# Design of an Anatomy Information System

**James F. Brinkley, Benjamin A. Wong, Kevin P. Hinshaw, and Cornelius Rosse**
*University of Washington, Seattle*

**The Digital Anatomist Project aims to "put anatomy on a computer." Both the content and the structure-based information system should have wide applicability outside the initial implementation.**

Biology and medicine rely fundamentally on anatomy. Not only do you need anatomical knowledge to understand normal and abnormal function, anatomy also provides a framework for organizing other kinds of biomedical data. That's why medical and other health sciences students take anatomy as one of their first courses.

The Digital Anatomist Project undertaken by members of the University of Washington Structural Informatics Group aims to "put anatomy on a computer" in such a way that anatomical information becomes as fundamental to biomedical information management as the study of anatomy is to medical students. To do this we need to develop methods for representing anatomical information, accessing it, and reusing it in multiple applications ranging from education to clinical practice.

This development process engenders many of the core research areas in biological structural informatics, which we have defined as a subfield of medical informatics dealing with information about the physical organization of the body.[1] By its nature, structural information proves highly amenable to representation and visualization by computer graphics methods. In fact, computer graphics offers the first real breakthrough in anatomical knowledge representation since publication of the first scholarly anatomical treatise in 1546, in that it provides a means for capturing the 3D dynamic nature of the human body.

In this article we explain the nature of anatomical information and discuss the design of a system to organize and access it. Example applications show the potential for reusing the same information in contexts ranging from education to clinical medicine, as well as the role of graphics in visualizing and interacting with anatomical representations.

## Anatomical information

First we'll define anatomical information and discuss how to best represent it. Then we'll consider disseminating that information.

### Representation

Anatomy as a science studies the structures and spaces that make up the body, and the relationships among them. We can broadly classify representations of these entities and their relationships as either spatial or symbolic.[1] *Quantitative spatial information* derives primarily from images and includes, for example, annotated images, 3D models, 3D scenes, and animations. *Symbolic information*, on the other hand, includes metadata about the images and 3D models (file names, date of digitization, and so on). More importantly, it also includes symbolic knowledge describing the names and synonyms of anatomical structures, their classifications (the aorta is a blood vessel), their definitions, their partitive relationships (the left atrium is part of the heart), and their qualitative spatial relationships (the upper lobe of the left lung is superior to the lower lobe). Both spatial and symbolic information have equal importance in understanding anatomy, and both require visualization and graphics for optimal understanding.

To date, most computer-based educational programs in anatomy have concentrated on image-based spatial information in the form of digitized medical illustrations and clinical images, 3D image volumes, and 3D surface models. All these resources prove useful because of anatomy's highly visual nature. The availability of the Visible Human male and female data sets, which consist of image volumes acquired from frozen cadavers, has prompted many applications of 3D image processing and visualization (http://www.nlm.nih.gov/research/visible/visible_human.html). These applications include 3D slice browsers, segmentation tools, animations of 3D scenes such as those on our own videodiscs and CD-ROMs,[2] 3D fly-throughs, Virtual Reality Modeling Language (VRML) models, immersive displays, and surgical simulations. Because of the highly irregular, complex, and variable nature of biological structures, such appli-

 0272-1716/99/$10.00 

cations present particularly interesting and difficult challenges for computer graphics.

Symbolic representation and visualization in anatomy remain less developed, with most applications limited to textual descriptions or static labels of structures on images. We argue that symbolic information and, in particular, symbolic knowledge must parallel image-based spatial information because it provides the "intelligence" that gives meaning to the images.

Recently, we proposed a *foundational model of anatomy* as a basis for representing symbolic anatomical knowledge.[3] The foundational model provides a logic-based description of the body's physical organization. Specifically, we define the foundational model as

Fm = {Ao, ASA, ATA, Mk}, where

- Ao, an anatomical ontology, defines class and subclass relationships (left coronary artery *is-a* artery, which *is-a* blood vessel)
- ASA, an anatomical structural abstraction, defines the partitive and qualitative spatial relationships of structures (left lung *has-parts* upper lobe and lower lobe; upper lobe *is-superior-to* the lower lobe)
- ATA, an anatomical transformational abstraction, defines how structures develop into other structures, particularly during embryogenesis (proximal bulbus cordis *gives-rise-to* inflow part of the right ventricle)
- Mk stands for metaknowledge—principles and sets of rules, according to which we represent concepts and relationships

We believe that combining a foundational model of anatomy with image-based spatial information offers the best opportunity for maximizing the applicability of an anatomy resource. To date, very few efforts have taken this direction besides our own,[4,5] and none have taken as broad and comprehensive an approach as we attempt. Our ability to take this approach is a direct result of the long-term collaboration between anatomists and computer scientists in the University of Washington Structural Informatics Group.

### *Dissemination*

Once we have represented anatomical information in the computer, we must organize the information and make it accessible to users and to other computer programs. To this day, students obtain most anatomical information from lectures, cadaver dissection, textbooks, and hard-copy atlases. In the past five years or so CD-ROM has become the dominant format for computer-based anatomical information,[4,6] although the majority of people still see CDs as an optional supplement to traditional forms of anatomy education.

Ease of marketing and quick access of high-bandwidth images and models have made CDs the most popular computer-based access method. However, the World Wide Web offers great promise for delivery of anatomical information once we have solutions to some of the economic and bandwidth problems. These potential benefits include widespread access, essentially unlimited content, flexibility in presentation format, and an organizational framework for other Web-based biomedical information.

A simple Web search for "anatomy" brings up more than half a million pages, although most aren't relevant to anatomy education. We're currently classifying and surveying the most popular and useful of these sites (as defined by the number of sites that link to them, the availability of significant content, and their relevance to anatomy education). Preliminary results from our survey show that, in spite of the large number of pages, very few sites have significant content, and many simply advertise CD-ROMs or other offline offerings such as courses.

None of the sites (including our own) live up to what we see as the Web's true potential for delivering anatomical information. The Web makes it possible to combine a foundational model of anatomy with high-quality images and models, managed by an intelligent system that not only delivers customized anatomical information to diverse end users, but also serves as an organizing framework for other biomedical data and knowledge.

## The Digital Anatomist Information System

We have worked for the past 10 years to develop a distributed, Internet-based framework for organizing and delivering anatomical information.[7] This system comprises a series of spatial and symbolic anatomical information resources made available to authoring and end-user client programs through one or more structural information servers. Key features include

1. detailed anatomical visualizations,
2. symbolic knowledge and data,
3. distributed and Internet-accessible architecture, and
4. flexibility to adapt to new technologies, content, and methods of presentation.

