

Look Before You Click: A Relation Browser for Federal Statistics Websites

Gary Marchionini¹, Ben Brunk¹, Anita Komlodi², Fred Conrad³, & John Bosley³

Interaction Design Lab, University of North Carolina at Chapel Hill¹, University of Maryland² & Bureau of Labor Statistics³

[march, brunkb]@ils.unc.edu, komlodi@glue.umd.edu, [conrad_f, bosley_j]@bls.gov

Abstract

In this paper, we describe a tool that allows people to browse several characteristics of federal statistical data contained in two hundred web sites before committing to any specific site. The design process used and results from usability testing for the Relation Browser tool are presented and recommendations for comparable applications are made.

Introduction

Interfaces to large WWW spaces offer doubly-vexing challenges to designers. First, the large number of sites and pages relevant to even moderately general topics threatens to overload users whether they use analytical search strategies or informal selection strategies. Second, broad-based availability demands that user-centered approaches support very wide ranges of user capabilities and experiences. These challenges are especially problematic in federal government portals and websites that disseminate large volumes of information and aim to serve the entire citizenry from school children to world-class specialists. This paper reports the results of work that aimed to address these design challenges for the US government portal to federal statistics. Federal statistics have long been important to economic, health, and other professionals, but the WWW makes this data easily accessible to the general population as they make decisions about where to live, what to buy, and how to secure their future. Fedstats (www.fedstats.gov) directs people to the bulk of statistics collected and disseminated by 70 federal government agencies through 196 primary websites that each contain many different subsites and webpages. As part of ongoing work to understand citizen needs for statistics and to improve how statistics can be made more accessible and usable, a highly interactive tool for exploring the structure and content of federal statistical sites was developed and tested. This paper provides a brief framework for the problems of overload and diverse user communities, describes the design process used in developing a prototype solution, presents results from a usability test, and discusses the revised tool and implications for similar situations.

Overload: Overview and Preview Solutions

The scope of the WWW has induced the need to implement information finding aids. Search engines index web sites according to word occurrences and categorical directories partition sites by subject. These services are popular entry points to the web, but often return very large result sets. The well-known problems associated with search engines encourages the creation of specialized partitions that aggregate web pages and websites devoted to particular topics. Portal services such as Yahoo point people to these partitions through directory structures. New tools are emerging to assist people in creating personalized web collections (beyond the basic bookmarking facilities of today's browsers), and many organizations create "webographies" for specialized topics. Fedstats is a website that acts as a gateway for a partition of all federal government websites that disseminate statistical data produced by 70 federal agencies. Regardless of how people arrive at complex partitions of web space, they face challenges such as understanding what is and is not available; understanding the structure of what is available (how parts are related and organized); and understanding the control mechanisms for exploring and extracting information. The default mechanism is typically a series of hyperlinks that jump to specific websites or web pages in the partition—leading to the familiar drill-down behavior. As part of ongoing work to provide people with multiple, flexible, and easily controllable views of information spaces, we are developing the Agileview framework (Marchionini, et al., 2000). The ultimate goal is to integrate overviews, previews, reviews (histories), peripheral views (contexts), and shared views (collaborative) into a unified environment for information seeking. The bulk of our work has been devoted to WWW partitions (results of search or specific websites or collections of sites) and creating and testing overview and preview techniques and in this paper, we focus on the overview and preview aspects of agileviews.

Greene et al.,(2000) distinguish overviews that reveal information about collections from previews that reveal information about specific objects. Clearly, overviews of web space partitions and previews for specific websites and pages will be helpful if they minimize response time, are easily reversible, and visually couple different views—well-known characteristics of direct manipulation interfaces.

One common type of website overview is the sitemap. Sitemaps aim to provide overviews that allow people to observe what is available and infer what is not, as well as show the relationships among structural elements of the site. There are many different techniques for creating sitemaps, some use simple HTML text, while others rely on clickable image maps (with the aid of JavaScript) or fully-featured software tools written in Java. Our efforts focused on creating an alternative sitemap interface for Fedstats. Note that this was meant not to replace the existing sitemap but rather to provide an alternative that supports additional exploratory power for discovering relationships and understanding what Fedstats provides.

Sitemaps are not necessarily “maps” in the geographic sense but rather lists, hierarchical views, or visualizations that reveal the structure of pages within a website. Textual (HTML-based) sitemaps provide an outline of a web site’s content, while graphical sitemaps almost always depict a given web site as a hierarchy (minimum spanning tree) of nodes. For example, the NASA site index (<http://www.nasa.gov/siteindex.html>) is one that uses an outline written in HTML. The arrival of Java has facilitated the creation of more interactive sitemap tools such as Nation’s WebToc. WebToc uses a spider to explore the structure of a site, showing the results in an expandable tree view, allowing people to quickly explore hierarchical relationships and data types in the Library of Congress National Digital Library prototypes (Nation et. Al., 1997). Hyperbolic trees (Lamping & Rao, 1996) signify a more unusual approach to web navigation and overview creation. This kind of tool allows users to manipulate very large tree structures and find specific pages while maintaining an overview context, all using relatively little screen real-estate.

