Agileviews: A Human-Centered Framework for Interfaces to Information Spaces

Gary Marchionini, Gary Geisler & Ben Brunk
University of North Carolina at Chapel Hill
[march, geisg, brunkb]@ils.unc.edu

Abstract

A framework for interface design that provides people with flexible control over different views for an information space is presented. The agileviews framework defines overviews, previews, reviews, peripheral views, and shared views that help people make decisions about where they should focus attention during information seeking. In addition to the views themselves, control mechanisms that facilitate low-effort actions and strategies for coordinating the views are discussed. Agileviews are particularly useful when specific partitions of large information spaces such as the WWW have been identified. Examples of these views are provided from several different projects and suggestions for additional research and development are made.

Introduction

As information becomes increasingly pervasive in wider aspects of human endeavor, the search for information becomes correspondingly more important and complex. Information search is a type of problem solving that requires representation, inference, and action mechanisms that change the state of representation. Defining representations (e.g., queries, indexes), making inferences (e.g., judging relevance), and choosing actions (e.g., select, terminate) all require the searcher to make decisions at strategic (e.g., How to invest my time? Which approaches to use?) and tactical levels (e.g., Which choice to make next? Did I find the solution?). This problem solving interaction with an information space contained in an office, library, or the web is mediated by physical and conceptual interfaces that determine how representations are manifested, what action mechanisms are available, and the rules of engagement that support interaction and ultimately the inferences people make to solve their problem. Interaction designs that start with needs and processes to meet these needs, and then apply technology to serving those needs illustrate human-centered design.

Considering search from a human-centered perspective, we recognize that search is embedded in real life tasks. People are interested in accomplishing those tasks rather than executing searches. This is an instance of what Carroll and Rosson (1987) call the production paradox—people want to get their work done rather than learn or give attention to the interface that they must learn and use to do their work. This implies that interfaces for search should be easy to learn and use and as far as possible be transparent to the person working through them. In many respects, we have traditionally focused on technical models of information retrieval rather than the human-centered model of search. The technical model necessarily focuses on mechanisms that are abstracted from the needs, capabilities, and experiences of people. Information seeking research has demonstrated, however, that search is an interplay of analytical and interactive problem solving strategies. This is especially so in large information spaces where strategies to identify promising partitions of the information space (e.g., database/collection selections in DL, queries) are distinct from strategies to search and examine those partitions (e.g., result sets). Rather than nimble tools to leverage these different strategies easily, our systems apply clumsy mappings of human decisions and subtle inferences into a small set of mechanical mechanisms that compare, retrieve, and display information objects. The norm is a series of conscious attentions to these possible mappings rather than to the information need, thus hampering cognitive flow and interrupting our understanding of the problem and information we seek. In short, the interfaces get in the way.

As more of the world's population engages in electronic search, this perspective is changing to accommodate diverse human-centered approaches to search. First, our conceptions of design are maturing to ground systems in human capabilities and culture—preferences and beliefs. This is so because we recognize that not only does technological change alter what is possible, but also people and populations learn and change and global access increases the variances of skills, experience, and resources within these populations. Second, we have begun to develop interface environments that are more robust. These environments leverage multiple I/O devices, a variety of user control mechanisms beyond type and click, and rich sensory representations for information surrogates and objects. This paper presents a framework for a human-centered, multifaceted search environment that is based on flexible views. We term these views, *agile* to suggest the ease with which people can shift them to change focus and advance the search process. We illustrate the agileviews framework with examples of interfaces from different projects; and suggest directions for research and development to extend the framework.

Agileviews

About 70% of all the receptors in the human body are in the eye (Campbell, 1993). Vision thus provides the most significant source of information to sighted humans and plays major roles in information seeking tasks. Regardless of which psychological theory of image understanding one subscribes to, the role of vision to understand what is where in the space around us (Barlow, 1990 alluding to Marr) is pervasive and powerful. This marvelous information processing capability that nature provides us is highly flexible to enable us to perceive a multiplicity of information forms and functions and control our perceptual intent with little conscious effort. The products of vision—views of the space around us--have multiple facets to accommodate the rich variations in the natural world, yet are easily controlled by even young children. Views thus serve as a universally applicable metaphor for interfaces to abstract information spaces. Keeping in mind the many limitations of metaphors, we propose guiding interface design with the metaphor of *agileviews*—multiple views of information objects that are easily controllable by information seekers to inform decision making and understanding.