This last feature has let us maintain the same basic architecture without significantly revising existing content, to deliver information through new methods with changes in technology (videodisc, CD-ROM, custom clients and servers, Web-based servers, Java, VRML), and to create new applications that access and present existing resources in new ways.

### *Architectural framework*

Figure 1 (next page) shows the current instantiation of this design. The symbolic and image-based spatial anatomical information resources appear in the bottom (fourth) row. The structural information servers shown in the third row deliver these resources. Currently, most of the authoring and end-user client programs, shown in the top row, directly access either these servers or the resources. However, we're developing the Digital Anatomist and Digital Neuroscientist applications, intelligent agents that know about the available resources, the expected uses, and the user types. Using this knowledge, the agents will generate customized interfaces to fit each particular situation. Eventually, the Digital Anatomist will behave somewhat like a human

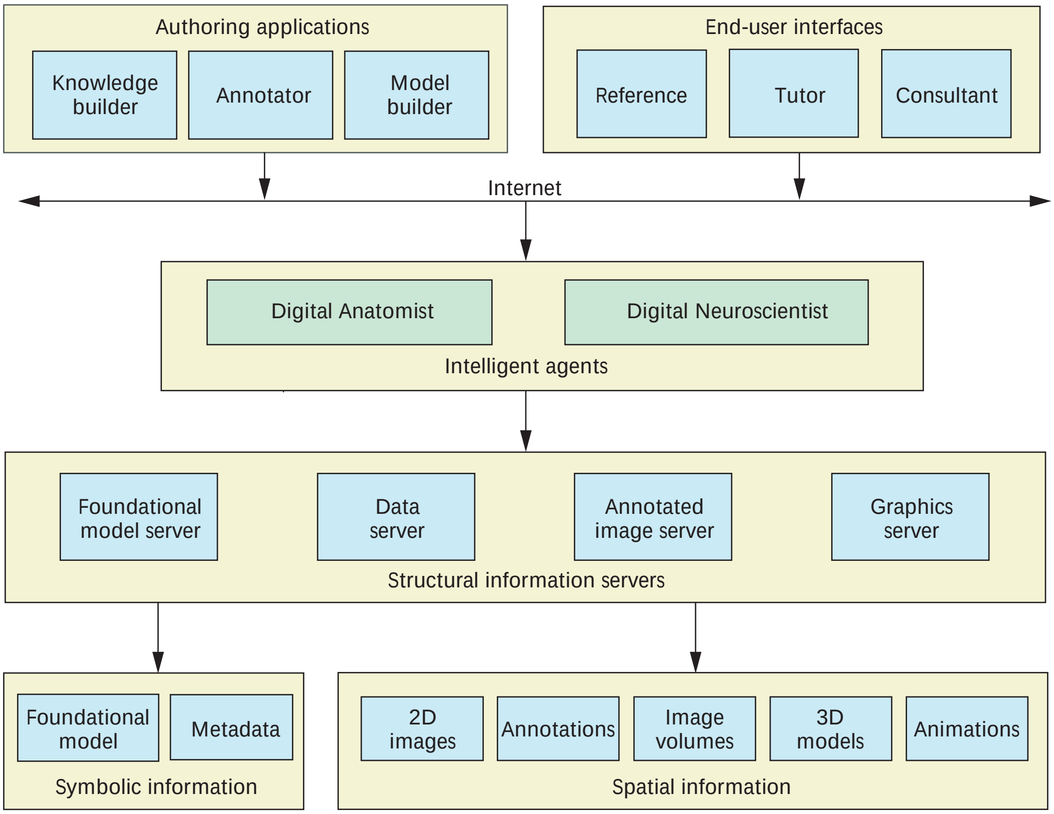


**1** Architectural framework. Authoring and end-user applications access spatial and symbolic resources by means of several structural information servers. Intelligent agents, such as the Digital Anatomist and the Digital Neuroscientist, generate customized interfaces depending on the user and the resources.

anatomist, and the Digital Neuroscientist like a human neuroscientist. We plan to develop other agents to access these and other resources (for example, a Digital Histologist or a Digital Pathologist).

## User interface

A crucial component of any information system is the user interface. Since we require the system to adapt to multiple user types and multiple modes of use, we can't specify a single, static user interface. Instead, we have classified the types of interface along three dimensions and created examples or mockups of each class. These examples then serve as templates for the development of the programs that generate them.

The three dimensions of interface classification are type of user, type of information, and mode of use:

- Type of user includes, for example, medical and other health science students, clinical specialists such as cardiac surgeons, researchers, and K-12 students as well as the lay public.
- Type of information ranges from image-based spatial to symbolic.
- Mode of use includes reference, tutorial, and consultant.

*Reference* mode lets the user freely browse the available information resources, with very little preprogrammed navigation other than indices and tables of contents. This mode resembles that encountered by library users as they browse the bookshelves.

*Tutorial* mode builds on reference mode to provide guided tours of the available resources. Examples would include a teaching module on the heart or a module tailored for cardiac catheterization.

*Consultant* mode acts like an expert anatomist who answers questions in response to specific queries. For example, radiation oncologists specify the optimal configuration of radiation beams to deliver maximal dosage to a tumor while sparing surrounding critical tissues. To do this, they might want to ask an anatomy information system to identify the critical radiation-sensitive structures in the vicinity of a tumor and the regions where the tumor would likely spread.

Figure 2 shows a mockup of an interface that might be generated for a medical student using the system in reference mode. In this case symbolic and image-based spatial information appear to the Web user as two frames synchronized to each other by the Digital Anatomist agent in Figure 1.

The left-hand (symbolic) frame of Figure 2a shows parts of the foundational model for the lower respiratory system, which includes the lungs and tracheobronchial tree. The presentation is modeled on "anatomy templates" developed by Rosse over 30 years of teaching anatomy. These templates include, for example, the class from the anatomy ontology (organ system subdivision) and partitive relationships from the anatomical structural abstraction. Many other attributes would be available, but don't appear in this figure.

