A visualization interface for interactive search refinement

Fernando Figueira Filho*
Institute of Computing
State University of Campinas (Unicamp), Brazil
fernando@las.ic.unicamp.br

João Porto de Albuquerque
School of Arts, Sciences and Humanities
University of Sao Paulo, Brazil
joao.porto@usp.br

André Resende
Institute of Computing
State University of Campinas (Unicamp), Brazil
resende@las.ic.unicamp.br

Paulo Lício de Geus
Institute of Computing
State University of Campinas (Unicamp), Brazil
paulo@las.ic.unicamp.br

Gary M. Olson
Bren School of Information and Computer Sciences
University of California, Irvine, USA
gary.olson@uci.edu

ABSTRACT

It is common practice nowadays to find, assess and explore the Web by groping scattered information presented through many search results. Browsing interfaces and query suggestion techniques attempt to guide the user by providing term recommendations and query phrases. In this paper, we introduce the browsing interface of Kolline, a community search engine under development. Two case studies are described and two distinct web browsing interfaces are analyzed. Based on this analysis, we present a new browsing interface, describing our design decisions and providing directions for future work.

Categories and Subject Descriptors

H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces—*web-based interaction, collaborative computing, organizational design*

General Terms

Design, Human factors

Keywords

Web 2.0, user-generated annotations, browsing interfaces

1. INTRODUCTION

In recent years, "Web 2.0" [3] applications have been employing tagging as a way for annotating published content. The greatest advantage of tagging systems is that they provide a means to gather the community vocabulary for further classification. Another important characteristic is that they carry different levels of specificity, ranging from very general, widely-used terms to domain-specific terms. This is especially useful in the case of on-line communities within which people are trying to interact and find shared content. In these environments, people may have different backgrounds and distinct areas of expertise, which leads to different perspectives on classification [1].

Regarding the user interface, there are two ways to take advantage of user-generated annotations when looking for information. First, the keyword-based search, which consists of a text box and a search button. The problem with this approach is that it assumes that the user knows how to formulate the query. This is especially hard for people who are trying to find information in different knowledge domains. A recently published article [5] points out that it is a common practice for researchers to find, assess, and exploit a range of information by scanning portions of many articles, instead of looking for a single article to read, in what the authors call "strategic reading". We also have seen this behavior within the open-source community. In order to solve a technical problem, sometimes one need portions of information that may be scattered throughout a series of different postings within a web forum. In these cases, finding the correct keywords which will lead to relevant results can be a time-consuming task. Term suggestion techniques attempt to address this issue, but still depend on an initial query in order to provide further suggestions. This issue is particularly relevant when exploring information in knowledge domains within which the user does not have a strong background, e.g. novice users searching for problem solutions in a web forum.

Second, many websites provide tag clouds or weighted lists, which consist of a visual depiction of user-generated annotations. In this approach, the criteria to show a given term is its use frequency, i.e. how many times users applied that term to annotate content. However, there are some problems with this type of visualization. On one hand, usually only popular terms are depicted, which might not be useful to a user who is searching within one or more specific

*Work done while at the University of California, Irvine.

topics [8]. On the other hand, a user can choose only one term at a time and very often one need a conjunction of terms to suitably express the search task. Consequently, it is impossible to refine a search by only using the tag cloud.

With the aim of addressing those issues, this work presents the navigation interface of Kolline, a community search engine currently under development. It features a term recommendation tool, which suggests terms based on the user's previous interactions. The tool does not require that the user provides an initial query and the interactions can be done solely by clicking on the recommendation tool. In addition, users can refine their search context by choosing new terms which are semantically-related using an underlying ontology. The tool recommends terms which hold subsumption relationships, so a user can refine the search by clicking on general terms at first, and then narrow down the search context by choosing more specific terms. A text box is also provided, so the user can add, remove or modify terms during the interaction.

The paper is organized as follows: Section 2 describes the two cases we studied before designing our application. Section 3 explores some concepts and examples of browsing interfaces for searching, finally introducing our solution. The paper finishes with Section 4 providing an overview of our evaluation plan for the future and drawing conclusions.

2. CASE STUDIES

To characterize the problem, we explored two cases of visualization interfaces for user-generated annotations. First, the case of Ubuntu Forums[1], which represents the major source of information in the Web about this Linux distribution and unites an open-source community of developers and users interested in sharing information about troubleshooting, new features, and other related content. Users exchange information by adding new posts that are shown in the form of threaded discussions. Some users annotate these threads with tags, which helps to categorize content by using a community-oriented, non-controlled vocabulary. Fig. 1 shows a typical tag cloud containing the most frequent terms associated with threaded discussions in the forum.

