

**An Alternative Site Map Tool for the Fedstats Statistical
Website**

Gary Marchionini, PhD

**Final Report to the Bureau of Labor Statistics
June 30, 1999**

An Alternative Site Map Tool for the Fedstats Statistical Website

Context and Problem Statement

For the past three years, a team of information specialists has been working with the Bureau of Labor Statistics to improve outreach to constituent groups by means of the World Wide Web (WWW). In our preliminary work, we conducted investigations of user needs and tasks and completed transaction log analyses of user behavior at the BLS website (<http://ils.unc.edu/~march/blsreport/mainbls.html>), resulting in a user task-type taxonomy. In the second year, additional user needs analyses were done (Hert), transaction log analyses were replicated, an interface prototype that aimed to provide alternative entry points for different user needs and types was built and tested, and recommendations for short and long term design strategies were made (http://ils.unc.edu/~march/blsreport98/final_report.html). Within this context, the 1998-99 project aimed to extend these efforts to the Fedstats website that includes the bulk of statistics collected and disseminated by 70 federal government agencies. Fedstats is a portal service (also called a locator service or a metasite) meant to serve as a gateway to all federal government statistics (www.fedstats.gov).

The work of the past two years clearly demonstrates 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 previous year's work took this approach with the BLS website. 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. 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. In this year's work, we focused 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.

Several challenges beyond the intricacies of statistical abstractions face people coming to such a complex site, for example, understanding what is and is not available; understanding the structure of what is available; and understanding the control mechanisms for exploring and extracting information. A common aid to such understanding is a site map. Site maps aim to provide overviews that allow people to observe what is available and infer what is not, and show the relationships among structural elements of the site. Site maps typically use well-known metaphors (e.g., geographic maps, table of contents) and popular control mechanisms (e.g., simple mouse actions to expand/contract) to facilitate use. We focused our efforts on creating an alternative site map interface for Fedstats. Note that this was meant not to replace the existing site map but rather to provide an alternative that supports additional exploratory power for discovering relationships and understanding what Fedstats provides.

Site maps are typically not maps at all but rather tables of contents that show the hierarchical structures of pages within a website. This is clearly an important function and some site mapping techniques add graphical visualization to demonstrate these hierarchical relationships. For example, hyperbolic trees (Lamping & Rao, 1996) allow users to manipulate large hierarchical structures to see overall structure and find specific pages (www.inxight.com/Content/7.html); Bedersen (1994) has been developing the Pad++ system (the Java version under development is called Jazz) of zooming interfaces that can be used to zoom in and out of hierarchical structures with continuous motion effects (<http://www.cs.umd.edu/hcil/jazz/>); and Nation created the WebToc tool to allow people to quickly explore hierarchical relationships and data types in the Library of Congress National Digital Library prototypes (Nation et al; 1997).

Many relationships within a complex site are not hierarchical but rather link distinct instances of different attribute sets. By attribute set, we mean a set of labels applied to an entity. For example, for the attribute set "topic" one attribute is "health." Note this is more restrictive than the use of attribute as a field in a table in database parlance. 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. Moreover, 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 (Lin, 1997). Robertson et al (1997) 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. 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 coverage.

After these objectives were identified and clarified, the work progressed in several phases that are detailed in the sections below: identifying and specifying the most salient attributes; creating paper mockups and an electronic prototype and getting feedback on feasibility and usability; creating a full electronic prototype and collecting underlying data to instantiate the prototype; conducting a usability test for the prototype and analyzing results; revising the prototype based on results of the usability test; and writing results and making recommendations for implementation and further development.

Attribute Specification and Preliminary Designs

The first goal was to determine what attributes were most useful to include as organizing schemes for federal statistics. Three input streams provided guidance: user needs assessments, actual data available in federal websites and underlying databases, and the existing structure of the various federal statistical websites and the Fedstats portal. Because Fedstats crosses all federal agencies, it was important to identify attributes that applied to as many types of statistics as possible. The challenge was to find attribute sets that were common to as many websites as possible but still provided some level of discrimination power for site map users.

From our previous work, we know that novice users of statistical websites are task and topic oriented rather than agency or data oriented. On the other hand, expert users can focus on specific data attributes such as source (survey) and variable technical name. For the site map, we decided that specific attribute sets like agency or survey, although they may apply across all websites, required too much knowledge about federal statistics to be helpful to novice or casual users. The first step was to investigate "topic" as the primary attribute set for the site map. The Fedstats and White House websites were studied since they presumably reflect the data available as well as designer expectations about user preferences for topical entry points.

Starting with the Fedstats sites, the fifteen agencies represented on the Interagency Council on Statistical Policy were examined to see what types of statistics they provided. Table 1 lists the different types of statistics provided by these agencies in the Fedstats site and gives a sense of the topics that might be of interest to citizens. Another way to look at the data is to consider which agencies provide information on the fourteen different programs on the Fedstats site. Table 2 shows this organization of data and agencies. To see how these programs are tied to agencies, consider that Income and National Accounts are both types of economic concepts but both have separate program listings each with a single agency providing data (Internal Revenue Service and Bureau of Economic Analysis, respectively). Moreover, only Census, BLS, US Geological Survey, NSF, and the Army Corps of Engineers are listed in more than one program area. Clearly, the Fedstats program areas (topics) are tightly tied to the agencies that provide statistical data.

Table 1. Data Provided by Interagency Council on Statistical Policy

Bureau of Economic Analysis provides statistics on:

balance of payments
 foreign direct investment
 gross domestic product (GDP)
 Gross State product
 industry product
 International trade
 national income and product accts (NIPAs)
 personal income
 personal income, by state

locality pay
 longitudinal surveys
 occupational projections
 Producer Price Index
 productivity
 real earnings
 state and area employment
 unemployment
 union members
 wages
 weekly earnings

Bureau of Labor Statistics provides statistics on:

compensation
 consumer expenditures
 Consumer Price Index
 contingent workers
 displaced workers
 employee benefits
 employer provided training
 employment
 employment cost trends
 employment projections
 foreign labor
 import-export prices
 industry employment
 job injuries
 labor force

Bureau of Justice Statistics provides statistics on:

Crime and Victims
 Drugs and Crime
 Criminal Offenders
 Law Enforcement
 Prosecution
 Courts and Sentencing
 Corrections
 Expenditure and Employment
 Criminal Record Systems
 Criminal Justice, general
 Firearms and Crime

Bureau of Transportation Statistics provides statistics on:

Aviation
Commodity Flow
Fatal Accidents
International Travel and Transportation
Motor Carriers
1994 Highway Statistics
1995 Highway Statistics
1996 Highway Statistics

Bureau of the Census provides statistics on:

agriculture
construction
governments
international trade
income and poverty
manufacturing
population estimates
population projections
social and economic characteristics
retail-wholesale trade
1990 data
census 2000
business ownership

