COMPUTERS AND BIOMEDICAL RESEARCH **30,** 472–503 (1997)
ARTICLE NO. CO971461

The Digital Anatomist Information System and Its Use in the Generation and Delivery of Web-Based Anatomy Atlases

James F. Brinkley,*,† Scott W. Bradley,† John W. Sundsten,* and Cornelius Rosse*

**Department of Biological Structure and †Department of Computer Science and Engineering, University of Washington, Seattle, Washington, 98195*

Received July 31, 1997

Advances in network and imaging technology, coupled with the availability of 3-D datasets such as the Visible Human, provide a unique opportunity for developing information systems in anatomy that can deliver relevant knowledge directly to the clinician, researcher or educator. A software framework is described for developing such a system within a distributed architecture that includes spatial and symbolic anatomy information resources, Web and custom servers, and authoring and end-user client programs. The authoring tools have been used to create 3-D atlases of the brain, knee and thorax that are used both locally and throughout the world. For the one and a half year period from June 1995–January 1997, the on-line atlases were accessed by over 33,000 sites from 94 countries, with an average of over 4000 ''hits'' per day, and 25,000 hits per day during peak exam periods. The atlases have been linked to by over 500 sites, and have received at least six unsolicited awards by outside rating institutions. The flexibility of the software framework has allowed the information system to evolve with advances in technology and representation methods. Possible new features include knowledge-based image retrieval and tutoring, dynamic generation of 3-D scenes, and eventually, real-time virtual reality navigation through the body. Such features, when coupled with other on-line biomedical information resources, should lead to interesting new ways for managing and accessing structural information in medicine.

CONTENTS

0010-4809/97 $25.00

1. INTRODUCTION

Anatomy is one of the most fundamental of the basic medical sciences, not only because it is essential for clinical medicine, but also because it provides an organizational framework for other medical disciplines (*1*, *2*). For this reason anatomy is among the first subjects studied by health sciences students.

Anatomy is primarily concerned with the physical objects and spaces that constitute the human body and with various attributes (e.g. semantic, spatial, functional, developmental) that are associated with these entities. The representation and integration of spatial and symbolic information pertaining to anatomical entities presents considerable challenges, which largely account for the difficulties in teaching and learning anatomy (*2*). Recent methods for imaging the living human body have extended the need for anatomical information from the classroom and anatomy laboratory to clinical settings.

The anatomy atlas, containing predominantly annotated images rather than text, has traditionally served in all these situations as the primary source of anatomical information. Such hard-copy atlases, however, have several limitations:

(1) They may not be available at the point of need; (2) the image needed to answer a particular clinical anatomy question may be hard to find; (3) 3-D information is hard to deduce from the 2-D illustrations; (4) the images cannot be manipulated to show views appropriate to a given query; (5) the atlases are expensive, yet only a small part of any atlas may be relevant to the problem at hand; and (6) the content is fixed, requiring a new edition when the information needs to be updated or expanded.

Computer-based methods have the potential to alleviate these deficiencies, and have therefore raised expectations for improving both the representation of, and access to, anatomical information. What is needed is an information system that delivers relevant 3-D anatomical images and text to the student's desk, the classroom and anatomy lab, the hospital bedside, the physician's office, and the patient's home, in response to specific queries from these diverse users. Such a system would be somewhat like an expert anatomist who is available where needed, who knows the available information sources, and who can answer specific questions by showing the relevant images or text from one or more of these information resources. Although no system meets such a need today, network and imaging technology have now progressed to the point where it is possible to make significant progress towards this goal.

An important development is the establishment of computer-based knowledge sources, for both spatial and symbolic anatomical information, that are qualitatively distinct from traditional methods. These resources include the National Library of Medicine's Visible Human (*3*), and less comprehensive anatomical spatial data sets such as the Digital Anatomist (*4*) and Voxelman (*5*). These strictly anatomical resources are augmented by growing repositories of clinical images with relevance to normal and pathological anatomy, e.g., the University of Washington Radiology Web server (http://www.rad.washington.edu) or the

University of Florida Interactive Tutorial on Normal Radiology (*6*). In addition, structured medical vocabulary projects, which target computerization of the medical record, incorporate symbolic information relevant to anatomy (*7–10*). An anatomy information system should be able to access and integrate these resources.

Currently, the most popular computer-based delivery method for anatomical information is the CD-ROM. Most of the CD-ROM anatomy atlases run on a Macintosh or Intel-based personal computer and are accessed using a combination of images and menus (*11–18*). Although these CD-ROMs allow nonlinear access to their contents, they are generally based on traditional hardcopy sources (published atlases, artist conceptions, photographs of anatomical specimens, 2-D radiological images), and have therefore not been shown to be more useful than hardcopy atlases (*19*).

There are only a few CD-ROM atlases that allow interaction with 3-D anatomical datasets. Among the first of these were the Digital Anatomist Interactive Atlases (*20–22*) in which the user can activate prerecorded animations, which assemble or disassemble 3-D computer graphics models derived from cadaver specimens. More ambitious is Voxelman (*5*), in which the user is able to interactively dissect a labeled 3-D volume, and to interface this volume with terms stored in a knowledge base. Both the representation and manipulation of anatomical objects approximate actual cadaver dissection, which the added advantages that the cadaver can be put back together, and the structures are labeled. However, like the other CD-ROM atlases, Voxelman still presents fixed information and cannot be interfaced with other resources. It is also expensive, limited to the UNIX platform, unavailable at many points of need, and too slow for the dynamic generation of images and real-time interaction.

The World Wide Web [23] has the potential to alleviate many of the problems with CD-ROM, and has been used to create several on-line hypertext systems in anatomy. However, these systems are mostly advertisements for stand-alone CD-ROMs available for purchase. The Visible Human has provided a source of image data which has led to some interesting Web-based applications. Examples include, (1) on-line slice browsers (http://www.npac.syr.edu/projects/vishuman/VisibleHuman.html), (2) images or movies of 3-D reconstructions, (http://www.crd.ge.com/esl/cgsp/projects/medical/), and (3) interactive atlases using in-line image maps (http://www.npac.syr.edu./projects/vishuman/VisibleHuman.html) or Java (http://www.remsoftware.com.skull/skull-anatomy.html). None of the Web-based Visible Human atlases currently has significant content, and none are yet integrated with the kind of symbolic information that would let them be used in an expert anatomy information system. However, many other applications are in development, and content continues to be added to these and other sites.

Recognizing the complexity of the problems involved, we have taken a rather different approach to the representation and distribution of anatomical information than initiatives that have focused on the generation and distribution of stand-alone products. The distribution framework we have established in the Digital Anatomist program has allowed us to take advantage of evolving technologies and provide access to expanding and reusable anatomical information resources. Our objective in this paper is to illustrate how the Digital Anatomist framework has evolved into a prototype for an anatomical information system which is well-suited for incorporating new and expanding knowledge sources as well as advances in communication technologies.

The main characteristics of our approach are: (1) information is distributed among many applications interacting over the Internet; (2) spatial and symbolic information are represented as distinct resources and are then integrated; (3) Web-based access allows the information to be delivered to any location with a Web browser; (4) the framework is independent of the particular anatomy content; and (5) a significant amount of content is available on-line. These characteristics have resulted in the widespread use of Digital Anatomist interactive atlases, both in local courses and throughout the world.

Section 2 describes the architecture and operation of the current software components, section 3 describes our evaluation of the atlas client programs, and section 4 discusses how this approach can be extended to an expert information system in anatomy.

2. ARCHITECTURE

The on-line anatomy information system we are developing is a project within the University of Washington Structural Informatics Group (*1*), the long term goal of which is to foster the development of a structure-based framework for biomedical information (*2*, *24*, *25*). The organizational framework is implemented in a distributed client-server system (*4*, *24*, *25*) in which various authoring and end-user client programs update and access an evolving set of anatomical information resources. These resources are made available over the network by means of one or more information servers. Different applications access these resources through various servers, and all components evolve as we respond to the needs of both the authors and users of the information.

The advantages of this distributed architecture are: (1) each component of the framework can be developed somewhat independently while still allowing stable versions of framework components to interact with each other; (2) partial but useful solutions can be implemented by those components that are best suited to each task; (3) information and software tools can be widely distributed; and (4) component programs of the framework can be replaced or upgraded as knowledge and technology advance, without requiring the entire system to be changed.

The distributed approach has allowed us to implement a research and development strategy in which advances in anatomical knowledge representation are incrementally combined with advances in technology, progressively leading to more sophisticated solutions to practical problems that require anatomical information. For example, the initial design of the framework was conceived in 1989 before the Web became popular (*24*), and therefore the systems included customized client and server programs. With the advent of the Web, it was relatively

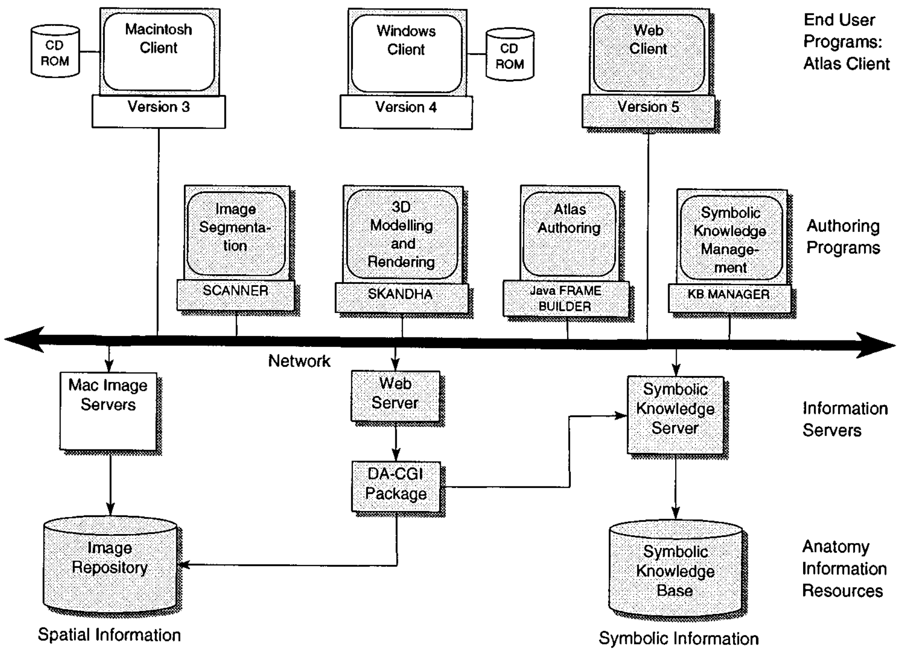

FIG. 1. Distributed software framework for the Digital anatomist Information System. Authoring programs shown in the second row are used to create the anatomy information resources shown in the bottom row. End user programs, in this case the Atlas client, access these resources by means of one or more Information Servers. Nonshaded boxes are specific for the earlier Macintosh and Windows clients, versions 3 and 4. All other components of the framework are currently in active use and development. See text.

easy to replace one or two components with Web servers and clients, a process that can be repeated as Web-based advances continue to develop.

Figure 1 shows the current implementation of this framework for managing and accessing on-line atlases of anatomy. The main components are spatial and symbolic anatomic information resources, a set of information servers that access these resources, authoring client programs for entering information into the resources, and end-user client programs for accessing these resources. A different implementation of this framework, not shown in Fig. 1, is the use of anatomy as the basis for functional brain mapping (*26*).

2.1. Evolution and Operation of the Atlas Client

Throughout this paper we distinguish between an anatomy *content atlas*, which is a collection of annotated images pertaining to a body part or region, and an anatomy *atlas client*, which is a computer program that provides access to these images. Most of this paper concerns the evolution and operation of the *atlas client*, rather than the specific contents of each atlas.

The atlas client, shown in the top row in Fig. 1, has evolved through several versions. Version 1 was a stand-alone Mac Supercard application. Version 2 was a Supercard client program that accessed information stored locally or over the network (*25*). Performance evaluation of this system (*27*) showed that delays were as much due to the client software as the network. We therefore re-coded the client in C, resulting in version 3, which is currently used not only for on-line access via the custom *Mac image servers* shown in Fig. 1, but also for accessing atlases off-loaded onto CD-ROM. Version 4 of the atlas client is written for Windows and can only access atlases off-loaded to CD-ROM. Version 5 implements Web-based access to the same images accessed by version 3. All versions of the atlas client have similar behavior, retrieving stored animations of 3-D scenes, as well as annotated 2-D images.

