

Visual Information Seeking in Digital Image Libraries: The Visible Human Explorer

Chris North[†], Ben Shneiderman^{†*}, and Catherine Plaisant

Human-Computer Interaction Laboratory

[†]Department of Computer Science

^{*}Institute for Systems Research

University of Maryland, College Park, MD 20742 USA

{north, ben, plaisant}@cs.umd.edu

<http://www.cs.umd.edu/projects/hcil/>

Footnote: This chapter uses extracts from the following paper, and is expanded to report on recent work:

C. North, B. Shneiderman, and C. Plaisant. User-Controlled Overviews of an Image Library: A Case Study of the Visible Human. *Proc of ACM DL96 - Conference on Digital Libraries*, pg 74-82, March 1996.

ABSTRACT

This chapter presents the Visible Human Explorer user interface, developed at the Human-Computer Interaction Lab of the University of Maryland at College Park, for remotely accessing the National Library of Medicine's Visible Human digital image library. With the interface, users can visualize the library, browse contents, locate data of interest, and retrieve and zoom on desired image details. The interface presents a pair of tightly coupled views of library data: an overview of the overall search space, and a preview of high-resolution images available for retrieval. To explore, the user sweeps the views through the search space and receives smooth, rapid, visual feedback of contents. Desired details are automatically downloaded over the internet from the library. The interface software is completely functional (runs on Sun Workstations) and freely available for public use, at: <http://www.nlm.nih.gov/>.

We also present several human-computer interaction design principles used to create the Visible Human Explorer interface, describe how these were applied to the design, and discuss issues in employing these techniques in user interfaces for other information collections. These principles are direct manipulation, visual information seeking, query previews, and multiple tightly coupled views. We illustrate these concepts with a plethora of pictures of user interface screens. Please also check the included CD-ROM for additional illustration media.

INTRODUCTION

As digital libraries quickly increase in number and size, researchers are continually developing improved public access methods for the many different types of library information. The primary user interface approach, especially for libraries containing unstructured collections of heterogeneous documents, is the keyword query method. In this case, Information Retrieval systems preprocess documents in the library, usually using automated algorithms, to generate

comprehensive indices (meta-data). Then, when a user enters a query, the system can quickly generate a list of matching documents. Since many of these libraries only contain textual documents and use strictly textual queries, interfaces are often text-based for simple WWW or teletype access. Image libraries can also employ query-by-image-content techniques. In this case, interfaces allow users to specify example images, interactively sketch objects, or use an image feature description language to create a query. Digital video libraries can expand the query method one step further, using speech recognition on the audio tracks of videos for additional textual indexing. While the query method may be an effective means for retrieving documents based on keywords or image features, it has several shortcomings:

- Requires recall. The user must supply a query, which may be difficult if the user is not sure what to ask for.
- Trial and error process. When a query returns a set of undesired hits, the user must blindly reformulate the query and try again. Users are frustrated by zero-hit or mega-hit query results.
- No model of library contents. What information is available in the library? What is NOT available? No context is provided for query hits, leaving the user unsure of how this information fits in with the rest of the library contents.

Hence, there is no capability to “explore” the library.

However, if the documents within a library are organized using some consistent, natural, comprehensible, orderly structuring method, then interfaces employing a browsing approach are also applicable. Advanced user interface techniques for visual information seeking can be used to explore the library structure, usually specified by meta-data, and its contents. This gives users a framework within which to learn about the library and locate interesting documents without relying on keyword search. Several common structuring methods exist. For example, hypertext systems such as the WWW organize documents in a hierarchy or other graph representation. With an interface for browsing graph-structured information, users can utilize the links, perhaps representing relationships between documents, to discover related documents that might not be found using a keyword query. Another structuring method organizes information by defined textual and numerical attributes, such as author name, publication date, etc. Then, visual dynamic-query interface techniques [DPS95] can help users to quickly learn overall contents, find trends, and locate documents of interest. Similarly, geographical systems use map interfaces [Pla93] to browse spatially organized information. All of these systems are examples of how advanced information-exploration user interface techniques can capitalize on the structure of an organized library to add a new dimension, namely the capability to visually browse, to the user’s tool set for the task of information retrieval.

The Visible Human Digital Library

The Visible Human digital library exemplifies another organizational method for structuring libraries. The National Library of Medicine (NLM), for its Visible Human Project, is in the process of creating a large digital library of anatomical images of the human body. It contains MRI and CT scans, as well as cryosection images (digital color photographs of cross-sections, Figure 1). Two cadavers, one male and one female, were carefully chosen as the subjects. The male dataset MRI and CT images were captured using the respective medical imaging scanners. The MRI images were captured at 4mm intervals throughout the body. Each is 256x256 pixels in resolution with 12-bit gray level encoding. The CT images were captured at 1mm intervals to correspond with the cryosections. Each is 512x512 pixels, also using 12-bit

gray level. To capture the cryosection images, the cadaver was first frozen solid inside a large block of blue gel. Then, 1mm thick slices were successively cut away from an axial cross-section (planar cut perpendicular to the longitudinal axis of the body), and digital color images were taken of each newly exposed cross-section. A total of 1878 cryosection images were taken, spanning the body from head to toe. Each is 24-bit color and has a resolution of 2048x1216 pixels. The total result is a 15-gigabyte image dataset of the male body. The female dataset is captured at 1/3rd mm slices and requires 40 gigabytes, nearly three times the storage space of the male dataset. In addition, 70mm film of cryosections may be scanned at higher resolutions for details of various anatomical structures.




Figure 1: A Visible Human cryosection image of the male at chest level. The actual images are full color and have 2048x1216 pixel resolution, much larger than most computer display screens.