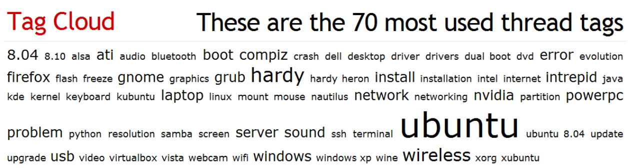

Figure 1: Ubuntu Forums' tag cloud

To better understand whether this type of visualization agregates any significant value to finding information, we refer to a user study which attempted to assess the usefulness of tag clouds in comparison to the traditional keyword-based search. [8] conducted an experiment, giving participants the option of using both approaches to answer various questions. They found that while some participants preferred to use the text search box exclusively, a significant proportion of participants used the tag cloud to find information. There were two scenarios in which the tag cloud's use outweighed the search box's use: (a) when the information-seeking task was broad and non-specific, such as "paste the title of an article you find amusing or interesting" and (b) when the tag cloud contained a term relevant to the question. The first case obviously does not apply to answering technical questions. Most people who visit a Linux distribution forum are looking for help from other community members on a specific topic. However, it could be scenario (b), i.e. the user finds a term in the tag cloud which is relevant to answering a technical problem. But, since the cloud shows only the most used tags, it is not clear if this is sufficient to fulfill a specific search task. In other words, the users will probably have to refine their search by adding other terms. This refinement phase is important for reducing the overall number of hits and excluding irrelevant information from search results. However, most tag clouds or weighted lists do not provide this functionality.

Our second case is a community of professors and researchers of the University of São Paulo. The School of Arts, Sciences and Humanities is an interdisciplinary institute where professors hold positions in a great variety of research areas. To stimulate scientific collaboration, an institutional website is under development, which will contain information about each researcher, organized by area of expertise. In a first phase, professors were asked to provide terms which would describe their research interests and current activities. The union of all colected terms is shown as a list, but because of the great diversity of topics, the result does not fit in one page.

The problems with this visualization approach are twofold. First, each professor uses a particular level of specificity to describe his/her research area. General terms such as "molecular biology" are separated from specific terms like "proteins", although both research areas may have a certain level of intersection. A weighted list approach which shows the most frequent terms in order to reduce the list size is not a suitable solution, because it would not show specific terms that are relevant to the researchers. Second, professors working in the same areas describe them differently, which is the synonymy problem commonly found in tagging systems [2]. For this reason, semantically-related terms end up in different positions on the list, so it is difficult to recognize inter-related subjects and research areas.

Although these cases are related to different communities, the practice of browsing and scanning many pieces of information to find relevant content is a very common issue. In both cases there are difficulties related to the query formulation, i.e. one only recognizes a relevant result when they go through it. In the case of the web forum, relevant results are posts, while in the case of the institute website, relevant results are professors or researchers with a shared goal or interest. We need a tool to visualize user-generated annotations which is able (a) to differentiate general and specific terms into different levels and (b) to provide a refinement mechanism which allows a user to browse horizontally, i.e. between different topics and, at the same time, vertically, i.e. doing an in-depth analysis and looking for specific terms.

3. BROWSING INTERFACES

Representing different levels of abstraction without polluting the interface is a challenging design task. A common way is to represent each level using indentation, e.g. the

[1] http://ubuntuforums.org

Clusty[2] search interface (Fig. 2a). The problem with this approach is that the user usually needs to scroll the page as he/she explores the structure, which requires an extra effort to keep the focus on a given abstraction level, i.e. a term and its proximate relationships. An efficient visualization technique which attempts to address this issue can be found in Google's Wonder Wheel[3] (Fig. 2b). It is a good example of a *focus & context* interface, which encompasses visualization techniques that allow a user to center his view on a part of the screen that is displayed in full detail (*focus*), while at the same time perceiving the wider screen surroundings in a less detailed manner (*context*). The major advantage of using these techniques is the improved space-time efficiency for the user, i.e. the information displayed per screen area unit is more useful and, consequently, the time required to find an item of interest is reduced as it is more likely to be already displayed [4].

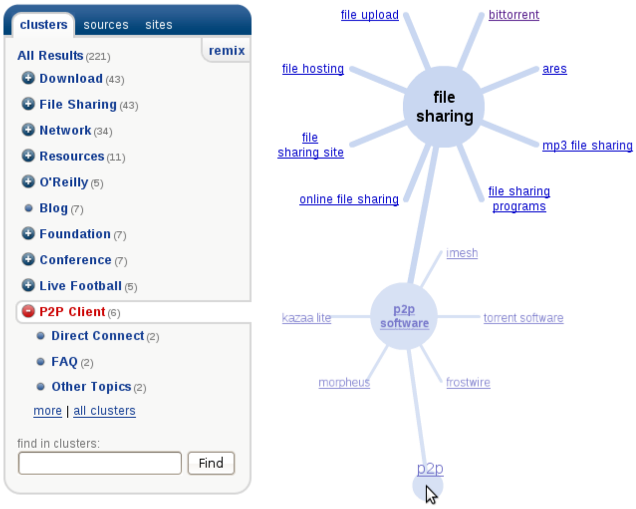

(a) (b)

Figure 2: (a) Clusty navigation menu and (b) Google Wonder Wheel.