The right-hand (spatial) frame of Figure 2a illustrates the lower respiratory system with a dynamically generated scene, created from stored 3D models. The user can manipulate this scene, causing the corresponding symbolic interface to update itself, or expand attributes in the symbolic interface to see more detail. Clicking on a template value brings up the template for that object. For example, Figure 2b shows the template for the left lung, one of the parts of the lower respiratory system shown in Figure 2a. Expanding the Inputs, Outputs, and Has-Parts attributes in Figure 2b reveals the names of the corresponding structures and causes the scene on the right to update accordingly. Similarly, clicking on one of the parts of the left lung, the upper lobe of the left lung, brings up the template and computed scene for the upper lobe in Figure 2c.

If the user were a K-12 student rather than a medical student, the interface might not show the lobes of the left lung, but would instead just show the whole lung in relationship to the heart. If the system were accessed in tutorial mode, the student might have to fill in the values for attributes in the symbolic display or click on the individual structures in the 3D scene. If the system were accessed in consultant mode, it might simply display the results for queries such as “Show me the lobes of the left lung.”

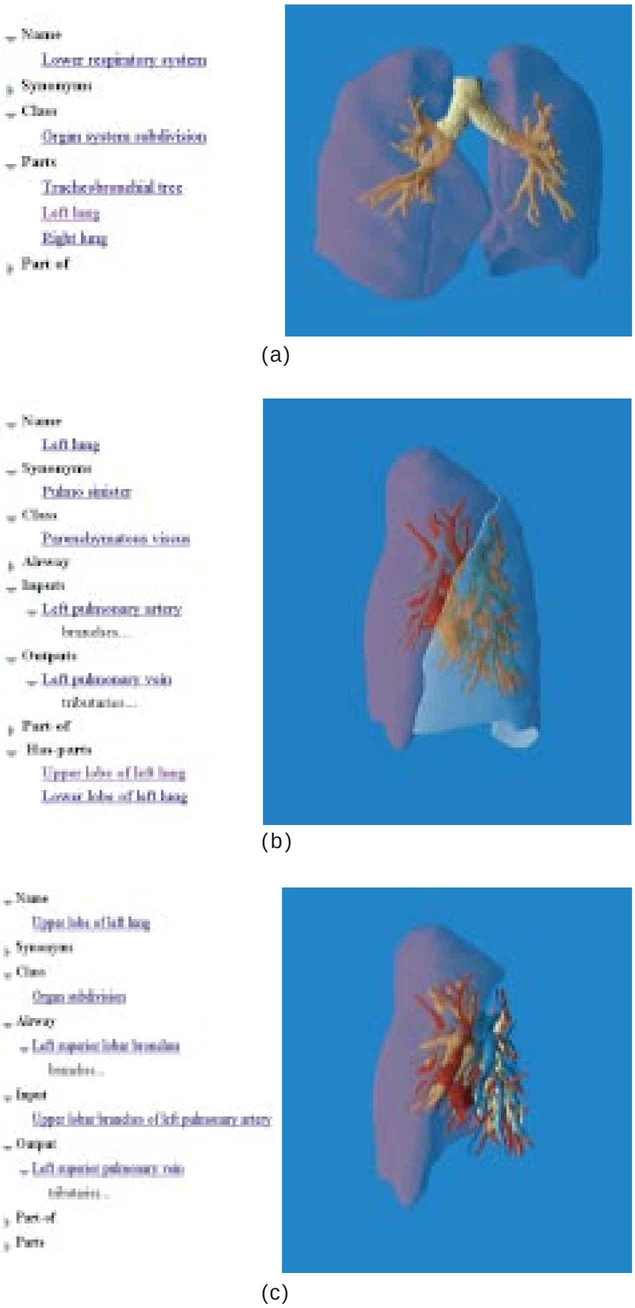


**2** **Mockup of a dynamically generated reference mode interface for specific structures. Symbolic information (on the left) and spatial information (on the right) are synchronized and displayed side-by-side. (a) Lower respiratory system. (b) Left lung. (c) Upper lobe of left lung. Working programs generate the symbolic and spatial interfaces, but the programs aren’t yet synchronized. David Conley created the 3D models using software developed by Jeff Prothero.**

### *Current status and implementation*

All the modules shown in Figure 1 work in one form or another, although many do not yet work together and all are in constant development. For example, individual working programs generate the interfaces shown in each window of Figure 2, but the windows aren’t yet synchronized. The Digital Anatomist and Digital Neuroscientist agents remain the least-developed modules.

Modules and resources are currently implemented in msql, Perl, C, Java, and Lisp. Servers and resources run under Silicon Graphics Irix or Linux. Authoring and end-user interfaces run in X, Java, or forms-based Hypertext Markup Language (HTML). The main constraint is that modules must communicate over the Internet. We implemented most of the graphics and image processing modules in Skandha Lisp, our hybrid Lisp-C development environment.[8] In this environment computationally intensive image and graphics operations are implemented in C as primitive Lisp functions. We can then combine these functions with other Lisp functions to generate diverse applications.

## Example anatomy authoring applications

Long-term, we want to create an intelligent agent that can dynamically generate new application interfaces for diverse users. Short-term, our implementation strategy concentrates on specific applications that drive the development of the overall system. We then integrate these individual applications so that they work together. This strategy has the benefit of generating useful programs as we move toward our long-term goals.

First we describe some of the authoring applications used to create resources for the anatomy information system. In the next section we describe some end-user applications.

Currently, we’ve represented the foundational model (lower left of Figure 1) as a semantic network stored in

**3 Knowledge builder used by anatomist authors to add content to the foundational model. The screen shows veins that are tributaries of the heart's right atrium. Software by Darren Stalder, content by Cornelius Rosse, Jose Mejino, and Augusto Agoncillo.**

a relational database. In its present form the model consists of more than 26,000 concepts, representing all structures visible to 1-mm resolution in the thorax, abdomen, pelvis, and perineum. We arranged the concepts in *is-a*, *part-of*, *branch-of*, and *tributary-of* hierarchies using more than 28,000 semantic links. Through a contract with the National Library of Medicine, the NLM incorporates successive versions of the foundational model in its Unified Medical Language System (UMLS) as part of an effort to develop controlled medical vocabularies.[9] We're developing other body parts under this contract as well.

Users or programs access the model through a foundational model server (third row in Figure 1) that provides a high-level interface to the model. Various client programs access this server. For example, the anatomist authors use a Java-based knowledge-builder applet (see Figure 3 and upper left of Figure 1) to add new terminology and relationships to the model. In addition, a set of common gateway interface (CGI) programs implement a forms-based interface to the foundational model server, like that shown in the left-hand side of Figure 2.