Figure 2 shows screen captures of two images displayed by the version 3 Mac client. Figure 2A shows a *control* image and Fig. 2B, a *content* image. Control images provide the most common form of navigation, and consist of small icons that direct the user to a desired content area within a content atlas or its chapters (sections). Once that area is reached, general navigation controls (shown in the upper right of Figs. 2A and 2B) assist in exploring the content domain. Control images either retrieve other control images depicting lists of subtopics, or retrieve content images. For example, clicking on the region labelled "Myocardium and blood supply: anterior view" in Fig. 2A retrieves the content image shown in Fig. 2B.

Content images may be accessed in two basic modes on the Mac or Windows atlas clients—*browse mode* and *quiz mode*. In *browse mode*, clicking on a structure shows the name and outline of the delineated structure ("anterior interventricular branch of left coronary artery" in Fig. 2B), whereas clicking on the "Play Movie" button causes a Quicktime movie to be brought up in a separate window. The movie shows a dynamic animation of the 3-D computer graphics model, of which the 2-D image presents but one static view. Another feature popular with students is "Show all Outlines," which outlines all identifiable structures in the image. Text definitions may also be associated with delineated structures in an image by selecting "Glossary" in the options pulldown menu.

In *quiz mode* the client program systematically displays each term one at a time, prompts the user to click on the area of the image that corresponds to it, and keeps track of the answers.

The following subsections describe the components of the current anatomy information system that support these operations. We also describe the detailed implementation of our current Web-based client (version 5), since the Web is rapidly becoming the major means for updating and accessing these information resources.

2.2. Hardware and Software Environment

Our hardware environment currently consists of Silicon Graphics workstations for image and graphics applications, NeXT computers for the database and knowledge base servers, and Macintoshes for the custom client programs. Commercial software includes the Sybase relational database (*http://www.sybase.com*), and AVS for volume visualization (*http://www.avs.com*).

Much of our custom software is written in *Slisp* (pronounced "Ess-Lisp"), a hybrid between C and a small version of Lisp (Xlisp, by David Betz (*28*)) that we developed for rapid prototyping and efficient execution (*29*). Computationally intensive data structures and functions are coded in C, then entered in the Lisp function table so that they become primitive functions at the Lisp level. Lisp level programming is then used to "glue" these primitives together in a rapid prototyping environment. Since Lisp is a complete language, arbitrarily complex Lisp programs can be created, parts of which can be re-coded in C if efficiency becomes an issue. Primitives have been added that allow any Slisp program to act as both a client and a server, with Lisp as the Application Program Interface (API) between programs. Even when Slisp is not used we generally use Lisp as the communication language between distributed applications unless a standard is already in place. Components of Fig. 1 that are written in Slisp include SCANNER, SKANDHA, and the Symbolic Knowledge Server. Other languages in use include Perl, C and Java. An advantage of the distributed framework is that individual components may be programmed in different languages as long as they can be made to communicate over the network.

2.3. Information Resources and Authoring Programs

The anatomical representations accessed by the atlas clients are implemented in the two information resources shown at the bottom of Fig. 1. Spatial information is represented in a repository of annotated 2-D images and 3-D animations, and symbolic information is represented as semantic networks and definitions in the symbolic knowledge base. These two resources are related to one another by means of anatomical terminology.

2.3.1. Symbolic Knowledge Base. The Symbolic knowledge base is being developed as an anatomical enhancement of the National Library of Medicine's Unified Medical Language System (UMLS) (30). Currently, it consists of a knowledge modeling ontology in which physical anatomical entities are assigned to classes according to defining attributes that they share (*30*, *31*). Anatomical terms denoting the concepts that constitute the ontology are linked to one another by the -is a- relationship, and are also grouped in additional semantic networks according to relationships such as -part of-, -branch of-, and -tributary of-, in order to

---

FIG. 2. Screen captures from Thorax content atlas displayed by the Version 3 Macintosh client. (A) Control image, in which active regions (indicated by thumbnail images) define commands to open other images. General navigation controls are a separate window in the upper right. Clicking on the region labeled "Mycoardium and blood supply: anterior view" (third column, second row) loads the bottom image. (B) Content image, in which active regions define structures that have been outlined and labeled. Associated Quicktime animations may also be played by clicking on "Play Movie."

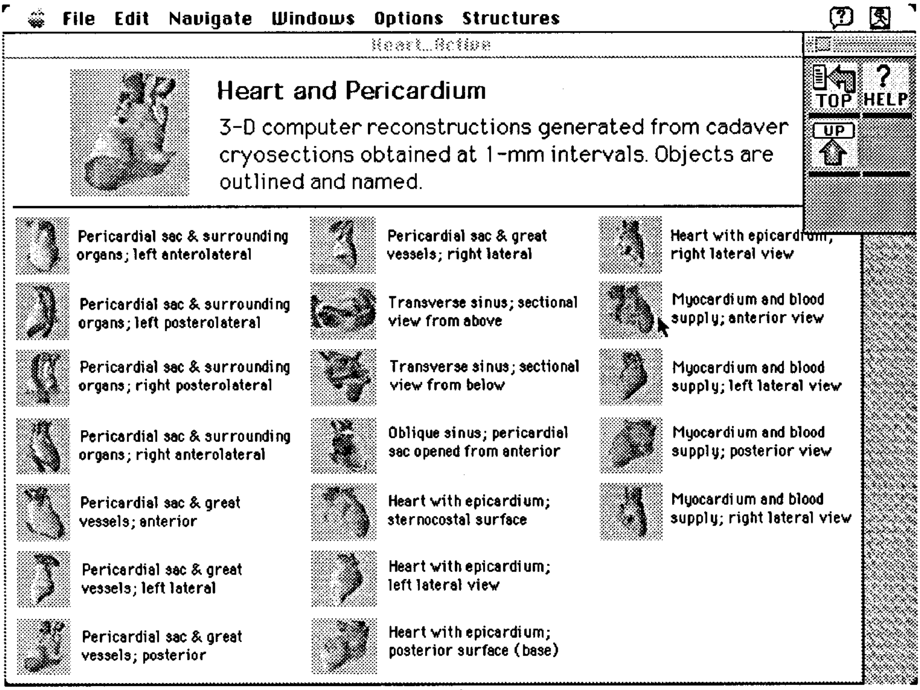
File Edit Navigate Windows Options Structures
Heart_Active
Heart and Pericardium
3-D computer reconstructions generated from cadaver cryosections obtained at 1-mm intervals. Objects are outlined and named.
TOP HELP
UP
Pericardial sac & surrounding organs; left anterolateral
Pericardial sac & surrounding organs; left posterolateral
Pericardial sac & surrounding organs; right posterolateral
Pericardial sac & surrounding organs; right anterolateral
Pericardial sac & great vessels; anterior
Pericardial sac & great vessels; left lateral
Pericardial sac & great vessels; posterior
Pericardial sac & great vessels; right lateral
Transverse sinus; sectional view from above
Transverse sinus; sectional view from below
Oblique sinus; pericardial sac opened from anterior
Heart with epicardium; sternocostal surface
Heart with epicardium; left lateral view
Heart with epicardium; posterior surface (base)
Heart with epicardium, right lateral view
Myocardium and blood supply; anterior view
Myocardium and blood supply; left lateral view
Myocardium and blood supply; posterior view
Myocardium and blood supply; right lateral view

A

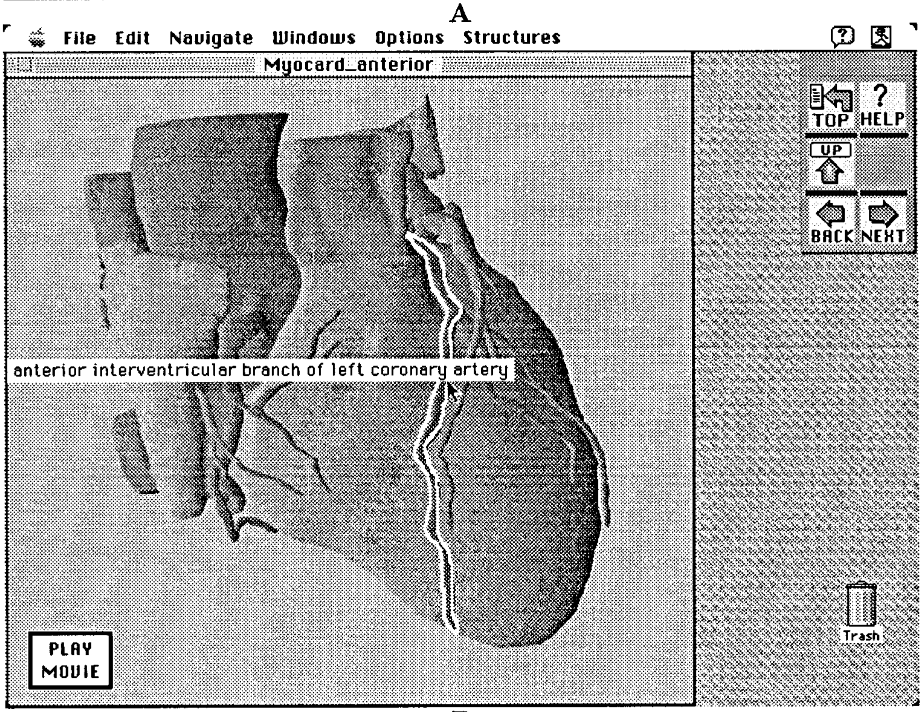
File Edit Navigate Windows Options Structures
Myocard_anterior
anterior interventricular branch of left coronary artery
PLAY MOVIE
TOP HELP
UP
BACK NEXT
Trash

B

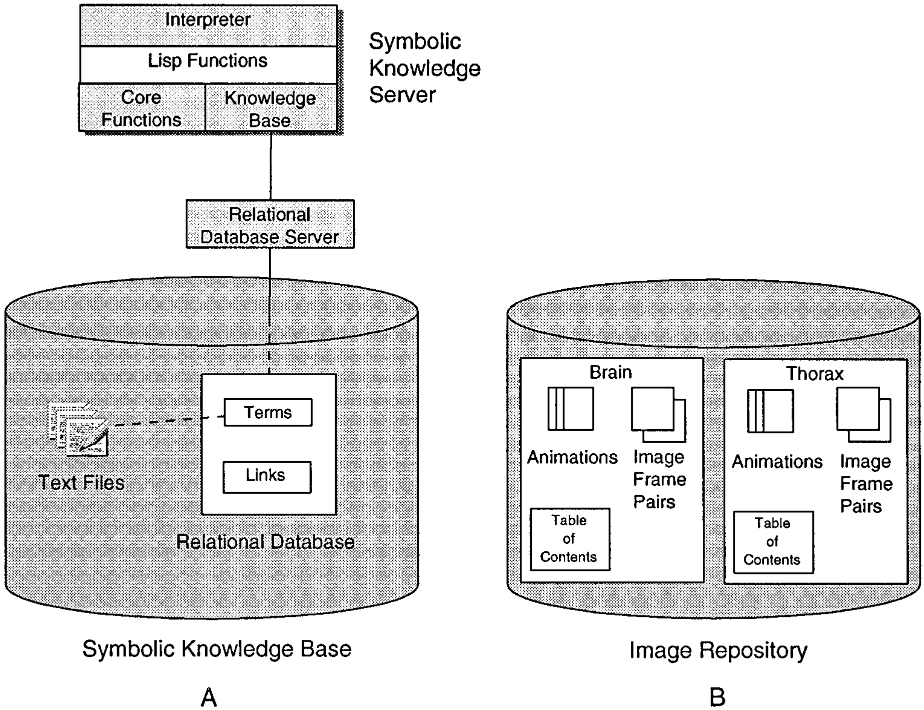

FIG. 3. Detailed view of the anatomy information resources shown at the bottom of Fig. 1. (A) Symbolic knowledge base and symbolic knowledge server. The knowledge base semantic network is implemented in two tables of a relational database. Text files contain associated definitions. The Symbolic Knowledge Server is an Slisp application, in which C-coded *Core Functions* and *Knowledge Base Functions* are called by Lisp-coded *Lisp Functions* which are in turn called by a C-coded Lisp *Interpreter*. Client programs send Lisp commands to the server. (B) Image Repository. Each content atlas (e.g., Brain or Thorax) is stored in a separate server directory, and consists of a set of animations, a set of image-frame pairs, and a derived table of contents.

model the physical organization of the body. To date we have entered 13,500 preferred terms and synonyms, representing 7,400 distinct structures visible within the thorax at 1-mm resolution.

