Finding and Understanding Government Statistical Information

Gary Marchionini, Stephanie W. Haas, Junliang Zhang, and Jonathan Elsas
University of North Carolina at Chapel Hill

There are many challenges to successful digital government operations, but serving large, diverse populations of people is one of the foremost. This is doubly difficult for information that is highly codified such as statistical data. Such data are important for citizen decision making, but low levels of numerical literacy limit widespread application and government statistical services aim to make such information more broadly accessible and understandable. The Govstat Project has been working with federal statistical agencies to make government statistical data findable and understandable (www.ils.unc.edu/govstat). Our approach is twofold: define a Statistical Knowledge Network (SKN) that can evolve over time as citizens and government agencies at all levels share and integrate statistical information; and create user interface prototypes that make finding and understanding easy for diverse populations. The overall project has several threads of work but this paper focuses on two specific user interface threads: interactive glossaries as examples of our general efforts to provide multiple layers of online help, and a family of highly interactive user interfaces for exploring data.

We start with a basic belief that a statistic, in isolation, is useless. If people are to find and understand statistical information, we need to surround statistical data with supportive information and means of manipulation that will reveal “the story behind the number,” or as one of our agency partner colleagues (John Bosley, Bureau of Labor Statistics) puts it, “no naked data.” Statistical experts depend on a variety of scaffolding to find and interpret federal statistics; such as the BLS Handbook of Methods or codebooks associated with surveys. We seek to selectively modify and augment the agencies’ scaffolding so that it will be useful to nonexperts as well.

Metadata describing the statistical information is one type of scaffolding we wish to incorporate into the SKN. It is crucial for understanding what a statistic represents, but is equally crucial for supporting the variety of actions users need to perform as part of accomplishing their tasks, such as searching and browsing, filtering, merging, and aggregating. For example, if someone wants information about individuals’ health, income, and education, they might seek to combine information from NCHS and the Bureau of Census. But he or she might not realize that levels of education are reported in different ways by different surveys: level of attainment (i.e., diplomas), and number of years of school attended. The metadata associated with the surveys holds that information: the challenge is how best to bring it to the users’ attention.

The life-cycle of a statistic provides several stages at which metadata is created, from the creation of the survey questions and variable definitions, to the row and column labels for a table drawn from the survey results. Ideally, the metadata should “travel” with the statistics from the point at which the raw data is produced, through aggregation or other statistical processes, so that it can be used by a visualization or manipulation tool, or be easily found by a user. Previous research (Denn et al. 2003) has shown that although agency staff may be able to find it (or know whom to ask), it isn’t always easy to find on the website, and sometimes isn’t there at all. Agencies are becoming more aware of the need to design their production processes to keep the metadata attached to the data itself. , but providing metadata is not yet a well-established part of the agency’s production cycles, and there have been few convincing demonstrations that there is real advantage in doing so. The GovStat prototypes are a step in that direction.

Layering Online Help: The Case of the Statistical Interactive Glossary

Government agencies provide enormous amounts of information via websites and a common experience is that the more information that is made available, the more email questions an agency gets as more people access this information. For example, from 1996 to 2000, the Bureau of Labor Statistics (BLS), saw a five-fold increase in email requests to the website help desk. Responding to these additional requests requires more time and effort by agency personnel and hiring more people to serve these needs is politically and economically difficult. Online help is one way to mitigate some of these pressures, however there is a significant rub. Online help has a bad reputation, often considered only as a last resort. But as more information, and more complex information, is placed online, improving the quality of help, including types of help mechanisms, its relevance to people's tasks, and how people know what is available, becomes more and more critical. It is too easy for someone to either abandon the task, or settle for less than good (or even erroneous) results, or send email asking for what they cannot find. In the commercial sector, people can be charged for such services; this is not politically feasible for government services where public funds are required. Thus, online help is particularly attractive for government agencies but it must be effective, efficient, and attractive so that people take advantage of it.