Many relationships within a complex site are not hierarchical but rather link distinct instances of different attribute sets. By attribute set, we mean a collection of content attributes. For example, for the attribute set "topic" one attribute is "health." In Fedstats, for example, a user may wish to see the relationship between a topic and an time period or between an topic and the types of data provided. These types of relationships are typically handled by database management system tables that support ad hoc user queries for arbitrary relationships. In site maps, the aim is to allow users to understand and explore arbitrary relationships without posing formal queries. Most site maps simply provide one specific type of relationship exploration (e.g., topic, subtopic) and serve mainly to give notice that an attribute exists in the website. As previously noted, most site maps are lists that may show hierarchical relationships within that list. Many other types of relationship may be useful to users. Graphical maps show spatial relationships, timelines show temporal relationships, and entity relationship diagrams show relationships between distinct entity sets contained in tables.

Some designers have demonstrated site maps that allow users to investigate non-hierarchical relationships. Lin used neural network techniques to produce semantic maps (<http://faculty.cis.drexel.edu/sitemap/index.html>) for document collections where proximity shows concept closeness and region size shows importance in the collection (1997). Robertson et al. (1998) used thumbnails of pages from different websites for visual bookmarking, allowing people to structure bookmarks spatially according to their interests; Marchionini et al (1997) used keyframes of video clips to give users overviews of multimedia databases; and Brunk (1999) is experimenting with animated thumbnails as sitemaps. In the WWW, all relationships, whether hierarchical or not, can be shown as direct links. In a site map, however, we would like to be able to explore links among various sets of attributes without literally following many links since clicking typically requires the user to wait for a new page to load and even with high-speed connections, the new page may overlay the existing focus. Moreover, not only should we be able to choose which relationships to explore, we should also expect more than only existence information about those relationships. We should be able to see some preview information (e.g., how large, language or format, etc.) information about the page(s) in the currently active relationship. We aim to provide new "look ahead" techniques that give additional information about result clusters and individual sites before users bear the costs of linking. A goal of the work reported here was to investigate techniques to show people the relationships among topics (Fedstats includes 14 different topical programs) and other attribute sets such as data types, region, and time period of published statistics, to allow users to explore those relationships easily without committing to link following, and to provide additional preview information such as the volume of information available.

Diverse users: The need for alternatives

Previous investigations of user characteristics and their needs for statistical information (Hert & Marchionini, 1998) demonstrate that government statistical websites must serve the entire range of citizens from school children to

seasoned researchers with the broadest range of experiences and statistical needs. Our perspective is that we must provide people with alternatives that best serve their capabilities and needs. There are different approaches to designing alternative interfaces. One approach is to provide completely different websites with highly specialized interfaces. This may be feasible for a small set of distinct user groups (e.g., medical researchers and the public) but is expensive and too coarse to assist the many variant needs and capabilities in large populations. Another approach is to model a small number of user needs and tasks and provide explicit entry points as well as generic entry points for cases that fall outside these models. In this approach, users are expected to choose which is most appropriate for their needs. The Department of Health and Human Services Healthfinder site illustrates this by offering entry in 11 categories related to age, expertise, gender, and language (<http://www.healthfinder.gov/>). The fact that users must characterize themselves into one of these groups is both a weakness when they are unable to do so, and a strength in that they are in control of the interface.

Another approach is to design interfaces that adapt to users' characteristics and behaviors. The user modeling community has long sought such solutions and a substantial experience base demonstrates the difficulties this approach raises. One variant is to allow people to create profiles (like preference settings in common desktop applications) that customize the website interface. My Yahoo is an example of this variant (<http://my.yahoo.com/>). This variant is useful for regular users who can amortize the time it takes to develop a profile over many uses. A more ambitious variation is to design interfaces that adapt to users automatically. The promise is to relieve the user from developing and maintaining profiles. Although this has long been a goal of interface designers, such interfaces have not been successful in practice (see the special issue of *New Review of Hypermedia and Multimedia* devoted to adaptivity and user modeling for examples and case studies, (Brusilovsky & Milosavljevic, 1998)). The complexities of human characteristics and needs mitigate attempts to create exhaustive mappings between user behavior and system response. More importantly, people prefer to be in control of their interactions and resist systems that behave inconsistently. Additionally, the time it takes to adapt to user behaviors makes such interfaces inappropriate for novice or casual users.