Fig. 2b shows an example of interaction. Let us suppose that the user is interested in downloading a peer-to-peer client, so he/she starts entering the query *"p2p"*. A new set of query suggestions appears and gains focus. Then, after selecting the *"p2p software"* query suggestion, the user is presented with new suggestions, among which is the query *"file sharing"*. Incidentally, the user may shift the search context and receive suggestions such as *"file hosting"* and *"file upload"*. If the purpose is to browse horizontally, the tool is very appropriate, leading the user to distinct domains with some level of specificity. However, the tool excludes the initial input term *"p2p"* from the query and eventually moves the users away from their search goal. Because the tool suggests related queries, it does not work as a query formulation tool. In other words, it does not necessarily keep all previously selected terms, suggesting queries that may not have a semantic intersection with the previous interactions.

[2]http://clusty.com

[3]At the time of writing this paper, one could reach the tool by selecting "Show options" in the Google's main page.

3.1 Kolline's interface

Our solution consists of an interactive tool for visualizing hierarchies of user-generated annotations. Terms are related using an ontology which is, in turn, derived semi-automatically by applying a probabilistic model similar to the one presented in [6]. In the case of Ubuntu Forums, we extracted both text corpora and user-generated tags from threaded discussions. As for the institute website, we expect to gather patterns of term co-occurrence from researchers' papers. The resulting ontology is a hierarchy, in such a way that the closer a term is to the root, the more general it is. The purpose of our interface is to allow the user to browse this hierarchy, at first selecting general terms and then refining the search context progressively by adding more specific terms. Fig. 3 depicts Kolline's interface and highlights the functionality of our query formulation tool.

The design of the query formulation tool is based on a colored pie and each slice represents a term in the ontology. The scheme was inspired by an electronic memory game popularized in the eighties called Simon. The main goal of this game was memorizing the sequence of colors displayed by the interface, adding to the sequence one color at a time. In our design, the colors have the purpose of enhancing the user's working memory. [9] shows that recognition memory is 5%–10% better on colored images in comparison to black & white images. Thus, one important design decision is based on the idea that colors may have an important role on helping the user to memorize previous steps when interacting with the interface.

Another important design decision is to avoid scrolling. [7] points out that this approach provides a better experience, especially for novice users. In both cases shown in Fig. 2, the structure grows vertically as the user browses the interface. As a result, scrolling eventually becomes a required effort during the interaction. To address this issue, our query formulation tool stays static and within a single, limited area of the screen, showing just the two previously selected levels as inner circles, i.e. context, and new term recommendations in the outer circle, i.e. focus. The path below the quadrant shows all previously selected terms and allows the user to go directly to a certain level. This has an important role in keeping the user's attention on the focus, without loosing the visual contact of the context.

The tool works as follows (Fig. 3). On selection of one of the general terms displayed by the interface, a transition changes the tool's shape. It becomes a quadrant through a smooth transition to transmit the idea of changing the focus. Each previous level of the hierarchy, i.e. inner circle, keeps the color of the previously selected term. At each new selection, new semantically-related terms are recommended in the outer side of the quadrant. The user can move the mouse over the inner circles to view the context, which causes the previously selected terms to be highlighted. Each new interaction with the tool changes the remaining parts of the interface. The search box is automatically updated with the effective expression resulting from the user's selection. Newer selections refine the search results which in turn gives an instant feedback, so the user can make a decision to continue refining the search context or to go back and browse horizontally over the ontology. To go back, the user can click: (a) on the back arrow displayed near the center of the quadrant; (b) on an inner circle or (c) on a previously selected term in the path below the quadrant.