Energy Information Administration provides statistics on:

coal
coal consumption
commercial consumption
electricity
energy: data | forecasts | statistical overview
environmental data
forecasts
gasoline
miles per gallon
natural gas
nuclear
oxygenates
petroleum: imports | consumption | production | prices
power plant emissions
prices, monthly all sources
prices, weekly diesel
prices, weekly gasoline
refinery operations
renewables
residential

Economic Research Service provides statistics on:

agricultural inputs and technology
banking and farm credit
farm sector economics
farm workers and farm employment
field crop economics
food consumption, prices, and expenditures

international agriculture
land, water and conservation
livestock, dairy, and poultry economics
Rural America
specialty agriculture
world agricultural trade
U.S. agricultural trade

Environmental Protection Agency provides statistics on:

Air Pollution
Emergency Response Notification System(ERNS)
Environmental Quality, Status, and Trends (CEIS)
Hazardous Waste Spills

Internal Revenue Service provides statistics on:

Corporation tax returns
Estate and Wealth tax returns
Excise taxes
Exempt organizations and Bond tax returns
Individual tax returns
International or Foreign Related tax returns
Sole proprietorships tax returns

National Agricultural Statistics Service provides statistics on:

Today's reports
Agricultural Statistics, 1995-96
agricultural statistics publications by Date
cattle
crops: state, historic | county | weekly weather
farm economics
field crops
fruits and nuts
hogs
livestock: state, historic | county
milk cows
poultry
sheep
trout and catfish
vegetables
USDA-Economics and Statistics System

National Center for Education Statistics provides statistics on:

Elementary and Secondary Education
Postsecondary Education
Longitudinal Studies
Educational Assessment
Finance
International Mathematics and Science Comparisons
Other Education Subjects

National Center for Health Statistics provides statistics on:

acute conditions (colds and influenza)

births
 breastfeeding
 chronic conditions (e.g., diabetes and asthma)
 deaths
 disabilities/Impairments
 diseases
 divorces
 hypertension
 infant mortality
 key national indicators of well-being
 leading causes of death
 life expectancy
 marriages
 overweight
 prenatal care
 serum cholesterol
 tobacco use
 unintended pregnancy

National Science Foundation provides statistics on:
 science and engineering education
 science and engineering work force
 R&D funding, expenditures, and infrastructure

Office of Management and Budget provides statistics on:
 budget of the U.S. Government

Social Security Administration provides statistics on:
 Highlights
 Fast Facts
 Health Care Programs
 Income of the Aged
 Old Age, Survivors and Disability Insurance Trust Funds
 Old Age, Survivors and Disability Insurance Covered Workers
 Old Age, Survivors and Disability Insurance Insured Workers
 Supplemental Security Income
 New Beneficiary Data System
 Other income support programs
 State Fact Sheets
 Social Welfare Expenditures

Table 2. Agencies by Fedstats Program

Agriculture	Census	SOI (IRS)
NASS	ITA	Labor
ERS	ESA	BLS
FAS	Corps Eng.	ESA
Crime	Customs	ETA
BJS	SBA	National Accounts
BOP	Education	BEA
DEA	NCES	Natural Resources
FBI	NSF	NRCS
INS	Energy	FS
Demographics	EIA	USGS
Census	MMS	NMFS (NOAA)
DIOR	USGS	FWS
ACF	Environment	Safety
AoA	EPA	BLS
FCS	NOAA	OSHA
ARS	NASA	MSHA
OASPE	USGS	CPSC
SSA	Health	Transportation
HUD	NCHS	BTS
VA	NIH	FAA
USAID	AHCPR	FHWA
EEOC	ATSDR	FTA
NSF	HCFA	MARAD
OSTI	HRSA	NHTSA
Economic	HIS	OST
Periodic	SAMHSA	FRA
Census	DOE OESH	Corps Eng
Current	Income	

To see how another group organized government statistics, the White House website was analyzed (www.whitehouse.gov). At the time, that site did not link to Fedstats but had two statistical briefing rooms: economics and social statistics. Thus, this scheme used a two-level hierarchy to organize federal statistical data. These briefing rooms provided data organized into the subcategories listed in Table 3. The two general categories are named with short terms aimed at the general public (note that the word "social" is used instead of the more technical term "demographics" in Fedstats). The subcategories, however, are less direct and it is curious that the social subcategories are presented in alphabetical order but the economic subcategories are not.

Taking these two organizational schemes as a starting point with the intention of coming up with a topical scheme that would be familiar and useful to a broad range of citizens, a variety of classifications were explored. One possible hierarchical scheme is illustrated in Table 4. An advantage of such a scheme from an interface point of view is that the three main categories could be spread out on a graphic with the various subcategories arrayed around them. Figure 1 illustrates one simple way to cluster these three categories with their respective subcategories. After considering these alternative schemes, we decided to focus on the existing Fedstats scheme as the basis for the topic attribute set and not cluster the topics in the three groups. This design decision supports interface comparisons and provides an easy transition if the alternative interface were to be implemented. The Fedstats team is encouraged to consider the hierarchical structure in Table 4 and layout in Figure 1 when a major redesign is undertaken.

Table 3. White House Statistics Briefing Room Organizations

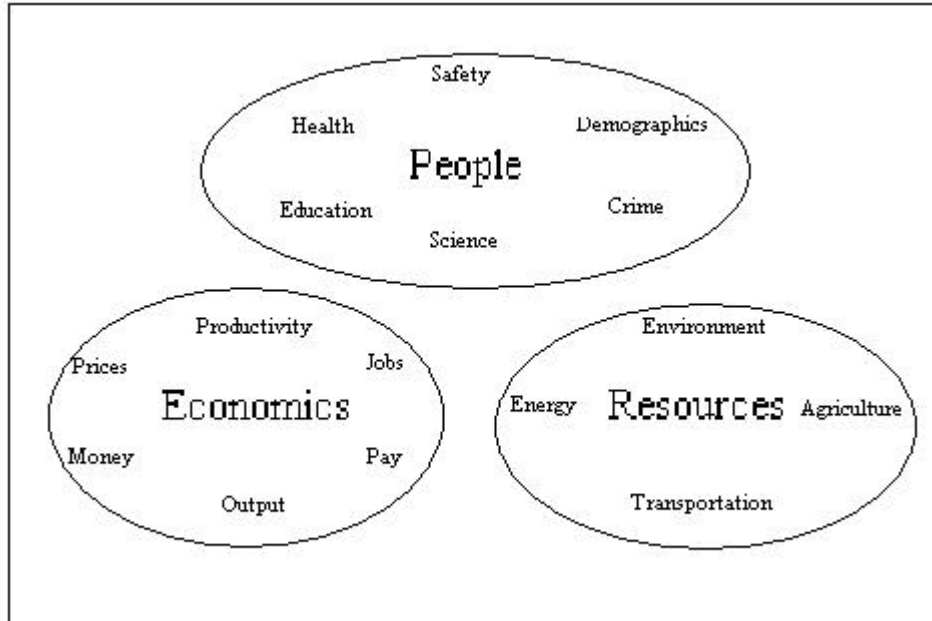
Economics Statistics	International Statistics
Production, Sales, Orders and Inventories	
Output	Social Statistics
Income, Expenditures, and Wealth	Crime Statistics
Employment, Unemployment, and Earnings	Demographic Statistics
Prices	Education Statistics
Money, Credit, and Interest Rates	Health Statistics
Transportation	