Yet another approach to this general design challenge is to create interface tools that support flexible interaction so that people can quickly and easily explore and learn what is available and control their own paths to meeting their information needs. Shneiderman (1998) and his colleagues have developed such interfaces under the rubric of "dynamic queries" that closely couple user actions to results and facilitate rapid exploration of databases. Such interfaces aim to give users quick overviews/understandings of website content and structure so that they can make informed choices about which paths will meet their particular information needs. This is the approach we took in the project reported here--focusing on the problem of giving people flexible control over the great variety of statistical information linked from the Fedstats website. More specifically, this work focused on the challenge of giving people easy ways to explore relationships among data in the 196 different websites available from the Fedstats site before they made commitments to access those sites. Our aim was to give citizens quick overviews of the federal statistics realm by allowing them to explore relationships among key data attributes.

Design Process

Based on our studies of user needs and meetings with citizen-support staff at the Bureau of Labor Statistics (BLS) and members of the Fedstats steering committee, we know that novice users of statistical websites are task and topic oriented rather than agency or data oriented. Based on discussions with citizen-support staff at the Bureau of Labor Statistics and members of the Interagency Council on Statistical Policy, we adopted the 14 topical descriptors used in Fedstats as one attribute set and data type, geographic region, and time period as three other attribute sets. Based on an analysis of which user control mechanisms might best be used in today's web environment, we developed a paper mockup that incorporated the attributes and specified the mouseover mechanisms keyed to these topics to trigger relationships for the other attribute sets. The first mockup on paper was sent to the two groups above via email for reactions. The design was meant to work as follows. As users mouse over a specific topic, related (available) items are highlighted in an attribute list in another window. The original intention was to provide a "supermenu" of four windows for the four attribute sets. If, for example, the cursor is over health as a topic, the number of websites available for health statistics for different geographic regions, time periods, and types types of information would be

displayed. The pertinent websites appear in a list box at the bottom of the screen. The important concept was to provide an exploration tool that allows users to explore links among two or more classes of attributes. Based on initial reactions by the BLS team to this paper mockup, work began on a Java applet that provided the functionality specified in the mockup.

Control Mechanism Mockups.

A Java applet was developed to instantiate a portion of the design so that the dynamics could be shown to the BLS and Fedstats teams. The aim was to demonstrate the dynamics between the topics and one other attribute set (data types). Both teams reacted favorably to the design and after lengthy discussion, five region names (values for the geographic region attribute) were agreed to (international, national, regional, state, and substate). There was considerable discussion about other ways to categorize the data, including a what/where/when/who/why/how scheme. Based on the results of this meeting and regular conference calls in the subsequent months, we developed a working Java prototype that was subsequently usability tested.

The Relation Browser Prototype.

In developing the prototype, we decided to focus on allowing users to explore pairwise relationships and also provide users with a sense of volume of data for these pairwise relationships. Based upon our previous user needs studies, we decided that linking topic and type of data was the most important relationship that people would use. Our aim to give the user some indication of the size or scope of information available for each topic-data type relationship creates a difficult challenge in the WWW for two reasons. First, boundaries are difficult to establish. Where does a website end? All internal pages? The entire underlying database or only the set returned for a query? Second, what metrics should be used for site volume? Bits are misleading for multimedia data and impossible for streaming data and on-the-fly computations.

To address the first question, we chose to treat the 196 different websites included in the Fedstats portal as distinct entities and provided simple counts of how many of these were appropriate for each pairwise topic-data type relationship. This approach avoided the need to deal with the second question since website was the unit of size and no attempts were made to distinguish large or small sites. We provided users with a visual approximation for the number of websites exhibiting the various attribute values. This aimed to give users a sense of how many chunks (where a website is a chunk) of data they might need to examine if they decide to continue down a query path. In addition to the mouse-over highlighting that showed relationships between topics and data types, a value bar showing an estimate of the number of websites that had such data was used to indicate the volume of information available for these relationships. This is not an ideal solution, but a first step toward giving users a sense of scope and size so that continued exploration is better informed about what lies ahead as they work in the website.

Additional paper mockups were first used to specify the interface. The prototype is shown in Figure 1. The left window contains the 14 Fedstats programs, which were labeled "topics" in the interface. The numbers in parentheses after each topic represent the number of websites that provide data for that topic. The second list contained seven data types that were identified by examining the types of data available from the Fedstats site. The data types: reports, tables, downloadable datasets, searchable datasets, and graphs/charts were clearly important to many websites. Maps and photos/videos were added in anticipation of their growing importance in the WWW environment and appropriateness to many of the statistical agencies. Given the limited screen real estate and the fact that there were just under 200 total websites linked from the Fedstats site map, we decided to use two lists, topic and data type, and make region and time period check box filters for the topic-data type relationships. The resulting prototype allowed users to explore the topic-data type pair while filtering on time and region. Before implementing the prototype, we examined all websites to collect all pertinent attribute information. A template was created based upon the specified attributes and each of the 196 websites linked from the Fedstats site was indexed by these attributes. Table 1 summarizes the data for the 14 topical areas. It is important to note that in many cases, additional links were followed from a specific website to make judgments about what regions, time periods, and data types were available.

