Visualizing Social Links in Exploratory Search

Justin Donaldson[1]* Michael Conover[1] Benjamin Markines[2,3]
Heather Roinestad[2] Filippo Menczer[1,2,3]

[1]Department of Informatics, Indiana University, Bloomington, Indiana, USA
[2]Department of Computer Science, Indiana University, Bloomington, Indiana, USA
[3]Complex Networks Lagrange Laboratory, Institute for Scientific Interchange Foundation, Torino, Italy

ABSTRACT

The visualization of results is a critical component in search engines, and the standard ranked list interface has been a consistently predominant model. The emergence of social media provides a new opportunity to investigate visualization techniques that expose socially derived links between objects to support their exploration. Here we introduce and evaluate network-based visualizations for facilitating the exploration of a Web knowledge space. We developed a force directed network interface to visualize the result sets provided by `GiveALink.org`, a social bookmarking site. The classifications and tags by users are aggregated to build a social similarity network between bookmarked resources. We administered a user study to evaluate the potential of leveraging such social links in an exploratory search task. During exploration, the similarity links are used to arrange the resources in a semantic layout. Users in our study prefer a hybrid interface combining a conventional ranked list and a two dimensional network map, allowing them to find the same amount of relevant information using fewer queries. This behavior is a direct result of the additional structural information present in the network visualization, which aids them in the exploration of the information space.

Categories and Subject Descriptors

H.5 [**Information Interfaces and Presentation**]: Hypertext/Hypermedia—*Navigation, User issues*;
H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval—*Retrieval models, Search process*

General Terms

Human Factors, Experimentation

Keywords

Interactive Network Visualization, Social Search, Social Links

*Corresponding author. Email: `jjdonald@indiana.edu`

HT'08, June 19-21, 2008, Pittsburgh, Pennsylvania, USA.

1. INTRODUCTION

The organization and presentation of result sets is an ongoing research focus in information retrieval (IR). A common convention is to present search results in an ordered list according to some measure of *relevance* to the query. This mode of interaction has been proven to be very useful for conventional information retrieval tasks involving the look up of answers to specific questions, or finding specific pieces of information. However, other search tasks may be better addressed by alternative result presentations.

Criticisms of conventional IR interfaces have been expressed in the context of *exploratory search.* Marchionini [18] gives an effective comparison and contrast between the information needs of conventional 'lookup' search versus 'learning' and 'investigative' search. Whereas he mentions that the information needs of 'lookup' search have been handled successfully by conventional query based list interfaces, he argues that the latter two methods, which he categorizes as 'exploratory search,' are not well served by such approaches. Methods of learning and investigative search often involve discovering abstract relationships between result items that go beyond similarities of content. We are focused on improving search result interfaces in this context.

Our project concerns itself with exploratory search leveraging user generated metadata in the form of bookmarks. This type of data can be used to uncover relationships between Web resources. We turn to social bookmarking [17] as a form of Web page indexing well suited for exploratory search. Social bookmarking provides users with a medium to annotate pages once controlled exclusively by Web authors.

Many participants of social Web (a.k.a. Web 2.0) applications engage in an activity called tagging. Tagging involves the use of social bookmarking tools to apply keywords or classifications to Web pages in which users are interested. Many online social sites also accept browser bookmarks from users. We can compute similarities between online resources by mining and aggregating individuals' tag annotations or the full hierarchical structure of their bookmarks [23, 19, 10, 5]. We use the terms 'social links' to refer to the semantic relationships derived from user generated metadata, and 'social search' for IR applications that exploit such links.

The similarities between Web pages defined by social bookmarking applications can be used to form an undirected weighted network. Strong relationships between pages in the collection indicate that they are commonly classified in the same or related bookmark categories, or that they are consistently tagged with shared labels. The user's rationale for the association of Web pages as bookmarks is disregarded in

such a network. While users may choose to associate pages for any number of diverse reasons, we assume that the associations are on average meaningful. The purpose of this paper is to evaluate the potential of exploiting such social links to visualize search results.