The NLM provides public access to the Visible Human digital library via the internet. The entire dataset resides on a high capacity storage system with a high-speed internet connection. A simple license agreement with NLM, intended for tracking data usage, allows any user with an internet capable computer to remotely access the full dataset. Each image is stored on the archive as a separate file with numerically indexed filenames ordered from head to toe. Using "ftp", a user remotely logs onto the archive machine with a special login ID and password provided with the license agreement, and downloads any number of image files by specifying the correct numerical filenames. The images can then be displayed locally using an appropriate image browsing software package. Also, several sample images are available to those without a license agreement on the Visible Human World-Wide-Web page at:

<http://www.nlm.nih.gov/research/visible/>

Because the library is available to the public, the user community is diverse. Users are of varying backgrounds and cover a wide range of expertise in human anatomy and computer skills. The most prominent user groups are medical professionals and students interested in using the images for education and modeling purposes, and computer science researchers and software developers creating medical imaging applications. Other users include elementary school teachers, curious internet surfers, and even martial arts experts! Some users wish to simply explore the dataset in an open-ended fashion, and download a sampling of images to browse. Other users, who desire a specific portion of the dataset, use a more directed search to locate and download images. Some desire only a single image or small set of images, perhaps to use as overheads while teaching an anatomy class. Others download an entire continuous block of images representing some portion of the body. In this case, the images are often used for 3D modeling of anatomical structures. For example, a medical expert might download all the cryosection images containing the heart, and load them into advanced software, such as Mayo Foundation's *ANALYZE*TM, to visualize ventricle chambers. Very few attempt to download the entire library, since it would require weeks of continuous internet activity as well as overwhelming storage capacity.

Unfortunately, using the Visible Human digital library as provided by NLM can be difficult and frustrating. The ftp interface merely presents users with a long list of numbered image filenames. Retrieving desired information requires a trial-and-error process of blindly guessing an image number, downloading the image (which typically takes several minutes), and displaying it locally to determine if it contains the desired data. A web page interface that displays a long list of thumbnails of the library images would allow the user to visually browse the contents of the library, but would require paging through hundreds of thumbnails that must be downloaded for each access. Hence, there is a great need for an effective user interface to assist users in remotely browsing the library and retrieving images.

The Visible Human employs an organizational method in which the medical imagery, although un-annotated, is concretely structured. That is, the images have a strict sequential ordering defined by the physical objects (the human body) represented. Hence, because of the previously described successes, we expect that a user interface, which utilizes this structure to provide a browsing environment for the library, would be an effective means for users to retrieve desired images.

THE VISIBLE HUMAN EXPLORER

We have developed a visual information-seeking user interface, called the Visible Human Explorer (VHE), for remotely exploring the Visible Human digital library and retrieving images. It allows users to rapidly browse, on their own machine, a miniaturized version of the Visible Human dataset and, based on that exploration, download desired full-resolution library images. The miniature dataset provides an overview of library content and acts as a preview mechanism for retrieval. The interface, along with the miniature dataset, is downloaded in advance and then used as an accessor to the archive. The payoff is quick, since downloading the VHE is equivalent to only about 3 image retrievals. With this paradigm, although implemented completely outside of any existing web browser, the interface can be thought of as an advanced, dynamic web-page application for browsing a medical image digital library.