Spatial information resources include 2D images, image annotations, image volumes, 3D models and scenes, and stored animations. All derive from medical images using the model builder and annotator tools shown in the upper left of Figure 1.

The model builder embodies the image processing and surface generation components of the standard graphics pipeline. Other modules then use these surfaces for scene generation and display. The 3D models shown in Figures 2 and 4 through 7 were derived from cadaver-based image slices obtained prior to the availability of the Visible Human. The models in these figures were created manually using a standard process of reconstruction from serial sections. Contours of structures were outlined on each slice, the contours were semiautomatically tiled to generate individual surface meshes, the meshes were combined to create 3D scenes, and the scenes were rendered as static 2D images or as saved animations.[2] Individual meshes were also saved as 3D "primitives," each indexed by a name from the foundational model, for use by the scene builder (described in the next section).

The anatomist authors use the Annotator authoring program to draw regions of interest on the images, label those regions with terms from the foundational model, and link the images together as an atlas depicting a region of the body. The annotations are saved in separate files stored with the images.

## Example end-user applications

In this section we describe three working end-user applications: online interactive atlases, dynamic scene generation, and brain mapping. All three use the Web as the primary end-user interface.

### *The Digital Anatomist interactive atlases*

Anatomist authors used various versions of the authoring tools to create a series of interactive anatomy atlases, each consisting of annotated images and QuickTime animations for different body regions. Atlases currently exist for the brain, the thorax, and the knee. All atlases rely primarily on 3D models created with the model builder, supplemented by 2D images such as photographs and radiographs.

Each atlas is saved independently from the programs that access it. This separation of program from content has allowed us to develop several versions of the interactive atlas client that change with advances in technology while maintaining the same contents. For example, we've created CD-ROMs in native Macintosh and Windows, and more recently in Java. We've also created a Web-accessible interface by means of a series of CGI programs that implement a forms-based interface. You'll find the Web atlases at http://www9.biostr.washington.edu/da.html.

On the Web, the server processes mouse clicks on the client browser to dynamically generate new HTML pages that emulate and extend the local interactions available on the CD-ROMs. The CGI programs, which constitute the annotated image server shown in the middle row of Figure 1, create a reference mode end-user interface.

In all the atlases users navigate through image thumbnails arranged in chapters and topics. Once the program retrieves an annotated image, the user can interact with the image in browse mode or quiz mode. In browse mode, shown in Figure 4a for the Web interface, the user can click on a region to get a structure's name, ask for all structures to be outlined, or ask for a pin-label diagram. The pin-label diagram (Figure 4b)—calculated on the fly—shows the structures' names in the margins, with lines leading to the centers of the regions.

In quiz mode, shown in Figure 4c, the computer asks the user to click on regions corresponding to structures in the annotation files. It keeps a score of the number of