Table 4. Hypothetical Organizational Scheme

- People
 - Crime
 - Demographics
 - Education
 - Health
 - Safety
- Economics
 - Jobs
 - Money
 - Output
 - Pay
 - Prices
 - Production

Resources
Agriculture
Energy
Environment
Transportation

Figure 1. Cluster Map for Federal Data



Based on discussions with the BLS team and the Fedstats team (September 24, 1998 meeting at BLS; October 2, 1998 conference call; November 5, conference call; Dec 3, conference call; various email exchanges) and examination of other websites that organize public information, consensus was reached that space (geographic region) and time (when the data were collected) were important attributes to include to characterize data. Because we felt it was important to help citizens quickly understand what types of data were and were not available, we aimed to include data type and data volume as attributes even though there was not consensus on their value or how they could be defined. These issues were discussed in detail at meetings at BLS (January 8, 1999) when the first prototype was presented. In sum, we settled on topic as the most crucial attribute set (using the 14 Fedstats program areas as topics), data type, region, and date as three other attribute sets, and data volume (number of websites) as a common property for all attributes.

The second goal of this phase of work was to consider different interaction mechanisms to control and represent the data attributes. There were several mechanisms considered, including:

- Hypertext links
- Expandable (cascading) hierarchical lists
- Tabbed views
- Zoomable space
- Magic lens
- Popups

- Mouse-over
- Coordinated lists (supermenus)
- Animated site

Hypertext links provide complete flexibility in specifying relationships and allow users to "drill down" in well-defined and familiar ways one step at a time. The well-known disadvantages of hypertext links are user disorientation, localization (losing a sense of the entire resource and/or problem), uncertainty about link destinations, and frustration with time-consuming clicking. Hypertext links require a strong mouse action--a click and subsequent load of another page. One goal of our interface designs was to minimize clicking and provide overviews and previews for content--important motivations for site maps. We thus aimed to reserve hypertext links for informed user decisions to retrieve a web page from another site.

Expandable hierarchical lists are popular in standalone applications (e.g., MS Windows and Macintosh directory managers) and interfaces such as WebToc (Nation et al) have extended this technique to the WWW environment. Expansions may be activated with a mouse click (typical) or mouseover. Users are thus familiar with these mechanisms, however several studies have shown that expandable lists yield poorer performance than simple lists as hierarchical depth increases (Tenissara, 1998; Zaphiris et al., 1999). Because we were interested in non-hierarchical relationships in this interface, we did not apply expandable lists in this design, although this may be a logical extension to add more details to the topic attribute set.

Tabbed views allow users to select alternative windows with a single click (or mouseover) and are familiar in standalone applications. They support arbitrary relationships since a distinct attribute set can be displayed. The displayed set typically hides the details of other attribute sets. Tabbed views have become viable in the WWW environment through Java. Our first designs aimed to display four attribute sets on a single screen and thus did not employ tabbed views. The final prototype compromised this objective for screen real estate and used a tabbed view to allow users to choose one of data type, region, or dates to pair with topic.

Zoom mechanisms offer a new form of control with the potential to reduce disorientation since users can continuously adjust their view of a window. The zooming metaphor lends itself best to well-ordered object attributes (e.g., time, space). Zooming is usually paired with panning capabilities and requires additional mouse options since button or motion actions must be mapped to the zooming and panning. Pad++ (Bederson et al., 1996) used zooming on high-end workstations and the Jazz descendent (Bederson, 1999) makes zooming viable in the WWW environment. We chose not to apply zooming in the relation browser because we were mainly concerned with unordered attribute sets, the user population (US citizens) is not yet familiar with zooming mechanisms, and the techniques require specialized Java classes and more substantial client machines.

Magic lens mechanisms (also known as movable filters and see-through tools; Stone et al., 1994; Bier et al., 1994) allow users to drag an object (e.g., a tool icon) over another object (e.g., a window or list) and immediately see alternative views of the underlying object (e.g., dragging an adjust for inflation tool over a table of prices would display the inflation-adjusted values as long as the tool remained over the table). Magic lens' offer users the ability to quickly manipulate or

view objects in arbitrary ways and to quickly undo those actions. Such tools require coordinated mouse down (to select tool), drag (to position tool), either continued depression of mouse or release to maintain activation, and drag out or click and drag out to dismiss tool action. We chose not to implement magic lens mechanisms because we aimed to support a small number of relationship views, the user population was not yet familiar with such mechanisms, and we were not aware of open source Java classes available for WWW applications.

Popup windows and menus were one of the first GUI mechanisms designers created to give users "two-and-a-half dimensions" of screen real estate. Popups are familiar to WWW users and allow subsequent freedom of action since all mouse events but one reserved for dismissal are potentially available. Popups may be activated by mouse clicks (mouse down and mouse up events), mouse down, or mouseover events (as well as keyboard events). The popup occludes what lies beneath it, must be properly placed on the screen to avoid occluding other pertinent information, may clutter the screen, and requires a dismiss action (dismiss button or move mouse out of popup). Additionally, designers must find ways to signal that popups are available. In the relation browser, we adopted the mouseover mechanism as an important part of the interface.

Mouseover events trigger special types of popups known as balloon help or tool tips. These mechanisms are popular in standalone and WWW environments and are used to provide additional information or help. In the relation browser, we used mouseover on a topic attribute as a cue to display additional information in the attribute set related to the topic (data type in the first prototype). Thus, mousing over a topic displayed a graphical bar estimate of how many web pages provided the different types of data for that topic. This mechanism allows users to quickly see what is available without committing to a mouse click and subsequent new window display or data transfer.

Coordinated lists (also known as supermenus) are multiple windows with distinct and arbitrary attribute sets. Supermenus are useful in search systems that allow users to limit queries to specific database fields. For example, in databases of research literature, searchers may wish to specify queries over certain ranges or values (e.g., date, language, methodology). This was exactly the case in the site map where we wanted people to be able to explore the relationships among statistical websites with respect to volume and name. Rather than providing four attribute sets (topic, data type, region, and dates) and allowing users to select from each to yield a result set (in effect executing a Boolean and for these parameters), we aimed to make the interaction more dynamic by showing relationships and partial results on-the-fly before users committed to the query.

Animated site maps were another mechanism considered. We have developed animated thumbnails for websites that give users a visual overview of a site (Brunk, in progress). Such a mechanism appeals to the perceptual system but even at high rates of display (e.g., 100 ms) does require time to cycle through large sites. We are working on incorporating textual information with the visual displays to provide coordinated conceptual and perceptual information to users. We chose not to include animation in the relation browser because the techniques are under development and require more substantial client machines to use effectively.