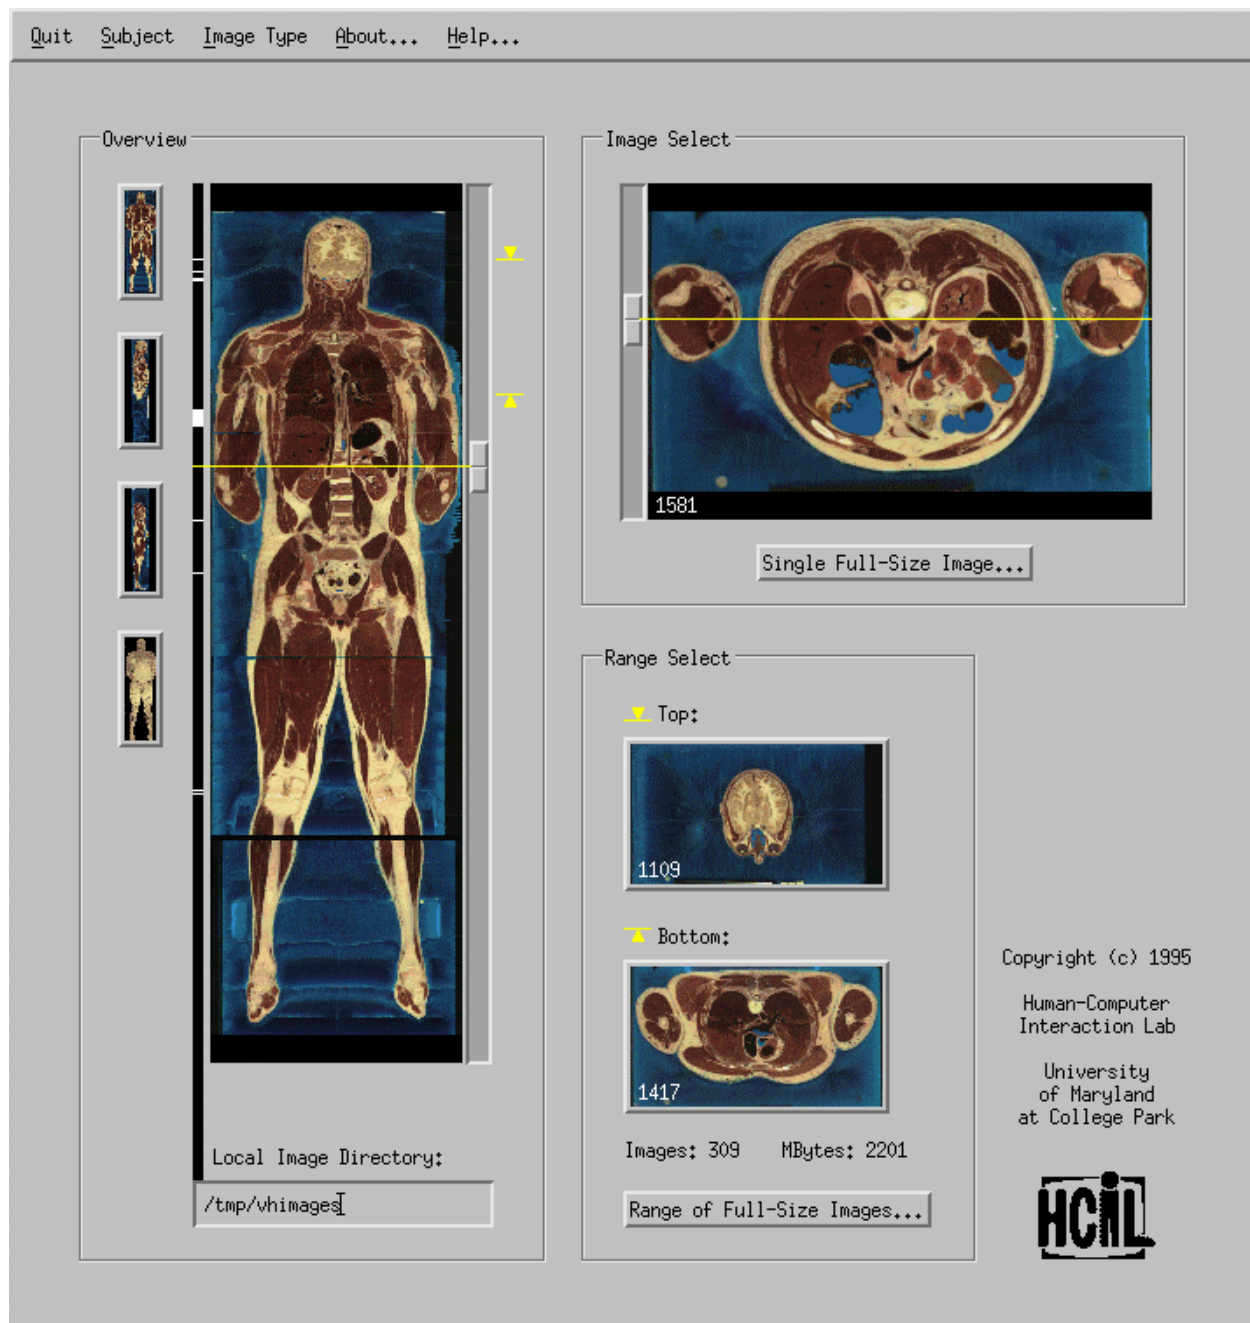


Figure 2: The Visible Human Explorer user interface, showing a digitally reconstructed coronal section overview and an axial preview image of the upper abdominal region. Dragging the sliders dynamically animates the cross-sections through the body in real-time.

The Interface

The VHE direct-manipulation interface (Figure 2) presents the user with a coordinated pair of orthogonal 2D cross-section views of the Visible Human body. The left view-window, or overview, displays a miniature coronal section, a front-view longitudinal cut of the body. It acts as an overview of the library, giving the user a general understanding of the contents of the body from head to toe. We digitally reconstructed the images for this view directly from the cross-section images, taking advantage of the structural organization of the library, by sampling a single row of pixels from successive images to simulate a perpendicular cut. The resolution is decreased to approximately 150x470. The right view-window, or preview, displays a miniature axial cross-section, the type of cut typically seen in medical atlas textbooks. The images for this view are reduced resolution thumbnail images (300x150) of the original library cross-sections and act as a preview of those higher resolution images. The magnification of this view is twice that of the overview, giving the user more detail of the library.

A horizontal indicator on the overview indicates the vertical position, in the body, of the axial section shown in the preview. The indicator is attached to a vertical slider widget spanning the height of the overview and can be dragged vertically across the body to sweep the cut, shown in the preview, through the body. As the user drags the indicator, the preview provides smooth, rapid feedback (<100msec) reflecting the axial cross-section at the sliding cut plane, resulting in a dynamic animated effect of motion through the body. This gives the user the ability to easily explore the contents of the entire body. Likewise, a sliding indicator in the preview controls the view shown in the overview. For example, the user can slide the overview to show a coronal section nearer to the front or back of the body. The result is a symmetric tight-coupling between the overview and preview. (See video report [NK96] for a demonstration of the dynamic interaction.)

While exploring, when the user locates an axial cross-section in the preview for which high resolution data is desired, pressing the Full-Size Image button opens the Retrieval dialog box (Figure 3). This dialog handles the details of accessing the Visible Human library archive. After users type their Visible Human login ID and password into the corresponding textboxes and press the Retrieve button, the system begins to download the requested full resolution cryosection image from the NLM digital library over the internet to the user's computer. Since retrieval may take several minutes, depending on internet traffic, a status meter provides feedback on the progress of the retrieval. Users have the option to cancel the downloading at any time. Once the meter reaches 100%, the download is complete and the image is then displayed in a large detail view.



Figure 3: The VHE Retrieval dialog box meters image downloading from the Visible Human digital library.

Since downloading over the internet can be a time consuming process, the system does have provisions for users with computers containing ample storage space. It can accumulate retrieved high-resolution images on the user's computer, thereby allowing the user to browse not only the remote library but also a local cache as well. A vertical bar aligned next to the overview image shows tic marks indicating which axial cryosections are contained in the local cache. When one of these cross-sections is previewed by placing the cut plane slider over a tic mark, a flag in the preview window highlights to indicate that the corresponding full-size image is in the cache. Clicking on the Display button will then display the image from the cache, saving the download time. Below the overview, a text box shows the name of the user selectable directory containing the local image cache. This feature is particularly useful for users who acquire their own copy of the Visible Human digital library, since they can use the VHE to browse it as well.

The VHE also provides the capability to download any user specified range of cross-section images into this cache at once. The user simply places range indicators next to the overview slider, and selects a button to begin retrieval using the Retrieval dialog box. Two smaller axial cross-section views display images of the top and bottom slices of the user specified range. Feedback on the number of images within the range and the total size, in megabytes, of those library images is also provided. This feature is useful for retrieving commonly examined portions of the body a priori, or for obtaining a continuous block of images for use in 3D modeling applications.