We define a view as a visual partitioning of the information space. The partitioning can be based on various metrics such as linear direction (objects are positioned in the information space), magnitude (where objects are ordered and thus may be near or far), spectrum (objects may be visible in selected radiations), and time (objects may have been active in the past, be active in the present, or be active in the future). Any view may be given human attention, and the view of interest at any given instant is called the *focus*. Visually, we may focus on objects in the background, thus making the foreground objects context (e.g., focus on a word on this page and the page view becomes context). In the online environment, the active window is the current focus. The mechanisms that change focus are crucial to how views are leveraged to make decisions and create understanding. Views must first be useful and appropriate, but this is not sufficient, they must also be easily controllable.

The physical book interface affords views of pages containing typographic units as well as the entire book as an object. These views have different possible expressions and functionalities that drive meaning (e.g., characters have font variability, pages may be narratives, indexes, and tables of contents). Additionally, the conceptual interface of the book relates the reader's experience, mood, and attention to the information organization and content the author designed for their use. When reading or searching in a text, we can shift focus easily among these views through fine-grained eye movements (e.g., saccades) and large muscle movements (e.g., head, hands). In the online physical interface, views are afforded by windows containing various bit patterns (including dynamic bits, and multiple layers of information representations), and the entire screen itself. Like books, each electronic view has many possible expressions and meanings. The conceptual interface in the online case likewise relates the user's experience, mood, and attention to the information organization and content that the provider designed for their use. Unlike the view of the whole book, which changes across books, the electronic interface has a single screen for all documents. Additionally, the screen has a fixed size to represent all possible views—the well-known limited screen real estate problem. When reading or searching online we can shift focus using fine-grained eye movements within the currently active window and whatever other views fit on the screen. At present, paper provides higher density resolutions to leverage fine-grain eye movements than screens. To change focus to another view requires a large-muscle movement to activate another window. The usual movement is a mouse click or key press. We suggest that a click is a radical act. First, it does change the window focus and unless the views are carefully coordinated and laid out, contextual information may be lost. Second, especially in a network environment, a click implies a wait for a new window to load. This interrupts thought, annoys people, and is one of the foremost complaints about the WWW. Fortunately, other control mechanisms are available when systems are able to distinguish and react to mouseover, mousedown, and mouseup events.

We believe that good design can provide simple mechanisms for using carefully coordinated views so that decision making and understanding are advanced and enhanced. In an ideal interface, these views are related in ways that are natural (easily recognizable) and easily controlled (agile) so that the interface fades into the context—disappears. The agileviews framework provides a step in this direction.

One way to consider how different views fit into the agileviews framework is to consider the popular metaphor of the information superhighway. Driving in an automobile, you are able to focus on the road or car immediately in front of you or to effortlessly shift your view to look at a distance to what lies ahead. These are previews and overviews of sorts, and if you stop to examine a map, you are using an overview at a greater scale (it is certainly feasible that heads up displays will eventually enable you to view the map overview without stopping). In addition, you may look in the rearview mirror and consult where you have just been, or at a higher level of abstraction, recall past driving experience that relates to the immediate drive at hand. These are examples of reviews. As you drive, you may observe the landscape rushing by or the general contours of the land around you. These are made possible by your peripheral view working in coordination with your macular view. Finally, you may be in a traffic situation where another motorist blows her horn at you to warn you of some impending danger that is outside of any of your personal views. This is an example of a shared view. Like all metaphors, this one has many limitations and flaws but it serves to suggest a physical analog to multiple, flexible, and effortlessly controlled views of information spaces.

Information seekers aim to focus their attention on an information object that serves to answer their need or advance their understanding (e.g., a document, image, map, etc.). These information objects may be considered to be the primary view of interest and bringing this (or these) view(s) into focus is a primary subgoal of search and the object of a good deal of decision making. Bringing these primary views into focus is expensive in several ways. First, and most importantly, processing these views (e.g., full text, high-resolution images, full video clips) demands considerable human attention time. Abstracts, summaries, and other surrogates are valuable because they lower attention time costs by informing decisions about whether to invest time in processing the full object. Additionally, regardless of increases in networking speed, in the WWW the primary views take time to load. The agileviews framework includes six genres of view that may be brought into focus during information seeking. Figure 1 depicts these different types of views.

