, users browse the web normally. Figure 1 shows an example of the extension at work. We arbitrarily chose a blog that mentioned a movie to illustrate the results returned from FilmTrust. After the page loads, the extension analyzes the page and communicates with the different services. When results come back, the terms in the page that have available social contextual information are highlighted. Moving the mouse over the highlighted term shows the information that a service returned.

Similarly, SocialBrowsing can provide information about the page as a whole. FOAF files can contain links to the blog of the person being described. If a user is browsing the web and comes across the blog of someone in the FOAF database, the FOAF service will generate a shortest path between the user and the blog author. This information is not connected to any particular item in the page, so it is accessed via a tooltip on the status bar icon.

Implementation

The current SocialBrowsing extension implementation takes advantage of many interfaces Firefox offers to modify the users' browsing experience. Written primarily in XML User Interface Language (XUL), Javascript, and Cascading Style Sheets (CSS), Firefox extensions can alter the webpage currently viewed as well as the browser's navigation pane itself.

Many of the algorithms that are executed to generate the contextual information displayed by SocialBrowsing are computationally expensive, as they involve searching large graphs. Furthermore, they all require access to centralized databases, often with private information. As such, the extension is configured to do only minor processing on the client side, and send relevant information to external services that process on the server side.

Registered Services

Browser

Extension

Step 1: The user accesses a page

Step 2: Extension Sends page info to registered services.

Step 3: Services process the page and send back relevant information

Step 4: Extension displays info from services in the browser window.

Figure 2: Architecture of the SocialBrowsing Extension.

The backend of the extension is composed of a module for each social network utilized by the SocialBrowsing extension. Each module has a parsing engine to pick out terms relevant to its respective social networks from a webpage. For example, the FilmTrust module identifies movie titles to query for additional information while the FOAF spider module uses email addresses.

Because the goal with this first version of the extension is to illustrate the concept of integrating social network information with browsing, we have adopted fairly simple methods of analyzing the page content. Currently FilmTrust movies are identified by parsing every link within a webpage using Regular

Expressions to find links to movies on the Internet Movie Database (IMDb). FOAF spider terms are identified similarly by looking for mailto: email links and querying for data pertinent to the given email address.

The mechanism for querying and receiving data to insert into tooltips is different for each module as well. Each social network that the SocialBrowsing extension gathers data from requires an interface for querying for terms and receiving a response. In the case of FilmTrust and FOAF spider, a server side script is used via HTTP get requests to ask about a particular term. The responses from these services are inserted directly into the tooltips. This response can be either HTML or just plain text to provide rich supplementary content for the user when viewing a webpage.

We eventually envision a web-service based architecture for connecting the extension to services, and a standardized format for services to indicate how the extension should process text in the web page.

Motivation and Related Work

There are several lines of work that relate to this project. In one sense, what we are proposing here is a complement to the study of online communities. Much of the research into online communities (well exemplified in [8]) addresses why users participate and how to motivate further participation. We begin with the premise that people are participating, and study how the information users input can be presented back to them in useful ways.

A project similar to our work is the Flock web browser[3]. Flock is a Mozilla-based browser that it's creators call a "social browser" because of its ability to interact with many social and collaborative websites. For example, Flock integrates the browser's bookmarks with del.icio.us and allows users to do batch uploads and tagging to photo sites like flickr. Our extension is in the same spirit of Flock, but has functionality not present in the browser. Flock is designed more to support interaction with a set of socially oriented websites while our extension uses data from those sites to add context to browsing.

Previous work has addressed augmenting web pages. Early work, such as [2], introduced navigation aids into web pages. More recently the Haystack project[9] and the follow-on work in [6] present mechanisms for extracting structured data from web-sites and using it to augment the browsing experience. Our work is similar in intent, but differs in approach and subject.

Conclusions and Future Work

The next step for social networks is to take information from them and use it to inform and enhance other activities. Inspired by the vast number of social network users, we created SocialBrowsing, a Firefox extension that integrates information from social networks into the user's browsing experience. We have completed the proof of concept implementation, and hope to evaluate it through experiments and study of user behavior with the tool installed.

In particular we would look into how the extension affects participation rates in social networking sites that provide context for our extension and compare. We would do this by having users sign up for new

social networks. After this we could observe two sets of people, those with socialbrowsing and those without. We theorize that those with the socialbrowsing extension would participate in these social networks with higher frequency.

Acknowledgements

Example citations

[1] Alexa Web Search - Top 500. http://www.alexa.com/site/ds/top_500.

[2] Barrett, R., P. Maglio, and D. Kellem. How to personalize the web, CHI 1997.

[3] Flock Web Browser. http://flock.com

[4] Jennifer Golbeck. The Dynamics of Web-based Social Networks: Membership, Relationships, and Change. MINDSWAP Tech Report TR-MS1285, 2006.

[5] Golbeck, Jennifer. Computing and Applying Trust in Web-based Social Networks. Ph.D. Dissertation, University of Maryland, College Park, 2005.

[6] Huynh, David F., Robert C. Miller, David R. Karger. Enabling Web Browsers to Augment Web Sites' Filtering and Sorting Functionalities. UIST 2006.

[7] Massa Paolo and Bobby Bhattacharjee. Using trust in recommender systems: an experimental analysis. Proc. of the 2nd International Conference on Trust Management, LNCS (2004).

[8] Preece, Jennifer Online Communities: Designing Usability, Supporting Sociability. Chichester, UK: John Wiley & Sons, 2002.

[9] Quan, D. A. and R. Karger. How to make a semantic web browser. Proc. WWW 2004. ACM Press(2004), 255–265.

[10] Ziegler, Cai-Nicolas. Towards Decentralized Recommender Systems. Ph.D. thesis, Albert-Ludwigs-Universitat Freiburg, Freiburg i.Br., Germany, June 2005.