Figure 4: Several different digitally reconstructed overview images: (a) a coronal section near the back of the body, (b) a sagittal section between the legs, (c) a front “skin” view.

Also, using the icon buttons adjacent to the overview, the user can select amongst different overview image types, including the coronal section, sagittal section (side view longitudinal section), and a simple front view of the body (no cut). This gives the user a variety of views of the overall search space. Menu options enable selecting between the male or female subjects, and between the MRI, CT, and cryosection medical image types.

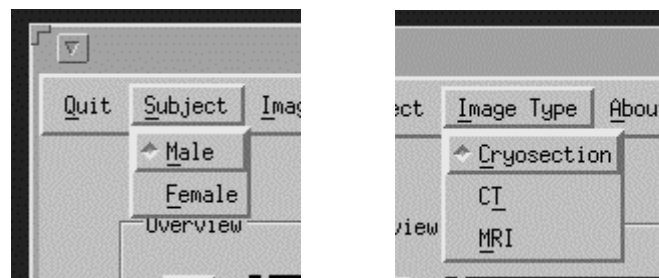


Figure 5: VHE menus for selecting among datasets.

Implementation and Performance

The VHE is implemented in the UNIX X/Motif environment on a Sun Sparcstation. The key to implementation is the use of meta-data, consisting of the miniature Visible Human thumbnail images, which acts as the library index. The benefit is two-fold: Firstly, the meta-data provides a manageable dataset for interactively browsing the library. Secondly, each axial cross-section thumbnail provides a preview of and an index to the corresponding full size library image for automated retrieval purposes.

For optimal browsing performance, the meta-data is stored locally on the disk of the user's machine and is preformatted in 8-bit XImage format. Hence, as a slider is dragged, new thumbnail images can be moved directly from disk to screen using X Windows shared memory. On an older Sparcstation 1+, images are updated at approximately 20 per second. As a result, we expect that an implementation on a PC or Mac with reasonable I/O capabilities could also yield acceptable interactive performance.

The meta-data is pre-generated and comes packaged with the VHE software. To create the meta-data, we developed several tools to automatically generate it from the library images. The thumbnail meta-data preformatted for the axial preview window is cropped, spatially sampled, and color quantized directly from library images, and requires less than 20 megabytes of storage space. The overview meta-data formatted as coronal images is spatially sampled and reconstructed from the axial meta-data, and requires less than 5 megabytes of space. The total cryosection meta-data is 1/600th of the total size of the full resolution cryosection library data. The meta-data size could be further reduced by trimming any empty space around the body (occupied by the blue gel) from each thumbnail image.

DESIGN PRINCIPLES

The VHE interface design was chosen to be in harmony with user tasks. The tightly coupled 2D views of the dataset, combined with rapid, dynamic user control of movement through the third dimension, provides a highly interactive interface yet avoids unnecessary complexities. The result is an elegant interface affording convenient user exploration of the image data, for both novice and expert users. To accomplish this, we have drawn on several major user interface design principles in the field of Human-Computer Interaction. These guiding principles are a result of extensive research in the field.

Direct Manipulation

First, direct manipulation, with compelling success in many information exploration interfaces, is a technique in which users directly interact with information objects in a visual manner. It is evidenced by these concepts: [Shn83, Shn92]

- Visual representation of the search space;
- Rapid, incremental, reversible exploratory actions;
- Pointing and selecting, instead of typing; and
- Immediate, continuous, visual feedback of action results.

With the VHE, users can quickly learn about the entire library by sweeping the views through the body, visually absorbing 20 megabytes of data in just a few seconds. Then, slower smaller movements to carefully examine interesting portions of the library help locate optimal images for downloading. Visual feedback is fast enough to engage the user. The smooth motion

is important in browsing interfaces because, when navigating from one point to another, the experience of moving to the destination helps users maintain context. Whereas, jumpy or discontinuous movement is very disorienting. Also, when dragging a slider, visualizing the resulting motion of the structural patterns in the cross-section thumbnails provides additional insight over, for example, simply viewing thumbnails side by side or in a click-and-wait incremental fashion (see the MPEG video on the CD-ROM). In addition, employing the direct manipulation technique, which users find very natural, essentially eliminates the need for user learning time.

Visual Information Seeking Mantra

At the HCIL, we have developed a mantra which embodies the process of users' information seeking activities in visual, direct-manipulation environments. It indicates an interface design approach which has evolved from many years of research and practice, and has proven to be very effective. It reads: "Overview first, then zoom and filter, and finally details on demand."

As has been discovered in exploration interfaces for many other types of data, a visual overview of the data space is extremely helpful [Shn94]. An overview provides a map from which users can construct a mental model of library contents. When users search for specific data, the VHE overview image quickly guides them to the desired location. It provides much-needed context for the axial preview thumbnails, some of which would otherwise be difficult to interpret exactly. It also encourages users to explore the library, by eliminating the penalty of possibly getting lost or not being able to return to a desired location in a timely fashion.

The VHE interface demonstrates an overviewing approach that is generalizable to other image sets in which images are sequenced in some meaningful progression. The Visible Human library contains spatially sequenced images, or volumetric data. Images can also be sequenced, for example, by time of capture, as in digital video. Figure 6 illustrates browsing a video clip of a cartoon using the VHE interface. The preview displays a single frame of the video, and the overview shows a cross-section of the clip across time. The overview image is digitally generated in the same fashion as for the Visible Human, by sampling a single row of pixels from each frame at the location of the yellow indicator line in the preview. Hence, time proceeds downward in the overview. As a result, it provides a visualization of the distinct scenes, their lengths, and overall motion patterns and trends of action in the video. At the beginning, there is a zoom in on a black object (the unicycle's seat). Then, the camera maintains focus on the unicycle as it rides past several lamp posts; we can see the posts moving off to the left and right in the overview. This action is followed by a definite scene break. The user can inspect the video by dragging the line in the overview to animate the frames in the preview. This interface would be especially helpful in video editing and sequencing tasks.