When people are learning a new skill or new sets of concepts, they generally welcome guidance from more knowledgeable sources (e. g., teachers). Receiving help is viewed as a normal part of learning. As they become more proficient at the new skills, less help is needed. The zone of proximal development (Vygotsky, 1978) refers to a stage where one can successfully accomplish a task with a little assistance that can not yet be done independently. As mastery is gained, less help is needed, and means of assistance should fade into the background (Carroll & Carrithers, 1984). One way of supporting different people who may be in different proximal zones of understanding is to design several layers of help. We aim to layer help with systems as well as help with the information content. In the case of help for a search tool, there are layers of functionality that guide layers of progressive disclosure of online assistance to minimize information overload (Kang, Plaisant & Shneiderman, 2003). This provides an easy way of getting over the common initial roadblock of "where do I start". In addition, it may pique the curiosity; having seen what can be accomplished with the first set of information, users may feel more confident in exploring on their own, or may continue on to the next level of help. From this framework, we derive design guidelines.

First, since many people visit statistical websites or seek statistical information only occasionally, expertise in using the sites and the information therein should not be the goal. People will not put a great deal of effort into learning to do a task they never expect to do again. Therefore, help information should be easily accessible.

Second, the amount of help information presented at any one time should be limited; rather than overwhelm the user with a full tutorial; just enough should be presented to enable the user to proceed with the next steps of the task at hand. In some cases, a discrete nugget of information such as the definition of a term, presented at the right point in a task, may provide sufficient guidance. We refer to this as "just-in-time, just-enough" help. In other situations, more comprehensive guidance is needed, layered to first present the information most likely to be needed. It is important to recognize that one kind of help is not appropriate for all types of information or user situations.

Third, since designing and developing good help requires time and effort, it makes sense to provide help for the tasks and information that people need most frequently. This requires a good understanding of the information that people most frequently seek, as well as the common problems they encounter. The statistical agency staff have been a rich source of information, augmented by user studies and query log analysis. It also suggests that efforts to streamline the development of help tools increase the likelihood of their eventual deployment.

The Statistical Interactive Glossary. The Statistical Interactive Glossary (SIG) provides just-in-time, just-enough explanations of statistical terms that people are likely to encounter on the agency websites. When someone needs to understand a column or row label, he or she needs the definition at the time he or she encounters the term, ideally without having to interrupt the task at hand. Interruption could mean leaving the current web page to search for and consult a glossary, or having to wade through a complete lesson about the term in order to find the specific nugget of information that is needed to complete the task. Either way, the train of thought has been broken, and cognitive effort will be needed to resume the task. Even worse, if looking up the term seems like too much effort, the person may just guess at the term's meaning, or just abandon the task.

A SIG explanation may be viewed by clicking on an icon associated with the term where it appears on a Web page. The scope of the SIG is limited to common statistical terms that users are likely to encounter on agency websites, and that our agency colleagues believe are important for people to understand in order to understand and use statistical information. The content of the SIG explanations was derived from existing agency glossaries and widely-used textbooks, and reviewed by our colleagues for accuracy and clarity. SIG explanations present declarative information about the term, rather than procedural information such as how to calculate a particular type of statistic. We have kept the explanations as brief and concise as possible.

A SIG explanation may be generally applicable or context-specific. A context-specific explanation incorporates entities or examples from the specific table or agency publication to which it is attached. For example, an explanation of *index* that is attached to a table showing the *Consumer Price Index* will incorporate the CPI into the explanation. Another version could use *Body Mass Index (BMI)* in its content, and be attached to occurrences of the term. The goal is to minimize the user's need to shift attention to entirely different topics and entities. It is not possible to create versions for every single type of index used by the statistical agencies; a general purpose explanation can be used in any context.

Each version (context-specific or general) of a SIG explanation is presented in three formats: text only, still graphics + text, and animated graphics + text. Research in multimedia teaching or tutorial materials provides mixed evidence for the effectiveness of various formats, which is not surprising given the many factors involved, such as learning style, fit between content and format (conceptual congruence), and dual coding in visual and audio channels (Blank et al., 2003; Morrison & Tversky, 2001; Paivio, 1986). But the agencies' wish to minimize the hardware and software requirements for their users, and their resulting reluctance to include audio was also an important consideration. Finally, we wanted people to find the presentations attractive; the best help in the world is useless if no one looks at it. To the greatest extent possible, we kept the content equivalent in all formats. Figure 1 illustrates an entry for the consumer price index.