The implementation of the symbolic knowledge base is shown in Fig. 3A. A relational database stores the terms and links in separate tables. The terms provide an index to associated files giving textual definitions for the concepts represented by the terms. The terms also provide linkage to images in the Image Repository (Fig. 3B).

The terms and semantic links are entered and arranged by means of the NeXT-based Knowledge Manager authoring program (Fig. 1), which provides a visual display of various semantic hierarchies. For example, Fig. 4 is a screen shot of a portion of the -branch of- hierarchy for the aorta. The structure outlined in

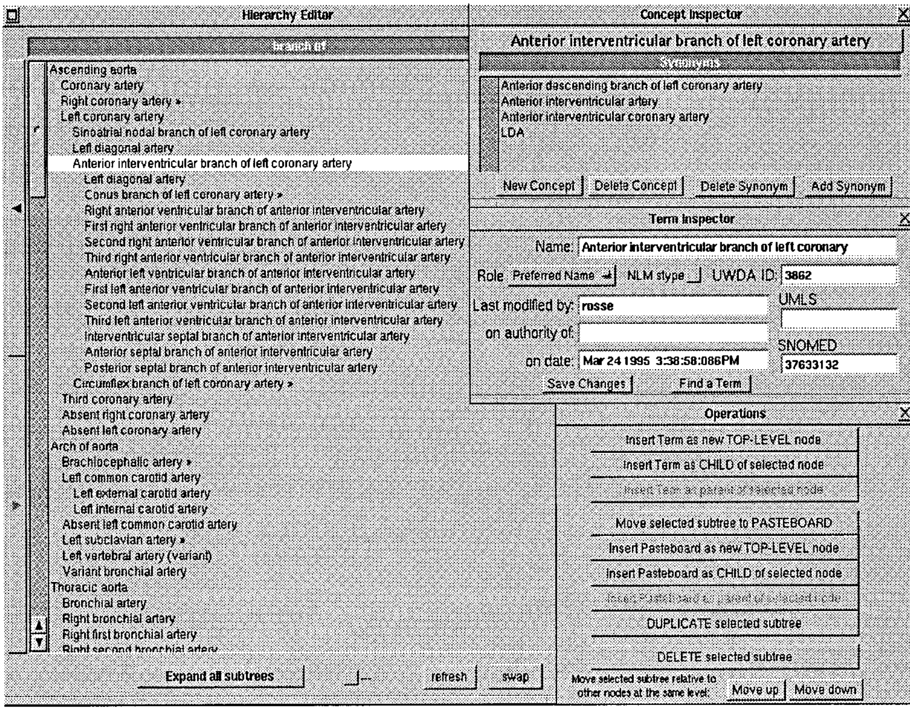

FIG. 4. Screen shot of the Knowledge Manager Program, which is used to add terminology and semantic relationships to the Symbolic Knowledge Base. The *Anterior interventricular branch of left coronary artery* is shown in a symbolic (or semantic) context, as compared with Fig. 2B, in which the same artery is shown in a spatial context. Tab indents indicate that this artery has a number of branches, and that it is a branch of the Left coronary artery, which is itself a branch of the Ascending aorta.

Fig. 2B, "anterior intraventricular branch of left coronary artery," is highlighted to show its semantic location in the -branch of- hierarchy.

The Knowledge manager communicates directly with the database by means of the relational database server, which is part of Sybase (Fig. 3A). Other application programs communicate by means of the symbolic knowledge base server discussed in section 2.4.2.

2.3.2. Image Repository. The 3-D animations and 2-D annotated images accessed by the various atlas clients are stored in the image repository shown in the lower left of Fig. 1, and in more detail in Fig. 3B. Each annotated image is stored in the repository as an image-frame pair, where a *frame* is an ASCII file containing delineated regions and structure names associated with the image. Note that our use of the term *frame* in this context is different from a Netscape frame.

Image-frame pairs and animations are organized into directories that pertain

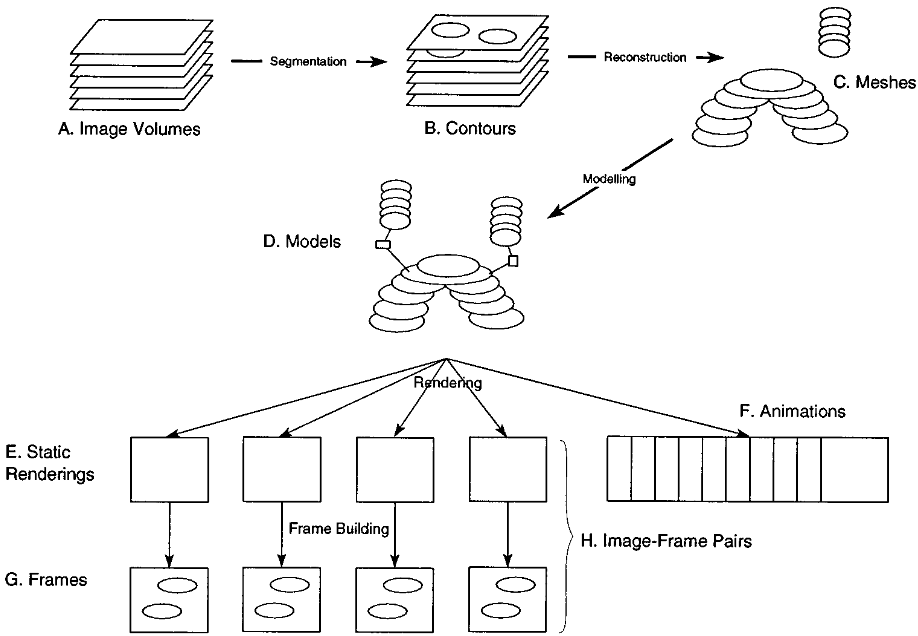

FIG. 5. Generation of 3-D renderings and animations. Image volumes (A) are segmented into contours (B), reconstructed into polyhedral meshes for each separate structure (C), then combined into models of structures with parts (D). Models are rendered to generate both static renderings (E) and animations (F). Structure regions on the renderings are delineated with the Frame Builder tool and saved as Frames (G). Animations and Image-Frame pairs (H) are accessed by the atlas client programs.

to body parts such as brain, thorax, or knee. We refer to such a directory as a *content atlas*, where a content atlas forms the subject domain in the navigation menus of the atlas client programs, and also forms the content of CD-ROM atlas publications (*20–22*). The majority of images in the image repository are static renderings of 3-D scenes created by the SKANDHA authoring program (Fig. 1), but radiographs, histological sections, photographs, or any other type of annotated images are also stored in the image repository.

The SCANNER and SKANDHA authoring programs (Fig. 1) support the process of 3-D reconstruction and modeling, shown schematically in Fig. 5. Input images are obtained from photographs of cryosectioned cadavers (*32*, *33*), or CT and MRI images of patients or volunteers.

The images are segmented with the SCANNER program, which also serves as a testbed for research on knowledge-based segmentation (*34*, *35*). For the most part SCANNER must be run in manual mode because the vast majority of structures can only be delineated by experienced anatomists.

The output of SCANNER is sets of contours, representing structures and spaces on adjacent slices. These contour sets are input to the SKANDHA program, where they are grouped into defined structure subparts, then tiled and interactively edited to generate 3-D polyhedral meshes, one for each distinct structure subpart. The meshes are edited and interactively combined to generate hierarchical 3-D models. The models are rendered by SKANDHA to produce both static 2-D snapshots of 3-D scenes and animated sequences saved as Quicktime files. The latter convey the 3-dimensional character for the rendering and also dynamically assemble and disassemble the renderings to demonstrate the spatial relationships of their constituent parts.

The static 2-D snapshots, plus other 2-D images such as X-rays or histological sections, are input to the Java-based Frame Builder tool (Fig. 1, second row, and Fig. 6), which is used by the anatomist authors to delineate named structures on the images, and to link frames together in a control hierarchy such as subject chapters organized by structure or function. The output of Frame Builder is a series of frame files, one file per image in the atlas. The frame file is written as a series of Lisp commands in order to maintain maximum flexibility as new features are added. Each region in the frame file is associated either with a named structure or with a Lisp command, such as (OPEN-FRAME *newframe*), that provides a link to another frame. Authors generally use OPEN-FRAME commands on *control* images, such as that shown in Fig. 2A. *Content* images, such as that shown in Fig. 2B, form the leaf nodes of the navigation tree, and generally associate structure names with delimited regions, although arbitrary Lisp commands may be associated with regions as well.

The animations and image-frame pairs for a given atlas are saved by Frame Builder within the relevant content atlas directory in the image repository. At periodic intervals an *update* program is run (by clicking a button on the Frame Builder Web page) which generates a cache of images with all structures outlined, and traverses the hierarchy implicit in the control frames to generate a tab-delineated Table of Contents. The files generated by the update program are used by the Web atlas client.

2.4. Web Atlas Client and Servers

The development of the Web atlas client was motivated by the spectacular rise of the World Wide Web. The advantages offered by the Web have freed us from having to maintain one-of-a-kind server applications, and we therefore have experimented with the Web for accessing the Digital Anatomist information resources.

Initially we made the version 3 Mac Atlas client program into a Web-helper application: the Web was used for navigation but the Mac client was called for interaction with a content frame. The advantage of this approach was that a particular content atlas could be made available for a much wider distribution through the Web, while interaction with the images could remain fast since the annotated images were accessed locally by the Mac client. However, the disadvantage was that it was still necessary to develop and maintain a custom helper application for each of the three major platforms (Windows, Mac and

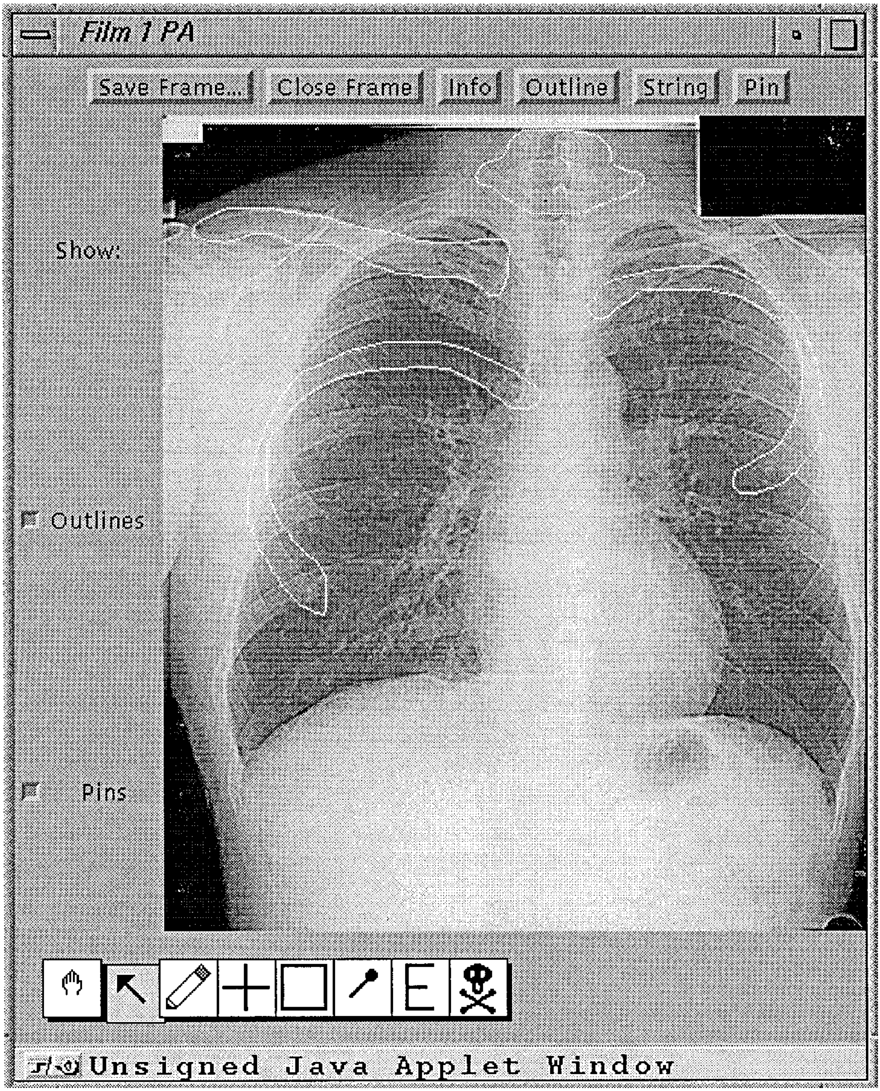

FIG. 6. Java-based FrameBuilder tool for annotating images. The author logs in to a given content atlas, uploads images from his or her local desktop, outlines and labels structures on the images, and links images together in a control hierarchy. The resulting image-frame pairs are saved in the relevant content atlas directory of the image repository, after which they become available to the Web-based atlas client, (version 5 in Fig. 1). The image shown is a chest xray being annotated for inclusion in the Thorax content atlas.