Figure 1. Agileviews Model

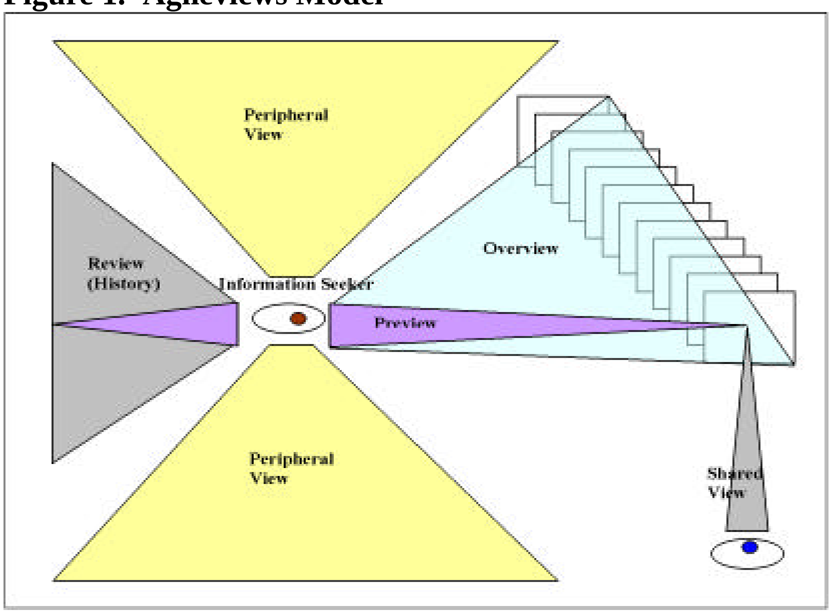

The **primary view** is represented by the stacked rectangles at the right. Overviews and previews provide information seekers with look ahead functionality to make judgments before primary views are displayed. Reviews provide historical information about the current or past searches. Shared views leverage the skill and knowledge of other people. Peripheral views provide the context for the current focus and includes any views in windows that are on the screen but not currently active.

Overviews aim to help people to understand the structure of an information space—what is and is not available, how the information objects are related, and what levels of granularity exist in these relationships. Overviews help people frame problems and determine points of attack (e.g., entry points, search strategies). There are many examples of overviews in common use in online environments today ranging from data dictionaries for databases to site maps on WWW sites. Many visualization techniques have been applied to create overviews, most notably the information visualizer suite of perspective walls and cone/cam trees from the Xerox PARC group (e.g., Card et al., 1991; Lamping & Rao, 1996), Lin's semantic maps (Lin, 1997), and the Human-Computer Interaction Laboratory's (HCIL) starfield displays (e.g., Shneiderman, 1998). Greene et al. (in press), characterize overviews as giving perspectives on collections of information objects, recognizing that objects may be considered from many levels of granularity. They provide examples that include the Visible Human Browser, NASA's Global Change Directory project, and the Library of Congress' American Memory digital library. New examples of overviews are discussed below.

Previews aim to give users quick glimpses of specific information objects, again recognizing the granularity issues related to objects (e.g., a web page is an object at one level but also a collection of associated objects such as texts, images, and link anchors). Previews focus on problem solutions—answers. They help us decide whether the information object has the answer or has a high probability of getting one closer to the answer. Previews facilitate a probability decision about getting closer to the problem solution. One type of preview that has had long-standing attention is a hypertext link anchor or tool tip that denotes link types (e.g., Nielsen 1998). There are many examples of previews in common use ranging from movie trailers to entice audiences to thumbnail images on web pages. Greene et al., (in press) provide examples of previews as well as overviews and new examples are discussed below.

Reviews can help searchers initiate and shape a search session or monitor progress as a search progresses. In the first case, reviews facilitate strategy selection by providing models from experience—we apply templates that have served us well in the past. In the second case, reviews assist the searcher to evaluate the overall search process and its relationship to current context and the larger global context. In this case, reviews facilitate a probability decision about the strategy in use or simply serve to recall recent intermediate results. Search histories in online systems and bookmarks and history lists in WWW browsers are prominent examples of review. There is considerable research on history mechanisms. For example, Abrams et al, (1998) provide a taxonomy of uses for bookmarks; Robertson et al., applied page thumbnails for personal web histories; and the HCIL has applied a variety of visual techniques such as lifelines and zoomable maps to histories in several contexts such as personal histories (Plaisant et al., 1997), legal databases (Harris et al., 1999), web traversal (Hightower et al., 1998), and education (Plaisant et al., in press).