Figure 6: Browsing a digital video clip with the Visible Human Explorer interface. The overview shows trends of action in the video and scene breaks.

However, for libraries containing heterogeneous collections of arbitrary, unrelated images or textual documents, the VHE overview is not useful (See Figure 7). In this case, the ability to quickly flip through many images or pages, like flipping through a book, is very helpful when exploring large libraries, but the overview window would need to be redesigned to show a meaningful representation of the library. For example, chronological timelines, geographical maps, starfield displays, arrays of tiny thumbnails, or even simple page numbers would provide additional understanding of library contents.

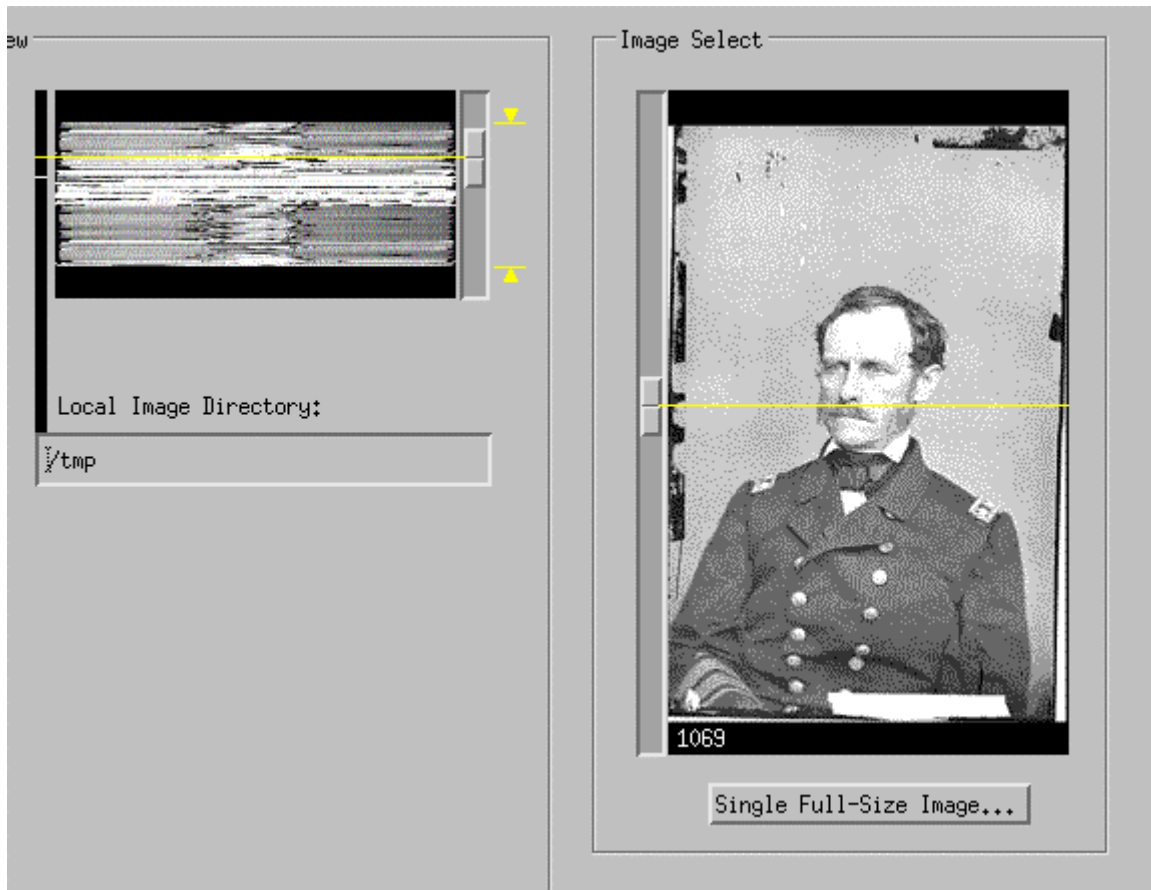


Figure 7: Browsing the Library of Congress Matthew Brady photography collection using the Visible Human Explorer interface. Since individual images are unrelated and unsequenced, a different overview technique is needed.

After starting with an overview, the next step in the visual information-seeking process is the zoom-and-filter stage. Users zero in on information of interest, while removing extraneous information from view. As users narrow the information scope, they can access the fine-grain details of the desired information on demand. In the VHE, the direct manipulation of the preview allows users to focus in on and select desired images, filtering out unwanted information. Then they can retrieve and zoom on the desired high-resolution cross-section images as needed.

As the Visible Human digital library advances to include the capability to retrieve 3D sub-volumes of the body at varying resolutions, we can better capture the more iterative visual information-seeking process, starting from overview (high breadth, low detail) and gradually zooming to detail view (low breadth, high detail). To accomplish this, we can augment the VHE interface to allow users to zoom on and browse higher resolution detailed data of a smaller portion of the body in the same fashion used to browse the initial miniature body. For example, the user could specify a 3D sub-region containing the heart. Then, the system would retrieve only that region in high resolution, and the interface would be reused for exploring that more detailed data (the same pair of tightly-coupled views, but showing high-resolution images of just

the heart). The user could continue to refine by zooming again, for example, on a ventricle valve (perhaps scanned from the super-high-resolution 70mm film).

