4

THE USE OF ANATOMICAL KNOWLEDGE IN MEDICAL IMAGING: AN OVERVIEW OF THE UNIVERSITY OF WASHINGTON STRUCTURAL INFORMATICS GROUP

James F. Brinkley*

Abstract: The purpose of medical imaging is to provide information about the physical organization of the body. Structural informatics is a subdiscipline of medical informatics concerned with the development of methods for representing, managing, visualizing and utilizing information about the physical organization of the body. Therefore, the tools of structural informatics should be useful for medical image management and analysis. In the University of Washington Structural Informatics Group we are attempting to develop a knowledge base of gross anatomy as one form of structural information, and to use this knowledge base to develop practical applications in clinical medicine, education and research. The applications we are currently working on include the Digital Anatomist, an information system in anatomy, and the UW Human Brain Project, an information system for brain mapping. These projects give rise to structural informatics research in the areas of image understanding, graphics, artificial intelligence, and multi-media databases. These projects and research areas are summarized, along with their relevance to medical image databases. We conclude that a knowledge base of anatomy will become an essential component, not only for medical image databases, but also for other forms of medical image analysis as well.

4.1 INTRODUCTION

In the past few decades developments in medical imaging technology have led to major advances in diagnosis and treatment. The very success of these technologies has resulted in the proliferation of both the number and variety of medical images. It is now becoming difficult, if not impossible, to manage

---

*Structural Informatics Group, Department of Biological Structure, University of Washington, Seattle, Washington 98195 and analyze the images by manual methods. The chapters in this book reflect this situation, and document the need for an informatics of medical imaging, (Kulikowski, 1997).

The most common method for organizing medical images, as well as all other clinical information, is to relate them to the individual patient record, which is generally retrieved according to the patient identifier. Although this approach facilitates individual patient tracking, it does not easily allow images from multiple patients to be combined together. Image retrieval from multiple patients is useful for comparing a current patient with similar cases, for developing general models of organ size and shape in both normal and disease states, and for other kinds of "data mining" tasks that might be applied to clinical image databases.

In order to relate multiple patients it is necessary to find the common elements among those patients. In the case of medical images a natural choice is anatomy, since the delineation of normal and abnormal anatomy (both static and dynamic) is the primary purpose of medical imaging. If individual images can be indexed according to the anatomic structures they delineate, then it will be possible to relate multiple images using a common anatomic framework.

The most basic requirement for anatomically-based image management is a well-defined and standardized set of anatomic terminology. Because of the large number of anatomic structures, this terminology must be organized so as to provide *symbolic knowledge* of anatomy in the form of attributes and semantic relationships among structures. If images are to be retrieved by content (that it, without requiring manual indexing with keywords) then the terminology must also be linked to *spatial knowledge* of anatomy that provides information about the shape and range of variation of anatomical objects, as well as their spatial relationships.

The implementation of these spatial and symbolic representations would constitute a knowledge base of anatomy that could be re-used for multiple applications, including the organization and retrieval of medical images. The development of these representations gives rise to some of the main research problems in *structural informatics*, which we have defined to be the development of representations and methods for organizing, managing and utilizing information about the physical organization of the body (Brinkley, 1991). Since structure provides a rationale foundation for understanding in biology, a structural information framework should provide a rational basis for organizing medical information, including medical images.

At the University of Washington we have formed a Structural Informatics Group for dealing with some of these problems. The goals of this group are: 1) to develop methods for representing both spatial and symbolic information about the physical organization of the body, 2) to develop Web-accessible computer programs that utilize these representations to solve practical problems in clinical medicine, research and education, and 3) to initially develop a knowledge base in gross anatomy, and to apply this knowledge base to the development of anatomy and brain map information systems.

The remainder of this chapter describes our conceptual framework for organizing and accessing anatomical knowledge and data, the current applications that drive the development of the representations, and the informatics research issues that arise from these applications. We end by discussing the impact that these techniques could have for medical image databases.

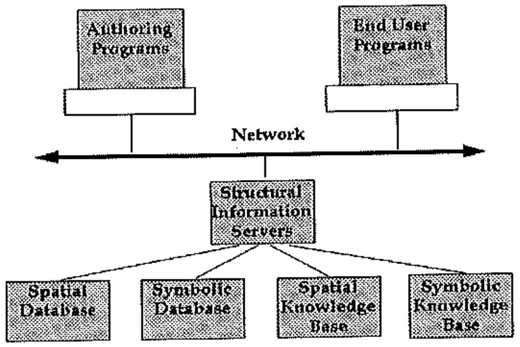

Figure 4.1 Conceptual architecture. End user and authoring programs access a set of re-usable structural information resources by means of one or more structural information servers.

4.2 CONCEPTUAL FRAMEWORK

Figure 4.1 provides a conceptual overview of the distributed architecture we are developing for representing and utilizing anatomical (or more broadly, *structural*) knowledge (Brinkley et al., 1989; Brinkley and Rosse, 1997).

In our conceptual framework, four kinds of representations are implemented in the structural information resources shown at the bottom of Figure 4.1. These four categories result from our classification of structural information along two different dimensions: spatial versus symbolic, and data versus knowledge (Brinkley et al., 1989). This classification has since been adopted by others, and it has also influenced the evolution of the National Library of Medicine's Visible Human Project (Ackerman, 1991).

Along the first dimension, we think of *spatial* information as information that has dimensionality, such as a molecular sequence or a medical image. We think of *symbolic* information as all the other kinds of structural information; for example, the names and semantic relationships of anatomical entities. In general, spatial information is conveyed by the illustrations in an anatomy textbook, whereas symbolic information is conveyed by text.

Along the second dimension, we think of structural *data* as information about a single individual: a set of CT images through the abdomen of a given patient, or a 3-D reconstruction of that patient's kidney. Structural *knowledge*, on the other hand, is usually expressed by models that capture information about classes of individuals: the class of all normal kidneys, including the range of variation, or the names and semantic relationships of anatomical objects in the body.

The four structural information resources are a result of these categorizations: a *spatial database*, consisting of quantitative structural information having 1 to 4 dimensions, a *symbolic database*, consisting of non-spatial data about individuals, a *spatial knowledge base*, consisting of quantitative spatial models about classes of anatomical objects, as well as the relationship between those objects, and a *symbolic knowledge base*, consisting of symbolic representations of the physical and conceptual entities that comprise anatomy.

The remainder of our conceptual framework consists of a set of structural information servers (including Web servers) that provide high level interfaces to these resources over the Internet, and both authoring and end user programs that constitute the applications. The development of the resources is driven by needs of the specific applications.

In the next section we describe our two driving applications, including their goals, current status and further work. We then describe the research issues that are raised by these projects, each of which is or will become an important component of medical image database research.

4.3 PROJECTS

Our current projects include the *Digital Anatomist* and the *UW Human Brain Project.*

The Digital Anatomist is an on-line information system in anatomy, the goal of which is to provide an "expert" anatomist that can answer queries from diverse Web users, both in the form of symbolic text and in the form of dynamically generated 3-D scenes. This application drives the development of spatial and symbolic representations for structural objects in the body.

The UW Human Brain Project is an on-line information system for managing and visualizing data about the human brain. In this case a structural model is used as a framework for organizing other information.

4.3.1 Digital Anatomist

The long term goal of the Digital Anatomist project is an anatomy information system that is available from any desktop computer on the network. The development of this system is driven by the needs of students learning anatomy, but the system has now evolved to a state where it is used by clinicians as well. A user of the information system should be able to query the knowledge base for specific anatomic questions, to retrieve dynamically generated 3-D scenes illustrating answers to the query, and to use the retrieved information as the basis for queries of related databases and image repositories available on the network. Such an information system requires many modules including visual databases, 3-D modelling, real-time rendering, virtual reality and others.

As in all our applications, our approach to meeting these requirements is an incremental one.

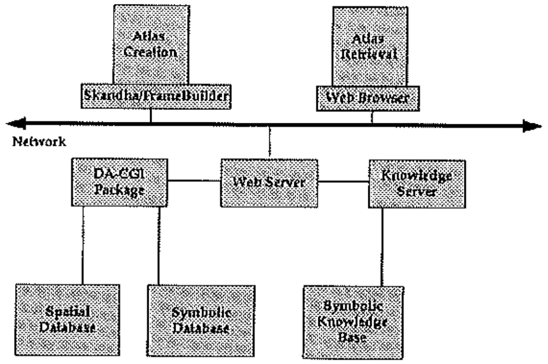

Figure 4.2 Anatomy information system, as an instance of the conceptual framework shown in figure 4.1. In this case the Spatial Database is an image repository containing a series of anatomy atlases. The atlas contents are generated by the Skandha and FrameBuilder programs, and are retrieved via the C-based DA-CGI Package.