We created the Content Development Template (CDT) (Wilbur & Haas, 2004) to automate some aspects of the development of the animations (e.g., placement and timing coordination of animations and captions). Designing effective graphics and animations is still a challenging job, but we wanted to help the agencies spend more effort on the cognitive and creative parts, and less on the mechanical parts.

We evaluated the effectiveness of SIG presentations in helping people answer questions about the meaning of statistical terms, as seen on agency websites. Details are provided in (Haas et al, 2005). Overall, study participants were able to answer more questions correctly after viewing presentations than before, and their confidence that they had chosen the correct answer increased. We found no significant difference in performance based on presentation format. In general, participants liked the presentations and found them helpful and attractive. But many participants preferred plain text or graphics + text, especially for "straight-forward" definitions; these versions seemed to be perceived as more efficient at delivering the information. Graphics + text and animations were seen as helpful in "showing how things worked". Animations were "really cool", but did not seem to have a real edge in the perceived effectiveness of the explanation in most cases, perhaps because they took longer to view. Although novelty can help users attend to a presentation, there is also a risk of having too much going on.

The combination of our findings, agencies' requirements and constraints, and the notion of conceptual convergence suggest that most definitions be presented in text and text + graphics. However, explanations of selected terms might benefit from animation. Some study participants found the animations especially easy to understand for *adjustment* and *index*, two terms that contain a semantic component expressing action or activity (as opposed to terms such as *mean* or *median*, for example). The animation did not seem like just an add-on, but rather an integral part of the explanation. Our experience in developing the animated versions provides support for this from another direction. Designing animations for *adjustment* and *index* was much easier than for other terms, such as *population* or *race*. It is extremely difficult to incorporate motion into an explanation of racial categories in any meaningful way. So if we can discover the characteristics of terms that are well-suited for animated presentations, the extra development effort will result in improved comprehension on the part of users. The agencies will maximize their effort by designing text and graphical versions for most terms, and limiting animation to those that need it.

Exploring and Understanding Information: The Case of the Relation Browser User Interface

As part of our ongoing efforts to design and deploy user interfaces that help people to easily find and understand information, we have developed a family of interfaces that we call Relation Browsers (Marchionini & Brunk, 2003). The key ideas are to facilitate exploration of the relationships between (among) different data facets, display alternative partitions of the database with mouse actions, and serve as an alternative to existing search and navigation tools. These particular design goals are rooted in the general principles that drive what Shneiderman (1994) calls dynamic query interfaces. In our interfaces we aim to: give users the ability to look ahead without penalty, minimize scrolling and clicking, provide alternative ways to slice and dice datasets, closely couple search, browse, and examine functions, offer continuous engagement, and bring data treasures to the surface. Our project partners at the University of Maryland have focused on applying these principles in interfaces that support exploration across several facets and with a range of specialized interface widgets specific to data types (e.g., maps for geographic data) and on novel interfaces for special populations such as an audio browser for maps. Our work has focused on the more basic dynamic query design that limits overall exploration but is very simple to use. The design concepts for a relation browser are to:

- Provide people with a small number of data facets such as topic, time, space, or data format; each of which is limited to a small number of attributes that will fit easily on the screen. These facets are displayed in columns or tabs and may be thought of as distinct ways of slicing the database. The attributes are displayed as rows within each column and may be thought of as subsets (selections) of the database. To promote understanding of the database and to cue searching, it is important that for each subset, the number of items in that subset is shown. In the RB, these numbers are shown absolutely as numeric values and relatively as graphical bars so that each column is visually displayed as a histogram of the number of items by attribute.
- Provide people with agile mechanisms for displaying the relationship between any single attribute and all attributes in the other facets. The control mechanism used in RB instances is the mouseover event and the visual display is to show the numbers and bar lengths on all foreign attributes (those in other facets) dynamically when the mouse is placed over any attribute.
- Allow people to select specific subsets (through a mouse click) and continue to explore relationships in the selected subset.
- Allow people to retrieve all records in a subset and search within it