Many methods of arranging search results have been attempted, using various graph visualization and interaction approaches. Herman *et al.* [9] and Kules *et al.* [12] provide good reviews of graph visualization techniques. These approaches have been proposed in a variety of different domains including music [21], medical and healthcare information [3], and general search [2, 4, 24, 11, 15, 6]. In particular, Hu *et al.* [11] and Allan, Leuski *et al.* [1, 16] provide direct comparisons of the effectiveness of list and visualization based interfaces in goal oriented IR tasks showing that hybrid visualization/list approaches improve retrieval task performance by reducing user query reformulations. Hu *et al.* [11] also reveal that modulating size and color of result nodes improves user performance. Here we build upon these findings by further exploring the use of network visualization in two directions: (i) the effectiveness of social links as the underlying relationships, and (ii) the application to exploratory search tasks.

The underlying data for conventional IR systems differ from those of social search systems. Nevertheless, the 'query plus ranked list of results' model is still the conventional interface. We hypothesize that by leveraging intuitive notions of similarity as proximity [14], users have a better chance of finding meaningful information.

A well defined set of pages, along with their social similarity links, allows the resulting network structure to be visualized with conventional network layout techniques. This arrangement of page nodes exposes more of the structure of the underlying network. Since one of the common goals of an exploratory search process is an analysis of the relationships between result items offered in the corpus content [18], the network visualization approach can be seen as offering a distinct advantage over list based visualizations, which only expose the relationship between items on a single ordinal relevance scale. Here we put these ideas to the test through a user study.

2. NETWORK VISUALIZATION

We used the social bookmarking site `GiveALink.org` as an environment in which to experiment with the visualization of social semantic networks in exploratory search. The *GiveALink* system, developed by our group for research purposes, provides us with a socially-derived similarity network between bookmarked Web pages (nodes), the details of which are outside the scope of this paper and can be found elsewhere [20, 19, 23]. Here we only sketch the salient features of the social search engine and its notion of social links. Bookmark annotations are aggregated across users so that if many users classify two pages in the same category, or tag them with the same keywords, then the two resources will have a high similarity.

The resulting similarity network in *GiveALink* supports the retrieval of pages related to a given URL query. *GiveALink* does not mine individual Web pages for their text or other content. To handle keyword queries, we parse and index annotation words — tags, categories, titles and descriptions — of the bookmarked resources. Results are retrieved by matching the query to the indexed keywords. *GiveALink*

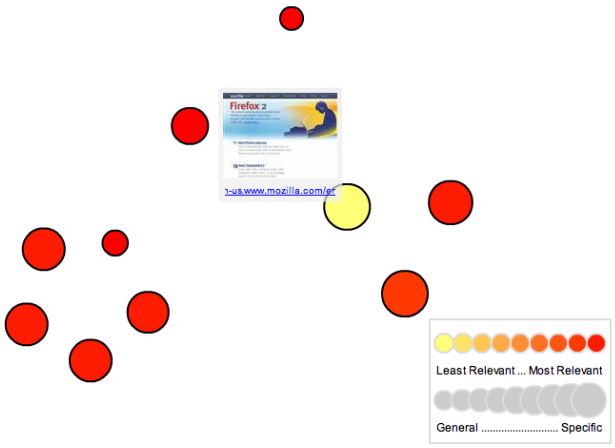

Figure 1: The network visualization applet.

then computes a *relevance* score by combining the number of keywords matching the query with the *strength* of each page. The strength of a node, defined as the sum of the similarities associated with its incident edges, is a measure of the node's centrality in the network. A central page tends to be of general interest, being related to many others. In our experiment we limit the result set to the top ten relevant resources per query.

We designed and developed a visualization applet in Flash to render search results as a network. Given a query, the applet retrieves an XML formatted result set from a *GiveALink* Web service and displays each document as a node in the map. The result set contains information about relevance scores, node strengths, and similarity among results. To visualize the social links between pages, their pairwise similarity is applied as a spring acting between the nodes according to a force directed graph placement method [7]. The applet calculates the position of each document as a real time physical simulation of attracting and repelling forces.

Each result node is colored according to the relevance of the page using a simplified 'heat' color gradient from yellow to red. The coloration was chosen to continuously alter both the hue and saturation so as to be detectable by individuals with color vision impairments. The size of each node is also scaled according to its centrality; more central nodes are smaller. This was done to emphasize nodes that are more specific to the query topic.

Finally, page previews were made available by hovering the mouse cursor over each node, allowing for a 'details on demand' approach recommended by Schneidermann [22]. This effect, along with the results of a typical queried Web page network, are visible in a screenshot of the visualization interface in Fig. 1.