Architecture. The current version of our anatomy information system is shown in Figure 4.2, and has evolved from several previous versions (Eno et al., 1991; Brinkley et al., 1993). Although this version does not satisfy all the long term requirements of an anatomy information system, it contains many components of that system, all within the conceptual framework shown in Figure 4.1.

In the current system, the spatial database is an image repository that consists of sets of annotated images and animations packaged into interactive atlases representing different body regions. Most of the images and animations are renderings of 3-D models created by our in house authoring program called Skandha, although any annotated 2-D images may be utilized. The atlases may be off-loaded to CD-ROM, where they are accessible by custom Mac and Windows clients. They also may be accessed directly by a custom Mac client (not shown) developed before the Web became popular (Brinkley et al., 1993; Brinkley et al., 1997a). In the past few years most of the access has been via a Web client implemented by a set of CGI programs called the DA-CGI package (Bradley et al., 1995; Brinkley and Rosse, 1997; Brinkley et al., 1997a). The DA-CGI package also provides links to the symbolic knowledge server described in section 4.4.2, and to a symbolic database that is used for searching atlases on the Web.

Atlas authoring. The construction of the animations and annotated images involves several image processing steps implemented by two main authoring client programs, Skandha and FrameBuilder. These steps give rise to research issues that are discussed in section 4.4.1.

The 3-D models are generated by a process of 3-D reconstruction from serial sections (section 4.4.1). These models are rendered by our locally-developed Skandha program (Prothero and Prothero, 1989; Brinkley and Prothero, 1997), either as static 2-D images or as Quicktime animations. The 2-D images are annotated by a Java-based software tool we call FrameBuilder, which allows the author to delineate regions on the images, and to label them either with the structure names or with commands to open other images. The annotations are saved in a separate file we call a frame. The combined animations and image-frame pairs are saved in the spatial database, one directory for each atlas.

Atlas retrieval. The atlases are accessed over the network by means of the Atlas Web client, which is implemented by a set of C programs that constitute the DA-CGI package. For a given atlas the user navigates through the set of images by clicking on small image icons that retrieve the full image, or by searching for images that depict a structure. Once an annotated image is retrieved it can be examined in either *browse* mode or *quiz* mode.

Figure 4.3 shows *browse* mode for an annotated rendering of a 3-D model in our atlas of Thoracic Viscera. Figure 4.4 shows *quiz* mode for an image from our brain atlas. In both these modes, state information is maintained by hidden form fields that are passed between the different DA-CGI programs.

Additional atlas client modes are described more fully elsewhere (Bradley et al., 1995; Brinkley et al., 1997a). These include a *Pin diagram* mode, reachable from the "Label All" button shown in Figure 4.3, in which the names of all structures are arranged in the margin, and lines are drawn to the centers of the regions depicting the structures. Other modes include a *table of contents* mode, which allows the user to see the navigation hierarchy implicit in the frames, a *knowledge base browse* mode, which displays the knowledge base hierarchy described in section 4.4.2, and a *search mode.* Search mode is entered either by typing a term name, or by clicking on a term in a knowledge base hierarchy. In either case, the search engine consults a separate relational database that contains the URLs for all images that contain the term. The user can then select from this list to retrieve the desired image. The search database is constructed by an "atlas crawler" program that searches all atlases on the Web that are known to it.

Evaluation. Various versions of the anatomy information system have been in use since 1992, and evaluations have been in the context of gross anatomy and neuroanatomy education. However, we envision that this kind of information will be of use in many areas of clinical medicine, research and education. Evaluations have primarily addressed the following questions: 1) How useful and general is the software framework for entering and delivering image-based anatomical content, and 2) How is the atlas used? We have also looked at the

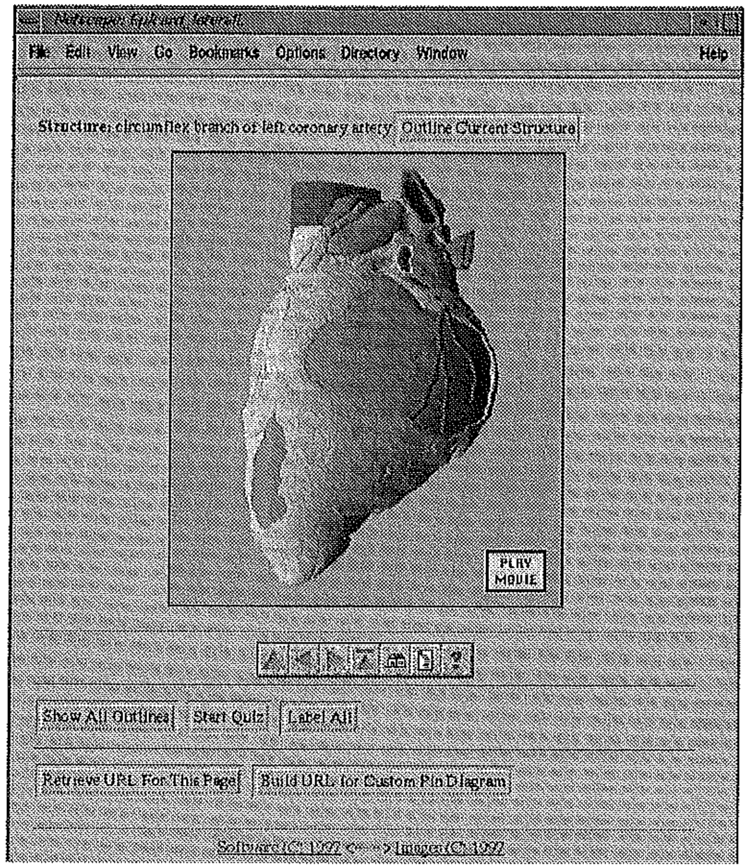

Structure: circumflex branch of left coronary artery

Figure 4.3 Browse mode for the Web atlas client, showing an annotated 3-D image from our Thoracic Viscera atlas. The user has clicked in the outlined region, which caused the Web server to re-send the image with the name of the structure shown along the top. The user then clicked "Outline Current Structure", causing the server to redraw the image with the current structure outlined. Other buttons include "Show All Outlines" to show the outlines for all structures, "Start Quiz", which initiates Quiz mode, and "Label All", which labels all the structures on the image. The icons just below the image are for navigation, help and a table of contents. The bottom buttons generate URLs that can be pasted into on-line tutorials or syllabi. All processing is done by the server DA-CGI package.

effect of network response time on the usefulness of the atlas (Dailey et al., 1993; Dailey et al., 1994; Dailey and Brinkley, 1996).

The *utility* of the distributed framework for entering and delivering image-based anatomical content is evidenced by 1) the number of different interactive

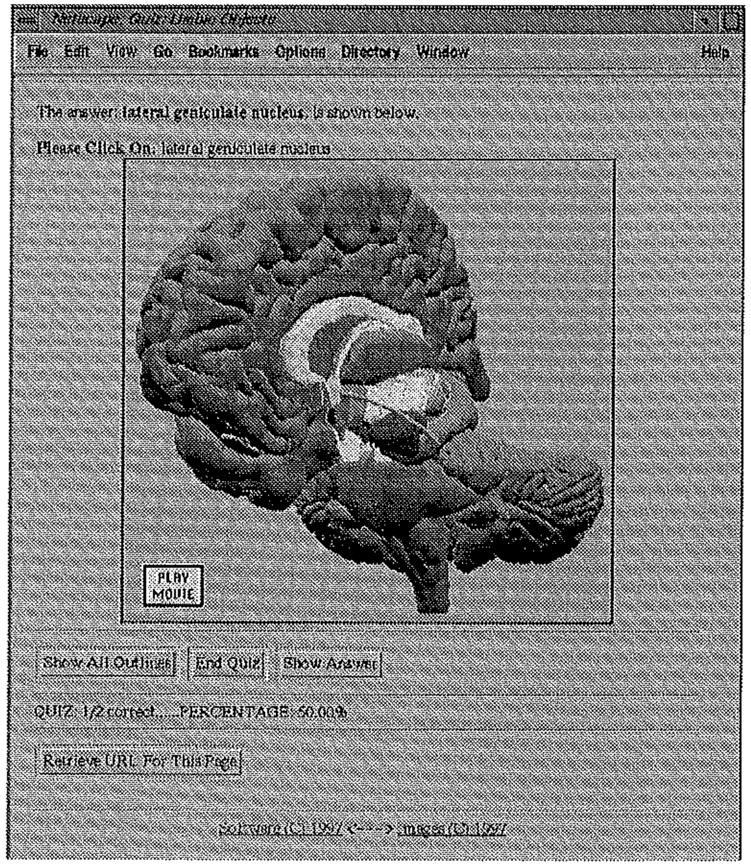

The answer: lateral geniculate nucleus, is shown below.