Figure 2 illustrates an instance of the RB for FedStats. Over time, we have implemented relation browser interfaces in several different ways and conducted user studies to examine their effectiveness. One early instance for FedStats data used a flat file backend design and was beta tested on the FedStats (www.fedstats.gov) site for nine months in 1999-2000. Users were satisfied with the dynamics and interaction style, but were unhappy to have to download the Java plug-in. Given that most people have Java-capable browsers today, this is not as large an issue. The current version uses a database backend and adds fast string search capabilities for results (see Marchionini and Brunk, 2003 for details on the early evolution, and Zhang & Marchionini, 2005 for details on recent instances). A recent user study with 17 people compared the effectiveness of the RB to a standard web-based interface (with the usual query box and selection lists) and demonstrated that the RB interface outperforms the standard web-based interface for queries that involve searching over multiple attributes or making inferences (Zhang & Marchionini, 2005). The RB was strongly preferred by users as well.

Challenges to Highly Dynamic Interfaces for Large Government Websites. Although the general goals are appealing and people who use these highly interactive interfaces find them appealing, there are several challenges and limitations that fall into three categories: scalability, metadata, and user expectations. Scalability is a challenge on two fronts. First, the RB family has inherent limitations for data that has a large number of attributes. If there are more than a dozen or so attributes for any facet, it will be difficult to display them on a single screen. Scrolling destroys the dynamics of such interfaces and thus fairly coarse or hierarchical layers of attribute sets of are required. The other scalability issue relates to getting the metadata for the corpus to the client side so that rapid dynamics are supported. We have experimented with reading the metadata from a database server in real time but even under optimal conditions (e.g., dedicated database server in same room), the latencies are too long to maintain the dynamic refreshes needed as people mouse over any attribute. The problem becomes getting the number of items in all possible subsets to the client side to insure dynamic interactions. Consider an e-commerce application with product category, price range, and manufacturer as facets, each with a dozen attributes. Transferring the number of items for all possible database selects must be delivered to the client, which for many facets and attributes can take several seconds even with broadband connections. In our implementations, we constrain this combinatorial explosion by limiting selections within a column to Boolean ‘ORs’ and across columns to Boolean ‘ANDs.’

In addition to sending these metadata to the client, there is the challenge of obtaining metadata for extant information services. Our experience with federal statistical agencies illustrates that most of these websites grew over time to enormous size without consistent information architecture standards within an agency (let alone across agencies). Thus, naming schemes, organizational structures, and data flow models vary enormously. For example, in our early work with FedStats, there were extensive discussions about how to best define a geography facet. Although international, national, and state attributes were easy to agree on, finer grain divisions were more difficult because some states do not have counties and cities and metropolitan area definitions vary by survey and data set. We eventually settled on ‘substate’ as an attribute category for all data specific to geographic regions smaller than a state. Of course, the more difficult challenge is to then classify every webpage or data item according to one or more of these attributes. Manually indexing hundreds of thousands of web pages on federal agency websites will take enormous resources to do, thus, automatic classification techniques must be leveraged.

Because topic is a facet in almost any RB application as well as a critical component in an information seeker’s need, we have investigated several ways to automatically extract topics from statistical agency websites and use those topics in RB instances. Our approach is to apply machine learning techniques to the statistical websites in order to discover a small set of topical clusters and then classify all the webpages into one of these distinct topics (Efron et al., 2003). . Over the past two years, we have experimented with a variety of techniques to build these site-wide document classifications. Our approach has evolved considerably over this time and the following section presents an overview of the present multi-phase system.

- **Acquire**. The documents from an agency web site are mirrored locally with a web-crawler. Currently this is limited to only HTML documents.
- **Build Representation**. The documents are converted from unstructured natural language text to structured numeric representations in order to be suitable for use in statistical machine learning algorithms. This is currently done by building word histograms with the text from the document titles, any available "keyword" or "subject heading" metadata, and link-text from incoming links.
- **Filter**. Common stopwords such as "and" and "the" are removed from the vocabulary, as well as infrequent terms that occur in fewer than approximately 10% of the documents. Additionally, terms specific to government statistical agencies that do not provide strong topical evidence are removed. These include the terms "bureau", "statistics" and "agency".