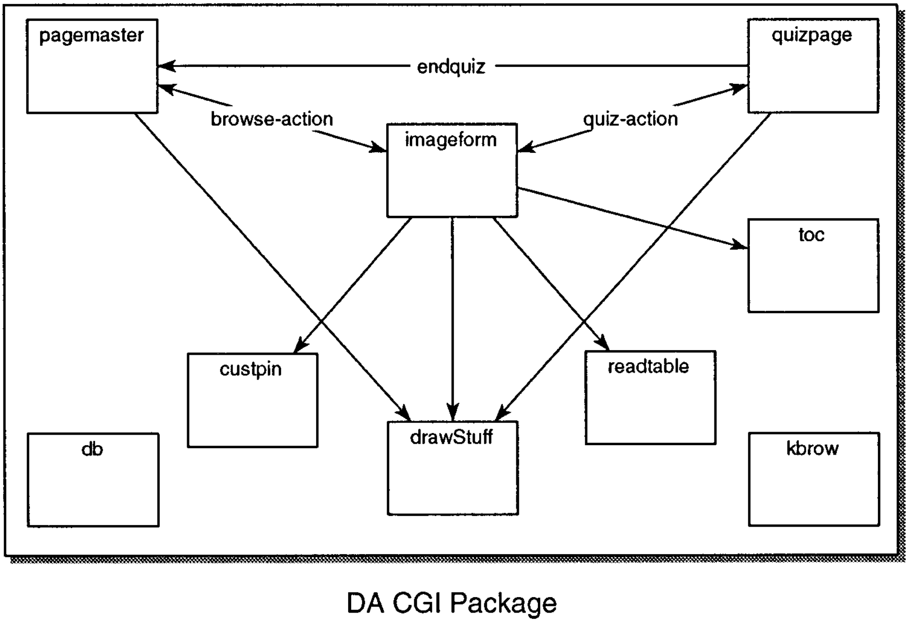

FIG. 7. Detailed view of the DA-CGI Package shown in Fig. 1. The DA-CGI Package consists of a set of C programs that interact via the Web client or via Unix system calls, and which access the anatomy information resources shown in Fig. 3. See text (section 2.4).

X). We decided, therefore, to explore the possibility of a Web-only atlas client, trusting that Web developments would soon allow us to replicate most of the features of the Mac client. The current Web client is the result of this effort.

In version 5 of the atlas client, the Mac client is replaced by a Web client, and the Mac image servers are replaced by a Web server and associated CGI (common gateway interface) package. All other components are essentially the same as they were before the Web became available. Thus the distributed approach has made it relatively easy to take advantage of advances in Web technology.

The primary mechanism for replicating the Mac features was the development of a set of relatively small CGI programs that are accessed by a standard Web server. The programs, which together comprise what we call the Digital Anatomist Common Gateway Interface Package (*36*), are shown as the box labelled *DA-CGI Package* in Fig. 1, and in more detail in Fig. 7. The programs are written in ANSI C, and to date have been compiled and tested on NeXT, SGI and DEC Alpha computers.

Since the current Web configuration does not maintain state it was necessary to develop mechanisms for passing state information between the user client and the various CGI programs, in order to create the appearance of an interactive atlas. The main design decision was the use of *forms* (http://www.ncsa.uiuc.edu/

SDG/Software/Mosaic/Docs/fill-out-forms/overview.html) in dynamically generated HTML documents, along with the use of hidden fields within the forms for passing state information from one program to another. The forms and CGI programs interact with the user in several modes of operation: browse and quiz modes (as in the Mac client), and additional new modes that are not available in the Mac client. As shown in Fig. 7, the central CGI program is IMAGEFORM, which is the CGI program, specified in the forms generated by both PAGEMASTER, which handles browse mode, and QUIZPAGE, which handles quiz mode.

2.4.1. Browse and Quiz Modes. An external link to the Digital Anatomist atlases usually specifies browse mode. For example, the URL *http://www9.biostr.washington.edu/cgi-bin/DA/PageMaster?atlas : Neuroanatomy+ffpathIndex: 3D^Object^Composites:Cerebellum^Plus+2* causes the Web server to run the PAGEMASTER CGI program, with the command line arguments *atlas:Neuroanatomy, ffpathIndex:3D^Object^Composites:Cerebellum^Plus*, and 2. The atlas argument specifies the name of the content atlas directory in the image repository. The ffpath argument specifies the path within the atlas directory to the frame file generated by Frame-Builder, where the ^s are placeholders for spaces in the frame name. The number 2, at the end of the argument list, specifies the number of "widgets" that control the appearance of the generated HTML page. The minimum widgets required are the name of the content atlas and the path to the frame file.

PAGEMASTER generates an HTML page like that shown in Fig. 8 (the exact format will change as the programs are developed). One of the fields of the form is an INPUT field of type *image*. The SRC URL in the image field is a direct reference to the GIF image specified in the Frame file. If the *Show All Outlines* button is on, the URL points to the outlined GIF image generated by the update program. If the outlined GIF image is older than the original GIF image (the author has not run the *update* program after uploading a new image), the image SRC URL is set to DRAWSTUFF (Fig. 7), a general purpose utility for drawing and labelling images. DRAWSTUFF then generates the updated outline image.

The *action* CGI program specified in the form generated by PAGEMASTER is IMAGEFORM. Therefore, this program is called by the Web server when the user clicks on an area of the form. IMAGEFORM is responsible for processing mouse clicks and for branching to other appropriate CGI programs, depending on the location of the click. For example, if the user clicks an area within the image, IMAGEFORM reads the framefile specified by hidden fields in the form, then determines which of the delineated regions was clicked. If a structure name is associated with the region in which the mouse click occurred, IMAGEFORM calls PAGEMASTER with the same framefile path and atlas initially specified as arguments to PAGEMASTER, but with an additional argument that specifies the structure name. PAGEMASTER then generates the same HTML page as before, adding the structure name at the top (*thalamus* in Fig. 8). This operation corresponds to browse mode on the Mac client, and is relatively fast for most Web browsers because the GIF image is generally available in the client cache.

If the delineated region specifies on OPEN-FRAME command, like those shown

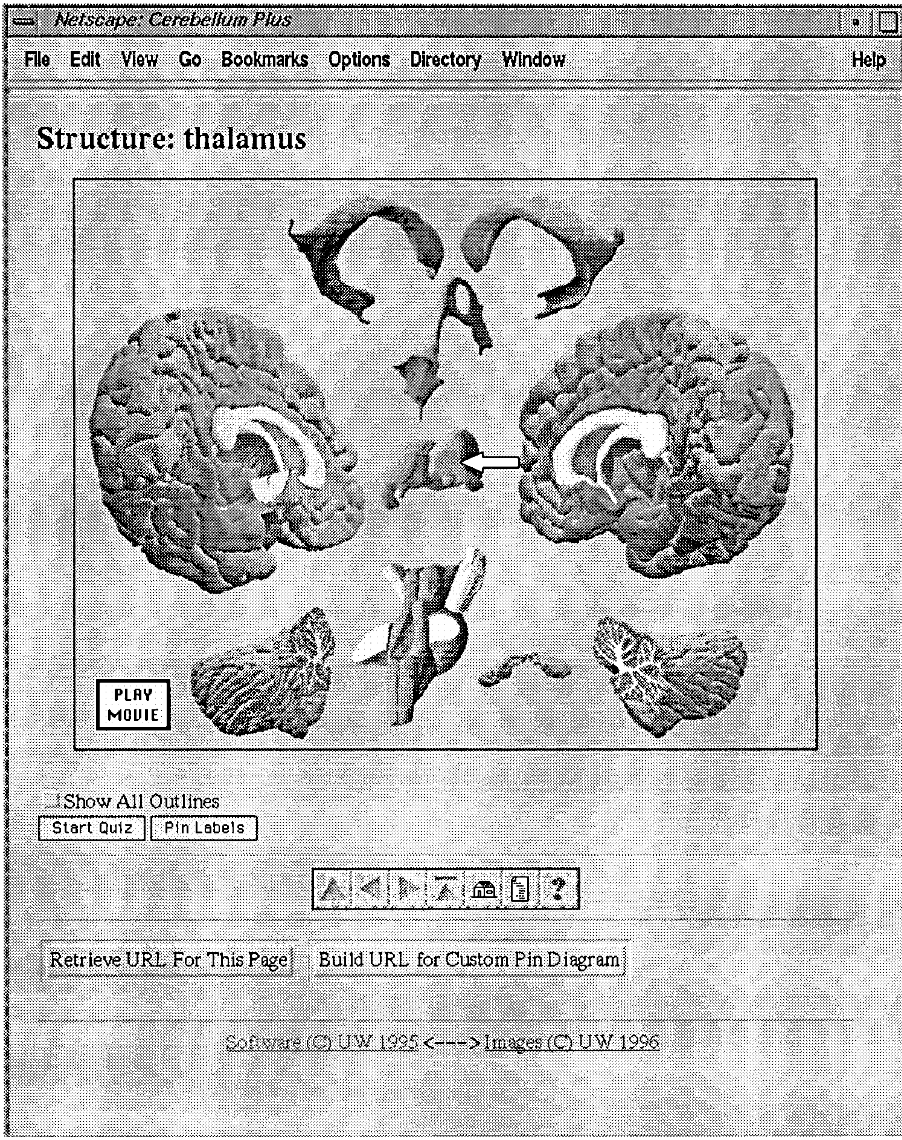

FIG. 8. Browse mode for the Brain content atlas. The HTML page initially generated by the PAGEMASTER CGI program (Fig. 7) is the same as this figure, except that the Structure field at the top of the page is blank. Once the user clicks on an annotated region in the image (as for example, that shown by the arrow), PAGEMASTER regenerates the same page with the addition of the term that identifies the structure, in this case *thalamus*. Other buttons can be clicked to enter quiz mode (Fig. 9), to show outlines of all annotated structures, to generate a custom pin diagram (Fig. 11), to navigate among frames (Fig. 10), and to generate URLs that can be used as links from Web-based tutorials or syllabi (Fig. 12).

in Fig. 2A, IMAGEFORM calls PAGEMASTER with the name of the frame that is to be opened rather than the frame by which it is called, thereby causing the Web client to load the new image.

Other actions are specified if the mouse is clicked in a different area of the HTML form. For example, if the user clicks in the navigation bar (the bottom row of icons in Fig. 8), IMAGEFORM calls the READTABLE CGI program. READTABLE loads the table of contents file generated by the offline update program (described in section 2.3.2), and finds the current frame name in this file. READTABLE then searches the table of contents for the parent, next sibling or previous sibling of current frame, depending on which button was pressed in the navigation bar. The name of this frame file is then passed to PAGEMASTER, which generates the new HTML form.

If the user clicks the START QUIZ button (Fig. 8) IMAGEFORM calls QUIZPAGE (Fig. 7) to handle quiz mode. QUIZPAGE is similar to PAGEMASTER, but formulates quiz page forms like that shown in Fig. 9. The forms are also processed by IMAGEFORM, which uses hidden items to keep track of the number of correct answers and the names of structures that the quiz has already presented. The user can request that the computer show the correct answer to a quiz question by outlining the relevant region on the image. In this case DRAWSTUFF is called to outline the structure. When the quiz is over, or when the user clicks the END QUIZ button, control is transferred back to PAGEMASTER.

2.4.2. Additional Modes. The ease of interface development with HTML has allowed us to add new features to the Web atlas client relatively quickly. Although similar features could be added to the Mac client, the effort would be much greater because Mac interface routines would need to be called from C. In addition to Browse and Quiz, six other modes are available on the Web atlas client.