Figure 1. The Relation Browser Prototype.

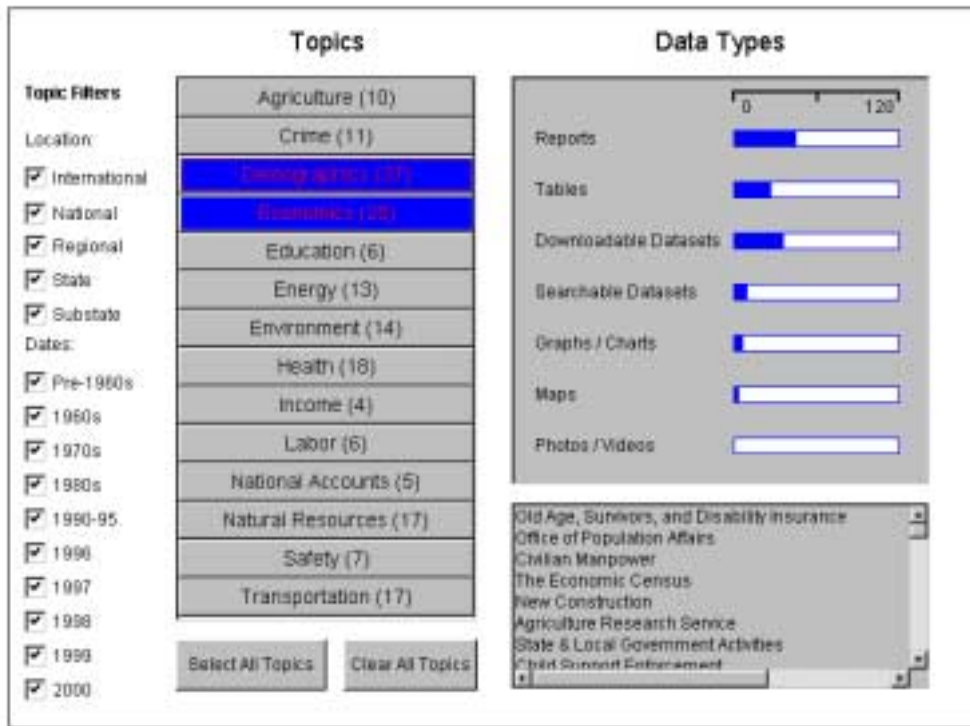


Table 1. Number of Websites with associated data types by topic

	Reports	Tables	Ddatasets	Sdatasets	Graphs/C	Maps	Photos/V
Agriculture (10)	10	5	5	4	4	2	0
Crime (11)	11	1	7	1	9	1	0
Demographics (39)	23	20	18	10	5	2	1
Economics (31)	24	10	10	3	3	1	1
Education (7)	2	3	1	3	2	0	0
Energy (14)	9	9	5	2	5	1	0
Environment (16)	15	1	9	1	5	6	4
Heath (17)	8	6	3	2	3	0	0
Income (4)	2	0	3	0	0	0	0
Labor (14)	6	4	2	2	1	0	0
National Accts (5)	3	3	1	0	0	0	0
Natural Res. (25)	8	5	0	1	3	6	1
Safety (8)	6	3	4	2	0	0	0
Transportation (18)	15	3	6	8	2	3	2

Note: parentheses indicate total number of websites for each topic.

Select all and clear all buttons were added to allow rapid selection and deselection of all topics. The data type list was also based on the paper mockup with the addition of a scale that allowed estimates for the number of websites containing that type of data. The big difference between the Java prototype and the paper mockup was the filters. Check boxes for location and date were used with all values set as the default (all values checked). For date, an ordered category scale was used to give preference to data from the 1990's. The websites associated with a given relationship are displayed in the window at the bottom right. As users mouse over the topics, the bars on the data type list change dynamically to reflect the number of websites associated with that topic and data type. Clicking on a topic freezes the bars and listed websites and allows continued exploration through mouseovers. In the figure, the user has selected both demographics and economics and the resulting number of data types available are shown in the data type list and the websites in the website window.

Usability Testing

To examine the efficacy of the prototype, a usability test was conducted with nine subjects. A testing protocol was drafted and discussed via email. The main goal was to determine how easy the relation browser (RB) was to use, how well it provided an overview of the Fedstats site and associated data, and how well it helped users to focus their exploration to promising websites and thus facilitate searching. Several specific interface issues were also tested to elicit user feedback on design solutions, e.g. filters, data types, and result box interaction.