- **Project**. The resulting matrix of term counts by documents is then reduced in size through Latent Semantic Analysis. This technique can arbitrarily reduce the size of a matrix wile retaining the underlying structure of the data. Currently, the term-space of the data matrix is reduced to between 50 and 100 dimensions.
- **Cluster**. The Expectation-Maximization (EM) clustering algorithm is then applied to the dimensionality-reduced matrix, resulting in a probabilistic assignment of documents to clusters. For some agencies we selectively chose documents to use in the training of the clusterer in order to take advantage of better-structured or more comprehensive portions of an agency's web site.
- **Name clusters**. The clusters produced in the previous step are only sets of related documents, and must therefore be labeled in order to be useful for the RB interface. This step was performed manually, looking at sets of highly associated terms with each cluster. The term-cluster associations built for this step include the most frequent terms and terms with the highest log-odds ratio for each cluster.
- **Assign pages to clusters**. Each page is then assigned membership to one or more clusters, based on a probability threshold. We have found that for the great majority of pages, there is an obvious single cluster membership.
- **Import to RB**. Once the pages were assigned to named topic clusters, the data was imported into the RB database scheme. Note that other facets such as geographic coverage and date were determined separately (e.g., we used knowledge-based heuristics to classify pages into geographic categories). The RB instance was then ready for use on the WWW.

We went through several major iterations of this process with different statistical agency websites (Bureau of Labor Statistics, Bureau of the Census, Energy Information Administration, National Agricultural Statistical Service, and Social Security Administration). It became clear that a fully automatic solution was not feasible and a more realistic way to approach the problem of discovering metadata for existing websites is to take a hybrid approach where people use computational processes to arrive at a satisfactory solution. Deep knowledge of the content and organization is important to all steps in the process except the projection, clustering, and importing steps, where good understanding of the computational tools is required. Thus, a team approach seems more prudent if these techniques are going to be useful in agency workplace settings.

It is difficult to evaluate classification solutions because a large collection can be 'sliced and diced' in many useful ways and assigning pages to a cluster (category) is often a subjective assignment. Thus, the best way to evaluate a particular solution is to have multiple human judgments--an expensive undertaking. For the BLS runs, we met with a small group of BLS personnel to critique one of our early solutions. We also worked with the FedStats group to install the BLS and FedStats RB instances on a FedStats test site and invited people within the various statistical agencies to try them out and provide feedback. A Relation Browser instance for the BLS is illustrated in Figure 3. To formalize the evaluation for this version a bit more, we asked BLS personnel to classify 50 randomly selected webpages from the BLS site into the ten categories. Seven people completed the web-based task of classifying the 50 pages and we used an inter-rater reliability statistic (Cohen's kappa) to compare the pair wise human-human and human-machine assignments. As expected, the overall human-human agreements were better than the human-machine agreements. For some categories, where there was strong human-human agreement, there also was strong human-machine agreement (e.g., Safety & Health, Prices), and the poorest human-human agreements also had the poorest human-machine agreements (e.g., Productivity, Occupation by Job). Although humans do better at classifying webpages, the machine classifications seem to parallel the same kind of 'fuzziness' trend as people when categories are less precise (e.g., occupation by job is easily confused with occupation by place) or more conceptual (e.g., Productivity).

Our final BLS results took advantage of particularly good training sets in the clustering step (we used the Editor's desk pages as the training set) and other agency results are similar at best. Whether a more ambiguous classification is satisfactory for an agency website is a policy decision. If the alternative is manual classification of tens of thousands of webpages, then using these hybrid human-machine techniques may be a good way to provide citizens with an alternative access mechanism to the existing search and browse websites. We argue that such alternatives will not only help people find what they are looking for, but also understand the overall scope of an agencies portfolio and contextualize and understand what they find as a result.

The third challenge to RB-like interfaces in government websites is that their deployment will increase user expectations about finding useful results. Improved service tends to lead to increased demand, which is welcome in business environments where costs may be passed on to consumers, but in fixed resource environments such as government agencies, increased demands may strain people and systems. One of the promises of IT is to increase productivity so this may be a non-issue over the long haul, but people who work in agencies and must adapt to these changes are clearly affected by these demands, in effect, having to run several parallel data dissemination operations while adding new services.