Table of Contents mode is reached by clicking on the Table of Contents icon in the navigation bar (Fig. 8). This mode calls the TOC CGI program (Fig. 7), which reads the table of contents file and displays portions of the contents depending on which bits are set in a hidden bit vector (Fig. 10). The user can expand or contract topics in the display by clicking on the small arrows, or can click on a frame name to call PAGEMASTER with the new frame.

A *Pin Diagram* is generated by clicking on the *Pin Labels* button in a content image (Fig. 8), which again calls DRAWSTUFF. In this case DRAWSTUFF reads the current image-frame pair and dynamically generates a new GIF image. In this image, structure names are placed in the margins, and lines are drawn between each name and the center of the corresponding delineated region on the image (Fig. 11). Heuristics are used to place the labels and the endpoints of the lines in reasonable locations within the structure outlines. The endpoints may also be explicitly specified in FRAMEBUILDER.

If the user clicks on *Retrieve URL for This Page* (Fig. 8), IMAGEFORM generates a URL that can be pasted into another document. If the user clicks *Build URL for Custom Pin Diagram*, the CUSTPIN CGI program (Fig. 7) is called to generate the page shown in Fig. 12, which allows the user to select which structures to

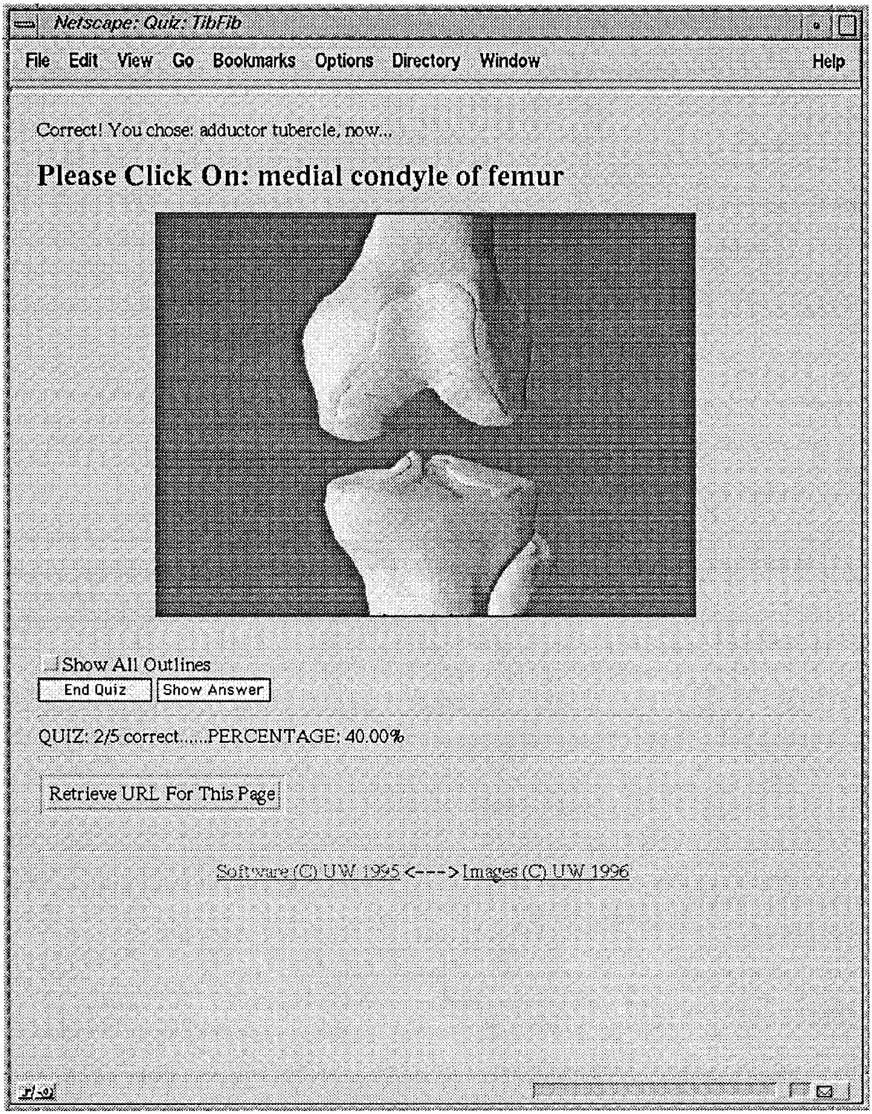

FIG. 9. Quiz mode for the Knee content atlas. The computer randomly asks the user to click structures listed in the frame file, in this case the "medial condyle of femur." At any time SHOW ANSWER can be clicked to see the outline of the correct structure. A score is kept of the percentage correct answers. Once all structures have been asked the user is returned to browse mode, or the user can END QUIZ at any time. Quiz mode, as well as all other modes, is automatically available for any image that has been annotated using Frame Builder.

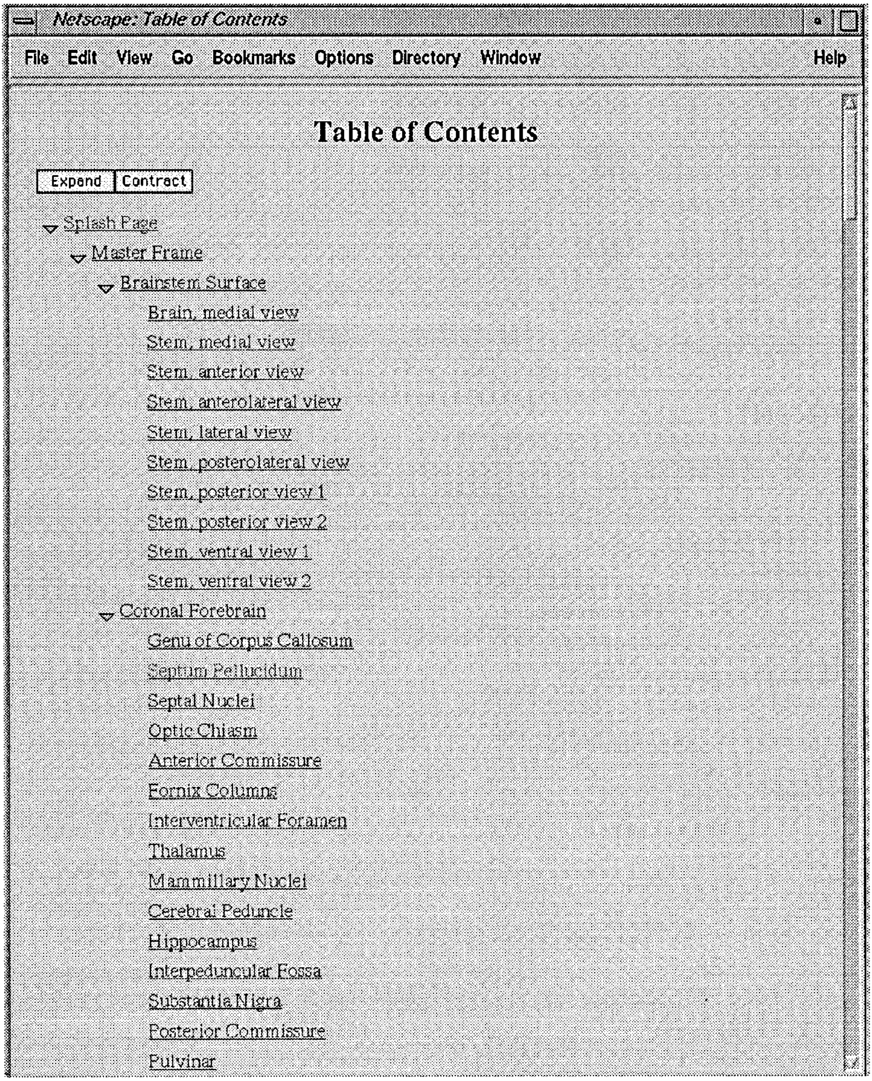

FIG. 10. Table of contents mode for the Brain content atlas. The user can expand and contract the hierarchy (by clicking on the small arrows), to show the names of images that are included within each subject area, then click on an image name to retrieve the image in browse mode.

show in a pin diagram, and to label those structures with names or numbers. The generated URL can then be pasted into an on-line quiz, in which the student is asked to type in the names of structures indicated by numbers. For example, Fig. 13 shows the image generated by the URL created in Fig. 12, which was then incorporated into a neuroanatomy course syllabus.

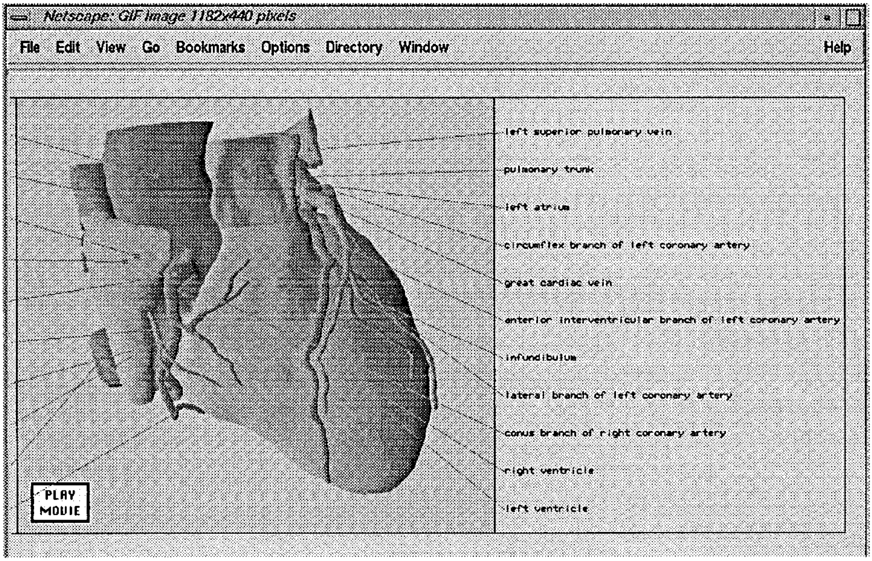

FIG. 11. Dynamically-generated "Pin" diagram for an annotated image from the Thorax content atlas. Note that this same image is available both from the custom Mac client (Fig. 2), and from the Web client.

Access to the Symbolic Knowledge Base is available via the KBROW program (Fig. 7), which initiates a connection to the Symbolic Knowledge Server (Fig. 3A). The knowledge server is an Slisp application that accesses the Sybase Relational Database Server via a set of *Knowledge Base Function* encoded in C as Lisp primitives. These functions are called by Lisp-level Functions which implement a high-level query language that hides the detailed Structured Query Language (SQL) (*37*) from the client programs. For example, the function (KB-GET-ATTRIBUTE termname definition) is called to get the textual definition of a term from the associated text file, and to display the text and any synonyms that are found in the knowledge base. If the function (KB-GET-CHILDREN *termname linktype*) is repeatedly called, the results can be used to generate a semantic hierarchy similar to the Knowledge Manager (Fig. 14).

The advantages of the intermediate Lisp-level query language are (1) the query language can be closer to the application domain than the detailed database language, (2) we are free to change both the user interface and the backend database since the query language remains the same, and (3) the query language may easily be extended to add new types of queries.