Figure 2. First Supermenu Mockup showing coordinated lists Phase One Mockups and Prototype

Topic	Region	Information Type
Economics Production, Sales, Orders & Inventories Output Income, Expenditures & Wealth Employment, Unemployment, & Earnings Prices Money, credit, & Interest Rates Transportation International Statistics Social Statistics Demographics Health Crime Education	USA State Local International	Tables Indices Time Series Data Sets News Releases Research Reports

National Center for Health Statistics (NCHS)	
National Institutes of Health (NIH)	
Agency for Health Care Policy and Research (AHCPR)	health care use and expenditures Medical Expenditure Panel Survey
Agency for Toxic Substances and Disease Registry (ATSDR)	
Health Care Financing Administration (HCFA)	Medicaid statistical data
Health Resources and Services Administration (HRSA)	health professions workforce organ transplants
Indian Health Service (IHS)	
Substance Abuse and Mental Health Services Administration (SAMHSA)	health problems
Department of Energy's	exposure to radiation

Based upon these analyses of mechanisms, we developed a paper mockup that incorporated the attributes and specified the mouseover and supermenu mechanisms. The first supermenu mockup on paper (see Figure 2) was sent to the group via email for reactions. The design was meant to work as follows. As users mouse over an item in any category (attribute set), related (available) items are highlighted. For example in Figure 2, the cursor is over health, and highlights on USA and International show that there is health data available for those regions and that there are tables, time series, news releases, and research reports available. The pertinent websites appear in hierarchical display in the window 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 (the coordinated columns are based on an underlying Relational table). Note that this technique could also be used as a thesaurus browser for non-hierarchical terms in the Fedstats A-Z list (the relationships across columns are semantic rather than topic-location-format) as well as for the query exploration example above.

Based on initial reactions by the team to this paper mockup, work began on a Java applet that provided the functionality specified in the mockup.

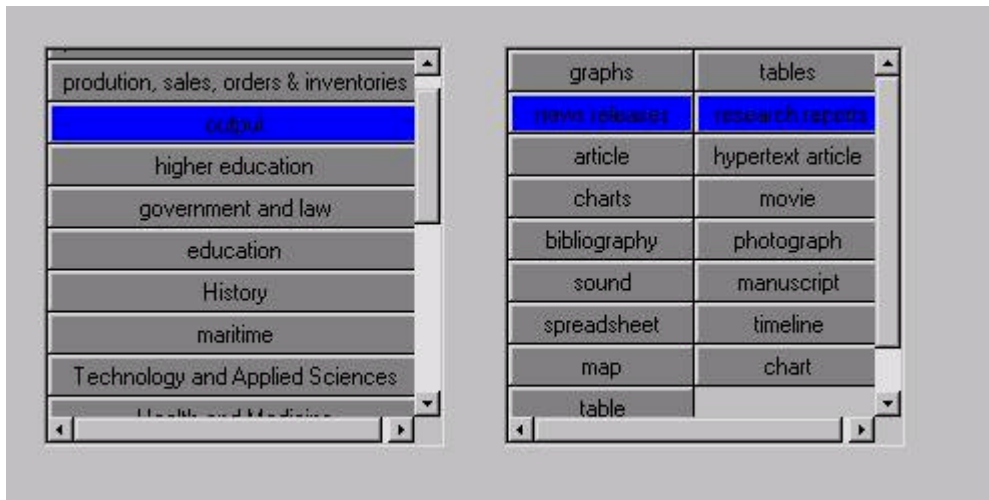
Relation Browser Mockups and Prototype

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 at meetings on January 8, 1999 at BLS. Figure 3 shows a screen dump for the prototype. The aim was to demonstrate the dynamics between the topics (e.g., output is highlighted in the left window topic list in the figure) and the data types (right window list). In the figure, we see that there are news releases and research reports available for the economic output topic. As the user rolls the mouse over other topics, the appropriate data

types are highlighted. This prototype was meant to show proof of concept and no actual data was included in the demo at the meetings in Washington.

Both teams reacted favorably to the design and after lengthy discussion, five region names 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 tested in the April-May period.

Figure 3. First Java Prototype



In developing the prototype, we decided to focus on allowing user 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 type of 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.

For the first question, we chose the 196 different websites as distinct entities and provided simple counts of how many were available for each pairwise topic-data type relationship. This approach avoided the issue of 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 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.

Paper mockups were first used to specify the interface. Figure 4 shows the paper mockup. Colored highlights were used on these mockups to show dynamics for the eventual Java prototype. The leftmost list 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. In the data collection effort undertaken to populate the prototype database, it became evident that a more detailed list is necessary (see data collection section below). The third list contained the five regions agreed upon by the Fedstats team. Given the limited screen real estate and the fact that that there were just under 200 total websites linked from the Fedstats site map, we decided that adding a fourth list for time period covered by the data was not feasible. Because we wanted to include as many attributes as possible, we decided to use two lists, topic and data type, and make region and time period 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 data on all pertinent attributes.

Figure 4. Paper Mockup

<ul style="list-style-type: none"> Agriculture (10) Crime (11) Demographics (39) Economics (31) Education (7) Energy (14) Environment (16) Health (17) Income (4) Labor (14) National Accounts (5) Natural Resources (25) Safety (8) Transportation (18) 	<ul style="list-style-type: none"> Reports <input type="text"/> Tables <input type="text"/> Downloadable Datasets <input type="text"/> Searchable Datasets <input type="text"/> Graphs/Charts <input type="text"/> Maps <input type="text"/> Photos/Video <input type="text"/> 	<div style="border: 1px solid black; padding: 5px;"> <input type="text"/> <input type="checkbox"/> International <input type="checkbox"/> National <input type="checkbox"/> Regional <input type="checkbox"/> State <input type="checkbox"/> Substate </div>
--	---	---

Data Collection.

A template was created based upon the attributes in the specifications. Figure 5 illustrates this form that was completed for each of the 196 websites linked from the Fedstats site map. 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. It is essential that a revised template (see suggestions below) be distributed to the individual webmasters responsible for these sites before full implementation is undertaken. Table 5 summarizes the data for the 14 topical areas.

Figure 5. Data collection template.

Name:
 URL:
 Program (topic):
 Agency:
 Region International____National____Regional____State____Substate____
 DataType Report____Table____Timeseries____Downloadable Dataset____Searchable
 sataseta Graphs/Charts____Maps____Photos/Video____Other____
 Time Period pre60s____60s____70s____80s____90-95____96____97____98____99____

Table 5. 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

The data collection effort suggests that reports be further distinguished as research reports, news releases, and periodic (e.g., annual) reports. Since many tables are timeseries, combining these seems sensible for novice users. Although photos/video instances were few at the time of collection, we encourage keeping this category to anticipate and suggest new data additions. Because this data collection effort is substantial, the Fedstats task force is encouraged to collect the data from webmasters electronically, adding the results to a database from which the interface can read data. This would also allow for updates and additions over time. In our prototype development, the data was collected on paper forms and then transferred to a Microsoft Access database. The Access database was exported as an ASCII text file that was

then used in the Java applet. A full implementation applet should read the data directly from the underlying database.