Conclusion

The challenge of helping diverse populations of people to find and understand statistical information is daunting. Academic research generates new ideas and techniques but transferring them to government practice takes time and patience. Academic researchers like to see their creations implemented and government agencies must work within constraints of policies that must be insensitive to short-term trends. When a government agency adopts a new user interface or help mechanism it must not only scale to very large websites, but to the full range of user populations. Interfaces must serve the broadest range of users and alternatives must be provided for those with special needs. Additionally, agencies do not work in a vacuum but must coordinate with parent organizations. All the statistical agencies in the US government are parts of larger units of government service that may have distinct requirements for incorporating new technologies into practice (e.g., the Bureau of Labor Statistics sits within the Department of Labor). Thus, implementing new technologies that have been proven effective in laboratory studies (let alone those that have not) may be impossible or take years rather than months because both technical and policy problems must be addressed.

Nonetheless, we do believe that theory and practice interaction is mutually beneficial and our experience working with statistical agencies illustrates some of these mutual benefits (e.g., Marchionini & Levi, 2003). Certainly, academic researchers benefit from working with agencies that provide large-scale data sets with real-world exceptions and special cases—beyond simply having 'lots of bits.' Likewise, working with applications that span enormous audiences of users and potential users and cover a wide range in interest, needs and abilities gives researchers much more leverage to examine design parameters. For example, webpages may contain many instances of a single kind of information (e.g., a table containing data cells), or combine text, forms, statistics, graphics, and maps into a complex structure. The scale puts real limits on "reasonable" solutions (e.g., limiting human effort), but challenges the imagination to make the information available to all. Although this is sometimes frustrating, it allows researchers to see the exceptions that may lead to new kinds of breakthroughs or alternative designs.

Our experience demonstrates that academic-government partnerships take time to develop and maintain. Trust must be developed on both sides so that different motivations and missions can come together to serve a larger need. Showing inspirational prototypes is helpful for considering the possibilities, but the more we can streamline, simplify, and explain what will give the best bang for the buck, the more likely our ideas will be adopted or adapted. Over the years, some of the user evaluation results and interface ideas we have developed for statistical agencies have been incorporated into agency websites and we hope that some of the ideas presented here will also be woven into efforts of government agencies to serve the larger needs of helping citizens find and understand statistical information.

Acknowledgements

This work has been supported by NSF Grant EIA 0131824 and by contracts from BLS, Census, EIA, SSA. and NASS. We would also like to acknowledge the contributions of the Govstat teams at UNC-CH and the University of Maryland, and the many statistical agency personnel who contribute their time and expertise to this project.

References

Blank, G., Roy, S., Sahasrabudhe, S., Pttenger, W. & Kessler, G. (2003). Adapting multimedia for diverse student learning styles. *Journal of Computing Sciences in Colleges*, 18(3).

Carroll, J. M. & Carrithers, C. (1984). Training wheels in a user interface. *Communications of the ACM*, 27(8), 800-806.

Denn, S. Haas, S. W. & Hert, C. A. (2003). Statistical metadata needs during integration tasks. In *CD-2003: Proceedings of the International DCMI Metadata Conference and Workshop*, 81090. htttp://www.siderean.com/dc2003/301_Paper50.pdf

Efron, M., Marchionini, G. and Zhang, J. (2003). Implications of the Recursive Representation problem for Automatic Concept Identification in On-line Governmental Information. ASIST SIG-CR Workshop (Long Beach, CA, October 18, 2003).

Kang, H., Plaisant, C., and Shneiderman, B. (2003). New Approaches to Help Users Get Started with Visual Interfaces: Multi-Layered Interfaces and Integrated Initial Guidance. *Proc. 2003 National Conference on Digital Government Research,* 141-146. http://www.dgrc.org/dgo2003/.

Marchionini, G. & Brunk, B. (2003). Toward a General Relation Browser: A GUI for Information Architects. *Journal of Digital Information*, 4(1), http://jodi.ecs.soton.ac.uk/Articles/v04/i01/Marchionini/

Marchionini, G. & Levi, M. (2003). Digital government services: The Bureau of Labor Statistics Case. *Interactions*: *New visions of human-computer interaction*, 10(4), July-August p. 18-27.