Please Click On: lateral geniculate nucleus

QUIZ: 1/2 correct.....PERCENTAGE: 50.00%

Figure 4.4 Quiz mode for the Web atlas client, showing an annotated 3-D image from our brain atlas. The DA-CGI programs systematically request the user to point to each annotated structure, keeping track of the number of correct answers. In this case the user has clicked "Show Answer", which causes the server to redraw the image with the currently-request structure outlined. Hidden form items record the score and the list of structures already asked.

atlases that have been created by different authors, and 2) our ability to deliver the same information via CD-ROM as well as via net-based client programs.

The software framework has to date been used by our group to create atlases of the brain from a cadaver and from MRI of a living volunteer, an atlas of the thoracic viscera, an atlas of the knee, and smaller atlases of the brainstem and chest radiology. The same material is available on videodisc, CD-ROM, and on the Internet.

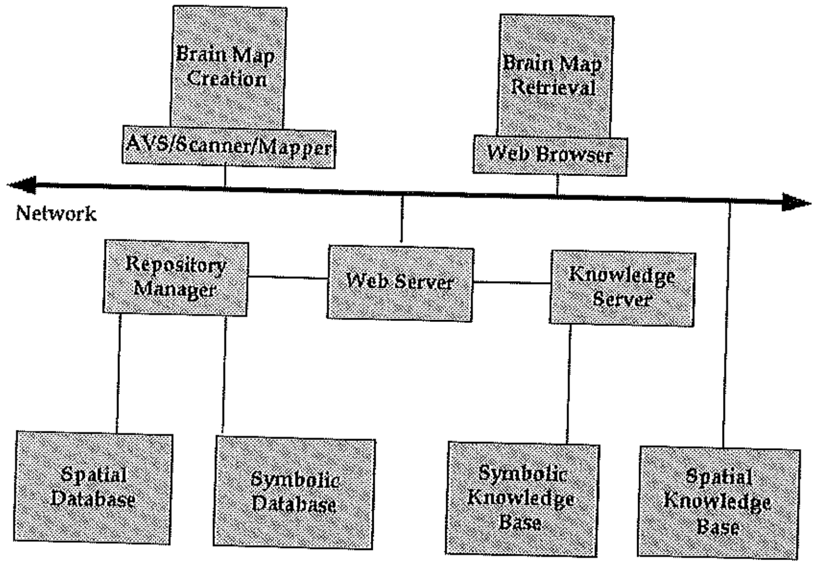

Figure 4.5 Brain map information system, as another instance of the conceptual framework shown in figure 4.1. In this case images and other spatial data are saved in a protected "File Storage Area" of the Spatial Database, which is indexed by relational tables in the Symbolic Database. The combined data repository is managed by a Web based Repository Manager CGI package. The brain maps are created by several programs: AVS, Scanner and Mapper. Mapper is shown in figure 4.6. The maps are retrieved by the Brain Map Retrieval Web client, screens of which are shown in figures 4.7 and 4.8.

The *usage* of the various atlases has been assessed by our local experience, and by examination of on-line log files and comments.

Various versions of the atlas client have been in use for neuroanatomy and gross anatomy classes at the University of Washington since 1992, and by other institutions over the network since 1994. The Web version was introduced in June, 1995, and has since received six Web awards by outside rating organizations. Since Winter 1994 we have only used net-based access for our local courses because of the convenience to the authors in updating content.

During a one year period from June 1995 through June 1996 the total number of sites accessing the Web atlas was over 13,000 from 81 countries. Usage peaked during midterm and final exams, and was very low during breaks. The average Web daily requests was 4073, with a peak of about 25,000 requests on one day. On-line comments have been very favorable; the major requests are for more material, for other ways to navigate through the information, and for links to other parts of the anatomy curriculum.

Further work. In addition to basic issues of 3-D modelling, much of the effort required to extend the anatomy information system will involve integration, particularly between spatial and symbolic information. For example, once a complete set of anatomic terminology has been entered as a result of work described in section 4.4.2, it will be necessary to ensure that all frames created by FrameBuilder are labelled with these terms.

Once all the frames have been indexed by terms in the knowledge base, the atlas search engine can be made more intelligent, since it can use semantic hierarchies in its searches (i.e., "Find all frames that show any branch-of the ascending aorta"). The atlas client can also use the knowledge hierarchies to dynamically change the level of annotation detail: high levels, such as "left ventricle" or "aorta", for K-12 students, and lower levels, such as "conus branch of right coronary artery", for interventional radiologists and cardiac surgeons.

A second major integration step will enable the Web atlas client to call the Skandha program as a server, so that it can dynamically generate 3-D scenes that are annotated with structure names. The scenes can either be rendered on the server as static annotated images, or they can be sent to a VRML client for 3-D interactive viewing. Skandha can already be called as a server; however, the 3-D models need to be indexed by terms in the knowledge base, and the models need to be saved in a more accessible database. The repository manager being developed for the brain mapping project (section 4.3.2) should serve as a very useful management system for atlas models as well.

4.3.2 UW Human Brain Project

The UW Human Brain Project (Brinkley et al., 1997b) is one project within the national Human Brain Project (Huerta et al., 1993), a multi-agency effort to develop informatics tools for managing the exploding amount of information that is accumulating about the human brain.

The objective of our Human Brain Project effort is to organize *functional* information about the brain around the *structural* framework that is conceptually illustrated in Figure 4.1. This application therefore extends the utility of the anatomy information system, described in the previous section, by using it to organize non-structural information.

The particular neuroscience problem we are addressing is the management, visualization and analysis of cortical language mapping data. In recent years, advances in imaging technology such as PET (Demonet et al., 1993) and functional MRI (Desmond et al., 1995) have allowed researchers to observe areas of the cortex that are activated when the subject performs language tasks. These advances have greatly accelerated the amount of data available about human language, but have also emphasized the need to organize and integrate the sometimes contradictory sources of data, in order to develop theories about language organization. Our hypothesis (and that of most Brain Project researchers) is that neuroanatomy is the common substrate on which the diverse kinds of data can be integrated.

In our initial prototype we are concentrating on electrical stimulation mapping data, obtained at the time of neurosurgery for intractable epilepsy. Once the neurosurgical data have been mapped to a patient-specific 3-D anatomical model of the brain, it will be possible to relate them to other kinds of data obtained from the same patient, such as functional MRI or PET. It will also be possible to develop methods for relating multiple patients, by mapping each patient-specific model (spatial *data*) to a generalized brain model that encapsulates ranges of cortical surface variation (spatial *knowledge*).

Patients enrolled in our initial study have epileptic foci in the left temporal region, requiring surgical excision of part of the cortex that may control language function. Previous electrical stimulation studies have shown that cortical language areas seem to be distributed as discrete, approximately $1cm^2$ regions on the surface of the left temporal cortex in right-handed individuals, and that the distribution of these sites varies from one person to the next (Ojemann et al., 1989; Haglund et al., 1993). It is necessary to map the location of these sites, in each individual before the surgical excision, in order to spare the language areas as much as possible.

For the patients included in our prototype information system, the language sites are determined by an object naming task performed by the awake patient following temporal craniotomy. Numbered tags are placed on the cortical surface (Figure 4.6, Top left). The patient is shown a series of slides of common objects, while the surgeon stimulates each numbered site with a small electrical current. A site is called an *essential* language site if the patient is unable to name the object while the current is applied to that site. The sites are called *essential* because previous studies have shown that surgical disruption of these sites causes post-surgical aphasia, whereas avoidance of the sites results in no language difficulties.

The language sites mapped in this way also provide a rich source of data for testing hypotheses about language organization in the brain. In one study Ojemann et. al. looked at 117 patients (Ojemann et al., 1989), and found that the distribution of language sites varies with both sex and verbal IQ, suggesting that there may be a relationship between anatomical location and language ability. It is therefore tempting to postulate that variations in cortical surface anatomy may relate to variations in language ability, since cortical surface anatomy is also highly variable from one individual to the next. By mapping the language sites to a patient-specific 3-D anatomical model, and by combining each patient in a generalized model of the cortex, it will be possible to test these sorts of hypotheses.