Because exploration is difficult to assess empirically but is entailed in a more easily definable task like searching, a set of six search tasks for the existing site map and the RB were devised to give test subjects some exploration guidance. In hindsight, this decision clearly biased users toward using the RB as a search tool rather than as an overview explorer. Participants were asked about the usefulness of the two interfaces both before and after searching them, and finally were prompted to compare the two. In addition to the search tasks, users were asked to find information about the websites from the RB itself, e.g. number of site related to a topic, or which websites have demographic data. Sessions lasted between 40 and 75 minutes depending on how verbal subjects were. In all cases, subjects used the Fedstats site map first (this sitemap provides a textual outline of the 196 sites clustered by topic) and then the RB since we wanted to establish a baseline for comparison rather than formally compare the two alternative site maps. The intention was to learn whether the RB was a viable additional alternative rather than a replacement for the site map.

To focus on the site map and RB features, we selected subjects who were experienced WWW users. Test participants were recruited by BLS staff. Six subjects with WWW and some statistical experience were selected from the BLS subject pool developed through ads in the Washington Post. Each of these subjects was paid \$25 for their participation. Three other subjects were recruited from BLS staff and were not paid for their participation. Six of the subjects were male, three female, all reported using computers on a daily basis, and all reported using the WWW at least once a week. Only one subject reported ever using Fedstats and that subject said he had used it once before. Three subjects reported using statistics on the WWW on occasion, and several subjects regularly worked with statistics on Intranets.

Testing took place in the BLS Usability Laboratory in Washington DC. The lab uses two video cameras, one ceiling mounted to capture keyboard and mouse activity and one wall mounted to capture user expressions, a microphone to capture audio, and a scan converter to capture screen activity. These video and audio signals are integrated in a Feral video mixer to produce a three window display that is recorded for later examination.

Results

Subjects had mixed success in using the Fedstats site map and RB to complete the six search tasks. One subject successfully completed all six tasks, and two subjects successfully completed no tasks. Subjects had equal success in the tasks using the Fedstats site map and the RB. Domain knowledge was clearly a factor in finding information.

One subject who worked with federal statistics regularly did not know that the CPI is produced by the Labor Department, but was able to use highly specialized knowledge about educational data to find education costs (searching for data on a specific city, knowing that the result would be compared to the national average). Another subject who was a medical student was able to find the health care cost value quickly, and noted in the debriefing that navigation was skillful "because I was knowledgeable about terms."

Overview browsing styles

In responding to the questions about Fedstats site map usefulness, subjects noted that it provides a good sense of which agencies produce statistics. Four of the subjects noted the importance of having a query capability as well as the site map, and one subject pointed out the usefulness of having multiple access points to the same information. One subject noted that it took too many clicks to get information. Subjects liked the short textual descriptions of websites provided on these pages. Two subjects mentioned that they would have liked more specific information about websites.

Initial responses to the relation browser before using it for the three search tasks varied. Overall, subjects talked about having more control, additional interaction, increased complexity, and getting a sense of how much data is available. One subject noted: "That's cool--it gives you your breakdown." (Sub 4). Another said that "the format looks more analytical--the other (Fedstats) looked less intimidating, but it looks like you can narrow down search." (Sub 6) Another subject said: "It tells you how much data is out there." This subject went on to note that it allows you to filter and is Windows oriented and very interactive. (Sub 7) Another subject liked the filters but did not like clicking off lots of check boxes; he also noted that it seemed more comprehensive and more detailed. (Sub 8) Another noted that it is more specific, "You can tailor it for what you need." (Sub 9)

After conducting three searches with the RB, subjects were asked a series of questions. When asked whether it was useful in gaining an overview of what statistics are available, several subjects noted that it provided a quick overview of the type and volume of data, allowed filtering, and provided rapid feedback. One subject noted: "Yes, it quickly tells you what is available and how much is available for each topic." (Sub 1). Another noted that it saves clicks and liked the filters. (Sub 4) One noted that you can see what kinds of information and how many websites are available (Sub 3). Another subject said it was more focused, having the number of sites per topic gives overall comparison of how much data is available. (Sub 6) One said: "The longer I spend with it, I like it more." This subject went on to say that "Being able to condense the 150 or so sites, you know it is not going to take forever." He also noted that it also provides easier and quicker access to information and that "If you are lost on a wrong site, you can get out quickly--everything moved quickly." (Sub 8) Another subject reinforced the notion that the RB provides look ahead information for next steps: "Yes, you can actually figure out what area you want to be in." (Sub 9)

When asked to compare the two site map tools, subjects generally said that the RB was more appropriate to expert users and liked the Fedstats descriptive texts for websites and its navigation familiarity. Subjects said: [Fedtstats site map] "gives a synopsis, this one [RB] shows general overview more quickly." (Sub 1). [RB] is more user friendly, [Fedstats site map] more simplified for novices." (Sub 4). One preferred the aesthetic of the Fedstats site map, and two others noted that the RB was more "focused" or "comprehensive." One subject noted that if he were looking for reports, he would prefer the Fedstats site map but the RB if he needed tables or data. (Sub 5) One subject noted: "With [Fedstats site map] I'd have to start clicking immediately, with the [RB] I could see things right in front of me." (Sub 8). Two subjects pointed out that the RB provided more specific information about the websites.