3. EXPERIMENT SETUP

Visualizing a set of search results as a network is not a trivial rearrangement of a list in two dimensions. In addition to selecting and ranking the results, search engines also incorporate the display of additional information about each result. This typically includes, in addition to the title and link to the result, a summary or 'snippet' of the result text:

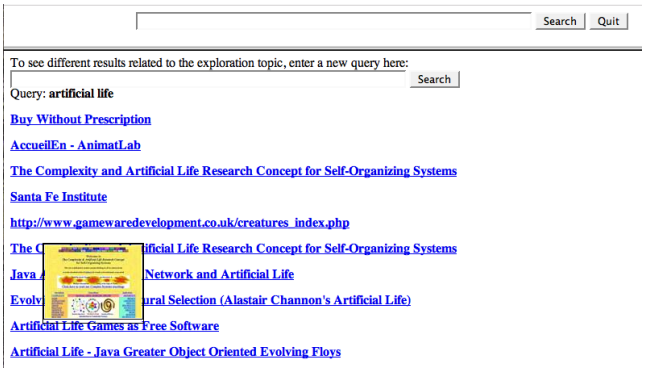

Figure 2: The list interface.

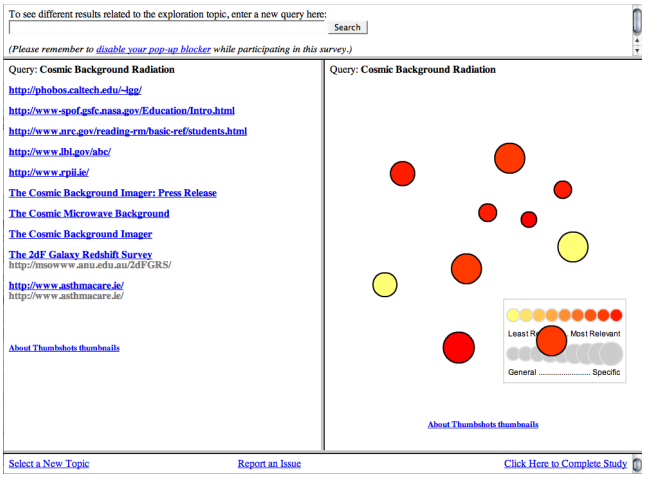

Figure 3: The hybrid interface.

a few descriptive sentences, or a relevant section of the page that relates to the query terms [13].

The inclusion of additional per-result information has become a standard feature for commercial search engines such as `google.com` and `yahoo.com`. Such a combination of ranked-list arrangement with extended result context has shaped the expectations of search engine users. However, the snippets of text limit novel arrangements of results, due to the difficulty of accommodating legible paragraphs of text in a dynamic two-dimensional layout. Furthermore, snippets have been shown to be helpful in informational tasks, but degrade performance for navigation tasks [8]. Therefore, rather than attempting a direct comparison between the network layout and the conventional search engine display of results, we made an experimental design decision to remove snippets from all interfaces. We naturally expect this to affect the perceived utility of the interfaces, but feel it necessary to better focus the study on the issue of structural positioning of results based on social links, independently of the textual context that any interface might add.

Study participants were solicited through a call for participation distributed to departmental and IR interest mailing lists.[1] One of three search interfaces was selected at random and assigned to each participant. The interfaces included a 'map' interface (using the network layout discussed above

[1]Indiana University IRB study #07-12006

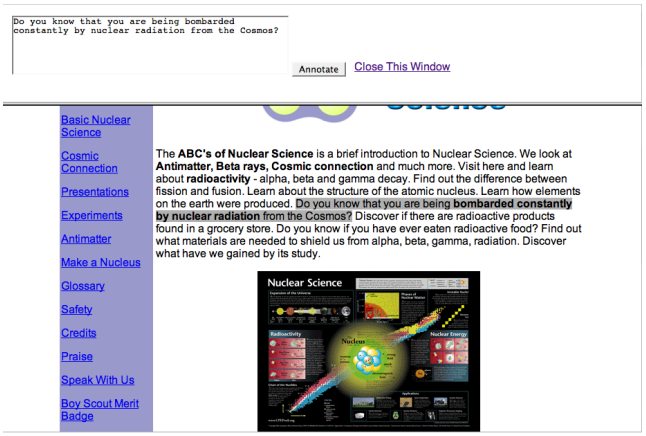

Figure 4: The annotation panel.