Architecture. Figure 4.5 shows our current brain map information system, as another instantiation of the conceptual framework shown in Figure 4.1. As in the anatomy information system, the brain map information system includes a spatial database, a set of authoring programs, and an end-user program. An important addition in this case is the repository manager, which implements a Web-based multimedia database system that manages both brain-map author-

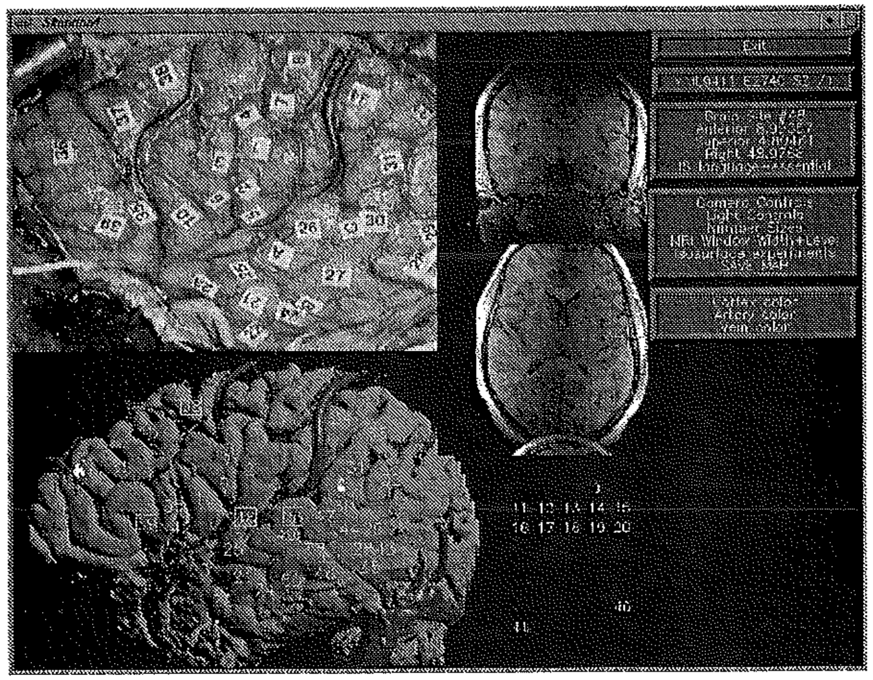

Figure 4.6 Brain Mapper. Top left: Intra-operative photo showing numbered stimulation sites. Bottom left: Surface-rendered left temporal surface with mapped stimulation sites; square boxes are essential for language. Top middle: Coronal and transverse slices through the MR volume corresponding to the location of the mouse click on the reconstruction. Bottom middle: Palette of numbers that are dragged over the reconstruction to perform the mapping. Right: Menu of operations.

ing and brain-map retrieval via Web clients. Such an approach will be very useful for the anatomy information system as well.

Brain Map Authoring. The first requirement for brain mapping is to relate the surgical sites to a patient-specific 3-D anatomical model. Our current approach is to create a 3-D surface model of the brain from MR scans taken prior to surgery, to render the surface and associated vessels, and to visually match the numbered tags seen on the photograph with the rendered image. The rendering shows cortical arteries and veins, as well as the surface anatomy, since vessels are important landmarks for the surgeon.

Given this approach, the first problem is the same as that of the atlas: to segment the cortex from the MR dataset. Unlike the the anatomy information system, it is not feasible to manually segment the brain for each patient, since we expect over one hundred patients to be included in the study. Therefore, our current segmentation approach in this case is semi-automatic: adaptive 3-D region growing, 3-D mathematical morphology, and volume or surface rendering (Modayur et al., 1996; Modayur et al., 1997c).

The process is implemented in six stages, numbered 1-6 in the following discussion. Prior to surgery, three sets of MR image volumes are acquired (1) within an interval that is short enough for the patient to remain motionless. One set is optimized to show cortical anatomy, one to show veins, and one to show arteries. The three image sets are transferred (2) over the network from the radiology MR machine to a database in the structural informatics group. They are then aligned (3) by registering and resampling the three datasets within the MR machine coordinate system, so that all voxels correspond and are the same size. The assumption in this case is that the patient does not move during the three sets of image acquisitions.

The cortical dataset is segmented (4) using an interactive 3-D region growing method (Myers and Brinkley, 1995) implemented in the commercial package AVS. The 3-D region is used as a mask for a standard marching cubes algorithm that extracts the cortex and surface vessels as polygonal meshes.

The resulting surface models are rendered (5) by a module of the *Mapper* program, a screen capture from which is shown in Figure 4.6. The intra-operative photo is shown in the the top left, a surface rendered image in the bottom left, and corresponding MR slices in the middle. The surgeon or technician visually matches the two images, then drags numbers from a palette to locations on the rendering that correspond to the photo. Essential sites are indicated by a double click, which causes a box to be drawn around the number. The resulting map (6) is saved in the spatial database as a set of 3-D coordinates with respect to the MR machine coordinate system.

Brain map retrieval. The brain map authoring tools are designed to be used by only one or two people. The retrieval tool, like the anatomy atlas client, is designed so that anyone with proper authorization can access the language maps, and can relate them to other brain map databases available on the net (Fox et al., 1994).

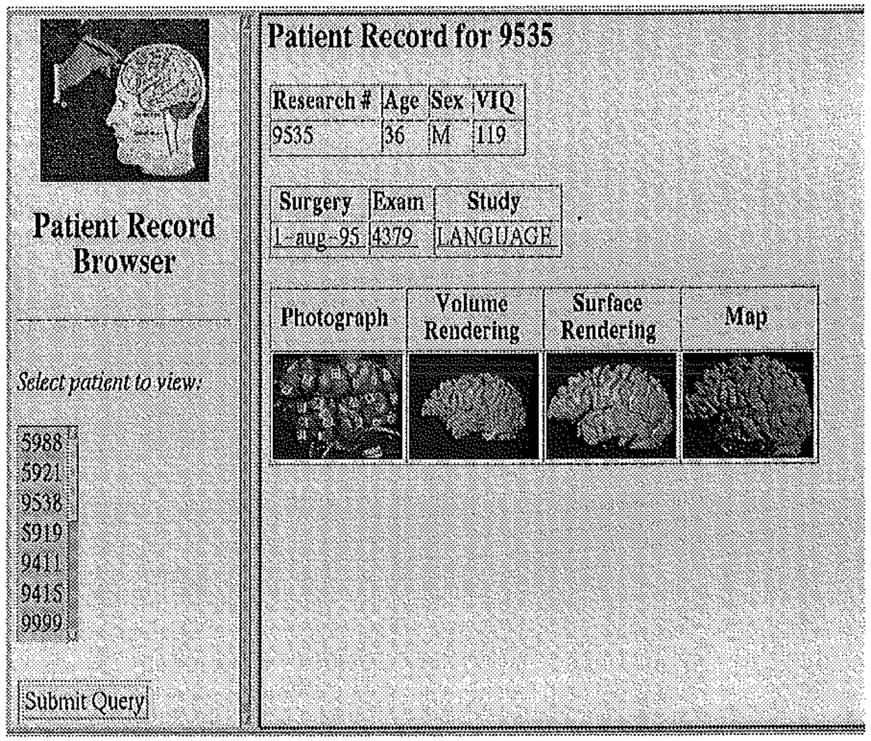

Patient Record Browser

Select patient to view:

- 5988
- 5921
- 9538
- 5919
- 9411
- 9415
- 9999

Submit Query

Patient Record for 9535

Research #	Age	Sex	VIQ
9535	36	M	119

Surgery	Exam	Study
1-aug-95	4379	LANGUAGE

Photograph	Volume Rendering	Surface Rendering	Map

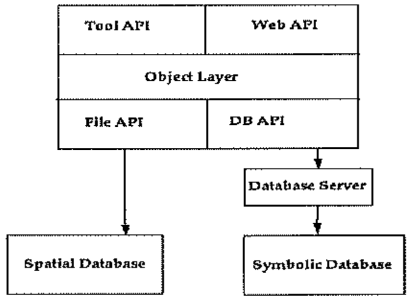

Figure 4.7 Brain Map Retrieval client (top) accessing the Web Repository Manager (bottom), a set of perl programs that implement API's to images and other multimedia data stored in the spatial database, and metadata stored in the symbolic database.. For brain map retrieval a frames-based Web interface is used to display a list of patient identifiers in the left-hand frame. The user selects the patient, then presses "Submit Query". The repository manager formulates a query to the symbolic database to retrieve the requested information, including the identifiers of image files in the spatial database. The results are packaged in the frame shown on the right, which in this case consists of symbolic and spatial data from the individual patient record.