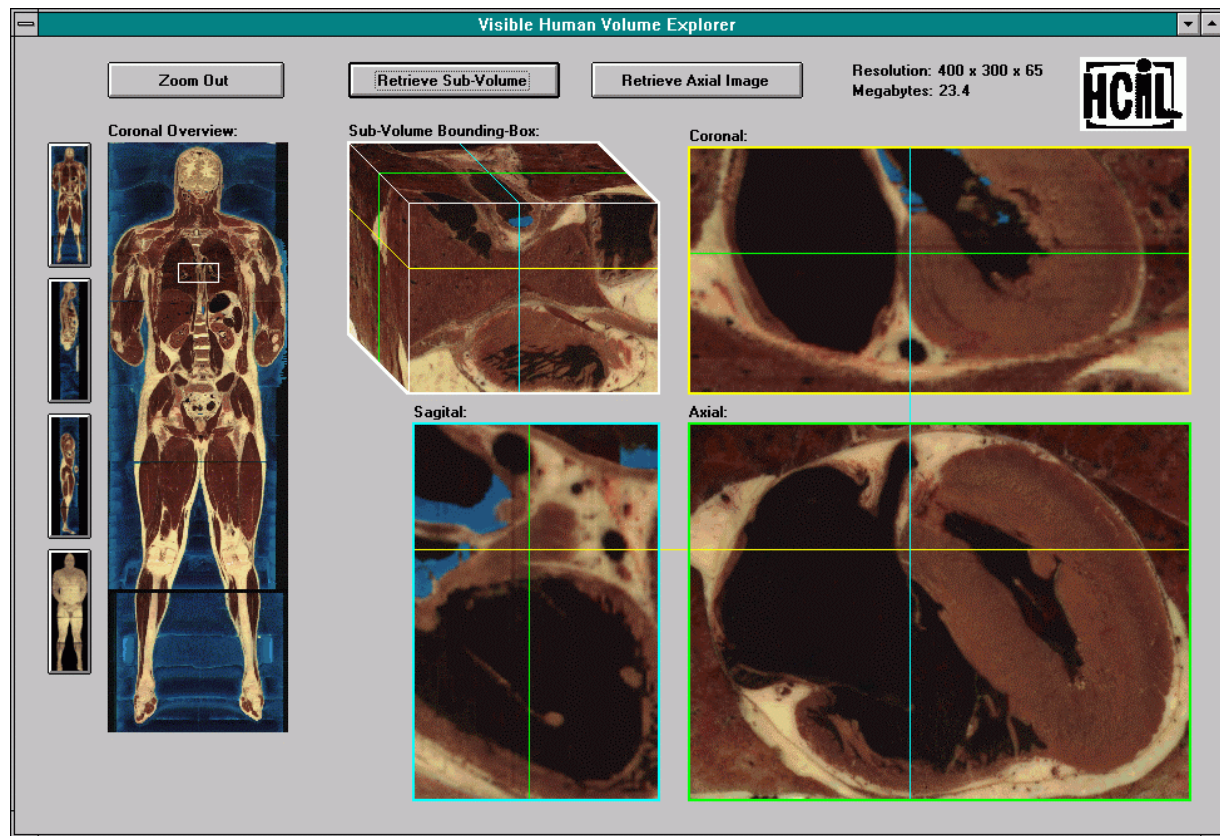


Figure 8: VHE Prototype for an improved Visible Human Library. The user has zoomed on a 3D sub-region of the body containing the heart. The interface presents three tightly coupled orthogonal views, a 3D bounding box of the sub-region, and the original coronal overview for context.

For insight on applying the Visual Information Seeking Mantra to textual and numeric data, see the chapter on Dynamic Queries.

Query Previews

In networked information retrieval environments, such as digital libraries, it is particularly difficult to provide direct manipulation style exploration interfaces. These interfaces attempt to provide real-time interactive browsing of library contents but, because contents are stored remotely, are hindered by comparably enormous network delays in accessing the contents. The solution is to provide a preview dataset (meta-data) that roughly approximates the original data, is reduced enough to be stored locally on the user's machine, and is appropriately preformatted to coincide with the type of queries issued by the interface. Then, a multi-stage visual-information-seeking interface approach can be applied in which users first browse the reduced dataset (in real time using direct manipulation techniques) to filter out the bulk of the

content, then retrieve details on only the remaining information which is now narrowed enough to be of manageable size for further detailed browsing. Hence, this technique avoids blind querying as well as slow delays in visual browsing.

The VHE accomplishes this by providing the thumbnail image set, which gives the user a rough preview of the contents of high-resolution retrievable library images. Hence, users can rapidly browse, make informed decisions about which images to download, avoid wasting time retrieving unwanted information, and, as a result, quickly obtain desired data.

The query previews technique can also be applied to scientific data. We have used query previews in conjunction with dynamic queries to design a user interface for browsing terabytes of highly distributed NASA environmental data records [DPS95]. The interface employs a two stage approach: First, users simply filter the number of record hits down to a reasonable quantity by using visual query widgets to specify rough ranges on a subset of attributes. Then, after the attribute details of those desired records are retrieved from the network, the user can use similar, but more detailed and complete visual query mechanisms to browse them. Finally, the user can then retrieve remaining storage-intensive details of individual records, such as satellite images, to complete the search task. The first stage uses a preview dataset, called a volume predictor, which simply indicates the number of records that match each possible query. Hence, contents of actual records are not available to users until the second stage, yet they still receive real-time feedback on number of hits while browsing in stage one.

Multiple Tightly Coupled Views

In many visual information-seeking environments, user interfaces employing multiple tightly coupled views of information can significantly enhance user performance. This technique provides users with several different views of the information, such that manipulating any one view has immediate visual effects in the other views. Multiple views enables users to simultaneously focus on several different aspects of the information, and browse from different points of view. Then, tight coupling, the glue that links the views together, enables users to navigate very large information spaces with simple direct manipulation techniques. Thereby, one of the perspectives on the library can be used to explore in the others.

For example, a standard overview and detail-view pair is helpful in browsing individual, very large, 2D images [CPH94, PCS95]. The overview presents a low-resolution global view of the entire image space, and the detail-view presents a high-resolution view of a small area of the image specified by a field-of-view indicator in the overview (see Figure 9). To navigate, users directly manipulate the field of view indicator in the overview, which is tightly coupled to the detail view to dynamically scroll its contents. The VHE applies this technique to 3D volumetric image spaces. A pair of tightly coupled 2D cut-plane views of the space provide several visual perspectives of the data. To explore, users rapidly sweep each view through the space by manipulating the guidance mechanisms in the other views, and receive continuous feedback of contents. Hence, the views act as maps for each other.