and shown in Fig. 1), a 'list' interface (visible in Fig. 2), and a hybrid interface combining the two (Fig. 3). We attempted to control the experiment as much as possible by providing the results in each of the interfaces with the same features: title text, hyperlink, and a pop-up preview of the result page. Relevance ranking was represented by the ordering in the list and by node color in the map. The title text appeared only upon hovering on the nodes in the map (along with the preview), providing for a weaker textual context than available in the list. On the other hand, the map had the additional context of the layout based on social links, and the result specificity represented by node size.

To evaluate these interfaces in the setting of an exploratory search task, each subject was asked to consider various topics from the following list:

- American presidential elections electoral colleges
- Alternative energy sources
- Artificial life
- Impressionism
- Partial differential equations
- Communism socialism fascism democracy
- Lung cancer
- Cosmic background radiation

The topics chosen reflect areas with adequate coverage in the corpus indexed by *GiveALink* and concerning domains deemed relatively unfamiliar to most participants.

We recorded relevant content discoveries as 'annotations.' We asked the subject to copy and paste helpful Web page content while they were viewing a resource linked from our interface. This content could be text, image links, or whatever they deemed to be important. Once a user navigated to a target Web page, two frames were presented, one containing the target Web page and the other a text box to record relevant information. Fig. 4 illustrates this interface. At any time, the user could enter a new query in the assigned navigation interface to retrieve another result set for the current topic. When done exploring one or more of the topics, the user answered a brief exit survey before ending the study.

4. RESULTS

We recorded each user's queries, URLs visited, and annotations made during topic exploration. We then compared

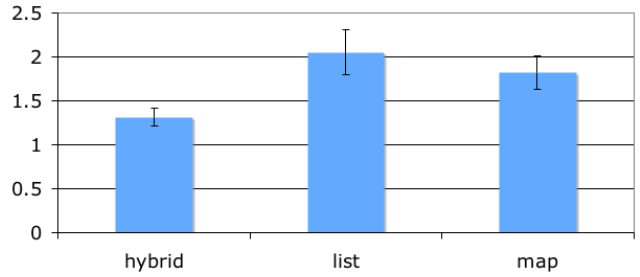

Figure 5: Number of queries submitted by the average user per topic, including the original topic query. Users who explored more than one topic contribute multiple data points. The error bars in this and the following charts represent ± 1 standard error around the mean. The number of queries for the hybrid group is significantly lower than for the list group with $p = 0.03$. We cannot statistically differentiate the other pairs ($p = 0.15$ for hybrid vs. map and $p = 0.40$ for map vs. list).

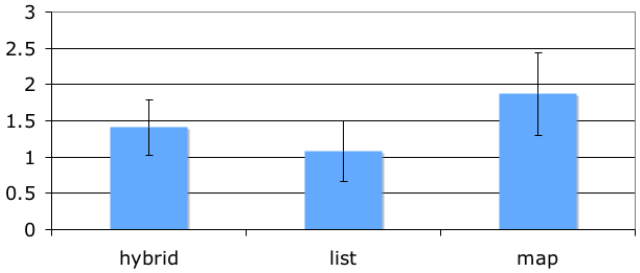

Figure 6: Number of annotations submitted per user, per topic. Users contribute a data point for each topic they explored. There are no statistically significant differences between the means ($p > 0.7$).

groups of users based on which interface they were assigned. Specifically we looked at the number of queries performed per topic, the number of search results visited per query, number of annotations per topic, and the number of topics explored by each participant. We also asked them to rate the usefulness of the interface and the quality of the information in the search results. Overall, we collected results from 65 participants, providing a total of 219 queries and 161 annotations.

The experimental results in Fig. 5 indicate that subjects using the hybrid navigation interface performed significantly fewer queries per topic than those using the list interface. On the other hand, no statistically significant differences were found in the number of annotations users provided per topic (Fig. 6). Similarly, users followed approximately the same number of results per query and explored the same number of topics, regardless of which interface they used.

Therefore users of the hybrid interface were able to gather similar amounts of relevant knowledge using fewer queries. These findings suggest that those who saw both views of the results side by side (their social links in addition to their rankings) explored the topics more efficiently.