Search mode can be entered either from the initial atlas Web page or after a term has been clicked in the Knowledge Browser (Fig. 14). Search mode is processed by the DB CGI program (Fig. 7), which queries a relational table via a standard Sybase relational database server (not shown, and different

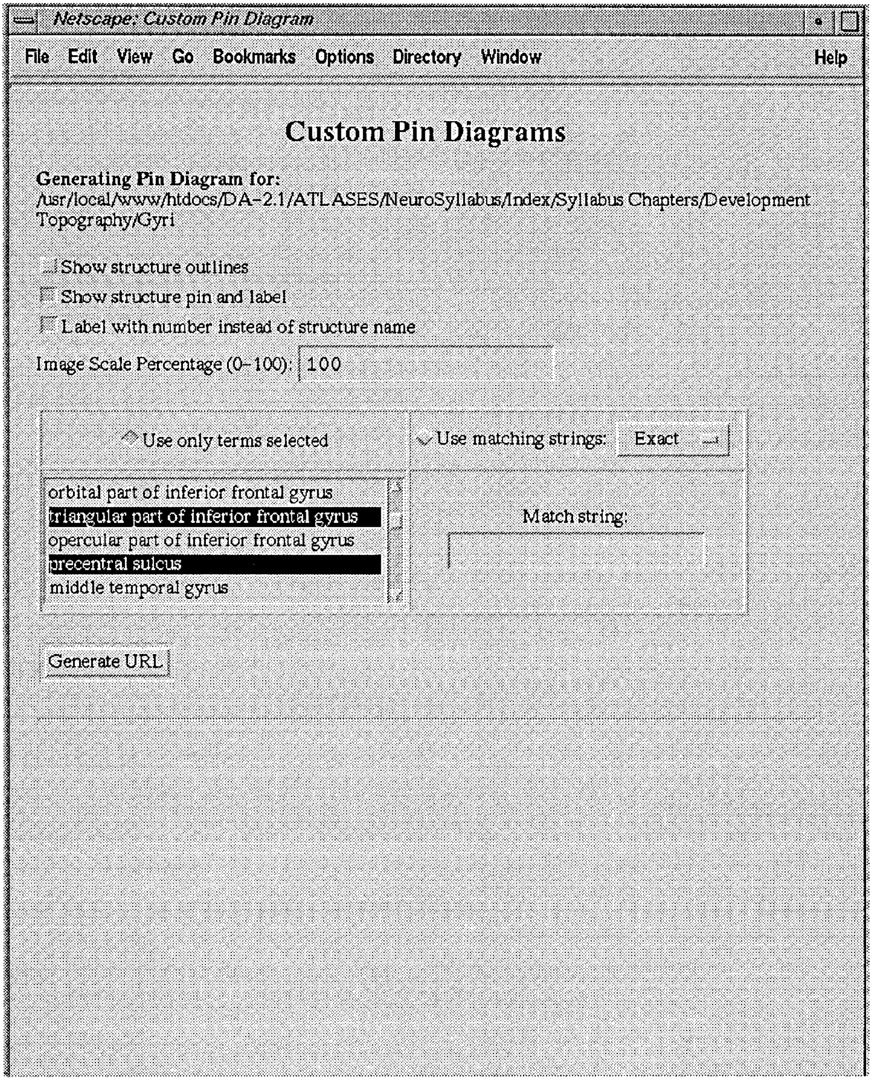

FIG. 12. Generation of URL to create a custom Pin diagram for inclusion in a Neuroanatomy course syllabus. The instructor selects which structures are to be labelled, and whether to use numbers or structure names, depending on whether the image is intended for a tutorial or a test.

from the one used to store the knowledge base in Fig. 3A). Each row in this table lists a content atlas, term and frame triple, along with the URL where the frame may be reached. For a given term, the DB program formats an HTML page that lists all frames that contain the term, in all content atlases known to the search engine (Fig. 15). The user may then click on

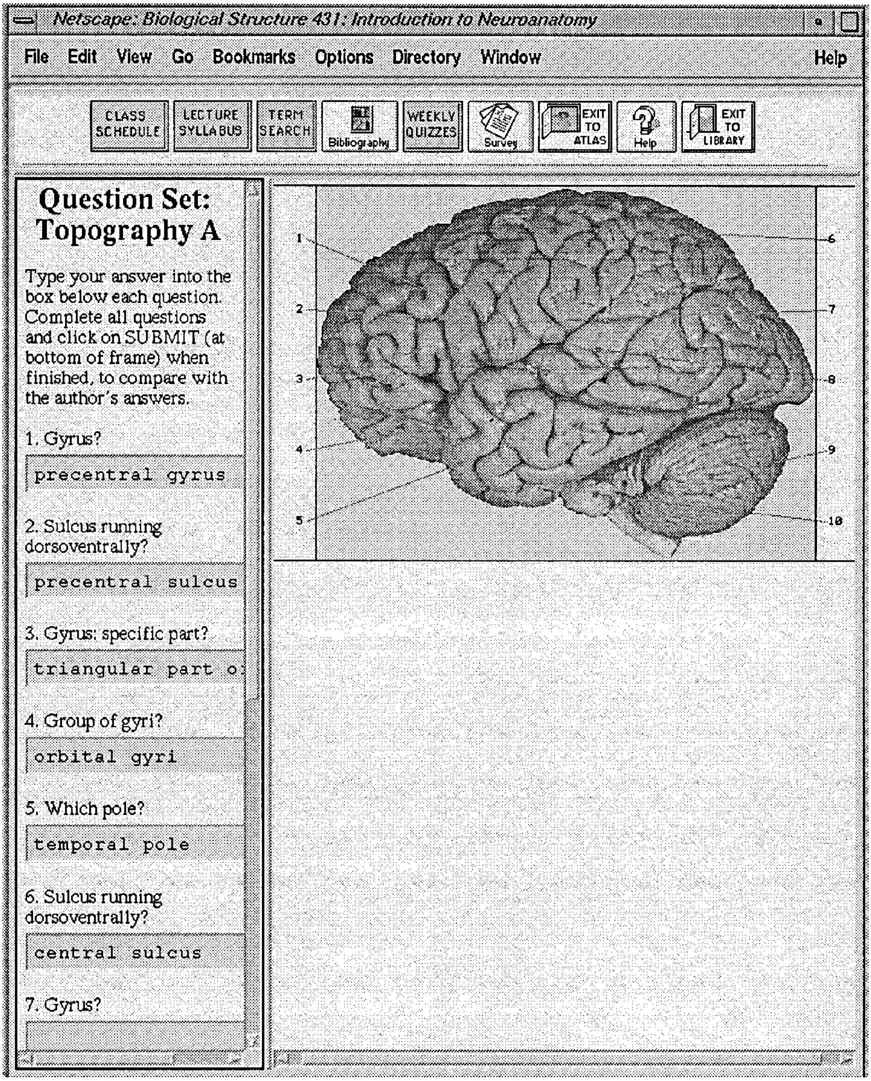

FIG. 13. Weekly quiz section in a Neuroanatomy tutorial, which uses the custom Pin URL generated as shown in Fig. 12, as the basis for a quiz shown on the left. The tutorial is written in a shell designed by Joan Robertson of the UW IAIMS program; the image is a photograph of a cadaver brain.

one of these frames to retrieve it. Search options permit the structure to be outlined and labelled on the retrieved frame. The list of atlas-term-frame triples is generated off-line by an "Atlas Crawler" program similar to the WebCrawler (*38*) which queries, via a build-in "back-door," all atlases on the Web that are known to it.

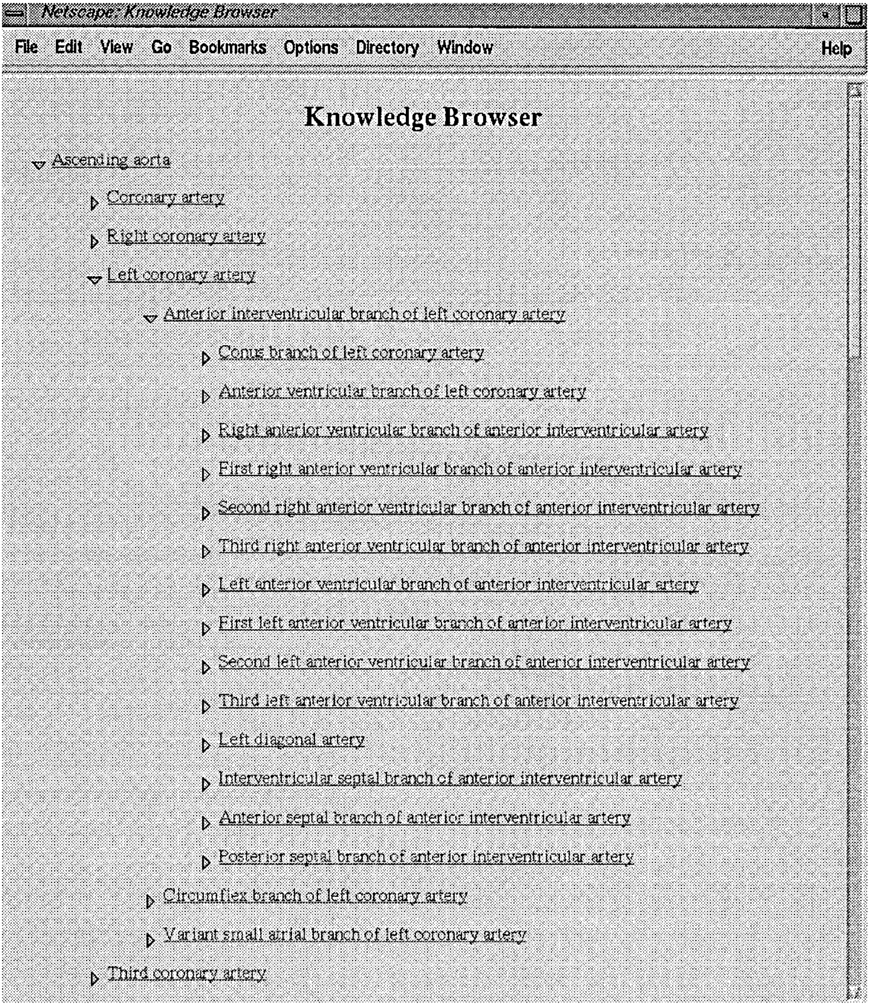

FIG. 14. Web knowledge browser, which displays the same knowledge base that is created by the Knowledge Manager shown in Fig. 4. Once a term is found a search can be initiated for all those frames that contain the term.

3. EVALUATION

Various versions of the Digital Anatomist interactive atlas have been in use since 1992 for gross anatomy and neuroanatomy education. Evaluations have addressed the following questions:

1. How useful and general is the software framework for entering and delivering image-based anatomical content?
2. How is the atlas used?
3. How do users like the atlas and what would they like changed?

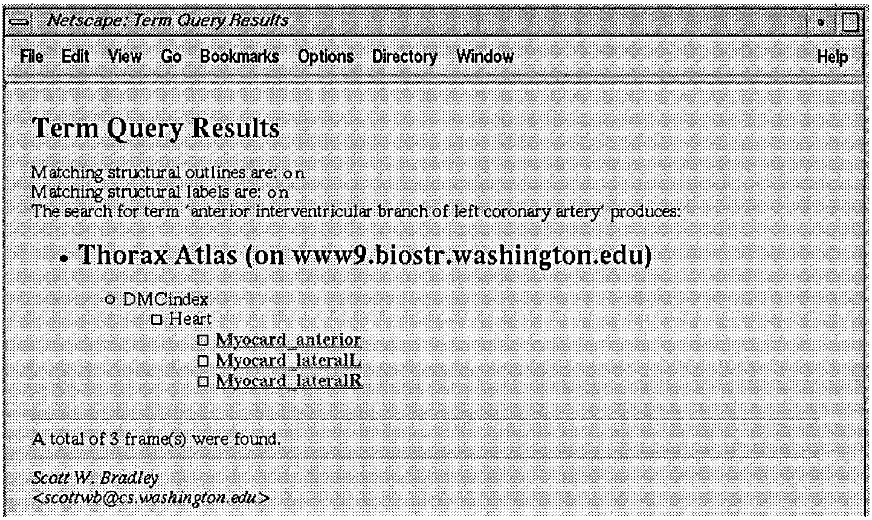

FIG. 15. Search mode for Web atlas client, showing results of a search initiated from the knowledge base browser (Fig. 14). All annotated images have been searched, in all atlases known to the search engine, that contain the search term, "anterior interventricular branch of left coronary artery." The user can click an image name to browse that image, and to request that the structures be outlined or labelled. Note that one of the images found, "*Myocard_anterior*," is the same as that shown in Figs. 2B and 11. The terms database is built by a program that periodically sends a special CGI request to each atlas Web server. This request causes all terms to be listed in a file that is returned to the search program, where it is added to a searchable database.

3.1. Utility of the Framework

The Digital Anatomist framework is both *distributed* and *independent of content*, features that are shared by no other anatomy information systems that we are aware of. The utility of these two features is evidenced by (1) the number of different content atlases that have been created by different authors, and (2) our ability to deliver the same information via videodisc and CD-ROM, as well as via three separate client programs.

To date, the software framework has been used by our group to create atlases of the brain (*22*, *39*) (Figs. 8, 9, 12, 13), thoracic viscera (*20*, *40*) (Figs. 2, 10, 11) and knee (*21*); others are under development. Much of the same material is available on videodisc (*39*, *40*), CD-ROM (*20–22*), and on the Internet in both Mac (*25*) and Web (*36*) versions.