The Prototype

The prototype was developed over the February-April period and was user tested in late April and early May. Ben Brunk was responsible for the Java programming. Figure 6 illustrates this prototype. The topic list was implemented as specified in the paper mockup with parenthetical values for the number of websites available for each topic. A select all and a clear all button 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 continual exploration through mouseovers. In the figure, the user has clicked on 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. See Appendix A for the Java class architecture. Because the alternative site map allows people to explore pairwise sets of relationships, the prototype was called the relation browser (RB).

Usability Testing

To examine the efficacy of the prototype, a usability test was conducted with nine subjects. A testing prototype was drafted and discussed via email. The main goal was to determine how easy the relation browser was to use and how well it provided an overview of the Fedstats site and associated data. Because exploration is difficult to assess empirically but is entailed in a more easily observable task like searching, a set of six search tasks for the existing site map and the relation browser were devised to give test subjects some exploration guidance. In hindsight, this decision clearly biased users toward using the relation browser as a search tool rather than as an overview explorer. The final testing protocol is provided in Appendix B.

To focus on the site map and relation browser features, we selected subjects who were experienced WWW users. Test participants were recruited by BLS staff. Six subjects 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 on April 30 and May 5, 1999. 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. Anita Komlodi conducted the tests and took written notes during the interaction. Sessions lasted between 40 and 75 minutes depending on how verbal subjects were. In all cases, subjects used the Fedstats site map first and then the relation browser 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 relation browser was a viable additional alternative rather than a replacement for the site map.

Subjects had mixed success in using the Fedstats site map and relation browser 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 Relation Browser. Domain knowledge was clearly a factor in finding information. One 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." Another 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).

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. One subject noted that it took too many clicks to get information.

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 Window 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 relation browser, subjects were asked a series of questions. When asked whether it was useful in gaining an overview of what statistics are available, 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 relation browser 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 relation browser was more appropriate to expert users and liked the Fedstats descriptive texts for websites and its navigation familiarity. Subjects said: [Fedstats 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 relation browser 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).

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:

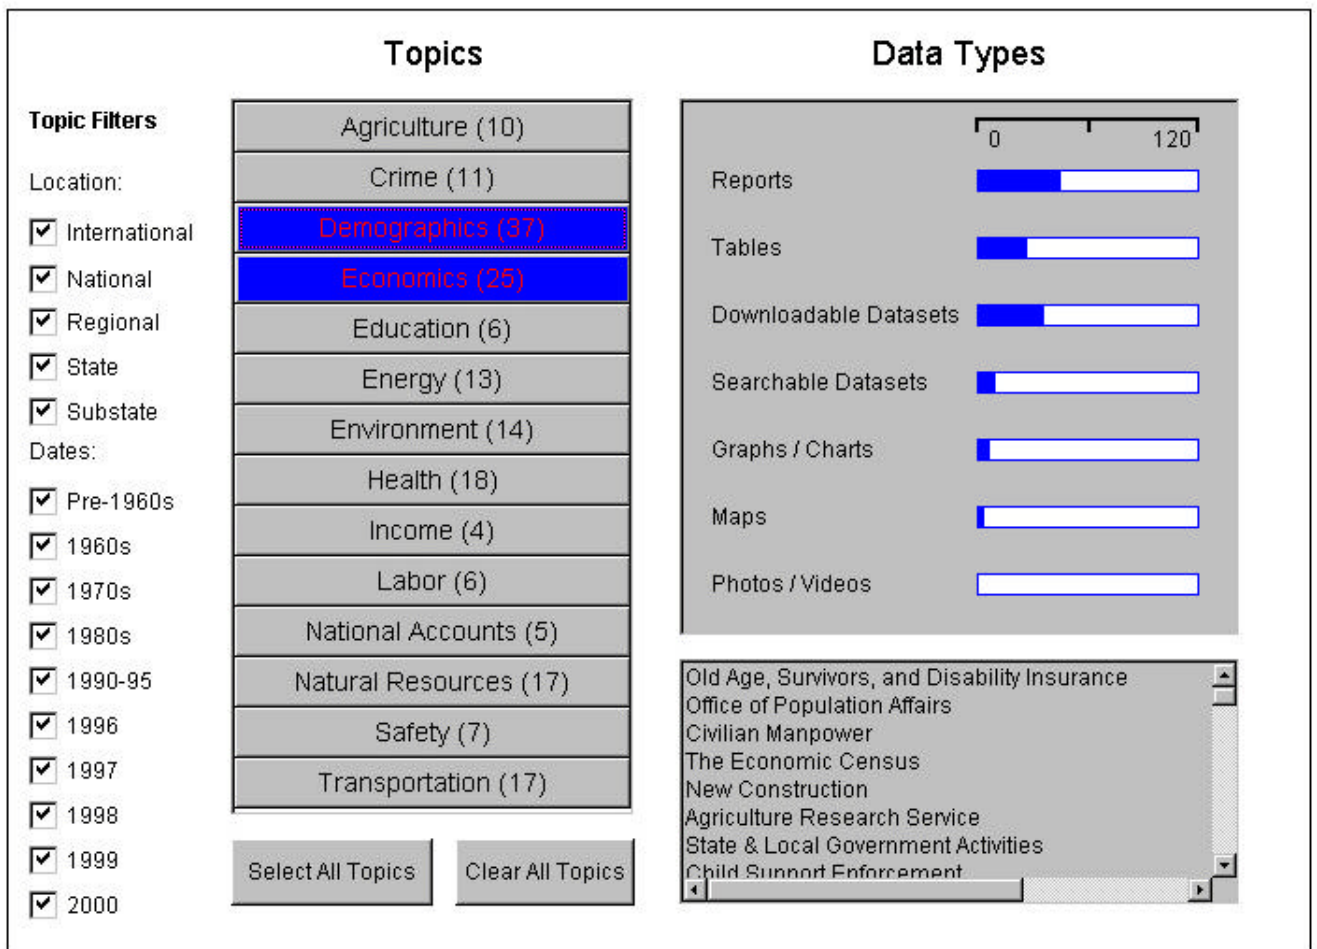
- Click on the URL to get a synopsis as in Fedstats.
- Label the 0-120 scale
- Use a better title than 'data types'
- Add search capability (two subjects)
- Show which websites have tables (in the result list)
- Add this [RB] as a link from the Fedstats site
- Change the highlight color to be different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

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.

Figure 6. The Relation Browser Prototype.



Revised Prototype.