In response to the prompt for best thing, subjects noted interaction speed; showing how much information is available; ease of going through topics without clicking; and filters (4 subjects). Thus, the RB seems to have accomplished the main mission of allowing people to quickly gain an overview of what and how much is available in the sites associated with the portal without having to directly link to those sites. The popularity of the filter is likely an artifact of the testing tasks that asked subjects to search for specific information. Although subjects liked the idea of the filters, they actually were not helpful in finding the required information.

Several of the suggestions for improvement resurfaced in the question about the worst thing. These included: estimating volume of data types; no search capability; not enough explanation of what is in the websites; not fitting on a single screen; square layout; and not knowing which dataset (website) has which data types. Many of these points were addressed in the revised prototype.

Specific interface design issues

The most common points of confusion caused by the RB were related to data types and volumes. The number of websites and graphical bar estimates were mentioned by three subjects and another noted that the term "data types" was confusing. Two others noted that the website names themselves were either confusing or not informative about what the site contained.

Subjects made several suggestions for improvement. These included:

- Click on the URL to get a synopsis as in Fedstats.
- Use a better title than 'data types'
- Label the 0-120 scale
- Add search capability (two subjects)
- Show which websites have tables (in the result list)
- Add this [RB] as a link from the Fedstats site
- Change the highlight color to be different from the color for selected topics since unclick does not change the color until the cursor is moved.
- Give numbers on the bars
- Provide subheadings for the websites
- Lump the 1990s more together

Participants in the usability study found the RB a viable alternative overview tool, especially for expert users. They stressed that it more control of the overview information, and provided more specific and detailed information quicker. They made several suggestions to make it more usable and less confusing. Many of these suggestions were adopted in the revised prototype discussed below.

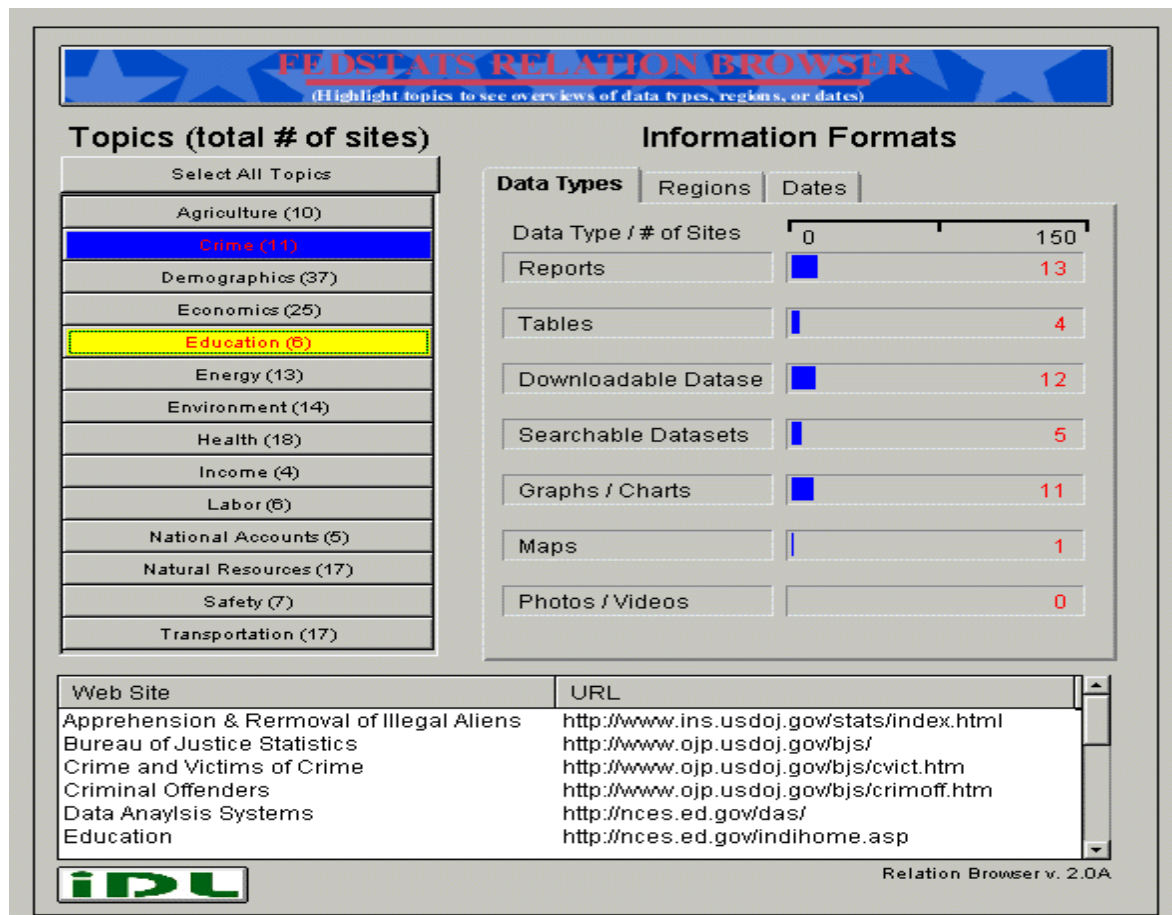
In response to the prompt for best thing, subjects noted interaction speed; showing how much information is available; ease of going through topics without clicking; and filters (4 subjects). Thus, the relation browser seems to have accomplished the main mission of allowing people to quickly gain an overview of what and how much is available in the sites associated with the portal without having to directly link to those sites. The popularity of the filter is likely an artifact of the testing tasks that asked subjects to search for specific information. Although subjects liked the idea of the filters, they actually were not helpful in finding the required information.