The brain map retrieval tool is generated by a Web-based repository manager (Jakobovits et al., 1996; Jakobovits and Brinkley, 1997), described in section 4.4.3. The top half of Figure 4.7 is one screen generated by this tool. The list of patient identifiers in shown in the left frame. Clicking on one of these identifiers retrieves the record for that patient from the symbolic database (a relational database), along with pointers to image files stored in the spatial database. The repository manager then dynamically generates a page showing the combined spatial and symbolic data in the right hand frame. Clicking on highlighted fields brings up additional information.

Evaluation. A total of 12 patients have been entered into the repository to date. All were mapped by our neurosurgeon collaborator (Ojemann), and by three non-experts in our group. Each observer had no knowledge of the results from the other observers. Repeatability analysis showed that all mapped sites (3-D location in MR machine coordinates) were located within 5 mm (2 SD) of the mean for the four observers (Modayur et al., 1997c; Modayur et al., 1997a). The surgical sites are only localized to within about $1cm^2$ because of the size of the stimulating electrodes. Therefore, in the absence of any other gold standard, the repeatability results show that the visual matching technique is a viable method for relating surgically-obtained language sites to a 3-D patient-specific model of the brain.

The repository manager has also been used in preliminary studies to look at correlations between language sites and behavioral factors, following earlier non-computer based analysis (Ojemann et al., 1989). In the left-hand frame of Figure 4.8 the user selects constraints that group the patients into two categories, in this case females with verbal IQ (VIQ) less than 99, and females with verbal IQ greater than 99. This query is sent to the repository manager, which retrieves those patients satisfying the criteria, summarizes the two groups, and determines if there is a significant difference between them. For this example, the dependent variable is the percentage of electrical stimulation sites that are associated with language. For this small sample, the low VIQ females have significantly more language sites than the high VIQ females (Modayur et al., 1997b). Although larger numbers of patients are needed, this result confirms earlier non-computer based results on larger numbers of patients (Ojemann et al., 1989). As new patients are added, only one or two button clicks will be needed to recalculate these kinds of correlations.

Further work. The ability to visually map the language sites to a 3-D model, and to store these results in a Web accessible multimedia database, is a prerequisite for other studies using the brain map information system. We are now in a position to integrate other forms of functional image data from the same patient, and to relate multiple patients. Accomplishment of these objectives will be facilitated by integrating components from our other projects.

As in many areas of image processing, a major bottleneck to all our applications remains segmentation. Although the 3-D region grower works, it is not robust enough to be run by a technician, nor is it generalizable to other

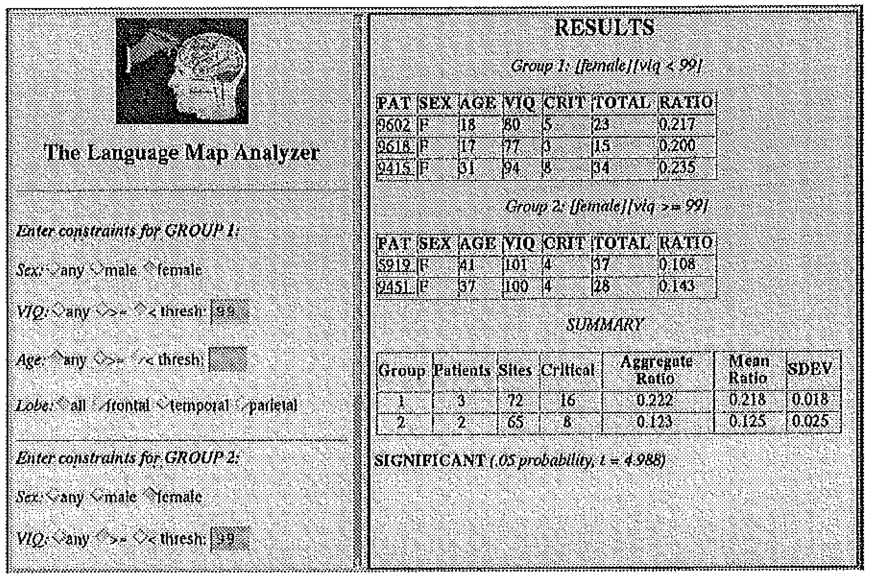

The Language Map Analyzer

Enter constraints for GROUP 1:

Sex: ◇any ◇male ◆female

VIQ: ◇any ◇>= ◆< thresh: 99

Age: ◆any ◇>= ◇< thresh:

Lobe: ◆all ◇frontal ◇temporal ◇parietal

Enter constraints for GROUP 2:

Sex: ◇any ◇male ◆female

VIQ: ◇any ◆>= ◇< thresh: 99

RESULTS

Group 1: [female][viq < 99]

PAT	SEX	AGE	VIQ	CRIT	TOTAL	RATIO
9602	F	18	80	5	23	0.217
9618	F	17	77	3	15	0.200
9415	F	31	94	8	34	0.235

Group 2: [female][viq >= 99]

PAT	SEX	AGE	VIQ	CRIT	TOTAL	RATIO
5919	F	41	101	4	37	0.108
9451	F	37	100	4	28	0.143

SUMMARY

Group	Patients	Sites	Critical	Aggregate Ratio	Mean Ratio	SDEV
1	3	72	16	0.222	0.218	0.018
2	2	65	8	0.123	0.125	0.025

SIGNIFICANT *(.05 probability, t = 4.988)*

Figure 4.8 Language Map Analyzer, accessing the same Web repository manager (using different perl routines) as that in figure 4.7. In the left-hand frame the experimenter chooses criteria that classify the patients into two groups. When the "Submit Query" button is pressed a query is sent to the symbolic database to retrieve those patients satisfying the criteria. The two groups are summarized, and the significance of the difference between the dependent variables ("Mean Ratio" of language sites to stimulation sites in this case) is determined. Individual patient records can be examined by clicking on the research number.

structures. Therefore, we are now in the process of replacing this approach by the 3-D shape based segmentation techniques described in section 4.4.1. This approach uses spatial knowledge of anatomy to guide low level image processing algorithms. The local coordinate system in this case is the Talairach system (Talairach and Tournoux, 1988), which will allow us to relate 3-D language maps from one patient to another, and to other language data on the network (Fox et al., 1994).

Symbolic indexing of the language data will be possible once the symbolic knowledge base has been augmented by brain terminology. This terminology will be based on Neuronames, a set of terms that were developed at the UW and are now part of the UMLS (Martin et al., 1990). Since the symbolic knowledge server is already on our local Intranet, the Mapper program will be able to call on the knowledge server, in order to display cortical surface terminology and use it for labelling. Terminology, such as that described by Ono for sulcal variation (Ono et al., 1990), can also be used to symbolically classify the variations in cortical anatomy for each patient, thereby allowing correlations to be established between language site variation and variation in cortical topology

From an information management point of view the common integration point for the brain project is the repository manager. Although the repository manager currently is only implemented for image access by the brain map retrieval Web interface, it will evolve to handle all the spatial data generated in the process of brain map creation, and in fact the data generated for the anatomy information system as well.

4.4 INFORMATICS RESEARCH

The projects described in the previous section give rise to many interesting and challenging research issues in structural informatics. These problems can be classified according to the kinds of structural information they deal with, and by the computer science and engineering fields which study them: 1) spatial data and knowledge (graphics and image understanding), 2) symbolic knowledge (artificial intelligence), and 3) spatial and symbolic data (multimedia databases).

4.4.1 Spatial Data and Knowledge (Graphics and Image Understanding)

A common problem in many of our projects is to extract anatomical objects from medical images, and to visualize these objects, as realistically and as efficiently as possible. Because of their rich and varied shape, biological objects pose many interesting and difficult challenges for researchers interested in image understanding and graphics.

The approaches we have taken can be divided into three categories: reconstruction from serial sections, 3-D region-growing, and shape-based surface models.

Reconstruction from Serial Sections. A common approach to extracting a 3-D object from medical image data is to draw the cross-section of the object on each of a series of parallel (or serial) image slices, then to "loft" or "tile" a surface over the adjacent contours. This approach remains the most accurate for defining models of the type required in the Digital Anatomist.