The prototype and preliminary testing results were presented and discussed at a meeting at BLS on May 10. 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 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

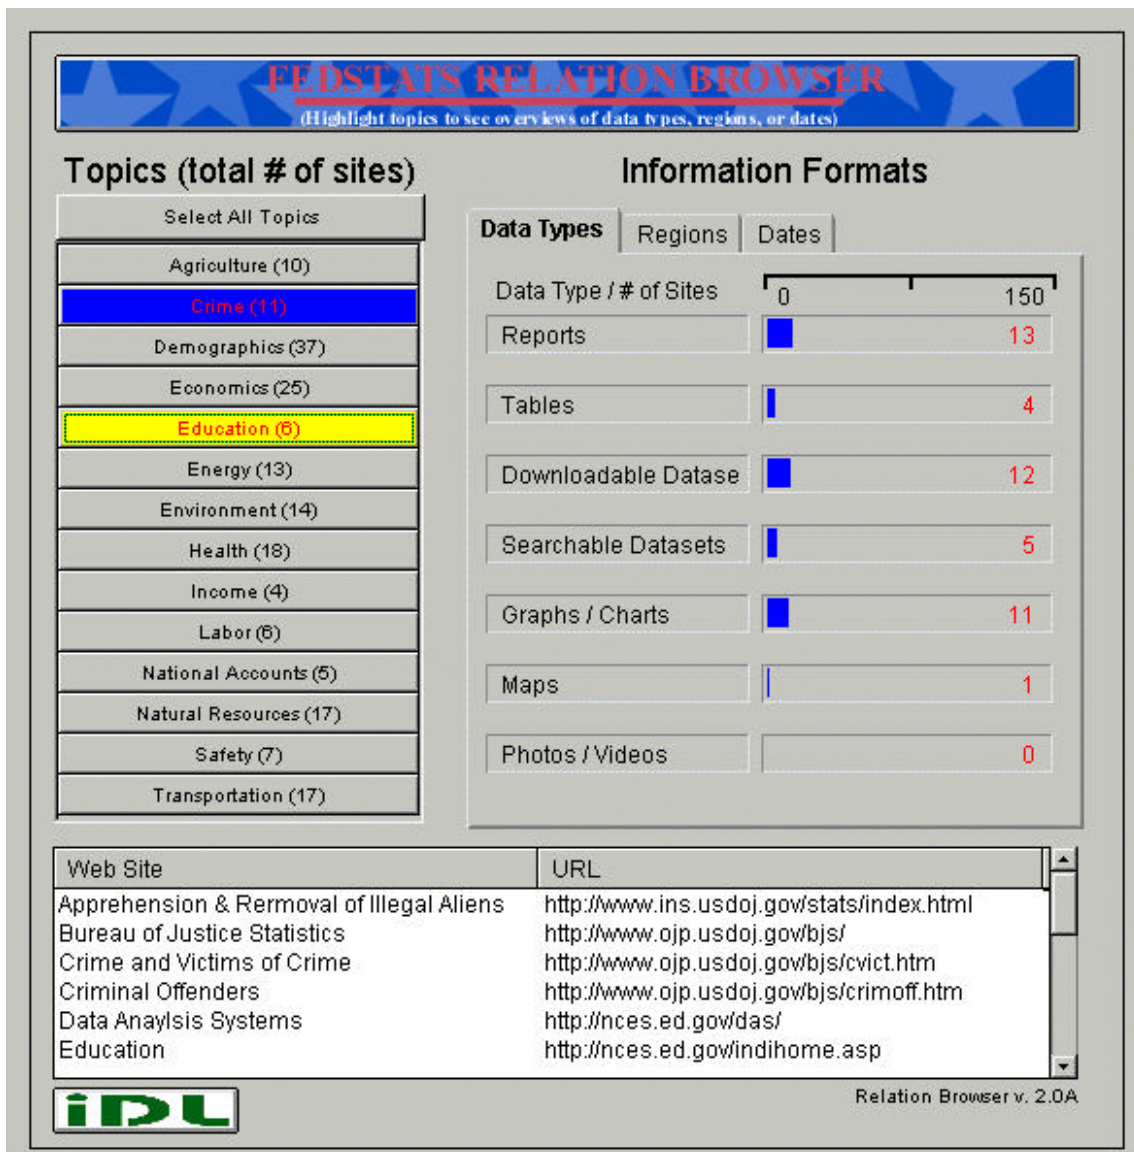
easy to overspecify a query and get no results. This is the classic "no hits" problem that has long plagued casual users of large online databases. Second, manipulating the two filter sets is awkward and confusing. Clicking off or on settings is awkward without additional buttons such as "select/deselect 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 yields less, whereas the linguistic connective AND yields 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 treated just as the data types and a tab mechanism added to allow users to relate topics to either data types, regions, or time. 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 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.
- Alt tags were added in the HTML to make the prototype Bobby¹ compliant.

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

¹ Bobby is a tool provided by the Center for Applied Special Technology (<http://www.cast.org/>) that analyzes web pages for their accessibility to people with disabilities.

Figure 7. Revised Prototype of the Relation Browser



Summary and Recommendations

This project focused on creating an alternative site map 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 relation browser 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 relation browser 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. 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. First, Fedstats should consider conducting a field test for the revised prototype. Such a field test would give users the option of using the relation browser in addition to the existing site map. Such a test could be time-bound to a few weeks and a feedback button linked to a brief questionnaire. Such a field test would inform decisions about whether to implement a full-scale relation browser for Fedstats. 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 distributed the template electronically to webmasters at the various federal statistics websites, collected the completed templates via email, piped results to a database, and used the database for populating the relation browser applet. A policy on updates would also be required so that new websites and changes to existing websites could be made automatically.

Other recommendations are not specific to the relation browser but to the organization of the current Fedstats site. First, links from Fedstats should not go to home pages of agencies but rather to the pages that give statistics. Sending a citizen to the home page of a large agency from

a statistical portal site is a disservice, yet the current site does this in several cases. Second, the overall Fedstats site and the site map should become more topic centered and less agency centered. The fourteen program areas reflect agency organization rather than topical organizations. The A-Z list in Fedstats is perhaps a more useful entry point in this regard. One alternative organizational scheme is presented in this report.

The primary promise of alternative site map tools 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 agencies as they are able to examine and compare the types and volume of information they provide through the Fedstats portal. Agencies that link to main home pages rather than statistical pages will appear to have little data in the RB and those that provide many pages for little data will become obvious to users and other agencies alike. These appearances may help resolve policy issues that recur in our work and in digital libraries in general. In addition to providing an alternative tool for user exploration, a tool like the RB can help agencies develop common agreements and highlight trends and gaps in their data.

References

Bederson, B. & Hollan, J. 1994. Pad++: A zooming graphical interface for exploring alternative interface physics. In *Proc. of ACM UIST '94*. 17-26.