Several of the suggestions for improvement resurfaced in the question about the worst thing. These included: estimating volume of data types; no search capability; not enough explanation of what is in the websites; not fitting on a single screen; square layout; and not knowing which dataset (website) has which data types. Many of these points were addressed in the revised prototype.

Revised Prototype.

The prototype and preliminary testing results were presented and discussed at a meeting at BLS in May of 1999. Based on the user testing, a number of changes were made to the prototype. Although subjects said they liked the filters, few were able to really take advantage of them in doing their search tasks--they liked the idea of filters for

Figure 2. Revised Prototype of the Relation Browser



searching but they were not effective in the assigned tasks of the usability study. This was so for two reasons. First, in a corpus of only 200 entries (note that websites were being filtered, not words or numbers in those websites), it is easy to overspecify a query and get no results. This is the classic "no hits" problem that has long plagued casual users of large online databases. Second, manipulating the two filter sets is awkward and confusing. Clicking off or on settings is awkward without additional buttons such as "select/deselect all." Confusion comes from the fact that checking on a filter yields more results--the same well-known problem people have with AND in Boolean queries (the logical connective AND entails less, whereas the linguistic connective AND entails more). Additionally, managing two sets of filters is much more complicated as users must understand that the filter settings for the two attribute sets--location and date-- are ANDed rather than Ored. Thus, by using the filters, people must understand they are specifying a logical query of the form ((L1 OR L2 OR...) AND (D1 OR D2 OR...)). To avoid the confusion and awkwardness in the revision, the filters were removed and extra tabs showing the other attribute sets were added to offset the loss and maximize the utility of the database. This revision was more in line with the original design. In addition to the major change of making the date and region attribute sets tab-selectable alternatives to data type (rather than filters), the following changes were also made:

- The website list was alphabetized.
- Labels were revised and instructions were added.
- The highlight color for mouseover was changed from the color used when topics were selected.
- The actual number of websites for each data type was added to the bars.

- The date ranges were collapsed for the 1990's into two 5-year increments.
- The website list was expanded to a wider window directly below the two attribute lists.
- The two select all and clear all buttons were combined into a toggle button.
- The size of the overall window was changed to fit on a single screen.

Figure 7 shows a screen display of the final prototype. The prototype is available at <http://squash.ils.unc.edu/bls/bls.html>

Summary and Recommendations

This project focused on creating an overview tool for the Fedstats portal to federal government statistics. This alternative tool, called the relation browser, aims to help people quickly get an overview of what and how much data is available by topic and a choice of data type, region, or date. This goal of getting an overview that informs subsequent clicks that retrieve other web pages is especially important in a portal site such as Fedstats. Because all citizens should be able to use it, the RB was designed to use well-known control mechanisms such as mouseovers and highlighted text and graphical bars. An analysis of the underlying data at the federal sites, organization of Fedstats itself, and previous user needs assessments led to four attribute sets (topic, data type, region, and date) as the basis for exploring the structure of the federal statistical corpus. Additional value was added by providing people with a sense of size and scope (number of websites) for different attribute pairings.

Several iterations of prototype and user testing were undertaken as paper and electronic mockups were discount tested with two project teams from BLS and the Fedstats Task Force and a subsequent prototype was laboratory tested with nine subjects. Based on these tests, a final revised prototype was implemented.