Among users who submitted only one query, ratings varied from the strongly negative to the strongly positive for all interfaces. Those who interacted more with the interfaces had more consistent responses. For evaluating the closing questionnaire we chose to include only results from users who performed more than one query. This ensures that the qualitative evaluations of the subjects were based on a baseline level of experience with the interface. The result of this decision is a reduction in the variance of likert-scored results, without a statistically significant change in mean.

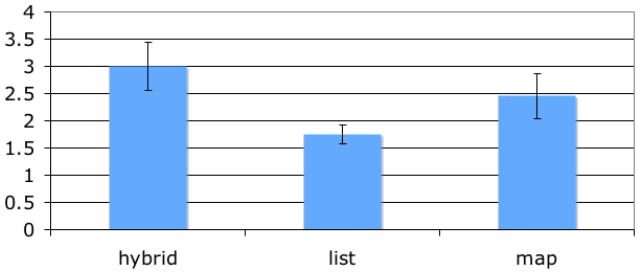

Figure 7: Ratings by group when asked about how much the interface helped exploration of the topics. The hybrid interface was rated more highly than the list interface ($p = 0.024$). The ratings for the map interface cannot be statistically differentiated from either the hybrid or the list interfaces ($p = 0.3$ and $p = 0.13$ respectively).

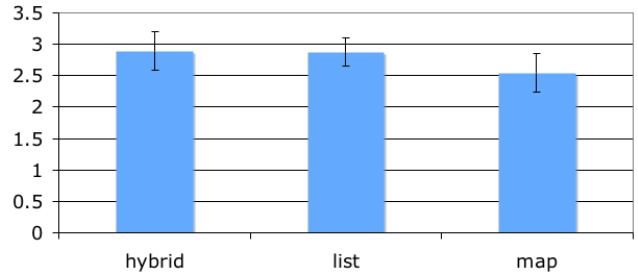

Figure 8: Ratings by group when asked about the usefulness of the data. The differences are not statistically significant ($p > 0.4$).

When asked about the usefulness of the interface for the task, users in the hybrid group gave significantly higher ratings than those in the list group (Fig. 7). Thus, it appears that the addition of the visualization of the social links between search results assists in the exploration task.

The rating for the quality of information is similar across all interfaces (Fig. 8). This indicates that the interfaces alone did not affect the user's perception of the search result quality — indeed the results presented were identical.

5. CONCLUSIONS

The preliminary results in this paper show that there is promise in the visualization of social similarity relationships between search results to complement the traditional ranked list arrangement. By comparing the interfaces solely on the basis of how they lay out the results, in the absence of text snippets, we find that participants prefer a hybrid approach. Furthermore, the hybrid method is more 'efficient' in that users generate the same number of annotations with significantly fewer queries. Our findings for the exploratory search

task are consistent with those of prior information visualization research applied to conventional 'lookup' retrieval tasks.

However, participant evaluation of the hybrid interface in its current incarnation was not overwhelmingly positive. This may have been a function of the perceived quality of the data, or it may have been due to the lack of familiar text summaries, which several participants mention in the open ended comments. While we made a conscious design decision to exclude summaries for reducing the number of free variables and improving experimental control, an important direction for future work is the design of network interfaces capable of incorporating all of the contextual information currently displayed in ranked lists, including snippets. However, the inclusion of summary text must be done in such a way that positioning and legibility requirements are met.

The primary purpose of the experiment described here was to evaluate the potential for network layouts of search results based on social links. Many potential additions to the interface could be investigated, as well as alternative layout algorithms and sources of similarity data. As an immediate follow-up of this work we will implement a prototype network navigation applet to be integrated into the `GiveALink.org` site, to expose this social interface to a broader community of users.

Acknowledgments.

Thanks to Dennis Groth for helpful discussions. We acknowledge generous support from the Indiana University School of Informatics and the ISI Foundation in Turin, Italy.

6. REFERENCES

[1] J. Allan, A. Leuski, R. Swan, and D. Byrd. Evaluating combinations of ranked lists and visualizations of inter-document similarity. *Information Processing and Management*, 37:435–458, 2001.

[2] R. Beale, R. McNab, and I. Witten. Visualising sequences of queries: a new tool for information retrieval. In *Proc. IEEE Conf. on Information Visualisation*. IEEE Computer Society, 1997.