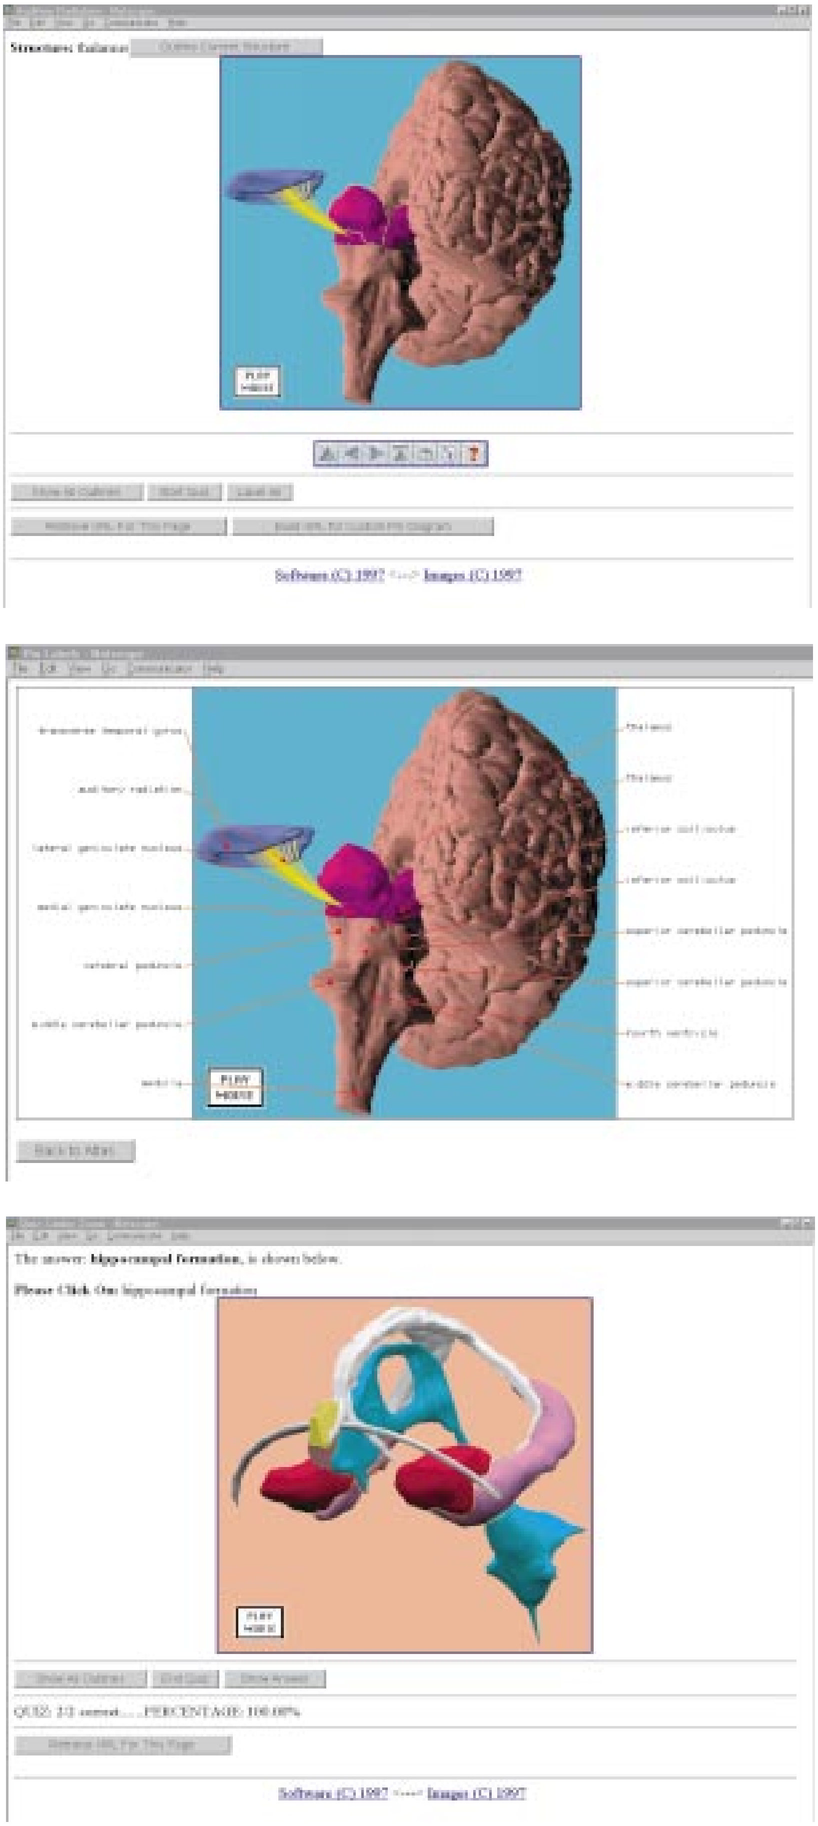


(a)

(b)

(c)

**4 Web-based interactive atlas of the brain. CGI programs at the server process mouse clicks, which consult annotation files associated with the images. (a) Browse mode. Clicking on a structure returns the name. Clicking "Outline Current Structure" generates a new image with the structure outlined. (b) Pin-label mode. Clicking "Label All" in (a) generates a new image with structure names placed in the margins and lines drawn to the regions' centers. (c) Quiz mode, which asks the user to click on structures in the annotation files. Clicking "Show Answer" outlines the structure. Quiz mode maintains a score and is automatically available for all annotated images. CGI programs by Scott W. Bradley, neuroanatomy 3D models by John Sundsten, from a serially sectioned cadaver brain.**

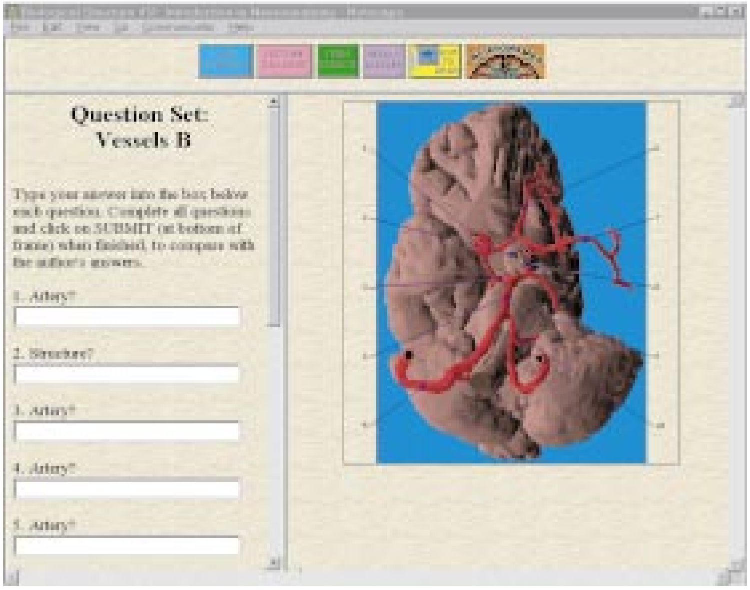


**5** **Accessing the interactive atlas from other Web pages. The atlas image is embedded as the right-hand frame in a larger Web page containing a weekly quiz that's part of a neuroanatomy tutorial. The numbers and lines are dynamically generated. Tutorial by John Sundsten, Joan Robertson, and Kate Mulligan. 3D models by John Sundsten from serial sections.**

```
> 200 FM Server 0.42 Fri Jan 15 20:44:53 1999.

Writing to "cache/posterior mediastinum-part of.str"
Adding "Thoracic part of esophagus":
        "/usr/local/data1/rausch/em65f/d/esophthoracic"
        (ESOPHAGUS-MATERIAL)
Adding "Right vagus nerve":
        "/usr/local/data1/rausch/em65f/d/vagusright"
        (NERVE-MATERIAL)
Adding "Left vagus nerve":
        "/usr/local/data1/rausch/em65f/d/vagusleft"
        (NERVE-MATERIAL)
Adding "Esophageal plexus":
        "/usr/local/data1/rausch/em65f/d/esophplexus"
        (NERVE-MATERIAL)
Adding "Thoracic duct":
        "/usr/local/data1/rausch/em65f/d/thoracicduct"
        (LYMPHATICS-MATERIAL)
Adding "Descending thoracic aorta":
        "/usr/local/data1/rausch/em65f/d/aortadescthoracic"
        (ARTERY-MATERIAL)
Adding "Hemiazygos vein":
        "/usr/local/data1/rausch/em65f/d/hemiazygosvein"
        (VEIN-MATERIAL)
Adding "Azygos vein":
        "/usr/local/data1/rausch/em65f/d/azygosveinmain"
        (VEIN-MATERIAL)

Finished writing to "cache/posterior mediastinum-part of.str".
```

(a)

(b)

**6** **Scene builder and navigator. (a) Scenes are automatically created for a given structure with a given semantic relationship, in this case the parts of the posterior mediastinum. (b) Four such scenes are combined and rendered by a graphics server, then sent over the Web, where the user can manipulate the combined scene. 3D models by David Conley, software by Ben Wong and Jeff Prothero.**

correct answers. At any time the user can ask that the structure be outlined.

We're converting our stand-alone Java-based CD-ROM client into an applet to provide faster access to the atlases over the Web. However, we'll continue to maintain the forms-based interface because

1. virtually all browsers support forms, whereas Java continues to evolve, and
2. users can access each CGI-generated page directly through a URL.

The anatomists have used this latter capability to generate Web-based anatomy quizzes that access specific atlas pages with only some of the structures labeled. For example, Figure 5 shows a neuroanatomy quiz that includes static HTML in the left-hand frame and an atlas page in the right-hand frame. In this case, the atlas URL calls up a pin diagram that labels only certain regions, with numbers rather than names on the pins. We've used similar methods to create Web-based tutorials that guide the user to specified atlas pages.

Knowledge-based search mode—one form of consultant mode interface—also uses the ability to specify every possible page as a URL. In this case, the forms-based interface to the foundational model, like that shown in the left of Figure 2, finds an anatomical term. Clicking on this term brings up a search form used to search all annotated images, on all atlases on the Web that contain the term. Options in the search engine allow dynamically outlining and labeling the structure, all by providing a particular URL.