Peripheral views bring context into the space—they contextualize overviews and previews by verifying the neighborhood of the current search. They enable a decision about whether the probabilities of solution in the overviews and previews are compatible/consistent with the overall need. When other views go out of focus but remain on the screen, they also serve as peripheral views. Query services that display the query along with search results maintain such context for searchers. Pop-up windows help maintain context by allowing the viewer to change focus while keeping the screen context below/around the pop-up. Within site consistencies of look and feel also support peripheral views. Iconized applications on status bar and the screen background also serve as peripheral views.

Shared views tap into current or historical decisions made by others. Considering what others are doing or seeking specific inputs from others engaged in the same or similar problem solving can aid decision making during search. Collaborative filtering as applied in recommender services for entertainment or consumer services leverage the products of shared views. New research on collaborative information retrieval also illustrate the recognition that searchers can benefit from what others see and do (e.g., http://www.ischool.washington.edu/people/facstaffdirectory.htm .

In all cases, people must understand the control mechanisms (e.g., mouse actions) for exploring, extracting, and using information. In the human vision sense, we simply refocus our eyes to change distance or move our head to change direction. We refer to artifacts (e.g., notes) to peer back and consult our memories to make inferences about the future. Human-computer interfaces are not so natural and easily controllable. Linking pop-up windows and other focus-changing actions to mouseover, mousedown, and mouseup actions is the current best control mechanism for managing views. Although eye gaze, gesture, voice commands, and other mechanisms have been proposed, mouse events are the most pervasive solution today.

It is important to consider the relationships and possible interactions among the different views. It is important that shifts between views be easy and seamless. This requires careful attention to the underlying system architecture such as metadata, control mechanisms, and consistent displays. In all the examples that follow, the underlying metadata and system features constrain what is displayable and displayed. If a view is to be useful, it must provide the most salient cues from the primary object and represent these cues effectively. It must also be linked to other views for that object and the object itself. For example, if previews are embedded within overviews, then window positioning must be carefully considered so that critical information in the overview is not obscured. This implies a robust data management structure behind the scenes and a compact way of transferring this data in the WWW environment—if there is too little metadata associated with the view it will be impoverished and not useful for decision making, if there is too much it will take too long to transfe/process (both electronically and mentally) and mitigate its usefulness. Additionally, techniques to automate the accumulation of metadata are needed if a rich variety of views is to be provided. For example, if number of webpage accesses is a part of shared view, ways to continually update these values without burdening the server with robot requests must be found (see Geisler, [2000] and Weinreich http://vsys-www.informatik.uni-hamburg.de/projects/hyperscout for approaches to this challenge.)

Examples

BLC Resource Explorer

Overviews of a multimedia database for middle school teachers is provided through a dynamic query interface called the Explorer that uses a barfield display to array various multimedia objects (Marchionini et al., 1998). Figure 2 shows the barfield display and the partition consisting of objects indexed as pertaining to ecology topics and earth science standards. Attributes such as media type (video clip, text, image, audio clip, website, instructional module), source (e.g., Discovery Channel, US Archives, etc.), instructional standard (various standards for science and social studies), and topic (science and social studies), can be mapped to the two dimensions of the barfield display, bar colors, and various sliders according to the desires of the teacher using the explorer. As users select an attribute set for a desired partition of the database, the results list (in the rightmost part of the figure) provides various types of previews that include bibliographic information for the object as well as media specific previews such as thumbnails for images and two alternative video previews—a slide show or storyboard of keyframes. In the figure, a slide show preview is playing and was captured midway between two keyframes. The user gets a simple peripheral view through the placement of the previews over the result list and barfield display. Most importantly, the overviews and previews are closely coupled and work together to give the user flexible control over the database and objects in it at different granularities before they commit to loading the full object. It should be noted that the Explorer itself is embedded in the larger instructional module (lesson plan) construction component so that users move across different granularity levels in seamless ways. It must be noted that although this system is fully functional in a high-speed network environment with the most recent WWW browsers, and teachers like the features it provides, load times and inconsistent client platforms in the schools have led teachers to opt to use simple database lookups in the underlying database rather than the explorer. They then paste those resources into the lesson construction component.   This illustrates the production paradox—even though these parts of the system are attractive and give more flexible options (e.g., provide the video previews), fast and reliable actions are preferred.