The input to this procedure is an image volume consisting of a set of serial sections. Two well-known examples of this kind of input are the Visible Human male and female (Ackerman, 1991), but clinical image volumes can also be used. The cross-section of each structure on each of these images is segmented, using a manual segmentation tool we developed called Morpho. The resulting stack of contours is input to our locally-developed Skandha program (Prothero and Prothero, 1989; Brinkley and Prothero, 1997). Skandha is then used to reconstruct the contours into a 3-D surface by semi-automatic tiling, to combine surfaces into 3-D models, and to render the models, either as static 2-D images or as Quicktime animations. The renderings and movies are input the the Frame Builder program for inclusion in the anatomy information system (section 4.3.1). Open research problems include segmentation, tiling, and multi-resolution smooth models for real-time rendering.

3-D Region Growing. Although accurate, the reconstruction technique is far too tedious for clinical applications like the UW Human Brain Project. It is therefore necessary to develop automatic or semi-automatic approaches in these cases.

Common low-level image segmentation approaches include edge following and region growing. In the UW Human Brain Project, the results from our first twelve patients were obtained using an adaptive 3-D region grower to segment the cortex from the input MR images (Myers and Brinkley, 1995; Modayur et al., 1997c; Modayur et al., 1997b; Modayur et al., 1997a).

The basic idea behind the adaptive region grower is that voxels belonging to the same tissue type and adjacent to each other will have fairly homogeneous grayscale properties. The algorithm thus attempts to find these initial or seed areas by looking for voxels whose neighborhood variance is below a user-specified threshold. Regions are then grown recursively from these seed voxels. As the regions begin to grow in size, the algorithm starts considering grayscale characteristics of the growing region as well as adjacency constraints of other voxels in the region.

The output of the region grower is a set of labeled 3D regions or volumes, from which the user selects the cortex. The extracted cortex volume is then input to a marching cubes algorithm (Lorensen and Cline, 1987) to extract the surface, and to 3-D morphological operators which expand the volume enough to include the surface arteries and veins. The expanded mask is used to help extract the vessels from the corresponding MR images. The resulting surfaces (veins, arteries and cortex) are then input to the Mapper program for surface rendering and mapping (Figure 4.6). Open research problems include automatic learning of parameters controlling the region grower, and exclusion of non-cortical regions with similar greyscale characteristics.

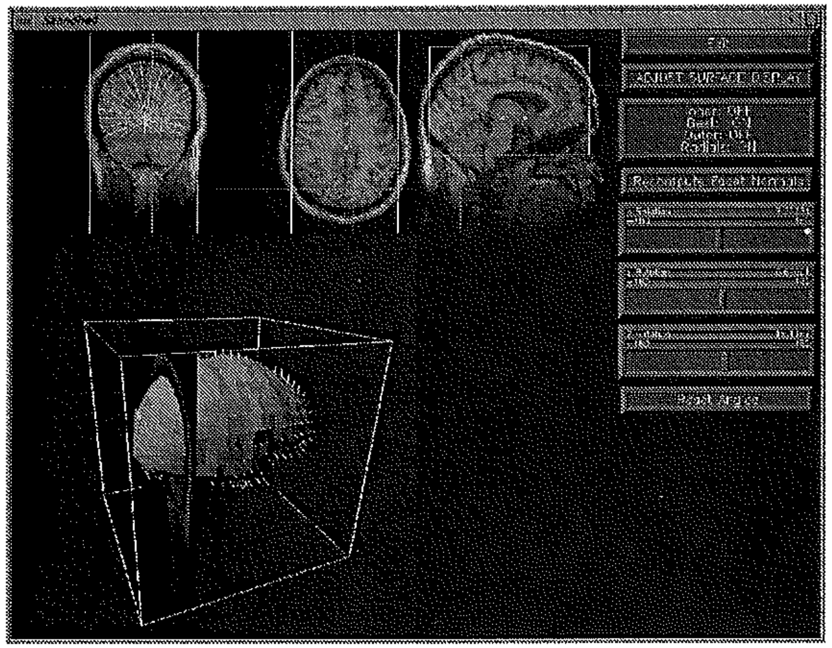

Figure 4.9 Scanner user interface for shape-based segmentation. The top three panels are orthogonal views through a volume MR dataset of the brain. The bottom panel shows an instantiated model of the cerebral cortex after the surface has been found. The location of the current slice is shown in relation to the model. This coarse model is then used as a guide for extraction of the detailed surface anatomy.

Shape-based surface models. Region growing and other low level segmentation techniques work well when the organ of interest has high contrast with respect to neighboring organs. However, these techniques often fail when the border between organs is not distinct, a common occurrence. Knowledge-based approaches attempt to compensate for this ambiguity in the data by giving the computer anatomic knowledge that can be used to guide the low level processes. One such class of knowledge is *spatial knowledge* of the shape and range of variation for various organ classes.

For both our image segmentation and earlier protein structure work we developed a representation, called geometric constraint networks (GCNs) (Brinkley, 1985; Brinkley, 1992), that attempts to model 3-D shape and range of variation by a set of interacting local shape constraints. The hypothesis behind this representation is that a collection of local constraints can interact to generate an "emergent" representation for the overall shape of the object. This representation was initially implemented and tested in a 2-D version of a GCN called a radial contour model (RCM) (Brinkley, 1993; Hinshaw et al., 1995).

More recently, we have extended this approach to 3-D, and have used it to extract a region of interest from the cortical MR dataset (Hinshaw and Brinkley, 1997a; Hinshaw and Brinkley, 1997b). This ROI is then input to the same kind of marching cubes algorithm used by the region grower, and results in comparable surface reconstructions. Figure 4.9 shows the implementation of this procedure in an interactive image segmentation program we call Scanner.

The main advantages of this method over the region grower are that it is much more intuitive for a non-expert to control, and it compensates for missing or ambiguous edge data. We are therefore in the process of replacing the region grower for the brain project with this method. However, the radial models are not as detailed as those obtained from other methods, so the models must be combined with other methods such as isosurface extraction. Issues that we are therefore continuing to address include generalizing the models to handle more shape classes, providing better methods to backtrack in case of errors by the low level methods, and representing the relationships among multiple structures.

The result of these efforts will eventually be a spatial knowledge base of anatomy that captures not only the shape, but also the normal range of variation of anatomic structures and their relationships.

4.4.2 Symbolic Knowledge (Artificial Intelligence)

Our work on symbolic representation of anatomical knowledge is currently driven by the need to develop a comprehensive set of terminology, as an anatomical enhancement of UMLS's Semantic Network and Metathesaurus (Lindberg et al., 1993; Rosse et al., 1995; Rosse et al., 1998). This terminology is necessary for indexing normal anatomic structures found on medical images.

To date, we have established an ontology (Tu et al., 1995) that classifies physical anatomical entities in an Aristotelian hierarchy based on the ISA relationship, and have instantiated the ontology for all static, macroscopic, anatomical entities that comprise the thorax (Rosse et al., 1998). We associate a unique term, the preferred name, as well as all its synonyms, with each concept and use this term for human-readable representations of the ontology.

Our current anatomical ontology consists of three top level classes of Physical anatomical entity shown in the top part of Figure 4.10 (Anatomical spatial entity, Body substance and Anatomical structure) . Each of these classes has multiple subclasses along the ISA hierarchy. We have also implemented PART-OF, BRANCH-OF and TRIBUTARY-OF links, and have used them to represent anatomical relationships of very high granularity (see, for instance, Figure 4.3). Although this representation has been fully instantiated as yet only for the thorax, we contend that each macroscopically identifiable physical entity in the body can be assigned to the subclasses we have defined.