CGI programs comprising one of several data servers we've implemented (third row of Figure 1) carry out the search. In this case the search engine consults symbolic metadata (lower left of Figure 1), stored in a relational database, that associates anatomical terms with specific atlas images.

The Web atlases, and an earlier custom Macintosh client, have been available on the Internet since 1992. These "production level" atlases currently don't include access to the foundational model, only to image-based navigation and retrieval of static annotated images. Nevertheless, they have become quite popular for anatomy education and review, receiving over 10,000 hits per day (as of fall 1998) from more than 95 countries. More than

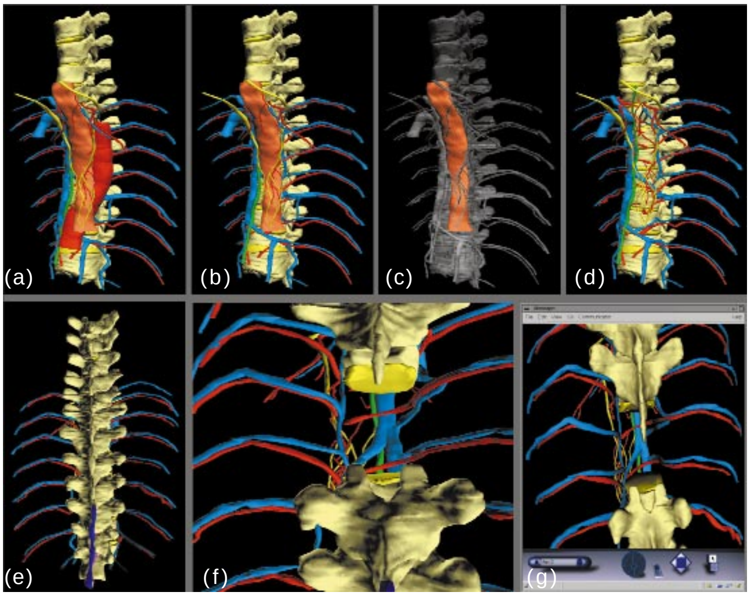


**7** An image sequence generated with the scene builder and navigator. (a) The full scene, depicting the parts of the posterior mediastinum. (b) The descending aorta is "dissected." (c) The user "highlights" the esophagus, causing the graphics server to render the other objects transparent and colorless. (d) The esophagus is removed. (e) The scene is rotated nearly 180 degrees to show the vertebral column from behind. (f) Two vertebrae are removed and the user zooms in for a closer look. (g) The scene is exported as VRML and sent to the client for further interaction. 3D models by David Conley, software by Ben Wong and Jeff Prothero.

1,000 sites have linked to them, and the atlases have received several awards. We also use them heavily in our local gross anatomy and neuroanatomy courses.

Users have made almost uniformly positive online comments, primarily expressing thanks for the large amount of content available, even in these limited regions of the body. Users also like the automatic quiz mode and the 3D scenes. They mainly request more content, additional connections with symbolic information, and dynamic interaction with 3D scenes.

### Scene generation

Partly in response to the users' desires to interact with the 3D scenes, we've created an experimental Web-based 3D scene builder and navigator. The scene builder is one of the modules within the evolving Digital Anatomist agent shown in the second row of Figure 1. The graphics server shown in the third row renders scenes.

Primarily forms-based, Web interaction occurs in what could broadly be called consultant mode. A Web user types the name of an anatomical structure (such as posterior mediastinum) and chooses a semantic relationship, (such as *part of*). The scene builder connects to the foundational model server to determine the parts of the posterior mediastinum, then consults a data server to determine the file names of those primitive 3D meshes that correspond to the parts. For each of these parts the scene builder sets the color depending on the class in the *is-a* ontology (for example, arteries are red, veins are blue), then saves the scene in a cache. Figure 6a show this process for the posterior mediastinum.

When the Web user requests display of the scene, the scene generator tells the graphics server to load the primitive models, render the scene, and save the rendering as a static 2D image. The image goes to the Web user as part of a forms-based interface that looks very similar to the static atlas pages shown in Figures 4 and 5, except that the user can now manipulate the scene.

Figure 6b shows four automatically generated scenes rendered together: parts of the posterior mediastinum, parts of the thoracic vertebral column, branches of the descending thoracic aorta, and tributaries of the azygos vein. This scene includes 89 automatically labeled 3D meshes, each label associating a name from the foundational model with a primitive 3D model.

In Figure 6b, the user has clicked on the third thoracic vertebra. The graphics server received the mouse's *XY* coordinates, simulated a 3D pick operation to determine the structure's name, then sent the name back to the Web client.

Figure 7 illustrates a sample session with the scene builder and navigator. Starting with the full scene (Figure 7a), the user clicks on the descending aorta and asks to "dissect" it. This request causes the graphics server to re-render the scene and send the new image back to the client (Figure 7b). Next the user selects the esophagus, "highlights" it (Figure 7c), and then removes it (Figure 7d). After rotating the remaining structures to see them from behind (Figure 7e), the user removes two vertebrae and zooms in for a closer look (Figure 7f). At any point, the scene can be saved as a VRML file, enabling the user to manipulate the models more interactively on their local machine (Figure 7g).

At the current stage the system only provides VRML as an output option rather than the primary mode of interaction. We opted for this approach because the forms interface is available to virtually anyone with a Web browser, whereas VRML continues to evolve. Moreover, our complex scenes tax all but high-performance clients. As technology evolves, we'll continue to examine the trade-offs between high-performance client workstations running VRML viewers and high-performance graphics servers connected to widely available but lower performance client workstations.

### Brain mapping

We hypothesize that anatomy provides a useful framework for organizing other biomedical information. We've begun to test this hypothesis in the context of the Human Brain Project, a national effort to build

**8** Brain mapping. (left) Photograph of the exposed brain taken during neurosurgery. Numbered tags indicate cortical locations tested for the presence of language function. (right) MRI-based reconstruction of the same patient's brain, rendered to match the photograph as closely as possible. The boxed area represents the portion of the cortex visible in the photograph. Veins appear in blue, arteries in red. Numbers interactively positioned on the reconstruction match the location of the tags in the photograph. Numbers surrounded by boxes indicate sites found to have language function. Software by Jeff Prothero and Kevin Hinshaw.