Figure 2. Baltimore Learning Community Overview and Previews

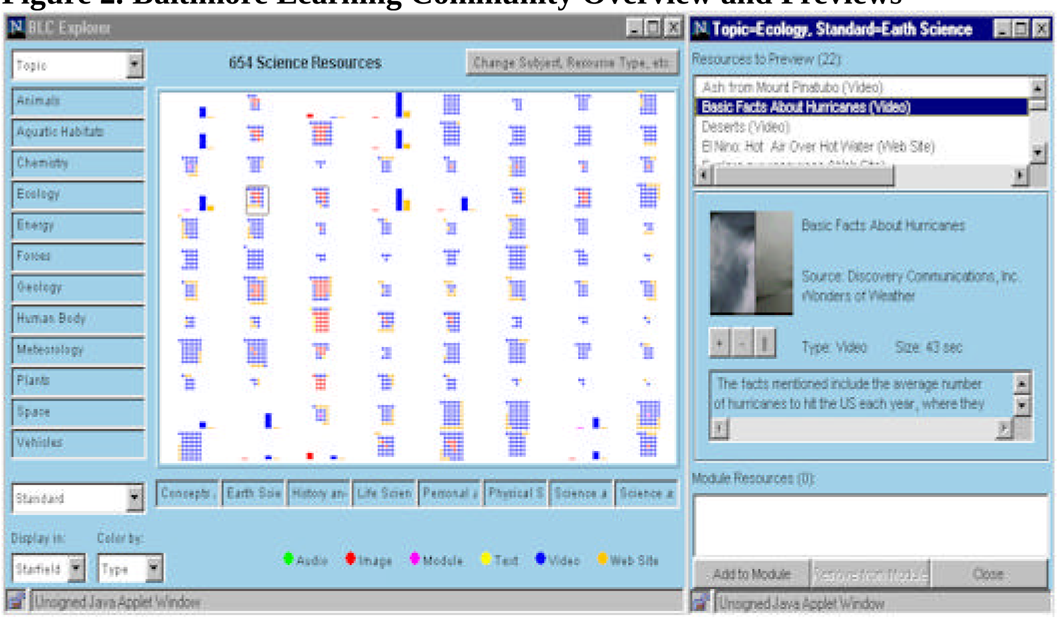

Federal Statistics Relation Browser

The Fedstats website provides access to the 200 websites and thousands of webpages containing statistical reports, datasets, tables, and other statistical elements produced by 70 federal government agencies. This site (www.fedstats.gov) acts as a portal to these websites and provides a sitemap as well as other entry points. We have been working with the Fedstats team to provide alternative access mechanisms and have developed an exploration tool called the Relation Browser (RB). The main idea is to provide the user with a mechanism to examine pair-wise relationships among attributes such as topic (economics, crime, education, etc.), data type (report, table, graph, downloadable dataset, multimedia, etc.), region (international, national, region, state, substate), and time coverage period. The RB aims to help people understand structure—the relationships among various attributes of the space. It facilitates a probability decision at the macro level about what strategies and tactics might be useful. In Figure 3, the RB tool shows a view of the Fedstats collection when the user has rolled the mouse over the environment topic. Note that the number of websites that have tables, reports, downloadable datasets and other information types are shown as blue bars and numerically so that user gains a sense of volume and type of data for this topic. The website titles and URLs appear below to provide rapid access to those websites. If the user changes the mouse position to another topic, the display instantly updates to show the corresponding data for that topic. By exploring the topics without committing to a mouse click, the user is able to gain a quick overview of the range of data in the 200 websites provided by these government agencies. Users may also choose to explore topics by geographic region or time period. We are currently working to give additional linkages to the data types in the result list to give users even more detailed looks ahead before they commit to a click. We have conducted several iterations of usability tests with the RB and are currently engaged in a field test of the tool which is accessible at the Fedstats website. See Marchionini et al. (in review) and see http://ils.unc.edu/~march/bls_final_report99.pdf for details of the system and testing.