Representation and implementation. The bottom part of Figure 4.10 shows the software components that support the current ontology, as an instance of the conceptual framework shown in Figure 4.1. In this case the symbolic knowledge base is represented by two tables in a relational database, and

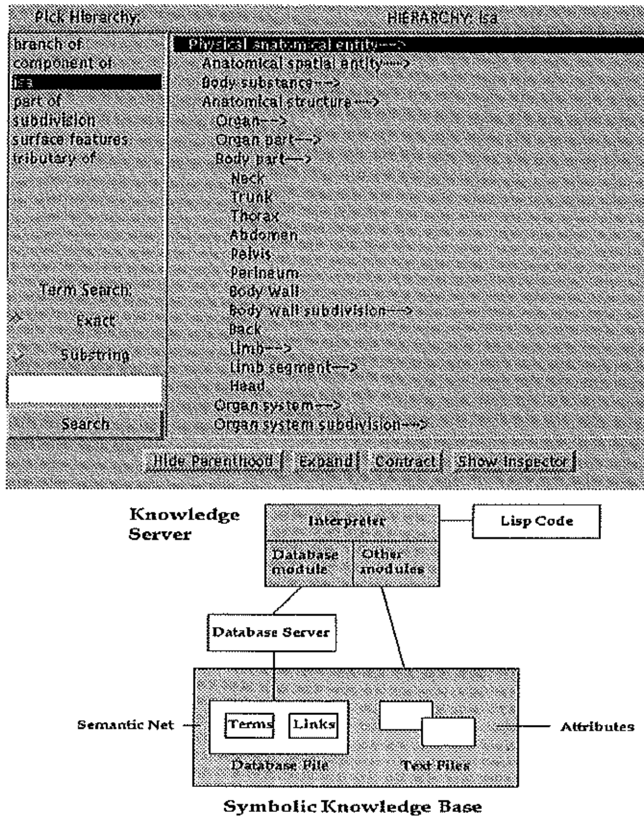

Figure 4.10 Java-based Knowledge Browser client, accessing Symbolic Knowledge Server. The client shows the top level classes, and some of the subclasses, in our anatomical ontology. The --> symbols indicate that the class has at least one subclass. For example, "Organ" has 19 first generation subclasses. The leaves of the tree are canonical instances, as for example "Thorax", which belongs to the class, "Body part". The knowledge server is a hybrid Lisp/C program that translates high level Lisp commands generated by the Java client, such as (KB-GET-CHILDREN "Body Part" "isa"), into SQL queries to a relational database.

by associated text files containing definitions and other textual attributes that provide the basis for later formalization. The terms table contains the names and synonyms for all anatomic concepts. Each concept is assigned a unique ID, and, as in the UMLS Metathesaurus, all synonyms refer to the same concept ID. The terms table also records the associated SNOMED and UMLS ID if the concept exists in these sources.

Terms are related by means of the links table, which records binary semantic relationships such as ISA and PART-OF. The relational database is accessed either by the Sybase commercial relational database server (http://www.sybase.com) or by the Msql free relational server. These servers are in turn accessed by the Symbolic Knowledge Server, a hybrid Lisp-C program that provides a high level Lisp-based query language for symbolic knowledge retrieval (Brinkley and Prothero, 1997).

The Knowledge Server includes a C-based interpreter that parses Lisp queries and dispatches those functions to various C routines that are encapsulated in separate modules. For example, the Lisp query (KB-GET-CHILDREN "Aorta" "part of") retrieves the children of the Aorta along the part-of semantic link. This query is implemented by a C routine in the Database Module that generates an SQL statement, sends the statement to the Sybase or Msql server, and packages the results as a Lisp list.

Knowledge acquisition is currently aided by means of a NeXT-based Knowledge Manager program, a tool we built for entering new knowledge in the knowledge base.

Experimental cgi scripts and Java applets have also been added to the Digital Anatomist Web atlases. These modules contact the Slisp knowledge server and generate html forms or Java-based interfaces (Figure 4.10) to the semantic network. Clicking on a term in these displays initiates a search of the Digital Anatomist atlases for those frames that contain the term. We are currently upgrading the Java interfaces to take over the knowledge acquisition task from the NeXT-based Knowledge Manager.

Evaluation and further work. To date we have entered a total of 11,944 terms, 7159 concepts, and 13,210 relationships, representing all distinct structures visible in the thorax to 1 mm resolution (Rosse et al., 1995; Rosse et al., 1998). Each distinct structure has been classified as an instance of one or more anatomic semantic types, where the anatomic semantic types are specializations of types defined in the UMLS semantic network. The correctness of the terms and relationships are currently being validated within our local group and at NLM, after which they will be integrated with the UMLS.

The ontologies we are developing form the basis of a symbolic knowledge base, but the current representations capture only certain aspects of anatomical knowledge. We are in the process of investigating alternative representations, and will likely port the knowledge base to an established frame-based system. However, we will need to integrate such a system with an object oriented database, since our knowledge base is, or will become, larger than many memory-based systems can handle.

4.4.3 Spatial and Symbolic Data (Multimedia Databases)

The management of structural information presents interesting challenges for database research, not only because there is a large amount of information, but also because the information is of complex and diverse types. A database for structural information needs to manage such diverse entities as molecular sequences, 3-D molecular reconstructions, images and image volumes, organ reconstructions, 2-D and 3-D image annotations, 3-D shape models, and symbolic terminology. The information needs to be efficiently represented and stored, and must be indexed so that it can be retrieved. These requirements are very general, and are studied in such fields as multimedia databases, knowledge-based database retrieval, content-based retrieval, and picture archiving and communications systems for radiology (PACS).

Our current database work is centered around the Web-based repository manager that manages brain map data (Figure 4.5, shown in more detail in Figure 4.7) (Jakobovits et al., 1996; Jakobovits and Brinkley, 1997). This work was influenced by our earlier work in NeXT-based distributed object-based databases for image segmentation (Brinkley, 1994).

Repository systems are becoming a popular approach to managing multimedia data (Bernstein and Dayal, 1994), in which data files are saved in a protected area that is indexed by meta-data about the files. Using the terminology defined in section 4.2, the protected data are stored in the spatial database, whereas the metadata are stored in the symbolic database.

The repository manager (bottom part of Figure 4.7) consists of a set of perl cgi routines that are callable from a Web browser. At the lowest level these routines implement a *File API* for accessing images and other files stored in the spatial database. A *DB API* provides access to the symbolic database, which in this case is implemented in the msql relational database. The database is accessed by the msql database server, which is in turn accessed by the perl routines constituting the DB API.

The perl *Object Layer* provides an object layer over the relational database, and maintains consistency between the spatial and symbolic data. The perl programmer generally calls routines at this level rather than routines at the level of the File and DB API's. At the top level are a set of *Tool APIs* for controlling external programs (one for each program), and a *Web API* for constructing Web forms.

As shown in Figures 4.7 and 4.8, these routines are used to manage the brain project data, and to do preliminary data analysis. As yet only one Tool API has been developed, for controlling an image retrieval program. This program connects to the Radiology MR machine, shows a list of patient imagesets on that machine, and retrieves the selected images over the network. We are in the process of developing additional Tool APIs for controlling the other programs used for brain map authoring.

Further work. The repository manager has already proved very useful for managing multimedia data for the brain project. The perl APIs make it easy to develop new Web interfaces, and facilities such as relational schema evolution (adding, modifying and deleting columns in tables) allow us to change the repository as we add more complex types. The fact that all components are free means that we will be able to distribute this to other users who need to manage personal or small group multimedia data.

We will continue to extend the Web-based repository manager to handle more complex datatypes, to integrate symbolic information for intelligent queries, and to interface to other programs. We will also generalize these methods and database models to handle a wide range of multimedia data, and will test these general models on other types of data. We are also considering an object oriented databases as a replacement for msql. In the longer term the database system, when coupled with the symbolic knowledge base, will become the backbone of our structural information framework.

4.5 CONCLUSIONS

In this chapter we have described the major projects and image-related research topics of the University of Washington Structural Informatics Group. Both of our major projects (the Digital Anatomist and the UW Brain Project) involve extensive use of medical images, and require at least some form of image database.

Our current database system is the Web-based repository manager. As we continue to develop this program for the UW Brain Project, it will become increasingly useful for the Digital Anatomist, and for other image-related projects as well.

Image management and analysis will be facilitated as anatomical knowledge is more closely integrated with the application programs. For example, once the symbolic knowledge base is accessible by authoring programs such as Frame Builder, Skandha and the Brain Mapper, the terminology will be useful for retrieving images and other multimedia data from the repository. The retrieval will become more "intelligent" as we add inference capabilities to the knowledge server.

Content-based retrieval will be facilitated as the spatial knowledge base evolves to the point where it can be used to help automatically guide segmentation of new images and image volumes. Eventually, using these and techniques such as deformable modelling, it will be possible to automatically deform a standardized model of anatomy, from the Visible Human or a similar source, to fit an individual patient imageset. This instantiated model will then provide a roadmap from which to search for abnormalities. Since the model will already be labelled by the correct terminology, both symbolic and content-based image retrieval will become possible without the need for manual indexing.

Given this scenario, the ultimate goal of our medical image database research becomes the same as that of medical image understanding: to instantiate a patient specific model of both anatomy and pathology, such that the model includes all the features of interest on the images. Once this model is available, the proper spatial and symbolic indexes become automatically available.