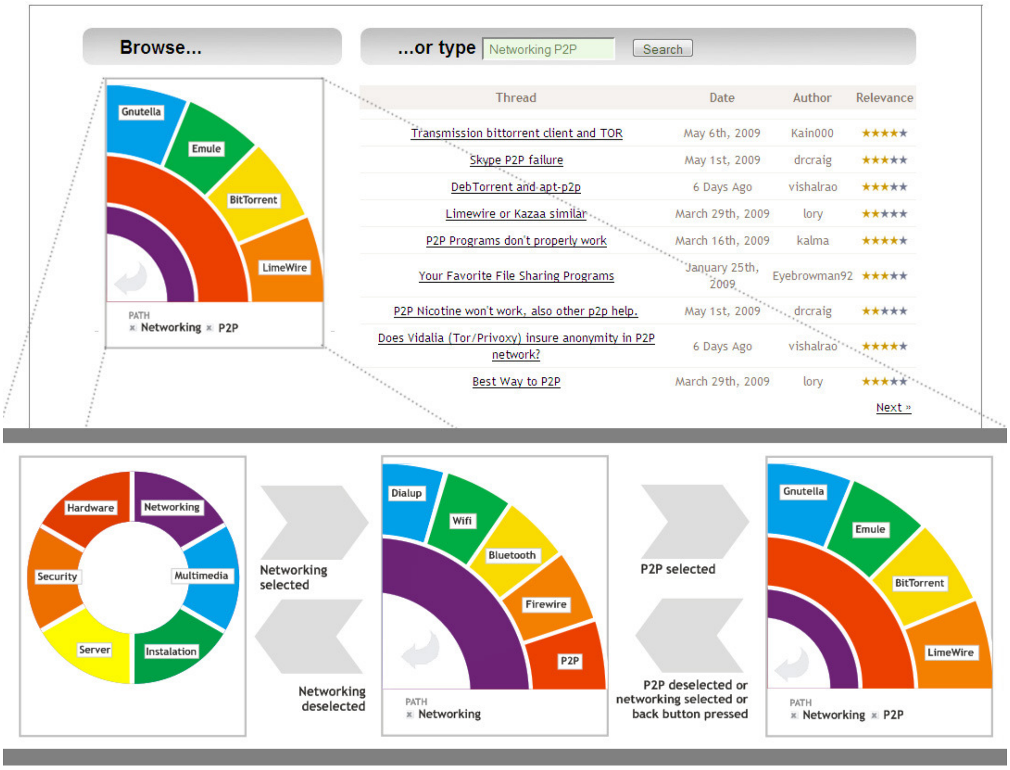

Figure 3: Kolline interface on top and a graphical representation of successive interactions on bottom.

4. FUTURE WORK AND CONCLUSION

As for the evaluation, we will conduct a user study in two phases. First, we want to identify the strategy used by users when performing search tasks and observe their browsing practices. We are particularly interested in better understanding the users' main difficulties when formulating queries and identifying relevant results. Subjects will be recruited to participate in individual, moderated sessions. A screen capturing software will record user activity and moderator will take notes. The aim of the second phase is to assess Kolline's effectiveness in comparison with the tools regularly used by users for searching. For this purpose, a comparison test will be conducted and a group of participants will be asked to perform a set of predefined tasks, in a between-subjects design.

This paper presented a query formulation tool which employs visualization techniques for browsing. We analyzed two cases which involve user-generated annotations to classify content and described two examples of browsing interfaces that attempt to provide assistance to the user in information-seeking tasks. Our design decisions are aimed at addressing the problems found in the case studies and at dealing with the issues identified in usual web browsing interfaces. Therefore, our interface differentiates general and specific terms into different levels and provides a refinement mechanism which allows a user to browse horizontally and vertically over large ontologies.

5. REFERENCES

[1] G. Bowker and S. Star. *Sorting Things Out: Classification and Its Consequences*. MIT Press, Cambridge, MA, 1999.

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

[3] T. O'Reilly. What is web 2.0 | o'reilly media. Available online: `http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html`. Last access: 8/21/2009, September 2005.

[4] J. Porto de Albuquerque, H. Isenberg, H. Krumm, and P. L. de Geus. Improving the configuration management of large network security systems. *DSOM'05: Proc. of the 16th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management*, pages 36–47, October 2005.

[5] A. H. Renear and C. L. Palmer. Strategic reading, ontologies, and the future of scientific publishing. *Science*, 325:828–832, August 2009.

[6] M. Sanderson and B. Croft. Deriving concept hierarchies from text. *SIGIR '99: Proc. of the 22nd international ACM SIGIR conf. on research and development in information retrieval*, pages 206–213, 1999.

[7] E. Schwarz, I. Beldie, and S. Pastoor. Comparison of paging and scrolling for changing screen contents by inexperienced users. *Human factors*, 25(3):279–282, 1983.

[8] J. Sinclair and M. Cardew-Hall. The folksonomy tag cloud: when is it useful? *Journal of Information Science*, 34(1):15, 2008.

[9] F. Wichmann, L. Sharpe, and K. Gegenfurtner. The contributions of color to recognition memory for natural scenes. *Learning, Memory*, 28(3):509–520, 2002.