Figure 3. Fedstats Relation Browser

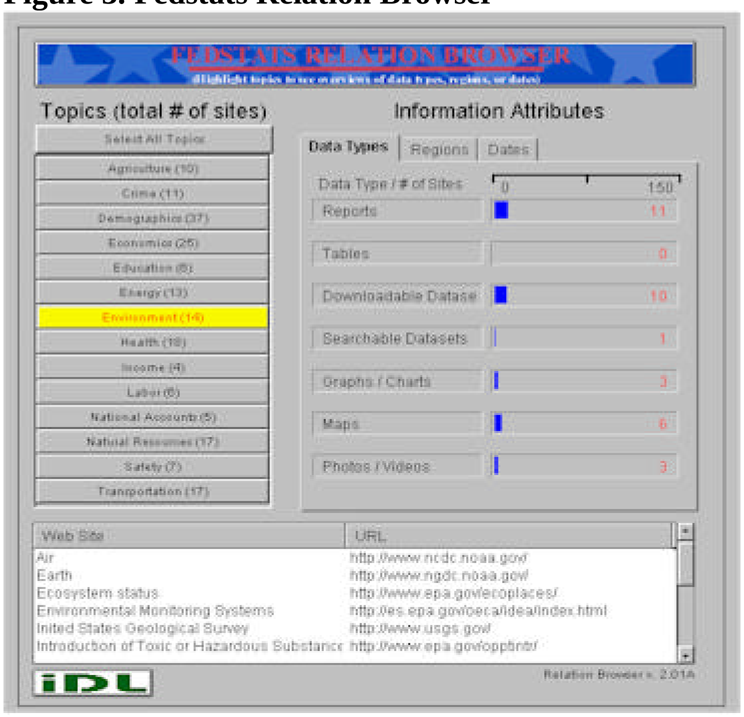

Enriched Links

An example of how several views can help a user navigate the Web is provided by the enriched links framework (Geisler et al., in review). The enriched links framework uses small pop-up visualizations—previews, overviews, and shared views—to provide the user with information about linked pages that can be used to evaluate the appropriateness of the pages before making a commitment to select the link and wait for the page to load. The preview information provides the most immediate look at the linked resource; it consists simply of a thumbnail image of the page and its file size, as shown in Figure 4. We believe that the visual cues in a page layout provide important clues that inform intelligent link following.

Figure 4. Enriched Link Preview

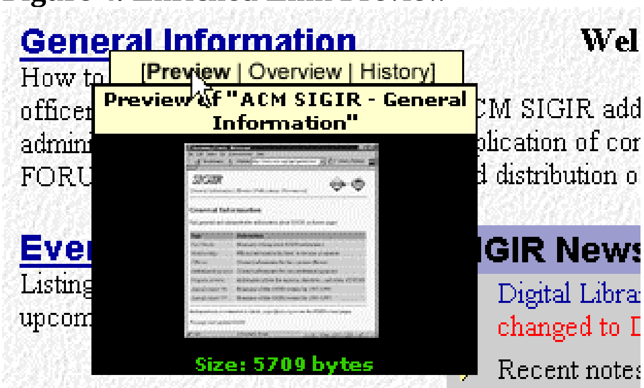

The Overview pop-up, shown in Figure 5, provides a more detailed summary of the objects that make up a linked page. A graphical representation shows the number of links, images, and other media contained by the linked page, each categorized by internal or external links, with internal links being resources that reside on the same domain as the parent page. In addition, a representation of dead links is provided to help identify pages that are out-dated or poorly maintained.

Figure 5. Enriched Link Overview

A variation of a shared view is provided by the History pop-up, which uses Web access log data to give the user an indication of how the page has been accessed by other users. As shown in Figure 6, this pop-up shows how recently the file was last updated, and displays a graphical representation of the hits to the page, both recently (the previous day, two days previous) and over time (the past month and year).

Figure 6. Enriched Link Shared View

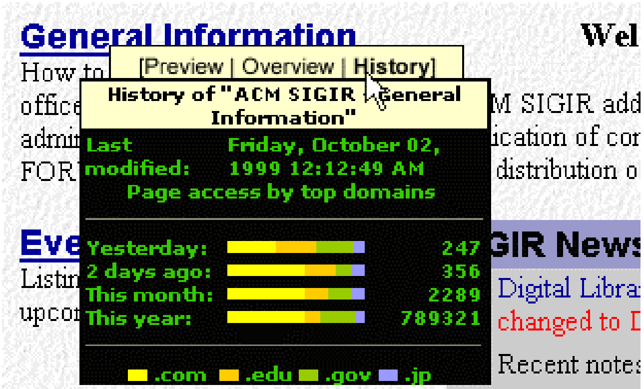

Other types of agileviews—reviews, peripheral views—could be integrated into the enriched links framework and provide the user with even more contextual information to improve the Web browsing task.