Bier, K., Stone, M., Baudel, T., Buxton, W., & Fishkin, K. (1994). A taxonomy of see-through tools. *Proceedings of ACM CHI '94* (Boston, MA, April 24-28, NY: ACM Press, 358-364.

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., 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 *Proceedings of Designing for the Web: Practices and Reflections* (3rd Conference on Human factors and the Web, Denver, June 12, 1997). <http://www.uswest.com/web-conference/proceedings/nation.html>

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.

Stone, M. Fishkin, K., & Bier, E. (1994) The moveable filter as a user interface tool. *Proceedings of ACM CHI '94* (Boston, MA, April 24-28, NY: ACM Press, 306-312.

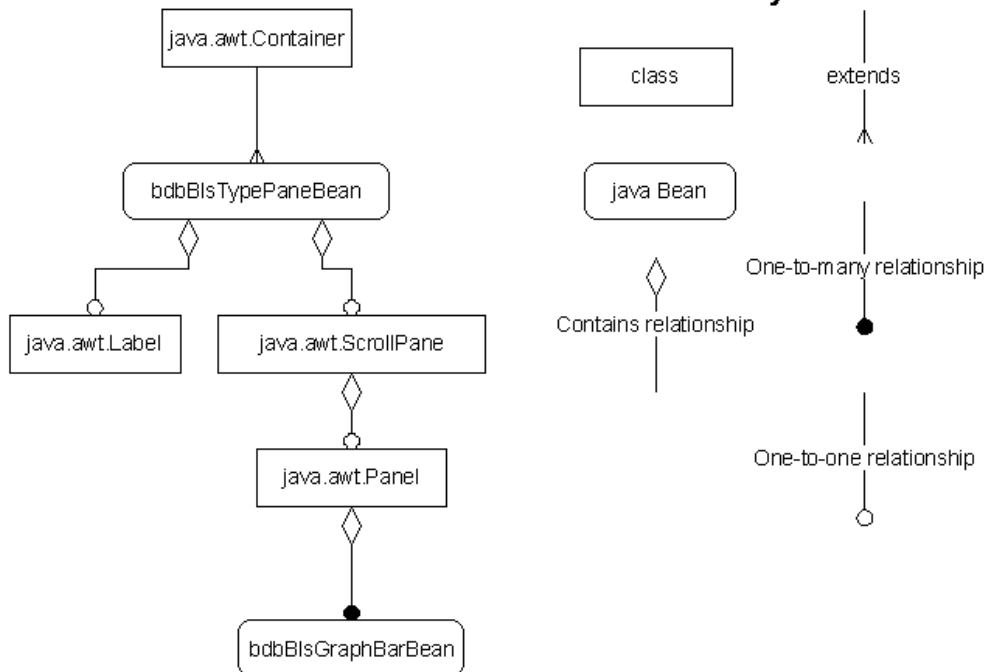
Tenissara, R. (1998). *The effectiveness of frames and expandable table of contents as user interface in hypertext information retrieval*. Unpublished doctoral dissertation, Indiana University.

Zaphiris, P., Shneiderman, B., & Norman, K. (1999). *Expandable indexes versus sequential menus for searching hierarchies on the World Wide Web*. Technical Report 99-15 University of Maryland Human Computer Interaction Laboratory.

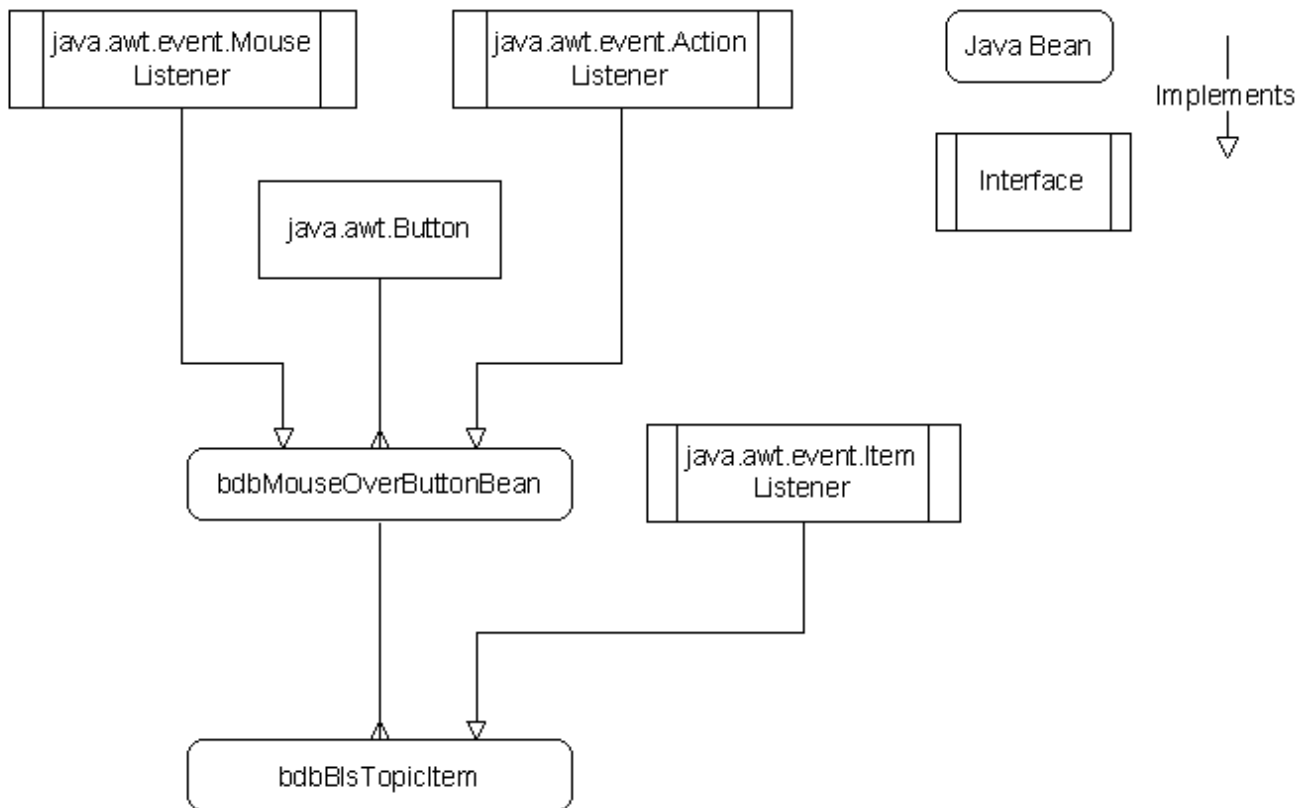
Appendix A. Java Architecture

The Java applet was developed using Symantic Café. There are three main Java Bean components, two of these are composed of several smaller Bean components. The three main Bean components are bdbBlsTypeBean, bdbBlsUrlBean, and bdbBlsScrollPaneBean. Other Beans include bdbMouseOverButtonBean, bdbGraphBarBean, and bdbAmountBarBean.