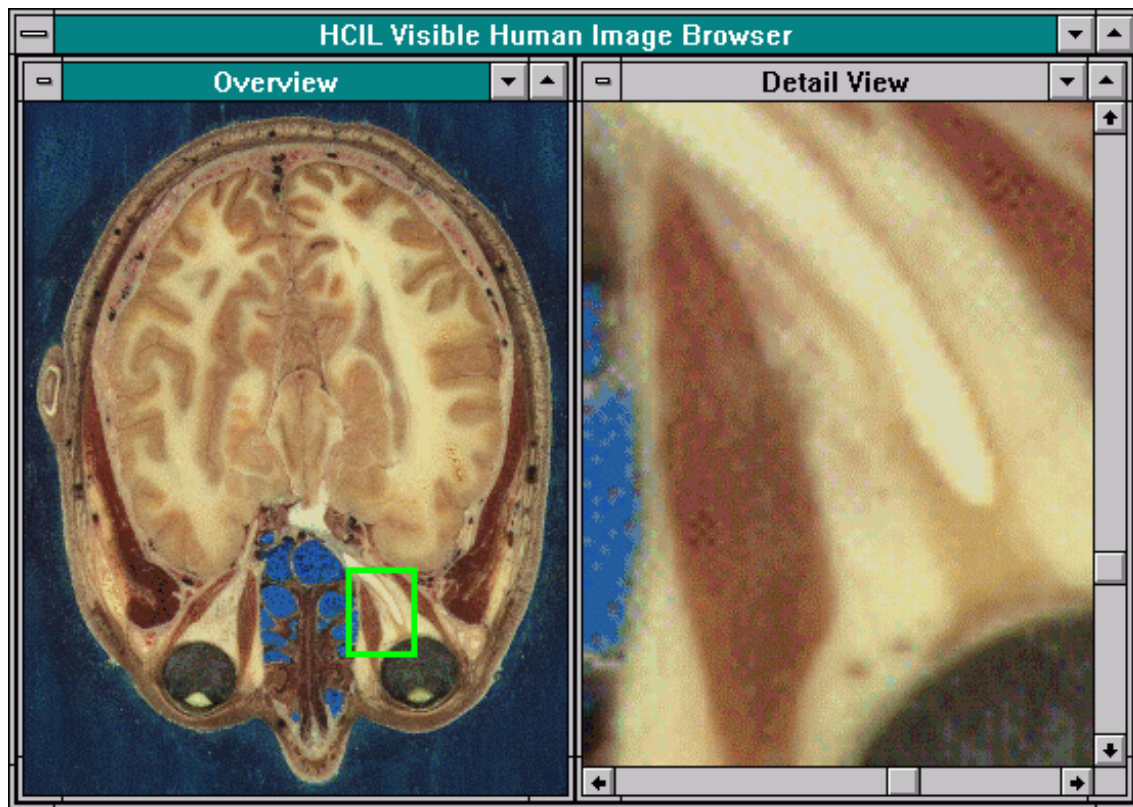


Figure 9: A standard overview and detail view pair for browsing single very-large 2D images. The field-of-view indicator in the overview is tightly coupled to the detail view.

Tight coupling can also be used to explore inter-relationships between information in the separate views. As the Visible Human digital library grows to include additional information, such as textual annotations, segmented images, attribute indexes, relationship graphs, and video animations, additional types of views can be integrated into the VHE interface. For example, integrating textual information such as the Medical Subject Headings tree (MeSH), a controlled-vocabulary medical terminology hierarchy published by NLM, gives the user another perspective on human anatomy [KS95]. Then, tight coupling between the textual and image views enables users to explore inter-relationships between terms and anatomical structures (see Figure 10). Selecting the term “Spine” highlights the spine in the image views and vice versa. Hence, users can browse with a combined strategy, utilizing their knowledge of both medical terminology and visual appearance of human anatomy, and extract additional structural knowledge of complex information. Generalizing this technique to tie many different types of information together will provide a rich and comprehensive browsing environment for an expanding Visible Human digital library and many other information collections.