Results demonstrated that the relation browser is a viable user-selectable alternative for the site map. The Fedstats site map has a familiar look and feel and allows people to click through/drill down in discrete steps. It also provides helpful descriptions for the main websites. The RB adds new capabilities including: the ability to explore quickly without clicking a commitment; provides data type, region, and date information not otherwise available in an overview; and provides a sense of volume/scope of information available for selected database partitions. These capabilities seem particularly useful for a portal site like Fedstats that points to other sites since making connections to other sites can be slow, especially in the home environment. The RB seems particularly helpful to get as many "advance organizer" cues as possible before committing to a new request.

Although this tool showed good promise for assisting people in understanding what federal statistics are available and finding such information in the various federal websites, there are many additional features that could be added. We obtained Office of Management and Budget approval to conduct a field trial of the RB and the trial was underway at the time this paper was written. The implementation is via a button on the Fedstats main webpage. When people select the button the RB is launched. A button requesting feedback is available on the RB and if people choose to provide feedback, a questionnaire with 11 items (and space for typed comments) is provided. Although preliminary results show substantial usage, a small fraction of the users choose to respond to the questionnaire (about 1 percent)

Additionally, much remains to be learned about how diverse populations of users can understand and use federal statistics. Additional work seems warranted in the following areas.

- Integrate RB into the larger Fedstats environment, for example, add a query function that limits queries to RB setting parameters.
- Develop techniques to link data type, time, and region to topic through mouseovers so that any of these attribute sets could be used to key explorations.
- Investigate ways to give more fine-grained information about volume of data.
- Investigate ways to support more detailed attribute sets, for example, topics and subtopics.
- Add text summaries for topics (as are given in the Fedstats site map).

- Provide topic, data type, time, and region data for websites in the results list.
- Provide additional results set capabilities such as sorting and graphical previews.

Several recommendations result from these experiences. A full implementation would require a revised template for data collection (based on the template used in this project) and a data flow and management plan that distributes the template electronically to webmasters at the various federal statistics websites, collects the completed templates via email, pipes results to a database that feeds the RB applet. A policy on updates would also be required so that new websites and changes to existing websites could be automatically reflected in the RB.

The primary promise of overview tools like the RB is to give better "look ahead" information to users so they can make more informed decisions about next steps in their data exploration. A secondary effect is to encourage better control across different websites as users are able to examine and compare the types and volume of information. The RB is appropriate to other web space partitions where the underlying data can be closely coupled to display representations and familiar control mechanisms.

Acknowledgements: This work has been partially supported by a contract from the Bureau of Labor Statistics.

REFERENCES

- Brunk, B. (1999). Overview and Preview Tools for Navigating the World-Wide Web, July 31, 1999. SILS Technical Report TR-1999-03, UNC.
- Brusilovsky, P. & Milosavljevic, M. (Eds.), 1998. Introduction to special issue on adaptivity and user modeling in hypermedia systems. *The New Review of Hypermedia and Multimedia*, Volume 4. 1-260.
- Greene, S., Marchionini, G., Plaisant, C., & Shneiderman, B. (2000). Previews and overviews in digital libraries: Designing surrogates to support visual information seeking. *Journal of the American Society for Information Science*. 51(4), 380-393.
- Hert, C. & Marchionini, G. (1998). Information seeking behavior on statistical websites: Theoretical and design implications. *Proceedings of the 61st Annual Meeting of the American Society for Information Science*. (Pittsburgh, PA, Oct. 25-29, 1998). P 303-314.
- Lamping, J. & Rao, R. (1996). Visualizing large trees using the hyperbolic browser. *Proceedings of ACM CHI '96 Conference Companion* (Vancouver, April 13-18, 1996). New York: ACM Press. 388-9.
- Lin, X. (1997). Map displays for information retrieval. *Journal of the American Society for Information Science*, 48(1), 40 - 54.
- Marchionini, G., Geisler, G., & Brunk, B.(2000). Agileviews: A human-centered framework for interfaces to information spaces. TR-2000-01 SILS, University of North Carolina at Chapel Hill Technical Report
- Marchionini, G., Nolet, V., Williams, H., Ding, W., Beale, J., Rose, A., Gordon, A., Enomoto, E., & Harbinson, L. (1997). Content+Connectivity=Community: Digital resources for a learning community. *Proceedings of ACM DL '97*. (Pittsburgh, PA, July 23-26, 1997), p 212-220.
- Nation, D. Plaisant, C., Marchionini, G., & Komlodi, A. (1997). Visualizing websites using a hierarchical table of contents browser: WebTOC. In *Proc. of Designing for the Web: Practices and Reflections* (3rd Conference on Human factors and the Web, Denver, June 12, 1997).
- Robertson, G., Czerwinski, M., Larson, K., Robbins, D., Thiel, D., & van Dantzich, M. (1998). Data mountain: Using spatial memory for document management. *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, Nov. 1-4, 1998). 153-162.
- Shneiderman, B. (1998). *Designing the user interface: Strategies for effective human-computer interaction* (3rd Ed.). Reading, MA: Addison-Wesley.