The success and generality of the software framework has prompted the Integrated Advanced Information Management System (IAIMS) program at the University of Washington to propose the Digital Anatomist Interactive Atlas client as a general tool that will be provided to faculty throughout the Health Sciences Center and the University of Washington, for the development of image-based atlases in various clinical and basic science disciplines. To that end we

TABLE 1
REMOTE UTILIZATION OF THE MAC AND WEB ATLAS CLIENTS FROM JUNE 1995 THROUGH JANUARY 1997 ($1\frac{1}{2}$ YEARS)

	Mac client			Web client		
Region	Sites	Requests	Percent	Sites	Requests	Percent
UW	65	89864	85	412	131807	5
US	177	14886	14	18710	1657664	66
World	26	284	<1	9104	567483	23
Unknown	7	173	<1	5033	137089	6
Total	275	105207	100	33259	2494043	100

Note. Sites are the number of unique client sites, based on Internet address. Requests are the number of requests to the server. These numbers for the Mac and Web clients are not strictly comparable since each mousclick is processed by the Web server, whereas only images are requested by the Mac client.

have developed a software package, combining Frame Builder and the DA-CGI programs, which is installable on any Web server that can run CGI scripts (see section 5). Since Frame Builder can be used to annotate any kind of 2-D image, we expect that it will soon be used by many authors to develop image-based atlases which can then be linked over the Web. In fact nothing in the atlas software requires the content to be medically related, so we anticipate that the atlas client will also be of use in areas other than medicine.

3.2. Usage

We argue that the most meaningful measure of utility is whether a system is actually used. Other measures, such as user satisfaction surveys, are useful for finding out how to improve the system so that it is used, but the critical question we have asked and will continue to ask is whether the atlas is actually used.

The use of the Mac client (versions 1, 2 and 3) was limited to the University of Washington until 1994, at which time the atlas client was made available to other institutions over the network. The Web client (version 5) was introduced in June, 1995. We have not done formal surveys of course usage outside the University of Washington. However, indirect measures indicate that both versions 3 and 5 of the online atlas client are widely used.

For example, Table 1 shows the number and distribution of remote sites that used the Mac or the Web client for a $1\frac{1}{2}$ year period. Sites specified as unknown in the Web logs were assigned to geographic regions by calling the Unix program *whois*, which resolved IP addresses to the domain name servers, from which the regions could be inferred. For the Mac client, this procedure reduced the number of unresolved sites from 275 to 7; for the Web client, the unresolved sites were reduced from 17230 to 5033.

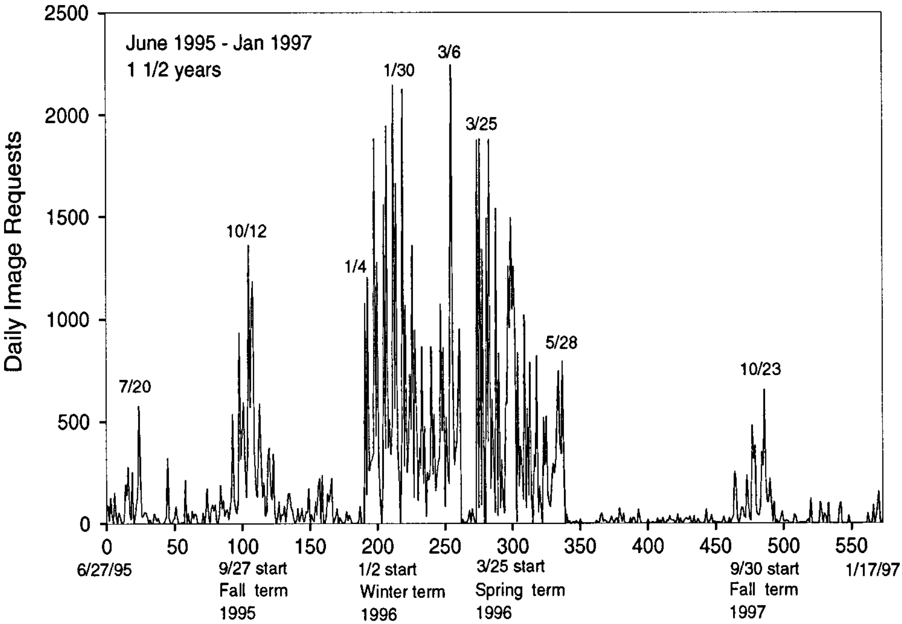

FIG. 16. Daily Mac server requests, for approximately $1\frac{1}{2}$ years from June 27, 1995 through January 17, 1997. Peak dates are indicated. Term start dates are also shown for University of Washington. Fall was thoracic anatomy, Winter and Spring were neuroanatomy. Mac usage is declining, but it is not yet zero.

The average number of daily server requests via the Mac client was 184, which reflects image retrieval alone, since interaction becomes local once the image is retrieved. Most of the Mac client requests were from the University of Washington (85%), reflecting the fact that local anatomy courses are taught by the authors of the content atlases, who are familiar with the Mac atlas client and find it faster for interactive use than the Web client. Of the 177 US sites outside the UW, 64 were from the Medical University of South Carolina, and 31 were from Washington University in St. Louis, both of whom have indicated to us that they are using the Mac atlas in their curricula.

The average number of daily server requests via the Web client was 4383. In this case, the number of requests includes image retrieval as well as interaction with the images, since all processing is done by the server. Nevertheless, even if there are an average of 10 Web requests for each image retrieved, overall usage of the Web client is higher than that of the Mac client. Most of the Web usage occurred outside the University of Washington (95% including unknown sites) with 66% in the US, and a significant amount (23%) outside the US.

Figure 16 shows the usage patterns of the Mac atlas client for the $1\frac{1}{2}$ year time period. Because of the similarity of the timing of anatomy courses and vacation periods at different institutions during the academic calender, it is informative

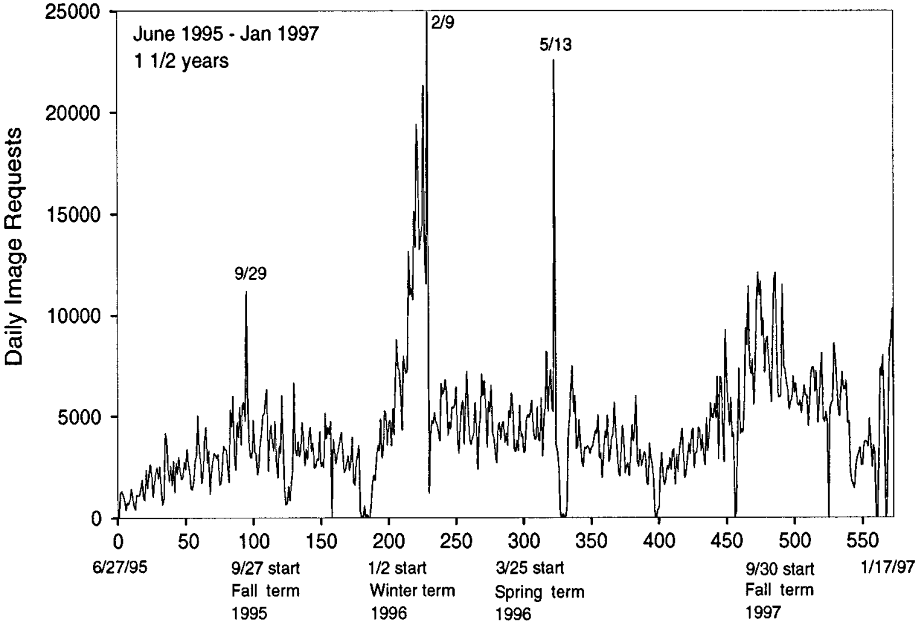

FIG. 17. Daily Web requests, from the time of introduction of the Web atlas on June 27, 1995, to January 17, 1997. Peak dates are indicated. Web usage is gradually increasing.

to correlate the usage patterns with anatomy courses and examinations at the University of Washington. For example, at the University of Washington, thoracic anatomy is taught during fall term, and neuroanatomy both during the winter and spring. Relevant peaks on the graph are labelled with the corresponding dates, and the start of the term is shown below the graph. A good correlation is revealed between usage patterns and the timing of the relevant courses and examinations. For example, usage dropped to near zero after the neuroanatomy final examination (March 12), stayed near zero during spring break (March 12–25), then rose again for spring quarter, falling again to near zero after the final exam (May 28). Overall, usage of the Mac is declining as a gradual transition is made to the Web, but it has not yet reached zero. An important advantage of the distributed approach is the feasibility of these kinds of gradual transitions to new technologies.

Figure 17 shows usage patterns of the Web atlas client for the same $1\frac{1}{2}$ year period. Although we do not have the dates for exams at other medical schools the fact that the Web pattern appears similar to the Mac pattern suggests that the Web atlas was predominantly used for courses rather than random browsing. Comments in the guest log confirm that many schools were using the Web atlas client for courses.

The Web usage pattern shows a gradual increase in average requests per day

TABLE 2

Responses in Web On-Line Guest Book from June 1995 through January 1997, Categorized According to the Reason for Accessing the Atlas, Based on a Selection List

Reason	Responses	With comments
Just for fun	525	135
Interested in server software for commercial use	10	7
Interested in server software for educational use	444	183
Using atlases for course	310	118
Interested in images (not software)	310	112
Other	168	104
Totals	1767	659

Note. Responses with added free-text comments are shown in the last column.

(4073 in June, 1996 versus 4383 in January 1997) with a large increase in the number of sites that have accessed the atlas (13,657 as of June 1996 versus 33,259 as of January 1997). The increase in sites is most likely a reflection of the explosive growth of the Internet within the past year; the fact that the average requests have not increased as much suggests that many of the sites were not "serious" users. However, an AltaVista *link*: search shows that over 500 sites have linked to the Web atlases, implying that many people find the atlases to be of enough interest to create a direct link to them.

3.3. User Satisfaction

Since its introduction in June 1995 the Web atlases have received six unsolicited awards (that we know of) by outside rating organizations: a ranking in the top 5 percent of all Web sites by Point Survey, an Exceptional Resource Award by the Multimedia Medical Reference Library, a designation as a "Cool Nursing Site of the Week," an "All Campus Best Site," a "Stanford MedWorld Best Medical Site," and an "Eisenhower National Clearinghouse Digital Dozen." In addition, we have provided an on-line "guest book" as well as an email address, either of which can be used for comments. Table 2 is a summary of the responses recorded in the on-line guest book as of January, 1997, categorized by "Reason for using," as selected from a menu. A total of 1767 responses were recorded in the guest book, of which 659 provided free-text comments. We have not done formal surveys to find out how many of the 310 respondents who said they were "Using Atlases for Course" were actually using them in a formal course, but the usage patterns shown in Fig. 17 suggests that many of them were.

Almost all the comments were positive. The most common suggestions were for more content. One comment suggested that structures depicted on the images were too detailed for K–12 education, a few comments requested alternate means for navigation, and a few requested linkages with functional and clinical information about the structures.

Example comments include: (from Tulane Medical School) "This is fantastic! . . . I hope it's all right if I share this with the other 150 students in our Neuroscience class," (from South Carolina State University) ". . . Your images are the best on the net . . .", (from an M.D. in Montreal) ". . . Wonderful site. Congratulations," (from University of Tennessee) "Wow. Wow. Wow. Why couldn't I have found you guys *before* I took gross anatomy and neuro?", and (from a sophomore in high school) "That is realy cool. I am currently a sophmore in high school, and this program, has helped me to study for my final exam that is comming up in the next few days. Thanks a lot!" (sic).

4. DISCUSSION

This paper had described an evolving information system in anatomy, and a Web-based client program that accesses the information to deliver on-line atlases of anatomy. Although the client program cannot yet substitute for an expert anatomist, the distributed framework on which the client program relies, and of which it is a component, is well-suited for evolving into an information system that can provide anatomical knowledge appropriate for solving diverse problems at the points of need. The utility of the Digital Anatomist framework is evidenced by (1) its adaptability and scalability to advancing technology such as Java and the Web; (2) its flexibility and suitability for handling different types of image data; (3) its capability to integrate spatial and symbolic information from on-line resources; and (4) the reusability of its information resources and other framework components for various clinical, as well as educational, applications.