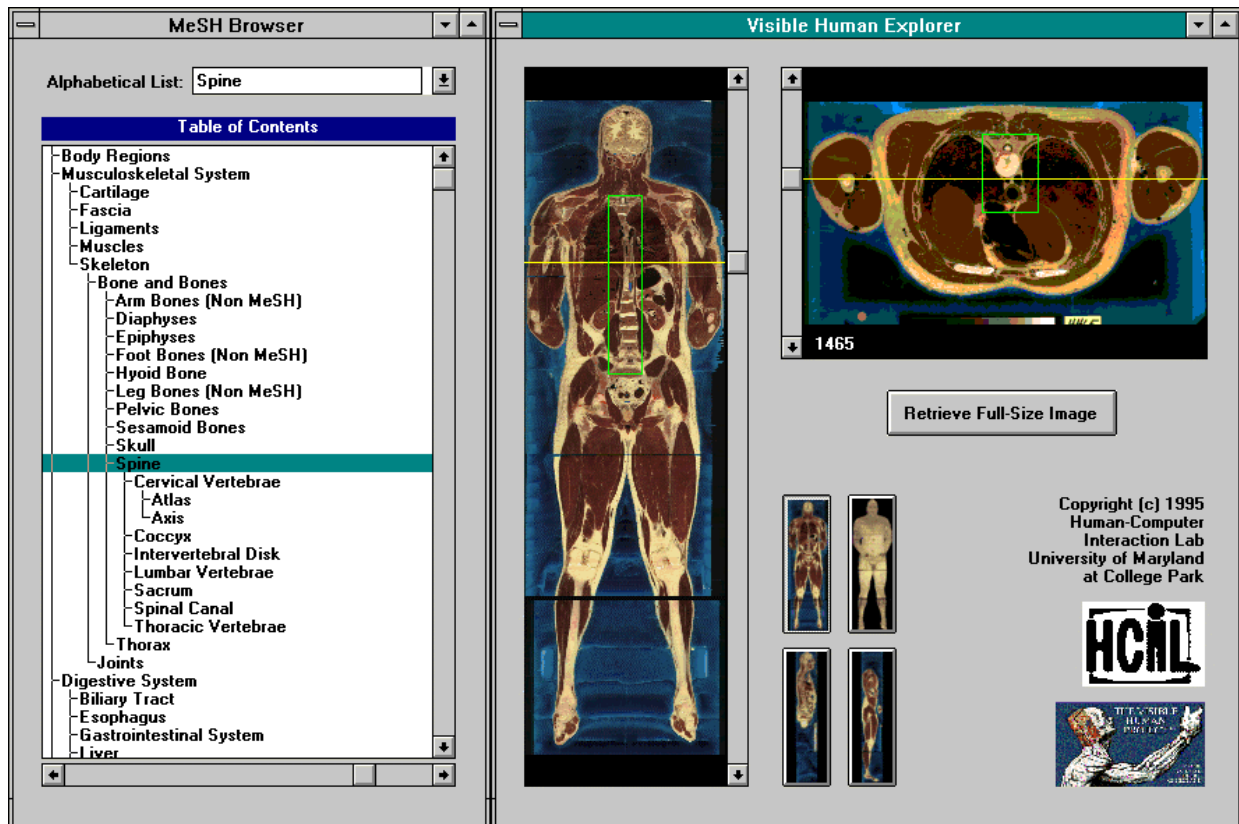


Figure 10: Tight coupling between image and textual views. Clicking on the term “Spine” in the medical terminology hierarchy reveals the anatomical structure in the image views.

CONCLUSION

The VHE interface software is fully functional and freely available for public use. We distribute it on our web site, enabling users to download and use it to access the Visible Human digital library. As a result, the VHE has gained widespread usage across the world, and we have received significant informal feedback about the interface from actual users. Many report that the VHE interface is very helpful for visualizing the library. One user summarized this well: “Now I have a far better understanding of how the whole data-set looks.” Users who need to download specific individual images were very pleased that, with the VHE interface, they were finally able to quickly locate and retrieve the proper library image. “Before using the VHE it was always a bit of a problem to get the right images.” Nearly all responses indicate the users’ excitement with the interface’s highly interactive nature, fast visual feedback, and quick learnability, all trademarks of a direct manipulation style interface. In addition, the VHE has been featured in several news and magazine articles, as well as a Japanese publication. Our distribution web page was also chosen among the top ten medical-related internet sites in *An Internet Guide for the Health Professional* [HH96].

The acclaim acquired by the VHE is a testament to the effectiveness of the design principles described in this chapter. These principles focus on providing users an information-abundant, visually intensive, dynamic, real-time, highly interactive browsing environment.

When these principles are successfully applied, the resulting interface enables quick mastery and long retention, speeds performance of tasks, eliminates the need for error conditions, and generates video-game-like enthusiasm among users.

SOFTWARE

The Visible Human Explorer interface software (for Sun Workstations), along with other information about our work on the Visible Human, including technical reports and videos, is available on our VHE distribution web page at:

<http://www.cs.umd.edu/projects/hcil/Research/1995/vhe.html>

or anonymous ftp to:

<ftp.cs.umd.edu> in `/pub/hcil/Demos/VHP`.

For information about the Human-Computer Interaction Lab (HCIL) at the University of Maryland at College Park, check out our homepage at: <http://www.cs.umd.edu/projects/hcil/>. Each of the papers in the references below is available on our web site.

ACKNOWLEDGMENTS

In support of this research, Chris North received a National Library of Medicine fellowship, administered by the Oak Ridge Institute for Science and Education. Additional funding was provided by a National Library of Medicine grant and National Science Foundation grant EEC 94-02384.

REFERENCES

- [CPH94] D. Carr, C. Plaisant, and H. Hasegawa. The Design of a Telepathology Workstation: Exploring Remote Images. University of Maryland, Dept. of Computer Science Technical Report, CS-TR-3270, 1994.
- [DPS95] K. Doan, C. Plaisant, and B. Shneiderman. Query Previews in Networked Information Systems. *Proc. of ADL96 - Forum on Research and Technology Advances in Digital Libraries*, IEEE Computer Society Press, pg 120-129, May 1996.
- [HH96] M. Hogarth, and D. Hutchinson. *An Internet Guide for the Health Professional: Second Edition*, 1996.
- [KS95] F. Korn and B. Shneiderman. Navigating Terminology Hierarchies to Access a Digital Library of Medical Images. University of Maryland, Dept. of Computer Science Technical Report, 1995.
- [NK96] C. North and F. Korn. Browsing Anatomical Image Databases: A Case Study of the Visible Human. *Conf. Companion and Video of ACM CHI96 - Human Factors in Computing Systems*, pg 414-415, April 1996.
- [PCS95] C. Plaisant, D. Carr, and B. Shneiderman. Image Browser Taxonomy and Guidelines for Designers. *IEEE Software*, vol 28, #3, pg 21-32, March 1995.

- [Pla93] C. Plaisant. Facilitating data exploration: Dynamic Queries on a Health Statistics Map. *Proc. of the Government Statistics Section, Annual Meeting of the American Statistical Assoc. Conf. Proc.*, pg 18-23, 1993.
- [Shn83] B. Shneiderman. Direct Manipulation: A Step Beyond Programming Languages. *IEEE Computer*, 16(8), pg 57-69, 1983.
- [Shn92] B. Shneiderman. *Designing the User Interface: Strategies for Effective Human-Computer Interaction: Second Edition*, Addison-Wesley Publ. Co, 1992.
- [Shn94] B. Shneiderman. Dynamic Queries for Visual Information Seeking. *IEEE Software*, pg 70-77, Nov 1994.