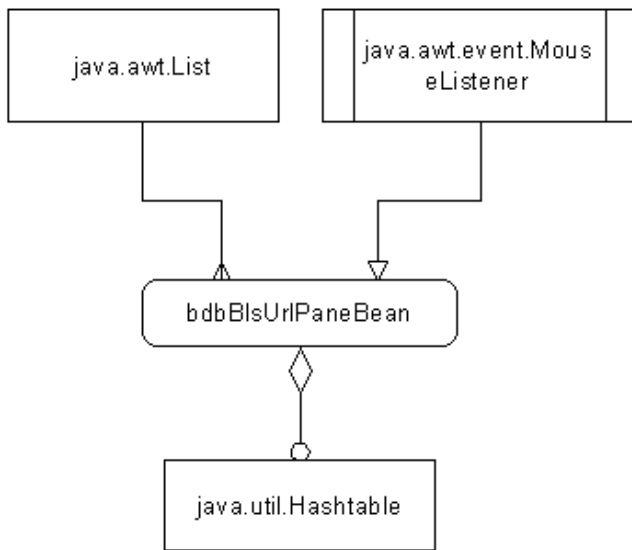
bdbBlsTypePaneBean



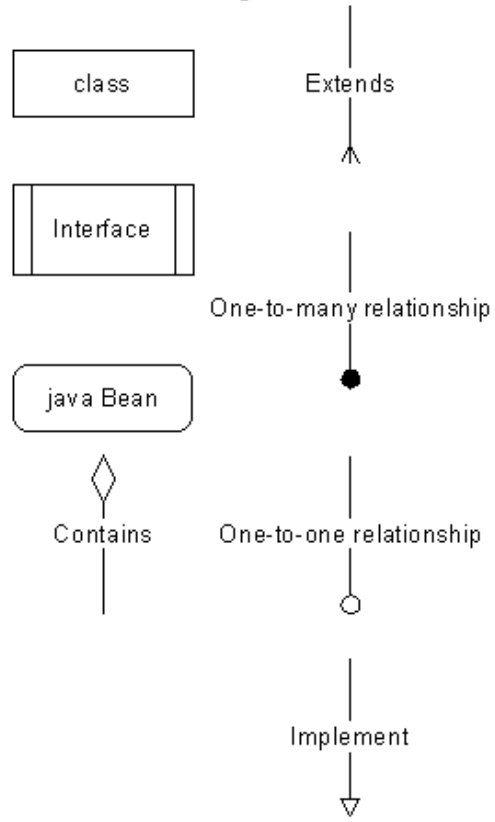
bdbBlsMouseOverButtonBean & bdbBlsTopicItem Inheritance



bdbBlsUrlPaneBean



Key:



Appendix B. Usability Testing Protocol

Fedstats Overview Usability Testing

A. Welcome

[subject invited to sit at computer, monitor off. Fedstats and RB are loaded with Fedstats in active window, full screen mode]

[be sure to clear cache before each subject arrives]

My name is Anita Komlodi and I'd like to thank you for helping us get preliminary reactions to a prototype interface tool. This tool is meant to help people get an overview of the different statistical products made available by 70 government agencies. This session will take no more than one hour and it is important that you understand that we are interested in assessing this system, not your knowledge or performance. We will videotape this session so your face may be linked with your actions and shown in research meetings outside BLS but none of the data that we collect will be attached to your name.

First, tell me a bit about your experience.

What type of work do you do?

How often do you use a computer?

How often do you use the WWW?

Do you regularly use federal government statistics?

Have you ever used the Bureau of Labor Statistics website?

Have you ever used the Fedstats website?

Other statistical websites?

B. Fedstats

[turn on monitor]

Look at this home page and please read the statement of purpose. [give time to read]

One way to get an overview of a website is through a site map--usually a table of contents for the site. Please click on site map.

Does this overview help you get a sense of what is available?

I'm going to ask you to start with the site map and find some specific information. We are more interested in how you would try to find these statistics than in actually locating them. [After 3 minutes, stop subject, thank them, and go on to next question]

1. Beginning with the Fedstats site map, find the average per capita health care costs for a US citizen. [\$3925 in table, just shy of \$4000 in text at HCFA]
2. How much coal was produced in California in 1997 [none from table]
3. What is the latest value of the consumer price index? [165 in March 99 at BLS]

[turn off monitor]

How useful is the site map for finding specific statistical information?

How useful is the site map for gaining an overview of the kinds of statistical information produced by the federal government?

C. Relation Browser

I am now going to ask you to try out an alternative tool for the site map you have just used. I'd like you to think about the pros and cons of such a tool as we work and at the end of the session we will discuss them.

[turn on monitor and activate RB window, ask subject not to touch the mouse until asked to do so]

Please study the screen. The topics are the same as those on the site map. What do you think the numbers in parentheses after each topic mean?

[After response, tell them this is the number of federal websites that have statistical information on that topic]

What do you think the data type column is for?

What do you think the filters mean?

Please use the mouse to roll over the topics but do not click the mouse.

What do you think is happening in the data type column?

Please click on a topic. What do you think is happening in the data type column? In the results box?

Please click on a website. Within that website, you will find statistical reports, tables, etc.

Close the website.

Click on the topic again. Note that it releases the highlighting and list of websites.

Click on another topic. What do the bars tell us here?

Click on another topic. Note the data type column and results box. What happens when more than one topic is selected?

Click again to release all the topics.

Suppose you only want international statistical data for the 1980s. The filters are meant to allow you to limit your browsing to certain regions and periods. Click off all boxes except international and 1980s.

Roll over the topics. How does this affect the number of web pages available?

[Turn off monitor]

Let's discuss this tool. Based upon this brief introduction, please tell me your first impressions about this tool. What might it be used for? How is it like the site map you used earlier? How does it differ from the site map?

I will now ask you to do a few tasks like the ones you did earlier.

[turn monitor back on]

1. Find the average K-12 public school expenditure in the U.S. [\$7163]
2. Find the most recent cost of a barrel of crude oil.
3. What is the latest producer price index value?

[Note for all the tasks, do not let users go very long--we are interested in how they start and use the site map and RB]

How many websites have downloadable data sets for health?

How many federal websites offer demographic statistics?

What do the bars represent?

What is the difference between mousing over a topic and clicking on the topic?

What happens when the filter check boxes are deactivated (e.g., if only substate and regional are checked)?

What is meant by substate regions?

What is meant by the time period filters?

D. Debriefing

Would you find this tool useful if you were trying to understand what statistics were available from the federal government; that is, does it provide an overview of the available statistics?

Would you find this tool useful if you were trying to find statistical information like the number of public high schools in the US?

How would you compare the relation browser to the Fedstats site map?

What was confusing about this tool?

What could be improved?

What could be added?

What is the best thing about this tool?

What is the worst thing about this tool?

Thank you for your help in testing this prototype tool!