Evaluation results show that the atlas client programs are widely used, both in local courses and throughout the United States and the world. This usage is due not only to the flexibility of the framework and the interactivity of the clients, but also to the large amount of high quality 2-D and 3-D anatomical content that they make available on-line. To our knowledge, there is no other on-line site that has sufficient content to support courses of study in anatomy.

The Digital Anatomist anatomy information system overcomes many of the limitations of hardcopy atlases, and is distinguished from other currently available computer-based anatomy atlases by the following characteristics: (1) it relies on an expandable knowledge base and database of structural information that can encompass domains ranging from macroscopic anatomy of living and decreased subjects to molecular structure: (2) these knowledge sources can be reused in multiple applications and can be integrated with other on-line information resources; (3) the software tools are independent of content, and therefore can be used for any kind of image-based knowledge domain; and (4) the tools and information resources are distributed over the Internet and can therefore be made available to the widest possible numbers of users. To our knowledge no other anatomical information system shares these characteristics.

Our plans for enhancing the information system will take advantage of evolving technology to deliver further developments in knowledge-based methods for representing anatomical information.

For example, now that Java is widely available we will update the Web client to be Java-based. This enhancement will make image interaction faster since processing of mouse clicks will occur at the client site. The current client will be retained, however, for those browsers that do not yet have robust Java interpreters. This update will be similar to the update we implemented from the Mac client (version 3) to the Web client (version 5), and will not require any fundamental changes to the content or the other software components.

More intelligent expert delivery of anatomical information will be made possible by taking increasing advantage of the evolving symbolic knowledge base. For example, semantic links will be used for (1) dynamically creating content atlases targeted for specific user groups (e.g., show only the major brain structures to K–12 users); (2) generating more intelligent quizzes (point to the nerve that supplies a particular structure); or (3) provide more sophisticated searching (retrieve all images that show any branch of the arch of the aorta).

Direct manipulation of 3-D scenes will be made possible by upgrading the SKANDHA program into a server, so that it can dynamically generate a 3-D scene in response to a query, then send the scene itself as a VRML description to a client-based 3-D viewer. Such a capability will replicate many of the features of Voxelman (*24*), but will also allow access to dynamic 3-D images by anyone with a Web and VRML browser.

Although the current atlas client programs do not yet provide all these capabilities, the general software framework on which they are based enables us to move closer to an on-line expert system in anatomy, including 3-D visual access to other on-line biomedical information. Such a visually based information system has the potential for fundamentally influencing both training programs for healthcare providers and the management of biomedical information.

5. AVAILABILITY

The Mac client (version 3) is free for downloading from ftp://ftp.biostr.washington.edu/pub/sig/atlas/mac/DA-demo.sea.bin. The Web DA-CGI package and Framebuilder (which allows authors to create their own content atlases on their own Web servers) is free for noncommercial use, and can be found at ftp://ftp.biostr.washington.edu/pub/sig/atlas/web/DA-5.2_1.tar.Z. The current Web content atlases can be reached at the URL http://www9.biostr.washington.edu/da.html. We encourage authors who create their own content atlases to provide us (digital_anatomist@biostr.washington.edu) with the URL so we may include them in this page and so that our search engine can access them. CD-ROMS of the brain, knee and thoracic viscera are available from the University of Washington Health Sciences Center for Educational Resources http://cer.hs.washington.edu/hscer/products/anatlst.htm. The CD-ROMS contain the same 2-D images as those accessible via the online atlas clients, but many more 3-D animations.

ACKNOWLEDGMENTS

This work was funded by National Library of Medicine Grants LM04925 and LM05982, National Library of Medicine Contract LM13506, National Cancer Institute Grant CA59070, Human Brain Project Grant DC/LM02310, cofunded by the National Library of Medicine and the National Institute on Deafness and Other Communication Disorders, and a Teaching Innovations grant from the University of Washington School of Medicine. We thank the following individuals, all from the University of Washington, for their help in various aspects of this work: David Conley (now at Washington State University) for the Thorax atlas; Peter Ratiu (now at Engineering Animations Associates) for the Knee atlas; John Bolles (Health Sciences Center for Educational Resources) for distribution of the atlas; Sherrilynne Fuller, Jim Barrett and Cliff Solomon (University of Washington IAIMS program) for assisting with the support of the Web atlas client; Kraig Eno (now in Kiev, Ukraine) for the Mac atlas client and servers; Roger Williams (University of Washington Department of Biological Structure) for the Windows atlas client; Jeff Prothero (now at Activision Corp.) for the 3-D graphics software; and Kevin Hinshaw, Rex Jakobovits and Bharath Modayur, for comments on various drafts of the manuscript.

REFERENCES

1. Brinkley, J. F. Structural informatics and its applications in medicine and biology. *Acad. Med.* **66,** 589 (1991).
2. Rosse, C. The potential of computerized representations of anatomy in the training of health care providers. *Acad. Med.* **70**(6), 499 (1995).
3. Ackerman, M. J. The visible human project. *J. Biocommun.* **18**(2), 14 (1991).
4. Brinkley, J. F., and Rosse, C. The digital anatomist distributed framework and its applications to knowledge based medical imaging. *JAMIA* **4**(3), 165 (1997).
5. Hohne, K. H., Pflesser, B. Riemer, M., Schiemann, Th., Schubert, R., and Tiede, U. A new representation of knowledge concerning human anatomy and function. *Nature Med.* **1**(6), 506 (1995).
6. Lanier, L., Rathe, R., and Seymour, J. An interactive tutorial on normal radiology. CD ROM, Gold Standard Multimedia, 1994.
7. Lindberg, D. A. B., Humphreys, B. L., and McCray, A. T. The unified medical language system. *Meth. Inform. Med.* **32**(4), (1993).
8. Cote, R. A., Rothwell, D. J., Palotay, L. J., Beckett, R. S., and Brochu, L. (Eds). "The Systematized Nomenclature of Medicine: SNOMED International. College of American Pathologists, Northfield, IL, 1993.
9. Schultz, E. B., Price, C., and Brown, P. J. B. Symbolic anatomic knowledge representation in the read codes version 3: Structure and application. *J. Am. Med. Inform. Assoc.* **4,** 38 (1997).
10. Rector, A. L., Glowinski, A. J., Nowlan, W. A., and Rossi-Mori, A. Medical-concept models and medical records: an approach based on galen and pen and pad. *J. Am. Med. Inform. Assoc.* **2,** (1995).
11. Chapman, C. M., Miller, J. G., Bush, L. C., Bruenger, J. A., Wysor, W. J., Meininger, E. T., Wolf, F. M., Fischer, T. V., Beaudoin, A. R., and Burket, W. E. Atlas-plus: multimedia instruction in embryology, gross anatomy, and histology. *In* "Proceedings, Symposium on Computer Applications in Medical Care," pp. 712–716. AMIA, 1992.
12. Freedman, S. J., and Chase, R. A. Electriccadaver: A dynamic book of human structure and function. *In* "Proceedings, Symposium on Computer Applications in Medical Care," pp. 1021–1023. AMIA, 1989.
13. ADAM Software. "Adam Scholar Series." CD-ROM, 1995.
14. The Parthenon Publishing Group. "The Anatomy Project." CD-ROM, 1996.
15. B. Hillen. "Temporal Bone and Posterior Cranial Fossa." CD-ROM, 1995.
16. Gold Standard Multimedia, Inc. "Human Anatomy." CD-ROM, 1996.
17. Mathers, L., and Chase, R. A. "C.l.a.s.s.-clinical anatomy study system." Mosby, 1994. CD-ROM.
18. Villarin, A., Schmidt, R., Newman, L., and Lewis, P., The gross project: Computer assessment

of gross anatomy through digitized cadaver photographs. *In* "Symposium on Computer Applications in Medical Care," pp. 1015–1017. AMIA, 1990.

19. Brinkley, J. F. Educational systems: synopsis. *In* (J. H. Van Bemmel and A. T. McCray, Eds.) "Yearbook of Medical Informatics," pp. 465–467. American Medical Informatics Association, The Netherlands, 1994.

20. Conley, D., and Rosse, C. "Interactive Atlas of Thoracic Viscera." *Available from* University of Washington Health Sciences Center for Educational Resources, 1994. CD-ROM.

21. Ratiu, P., and Rosse, C. "Animations of Knee Anatomy." *Available from* University of Washington Health Sciences Center for Educational Resources, 1994. CD-ROM.

22. Sundsten, J. W. "Interactive Atlas of the Brain." *Available from* University of Washington Health Sciences Center for Educational Resources, 1994. CD-ROM.

23. Berners-Lee, T., Cailliau, R., Groff, J. F., and Pollerman, B. World-wide web: the information universe. *Electron. Network.* **2,** 52 (1992).

24. Brinkley, J. F., Prothero, J. S., Prothero, J. W., and Rosse, C. A framework for the design of knowledge-based systems in structural biology. *In* "Proceedings, 15th Annual Symposium on Computer Applications in Medical Care," pp. 61–65, Baltimore, 1989.

25. Brinkley, J. F., Eno, K., and Sundsten, J. W. Knowledge-based client-server approach to structural information retrieval: the digital anatomist browser. *Comput. Meth. Programs Biomed.* **40,** 131 (1993).

26. Brinkley, J. F., Myers, L. M., Protero, J. S., Heil, G. H., Tsuruda, J. S., Maravilla, K. R., Ojemann, G. A., and Rosse, C. A structural information framework for brain mapping. *In* (S. H. Koslow and M. F. Huerta, Eds.), "Neuroinformatics: An Overview of the Human Brain Project," Chap. 9, pp. 309–334. Erlbaum, Mahwah, NJ, 1997.

27. Dailey, D. J., Eno, K., and Brinkley, J. F. Performance evaluation of a distance learning program. *In* "Proceedings, 18th Annual Symposium on Computer Applications in Medical Care," pp. 76–80, Washington, DC, 1994.

28. Betz, D. "Xlisp: An Object-Oriented Lisp." Unpublished reference manual for version 2.1, *available from* ftp.biostr.washington.edu and other Xlisp ftp sites, 1989.

29. Brinkley, J. F., and Prothero, J. S. Slisp: A flexible software toolkit for hybrid, embedded and distributed applications. *Software Pract. and Exper.* **27**(1), 33 (1997).

30. Rosse, C., Ben Said, M., Eno, K. R., and Brinkley, J. F. Enhancements of anatomical information in umls knowledge sources. *In* "Proceedings, 19th Annual Symposium on Computer Applications in Medical Care," pp. 873–877. New Orleans, Oct. 30–Nov. 1 1995. American Medical Informatics Association.

31. Rosse, C., Mejino, J. L., Jakobovits, R. M., Modayur, B. R., and Brinkley, J. F. Motivation and organizational principles for the digital anatomist symbolic knowledge base: an approach toward standards in anatomical knowledge representation. *JAMIA*. In Press (1998).

32. Conley, D. M., Kastella, K. G., Sundsten, J. W., Rausching, W., and Rosse, C. Computer-generated three-dimensional reconstruction of the mediastinum correlated with sectional and radiological anatomy. *Clin. Anat.* **5,** 1 1992.

33. Rauschning, W., and Gleen, W. W., "The knee: Mr-imaging, arthroscopy and anatomy correlations." Videodisc, 1986.

34. Brinkley, J. F. A flexible, generic model for anatomic shape: application to interactive two-dimensional medical image segmentation and matching. *Comput. Biomed. Res.* **26,** 121 1993.

35. Hinshaw, K. P., Altman, R. B., and Brinkley, J. F. Shape-based models for interactive segmentation of medical images. *In* "SPIE Medical Imaging 1995: Image Processing," pp. 771–780, San Diego, Feb 26–March 2, 1995.

36. Bradley, S. W., Rosse, C., and Brinkley, J. F. Web-based access to an online atlas of anatomy: the digital anatomist common gateway interface. *In* "19th Symposium on Computer Applications in Medical Care," pp. 512–516, New Orleans, Oct 30–Nov 1, 1995.

37. Date, C. J. "A Guide to INGRES." Addison–Wesley, Reading, MA, 1987.

38. Webcrawler, http://webcrawler.com/, 1996.

39. Sundsten, J. W., and Kastella, K. G., "Human Brain Animations." *Available from* University of Washington Health Sciences Center for Educational Resources, 1992. Videodisc.

40. Conley, D., and Rosse, C. "Animations of Thoracic Viscera." *Available from* University of Washington Health Sciences Center for Educational Resources, 1994. Videodisc.