Morrison, J. & Tversky, B. (2001). The (in)effectiveness of animation in instruction. *CHI '01 Extended Abstracts on Human Factors in Computing Systems*, 377-378.

Paivio, A. (1986). *Mental Representations: A Dual Coding Approach*. Oxford University Press: New York, NY.
Shneiderman, B. (1994). Dynamic queries for visual information seeking, IEEE Software 11(6), 70-77.

Vygotsky, L. S. (1978). *Mind in Society: The Development of Higher Psychological Processes*. Cambridge: Harvard University Press.

Wilbur, J. & Haas, S.W. (2004). The GovStat Content Delivery Template (CDT) for Creating Explanatory Animations. Proceedings of the dg.o2004, The National Conference on Digital Government Research, Seattle. Available at: http://dgrc.org/dgo2004/disc/demos/tuesdemos/wilbur.pdf (November 15, 2004).

Zhang, J. & Marchionini, G. (2005). Evaluation and evolution of a browse and search interface: Relation Browser++ *Proceedings of dg.o2005; The 6th National Conference on Digital Government Research.* Atlanta, GA: May15-18, 2005. NY: ACM Press. pp. 179-188.

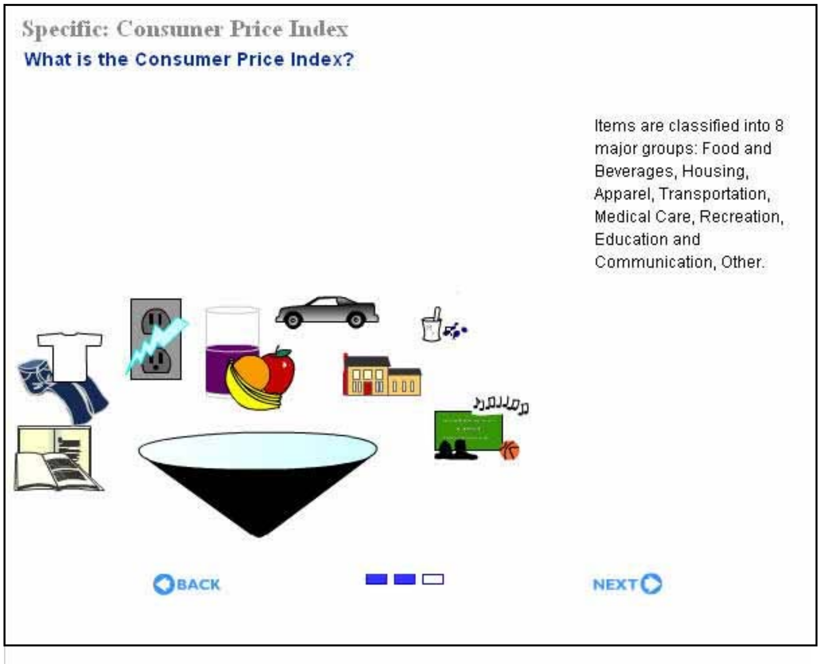

Figure 1: Screenshot from SIG animated presentation of the Consumer Price Index.