Animated Site Maps

A distinctly different approach to overviews of sites and previews of pages in a site is provided animated site maps. The BLC project used slide shows and story boards as visual previews to video clips and a number of associated empirical studies were conducted to evaluate how accurately and rapidly people can extract meaning from thumbnail views (e.g., Tse et al., 1998; Ding, 1999). Wittenburg et al., (1998; 1999) have also used these techniques in their information navigation systems. Figure 7 shows two frames of an animated site map for the School of Information and Library Science web site at North Carolina. The applet "plays" a slide show of salient (manually selected in this example) webpages for the site. The user can control the speed of the slide show with the slider and jump to pages by clicking on them. Mousing over the thumbnail also displays the page title. The idea is to provide a visual overview of the site that trades the time to view the slide show for the screen real estate taken by the thumbnail. We are experimenting with how small the animated view can be (ideally it could run in a tiny corner of the screen) and still be useful. It is easy to identify major visual distinctions with such a tool, for example the data entry form in the left part of the figure. Preliminary user reactions have been positive (Brunk, 1999).

Research Directions

These examples instantiate several aspects of the agileviews framework. Our goal in the coming year is to apply a fuller range of views to a table browser for statistics (NSF Grant: Citizen Access to Federal Statistics) and to a video repository for the video retrieval community (http://openvideo.dsi.internet2.edu/). The table browser will use peripheral views of the large realm of federal statistics to contextualize specific tables. One challenge in this regard is how to show context without complicating the screen display. Previews of the underlying metadata will also be provided for key elements of the tables. We have begun to incorporate enriched links into the interfaces to the Open Video project and will aim to develop shared views that are based on aggregated community access as well as concurrent access.