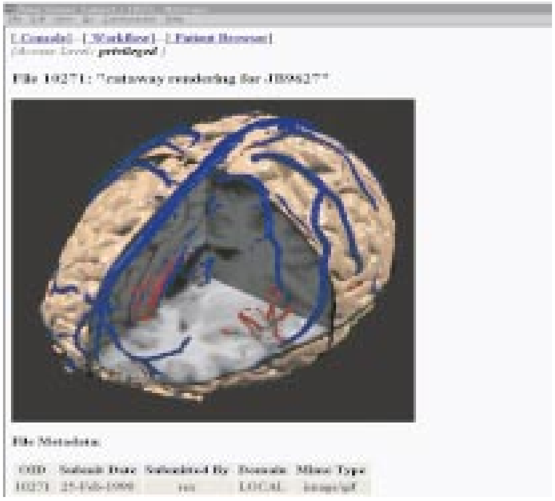

**9** Web-based repository manager. Users can retrieve images and other multimedia data from a Web-accessible multimedia database. The retrieved image shows a cut-away through an MRI-based reconstruction. Repository manager by Rex Jakobovits, visualization software by Andrew Poliakov.

information systems that can manage the massive amounts of data accumulating about the human brain. In the University of Washington Human Brain Project we're extending the Digital Anatomist information system to organize, visualize, analyze, and share cortical language maps obtained at the time of neurosurgery.

Previous studies have shown that language function occurs mostly in the left temporal area of the brain. However, surgical stimulation studies show that the specific location of language varies significantly from one patient to the next. Our project attempts to make sense of this variation by mapping the language sites to a 3D patient-specific reconstruction of each patient's brain, as obtained from magnetic resonance images (MRI). The patient-specific models must then be registered to each other to account as much as possible for anatomical variation among patients. We can then analyze the remaining language site variance in terms of its relationship to such factors as intelligence, gender, and age.

**Brain map authoring**. We generate the 3D patient-specific model from an image volume defined by the MRI slices, using a variation of the model builder program described earlier. Unlike the educational application, it's not feasible to use a purely manual method of reconstruction from serial sections, since the process would become too tedious for large numbers of patients. Instead, we developed a semiautomatic approach to reconstructing the brain surface that uses geometric spatial knowledge of anatomy. In this approach a training set of brain surfaces serve to create a low-resolution generic model of the brain that captures the training set's shape and range of variation. A model-based image segmentation program then uses the generic shape model to guide low-level image processing algorithms in their search for the surface.

The resulting low-resolution brain surface provides a volume mask that excludes nonbrain structures from the MR image volume. Processing the masked image volume with a standard isosurface algorithm obtains the cortical surface. We also use the mask to extract vein and artery surface meshes obtained from separate MR image volumes optimized for veins and arteries.

An interactive brain-mapping program enables the marking of points directly onto the three combined models. A user drags numbers onto the cortical surface such that they visually match the location of language sites visible on a photograph of the exposed brain at the time of neurosurgery (see Figure 8). The mapped language sites are saved as 3D coordinates with respect to a patient-specific coordinate system.

**Brain map retrieval**. The brain map database builds on a multimedia Web-accessible database system, called the Web Interfacing Repository Manager, which we wrote using all free components (msql and Perl).[10]

Multimedia objects such as images and 3D models reside in a protected file area, whereas metadata about those files occupy the msql relational database. Perl modules implement various programming APIs that maintain consistency between these two storage areas and facilitate creating Web interfaces through CGI calls. The collection of these routines constitutes our most complete instantiation of the data server shown in Figure 1.

In the brain map retrieval application, patient demographics, image metadata, 3D model file names, and language maps reside in the relational database, whereas images and 3D models occupy the protected file storage area. A forms-based Web interface allows different classes of users (privileged, collaborator, and public) access to the same information, but they can only see material relevant to their user class. For example, a privileged user might see the patient names, whereas a collaborator might see only research numbers. (In fact, for patient confidentiality reasons we don't include patient names in the database.)

The interface lets a user browse the database, drilling down to retrieve various items associated with the patient. For example, Figure 9 shows a screenshot of one of several renderings for a patient. In this case, the rendering shows a cut-away view of the cortical surface, with MR image intensities mapped onto the cut planes.

## URLs for Online Figures in this Article

Consult the following uniform resource locators (URLs) to view some of the images used in figures in this article. You'll find the Structural Informatics Group home page at http://sig.biostr.washington.edu.

Figure 4a: http://www9.biostr.washington.edu/cgi-bin/DA/PageMaster?atlas:Neuroanatomy+ffpathIndex/3D^Pathways/Auditory^Radiation+2.

Figure 4c: http://www9.biostr.washington.edu/cgi-bin/DA/quizPage?atlas:Neuroanatomy+ffpathIndex/3D^Pathways/Limbic^Zoom+2.

Figure 5: http://www9.biostr.washington.edu/cgi-bin/DA/drawStuff?NeuroSyllabus+Index/Syllabus^Chapters/Vessels^Ventricles^CSF/Circle^of^Willis+-lni+5:8:9:11:13:17:19:21:22:28. To find the names of these structures, try deleting the "n" in the –lni flag. To see fewer structures, try deleting some of the numbers at the end.

## Conclusions and future work

Ten years of experience with this anatomy information system architecture have convinced us of the usefulness of a distributed, network-based approach for reusing anatomical information in multiple applications.

Features that appear particularly useful include

1. detailed anatomical images and 3D models,
2. a foundational model of anatomy that gives meaning to the image-based information,
3. a significant amount of content entered by trained anatomists,
4. separation of content from access methods, letting us take advantage of changes in technology without significantly changing content,
5. distributed modules that permit applying the optimal software and hardware to a given task, and let software developers use the tools most comfortable for them, and
6. access through the Internet, allowing widespread dissemination at minimal cost.

Many different subfields of computer science prove relevant to an anatomy information system. Image processing, graphics, multimedia databases, artificial intelligence, and software engineering all come into play. For example, the generation and display of realistic 3D anatomical models gives rise to many daunting problems well known in the computer graphics field. Thus, this research is by necessity highly interdisciplinary, requiring expertise from many different areas.

The network-based approach lets other groups develop servers (as has begun to happen) so that we don't have to solve all these problems ourselves. Instead, we can access others' solutions. Network-based servers can also provide access to anatomical resources developed by others. In both these cases the problems of interoperability (such as common vocabularies and coordinate systems, common APIs and file formats) are important and not easily solved.