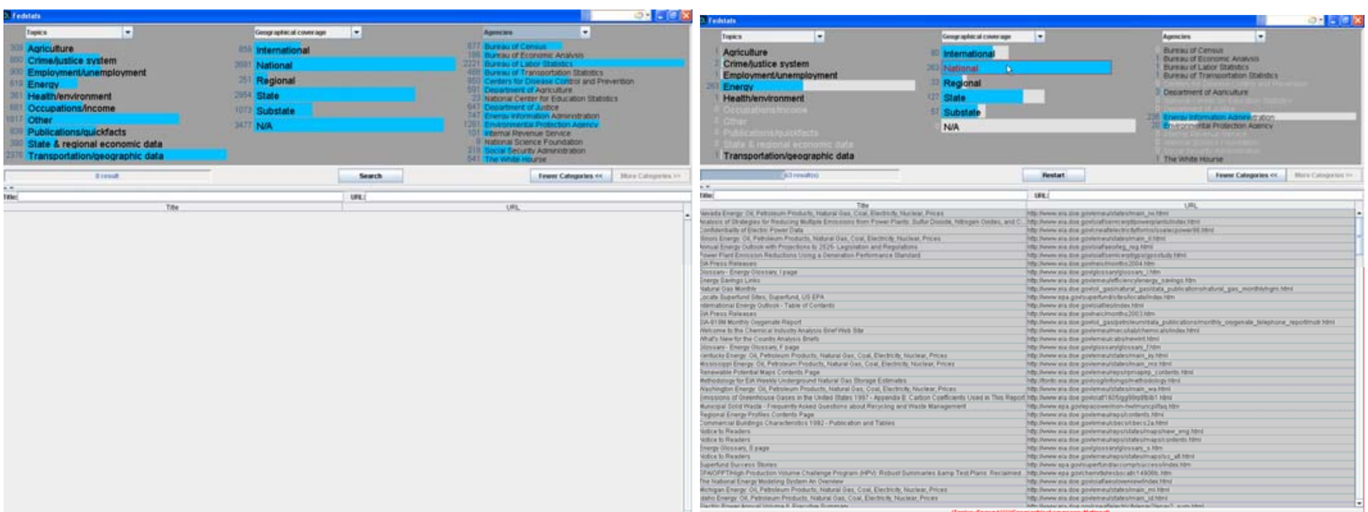

Figure 2. Relation Browser Screen Displays. Left panel shows initial screen for FedStats data. Right panel shows screen after energy selected, results returned, and user brushing over national facet.

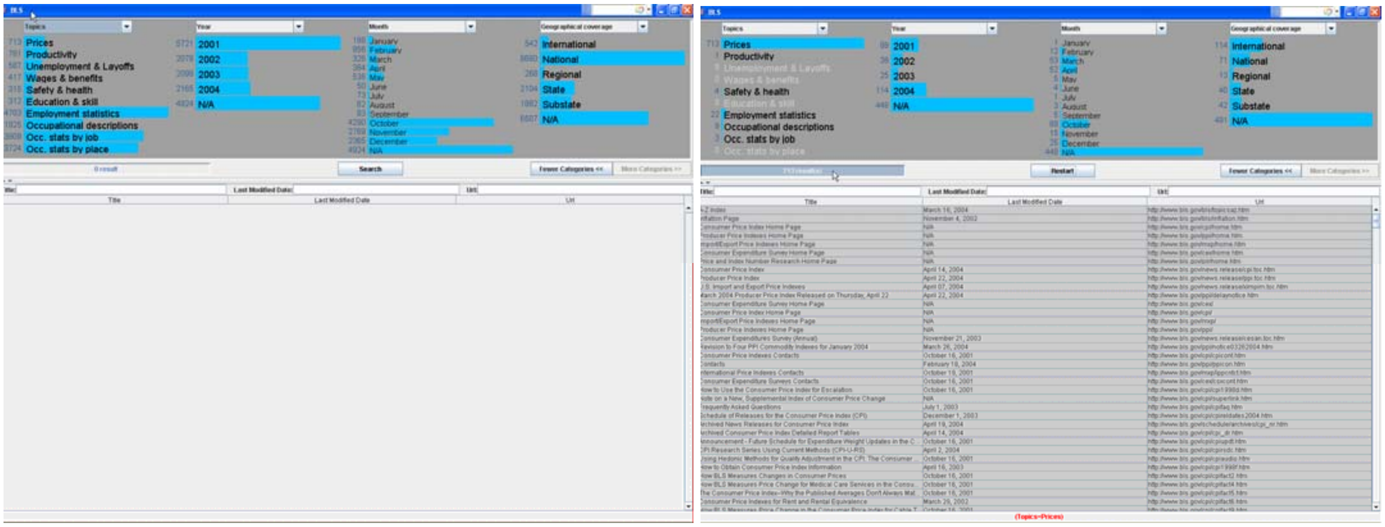

Figure 3. Relation Browser Screen Displays for BLS with Generated Topical Categories. Left panel shows initial screen for BLS data. Right panel shows screen after prices selected.