Figure 7. Animated Site Map Snapshots

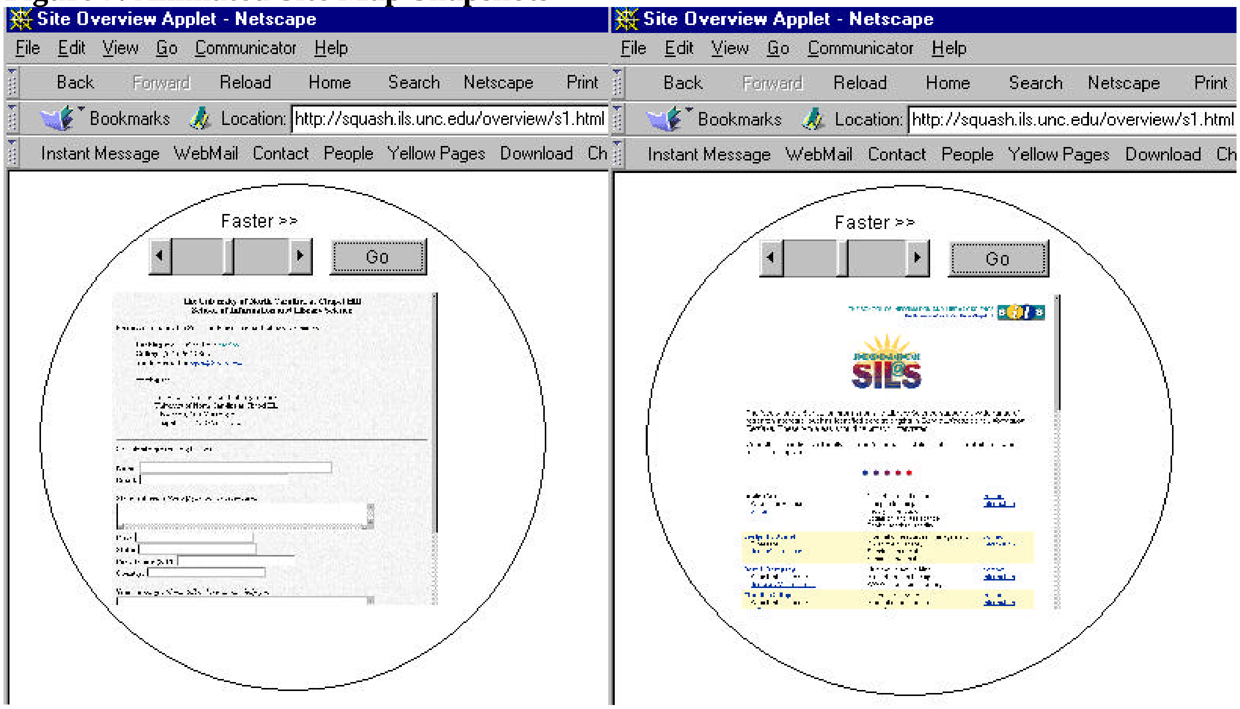

Metaphors are double-edged swords. On one hand, the view offers a universally recognizable anchor for working in information spaces. On the other hand, information space is not physical space and managing changes among one's mental focus, visual focus, and active window on a screen may be confusing rather than helpful. We will continue to conduct usability tests to understand the costs and benefits of the metaphor and the specific views based upon it. The WWW has demonstrated that people are impatient waiting for information to load and are easily disoriented by page after page of minimally connected information. We believe that alternative views that are carefully coordinated and activated through simple mouse events such as mouseover can go a long way toward minimizing these adverse affects. A click is a terrible thing to waste and agileviews may help get people more bang for their click.

References

Abrams, D., Baecker, R., & Chignell. M. Information archiving with bookmarks: Personal; web construction and organization. *Proceedings of ACM CHI 98 Conference on Human Factors in Computing Systems*. p. 41-48.

Barlow, H. (1990). What does the brain see? How does it understand? In H. Barlow, C. Blakemore, & M. Weston-Smith (Eds.). *Images and understanding*. NY: Cambridge U. Press.

Brunk, B. (1999). *Overview and Preview Tools for Navigating the World-Wide Web,* SILS Technical Report TR-1999-03. July 31, 1999

Campbell, N. (1993). *Biology* 3rd Ed. Redwood City, CA: Benjamin/Cummings Publishing. p. 1022.

Card, S., Robertson, G., & Mackinlay, J. (1991). The information visualizer: An information workspace. *Proceedings of ACM CHI '91 Conference on Human Factors in Computing Systems*. P. 181-188.

Ding, W. (1999). *Designing multimodal surrogates for video browsing and retrieval*. Unpublished doctoral dissertation (University of Maryland at College Park).

Geisler, G. Rekimoto, J. & Marchionini, G. (in review). Enriched Links. Submitted as short talk to ACM CHI 2000.

Greene, S., Marchionini, G., Plaisant, C., & Shneiderman, B. (in press). Previews and overviews in digital libraries: Designing surrogates to support visual information seeking. *Journal of the American Society for Information Science*.

Harris, C., Allen, R., Plaisant, C., Shneiderman, B. (June 1999). Temporal Visualization for Legal Case Histories. *ASIS'99 Proceedings of the 62nd Annual Meeting of the American Society for Information Science*, Conference October 31-November 4, 1999, Vol. 36, 271-279.

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., Brunk, B., Komlodi, A., Conrad, F., & Bosley, J. (in review). Look before you click: A relation browser for Federal statistics websites. Submitted to ASIS 2000 Annual meeting.

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.

Nielsen, J. (1998). Using Link Titles to Help Users Predict Where They Are Going. Alertbox #11. January 1998. http://www.useit.com/alertbox/980111.html

Plaisant, C., Rose, A., Rubloff, G., Salter, R., Shneiderman, B. (in press). The design of history mechanisms and their use in collaborative educational simulations. *ACM Transactions on Computer Human-Interaction*.

Plaisant, C., Shneiderman, B., Muhslin, R. (1997). An Information Architecture to Support the Visualization of Personal Histories. *Information Processing & Management*, 34, 5, pp. 581-597,

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.

Tse, T., Marchionini, G., Ding, W., Slaughter, L. & Komlodi, A. (1998). Dynamic key frame presentation techniques for augmenting video browsing. *Proceedings of the Working Conference on Advanced Visual Interfaces AVI '98* (L'Aquila, Italy May 24-27, 1998). NY: ACM Press. 185-194.

Wittenburg, K. Ali-Ahmad, W., LaLiberte, D., & Lanning, T. (1998). Rapid-fire image previews for information navigation. *Proceedings of the Working Conference on Advanced Visual Interfaces AVI '98* (L'Aquila, Italy May 24-27, 1998). NY: ACM Press. 76-82.

Wittenburg, K., Nicol, J., Paschetto, J., & Martin, C (1999) Browsing with Dynamic Key Frame Collages in Web-Based Entertainment Video Services. In *Proceedings of IEEE International Conference on Multimedia Computing and Systems*, (June 7-11, 1999, Florence, Italy).