[3] M. Boulos. The use of interactive graphical maps for browsing medical/health internet information resources. *International Journal of Health Geographics*, 2:1, 2003.

[4] S. Carriere and R. Kazman. Webquery: Searching and visualizing the Web through connectivity. *Computer Networks and ISDN Systems*, 29:1257–1267, 1997.

[5] C. Cattuto, A. Baldassarri, V. D. P. Servedio, and V. Loreto. Emergent community structure in social tagging systems. In *Proc. European Conf. on Complex Systems (ECCS)*, Dresden, 2007.

[6] F. Das-Neves, E. A. Fox, and X. Yu. Connecting topics in document collections with stepping stones and pathways. In *Proc. 14th ACM Intl. Conf. on Information and Knowledge Management (CIKM)*, pages 91–98, 2005.

[7] T. M. J. Fruchterman and E. M. Reingold. Graph drawing by force-directed placement. *Software-Practice and Experience*, 21:1129–1164, 1991.

[8] Z. Guan and E. Cutrell. An eye-tracking study of the effect of target rank on web search. In *Proceedings of CHI'07, Human Factors in Computing Systems*, pages 417–420. ACM Press, 2007.

[9] I. Herman, G. Melancon, and M. S. Marshall. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 6:24–43, 2000.

[10] A. Hotho, R. Jäschke, C. Schmitz, and G. Stumme. Information retrieval in folksonomies: Search and ranking. In Y. Sure and J. Domingue, editors, *The Semantic Web: Research and Applications*, volume 4011, pages 411–426, 2006.

[11] P. J.-H. Hu, P.-C. Ma, and P. Y. K. Chau. Evaluation of user interface designs for information retrieval systems: a computer-based experiment. *Decision Support Systems*, 27:125–143, Nov. 1999.

[12] B. Kules, M. Wilson, M. Schraefel, and B. Shneiderman. From keyword search to exploration: How result visualization aids discovery on the web. Human-Computer Interaction Lab Technical Report HCIL-2008-06, University of Maryland, 2008. http://www.cs.umd.edu/localphp/hcil/tech-reports-search.php?number=2008-06.

[13] S. Lawrence and C. L. Giles. Context and page analysis for improved web search. *Internet Computing, IEEE*, 2:38–46, 1998.

[14] A. Leuski and J. Allan. Improving interactive retrieval by combining ranked lists and clustering. *Proceedings of RIAO*, pages 665–681, 2000.

[15] A. Leuski and J. Allan. Lighthouse: Showing the way to relevant information. *Proceedings of the IEEE Symposium on Information Visualization (InfoVis)*, pages 125–130, 2000.

[16] A. Leuski and J. Allan. Interactive information retrieval using clustering and spatial proximity. *User Modeling and User-Adapted Interaction*, 14:259–288, 2004.

[17] B. Lund, T. Hammond, M. Flack, and T. Hannay. Social Bookmarking Tools (II): A Case Study - Connotea. *D-Lib Magazine*, 11(4), 2005.

[18] G. Marchionini. Exploratory search: from finding to understanding. *Commun. ACM*, 49:41–46, 2006.

[19] B. Markines, L. Stoilova, and F. Menczer. Bookmark hierarchies and collaborative recommendation. In *AAAI*. AAAI Press, 2006.

[20] B. Markines, L. Stoilova, and F. Menczer. Implicit tagging using donated bookmarks. In *Proc. WWW Workshop on Collaborative Web Tagging*, 2006.

[21] A. Rauber, E. Pampalk, and D. Merkl. The som-enhanced jukebox: Organization and visualization of music collections based on perceptual models. *Journal of New Music Research*, 32:193–210, 2003.

[22] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. *The Craft of Information Visualization: Readings and Reflections*, 2003.

[23] L. Stoilova, T. Holloway, B. Markines, A. G. Maguitman, and F. Menczer. Givealink: mining a semantic network of bookmarks for web search and recommendation. *Proc. 3rd Intl. Workshop on Link Discovery (LinkKDD)*, pages 66–73, 2005.

[24] L. Terveen, W. Hill, and B. Amento. Constructing, organizing, and visualizing collections of topically related web resources. *ACM Transactions on Computer-Human Interaction*, 6:67–94, 1999.