Our future work will continue to develop the individual components in Figure 1 or supplement them with components developed by others. We'll also continually look for ways to combine components. For example, we implemented both the graphics server and the brain mapper in Skandha Lisp. Since Skandha Lisp can establish network connections to other servers, it should prove relatively easy for the graphics server to contact the brain map database, then load and render 3D brain models and brain maps. A Digital Neuroscientist agent can then aid Web-based brain map researchers as they examine 3D patient-specific brain models, rotate them and dissect them, and overlay probability maps of likely language sites based on demographic characteristics that match a patient scheduled for neurosurgery. Such a map could help in surgical planning.

As we develop and integrate the individual components in Figure 1, many other applications will become possible in areas such as image segmentation, medical illustration, medical records, medical education, and therapy planning. As the structural information framework extends to finer levels, it might prove useful in managing embryological, cellular, and molecular data.

In addition, some of these techniques may find applicability outside of medicine. For example, the combination of symbolic ontologies, with 3D primitive models stored in a database, might benefit automated scene generation in such diverse areas as aircraft design or virtual worlds. Since the domain knowledge (anatomy) remains separate from the software modules that access and reason with it, applications outside of anatomy presumably can reuse the modules.

The fact that so many potential applications exist for an anatomy—and, more broadly, a structure-based—

information system, results from the fundamental nature of structure. If we can find better ways to represent structure and to access these representations in a reusable fashion, then we may come closer to a framework for organizing at least some of the information on the Web. ■

## Acknowledgments

This work was funded by National Library of Medicine grant LM06316 and contract LM43546, and Human Brain Project grant DC02310, funded by the National Institute of Deafness and Other Communication Disorders. In addition to the authors, the following individuals were involved in the development of the software tools and contents described in this article: Augusto Agoncillo, Scott W. Bradley, David Conley, Deepa Dinendra, Kraig Eno, Jill Gardner, Rex Jakobovits, Etore Lettich, William Lober, Kenneth Maravilla, Jose Mejino, Bharath Modayur, Kate Mulligan, Loyd Myers, Pamela Neal, George Ojemann, Andrew Poliakov, Jeff Prothero, Joan Robertson, Linda Shapiro, Darren Stalder, and John Sundsten.

### References

1. J.F. Brinkley, "Structural Informatics and Its Applications in Medicine and Biology," *Academic Medicine*, Vol. 66, No. 10, 1991, pp. 589-591.
2. D.M. Conley et al., "Computer-Generated Three-Dimensional Reconstruction of the Mediastinum Correlated With Sectional and Radiological Anatomy," *Clinical Anatomy*, Vol. 5, 1992, pp. 185-202.
3. C. Rosse, L.G. Shapiro, and J.F. Brinkley, "The Digital Anatomist Foundational Model: Principles for Defining and Structuring Its Concept Domain," *Proc. American Medical Informatics Association Fall Symp.*, Hanley and Belfus, Philadelphia, 1998, pp. 820-824.
4. K.H. Hohne et al., "A New Representation of Knowledge Concerning Human Anatomy and Function," *Nature Medicine*, Vol. 1, No. 6, 1995, pp. 506-510.
5. N. Wacholder et al., "An Ontology-Based Navigation System for Human Anatomy," *Proc. Second Visible Human Project Conf.*, M.J. Ackerman and R.A. Banvard, eds., National Library of Medicine, Bethesda, Md., 1998, pp. 95-96.
6. ADAM Software, "Anatomy CD-ROMs," Atlanta, Ga., 1999.
7. J.F. Brinkley et al., "The Digital Anatomist Information System and Its Use in the Generation and Delivery of Web-Based Anatomy Atlases," *Computers and Biomedical Research*, Vol. 30, 1997, pp. 472-503.
8. J.F. Brinkley and J.S. Prothero, "Slisp: A Flexible Software Toolkit for Hybrid, Embedded, and Distributed Applications," *Software—Practice and Experience*, Vol. 27, No. 1, 1997, pp. 33-48.
9. D.A.B. Lindberg, B.L. Humphreys, and A.T. McCray, "The Unified Medical Language System," *Methods of Information in Medicine*, Vol. 32, No. 4, 1993, pp. 281-291.
10. R.M. Jakobovits and J.F. Brinkley, "Managing Medical Research Data With a Web-Interfacing Repository Manager," *Proc. American Medical Informatics Association Fall Symp.*, Hanley and Belfus, Philadelphia, 1997, pp. 454-458.

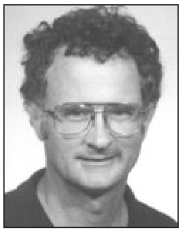

***James F. Brinkley*** *is a Research Associate Professor in the Department of Biological Structure at the University of Washington, with adjunct appointments in the Department of Computer Science and Engineering, and Division of Biomedical Informatics. He defined the field of structural informatics in 1991 and currently directs the UW Structural Informatics Group. He received an MD in 1974 from the University of Washington and a PhD in medical informatics from Stanford University in 1984.*

***Benjamin A. Wong*** *is a hacker. A recent graduate of the University of Washington, he received a BA in philosophy in 1999 and a BS in computer science. His current plans include pursuing a PhD in computer science.*

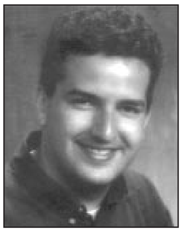

***Kevin Hinshaw*** *is a doctoral candidate in the Department of Computer Science and Engineering at the University of Washington. His main research interest is making computers more perceptive, expressive, and creative. He obtained an MS from the University of Washington in 1994 and a BS in computer science and math from the University of Puget Sound in 1992.*

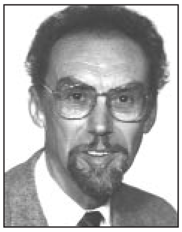

***Cornelius Rosse*** *is a professor in the Department of Biological Structure at the University of Washington, with an adjunct appointment in the Division of Biomedical Informatics. He has received numerous national teaching awards in anatomy and is a fellow in the American College of Medical Informatics. His current research interests are in computer-based anatomical knowledge representation. He received his medical degree in 1964 and Doctor of Science degree in 1983 from the University of Bristol, UK, for contributions in lymphocyte and cancer biology.*

*Contact Brinkley at the Structural Informatics Group, Dept. of Biological Structure, Box 357420, University of Washington, Seattle, WA 98195, e-mail brinkley@u.washington.edu, http://sig.biostr.washington.edu.*