In order to create such a model we argue that a knowledge base of anatomy is essential. The development of such a knowledge base gives rise to challenging issues that are at the core of structural informatics research.

Acknowledgments

This work was supported by National Library of Medicine awards R01LM04926, F37LM00017, N01LM13506, N01LM43546, National Cancer Institute award R29CA59070, Human Brain Project award R01DC/LM02310, jointly sponsored by National Library of Medicine and the National Institute on Deafness and Other Communication Disorders, National Science Foundation award IRI9116809, Achievement Awards for College Scientists, the University of Washington IAIMS program under the direction of Sherrilynne Fuller, the University of Washington School of Medicine, and equipment grants from IBM and the Murdock Foundation.

This chapter reviews work done by members of the Structural Informatics Group (formerly called the Digital Anatomist Program) since 1988. Other than the author, the following individuals contributed to the work reported here, all of whom were at the University of Washington. Software development: Bill Barker, Scott W. Bradley, Kraig Eno, Kevin Hinshaw, Rex Jakobovits, Bharath Modayur, Loyd Myers and Jeff Prothero. Anatomy content: David Conley, Jose Mejino, Kate Mulligan, John Prothero, Peter Ratiu, Cornelius Rosse, and John Sundsten. Neurosurgery and radiology content: Etorre Lettich, Ken Maravilla, George Ojemann, and Jay Tsuruda.

Further information can be found on our Web site http://www1.biostr/washington/edu/sig, including a list of CD-ROM products for sale, and a list of software tools available for downloading.

References

Ackerman, M. (1991). The visible human project. *Journal of Biocommunication*, 18(2):14.

Bernstein, P. and Dayal, U. (1994). An overview of repository technology. In *Proceedings of the 20th VLDB Conference*, Santiago, Chile.

Bradley, S., Rosse, C., and Brinkley, J. (1995). Web-based access to an online atlas of anatomy: the digital anatomist common gateway interface. In *19th Symposium on Computer Applications in Medical Care*, pages 512–516, New Orleans.

Brinkley, J. (1985). Knowledge-driven ultrasonic three-dimensional organ modelling. *PAMI*, PAMI-7(4):431–441. Also available as Stanford Technical Report KSL-85-33.

Brinkley, J. (1991). Structural informatics and its applications in medicine and biology. *Academic Medicine*, 66:589–591.

Brinkley, J. (1992). Hierarchical geometric constraint networks as a representation for spatial structural knowledge. In *Proceedings, 16th Annual Symposium on Computer Applications in Medical Care*, pages 140–144. SCAMC.

Brinkley, J. (1993). A flexible, generic model for anatomic shape: application to interactive two-dimensional medical image segmentation and matching. *Computers and Biomedical Research*, 26:121–142.

Brinkley, J. (1994). A distributed, object-oriented framework for medical image management and analysis: application to evaluation of medical image segmentation techniques. In *Proceedings, IEEE Workshop on Biomedical Image Analysis*, pages 194–203, Seattle. IEEE Computer Society Press.

Brinkley, J., Bradley, S., Sundsten, J., and Rosse, C. (1997a). The digital anatomist information system and its use in the generation and delivery of web-based anatomy atlases. *Computers and Biomedical Research.* In Press.

Brinkley, J., Eno, K., and Sundsten, J. (1993). Knowledge-based client-server approach to structural information retrieval: the digital anatomist browser. *Computer Methods and Programs in Biomedicine*, 40:131–145.

Brinkley, J., Myers, L., Prothero, J., Heil, G., Tsuruda, J., Maravilla, K., Ojemann, G., and Rosse, C. (1997b). A structural information framework for brain mapping. In Koslow, S. and Huerta, M., editors, *Neuroinformatics: An Overview of the Human Brain Project*, chapter 9, pages 309–334. Lawrence Erlbaum, Mahwah, New Jersey.

Brinkley, J. and Prothero, J. (1997). Slisp: A flexible software toolkit for hybrid, embedded and distributed applications. *Software – Practice and Experience*, 27(1):33–48.

Brinkley, J., Prothero, J., Prothero, J., and Rosse, C. (1989). A framework for the design of knowledge-based systems in structural biology. In *Proceedings, 15th Annual Symposium on Computer Applications in Medical Care*, pages 61–65, Baltimore, Maryland.

Brinkley, J. and Rosse, C. (1997). The digital anatomist distributed framework and its applications to knowledge based medical imaging. *Journal of the American Medical Informatics Association*, 4(3):165–183.

Dailey, D. and Brinkley, J. (1996). A metric for quantifying response time in a browser application. *IEEE Trans. Systems, Man and Cybernetics, Part A (Systems and Humans)*, 26(2):271–275.

Dailey, D., Eno, K., and Brinkley, J. (1994). Performance evaluation of a distance learning program. In *Proceedings, 18th Annual Symposium on Computer Applications in Medical Care*, pages 76–80, Washington, D.C.

Dailey, D., Eno, K., Zick, G., and Brinkley, J. (1993). A network model for wide area access to structural information. In *17th Symposium on Computer Applications in Medical Care*, pages 497–501. SCAMC.

Demonet, J., Wise, R., and Frackowiak, R. (1993). Language functions explored in normal subjects by positron emission topography: A critical review. *Human Brain Mapping*, 1:39–47.

Desmond, J., Sum, J., Wagner, A., Demb, J., Shear, P., Glover, G., Gabrielli, J., and Morell, M. (1995). Functional mri measurement of language lateralization in wada-tested patients. *Brain*, 118:1411–1419.

Eno, K., Sundsten, J., and Brinkley, J. (1991). A multimedia anatomy browser incorporating a knowledge base and 3-d images. In *Symposium on Computer Applications in Medical Care*, pages 727–731. SCAMC.

Fox, P., Mikiten, S., Davis, G., and Lancaster, J. (1994). Brainmap: A database of human functional brain mapping. In *Functional Neuroimaging*, chapter 9, pages 95–105. Academic Press, San Diego.

Haglund, M., Berger, M., Shamseldin, M., Lettich, E., and Ojemann, G. (1993). Cortical localization of temporal lobe language sites in patients with gliomas. *Neurosurgery*, 34(4):567–576.

Hinshaw, K., Altman, R., and Brinkley, J. (1995). Shape-based models for interactive segmentation of medical images. In *SPIE Medical Imaging 1995: Image Processing*, pages 771–780, San Diego.

Hinshaw, K. and Brinkley, J. (1997a). Shape-based interactive three-dimensional medical image segmentation. In Hanson, K., editor, *SPIE Medical Imaging: Image Processing, Volume 3034*, pages 236–242, Newport Beach, CA.

Hinshaw, K. and Brinkley, J. (1997b). Using 3-d shape models to guide segmentation of mr brain images. In *Proceedings, Annual Fall Symposium*, pages 469–473, Nashville, Tennesee. American Medical Informatics Association. Awarded second place, student paper competition.

Huerta, M., Koslow, S., and Leshner, A. I. (1993). The human brain project: an international resource. *Trends in Neurosciences*, 16(11):436–438.

Jakobovits, R. and Brinkley, J. (1997). Managing medical research data with a web-interfacing repository manager. In *Proceedings, Annual Fall Symposium*, pages 454–458, Nashville. American Medical Informatics Association.

Jakobovits, R., Modayur, B., and Brinkley, J. (1996). A web-based repository manager for brain mapping data. In *Proceedings, AMIA Fall Symposium*, pages 309–313, Washington, D.C. American Medical Informatics Association.

Kulikowski, C. (1997). Medical imaging informatics: challenges of definition and integration. *Journal of the American Medical Informatics Association*, 4(3):252–253.

Lindberg, D., Humphreys, B., and McCray, A. (1993). The unified medical language system. *Methods of Information in Medicine*, 32(4):281–291.

Lorensen, W. and Cline, H. (1987). Marching cubes: a high resolution 3-d surface construction algorithm. *ACM Computer Graphics*, 21(4):163–169.

Martin, R., Dubach, J., and Bowden, D. (1990). Neuronames: Human/macaque neuroanatomical nomenclature. In *Proceedings, 14th Annual Symposium on Computer Applications in Medical Care*, pages 1018–1019. American Medical Informatics Association.

Modayur, B., Jakobovits, R., Maravilla, K., Ojemann, G., and Brinkley, J. (1997a). Evaluation of a visualization-based approach to functional brain mapping. In *Proceedings, Annual Fall Symposium*, pages 429–433, Nashville. American Medical Informatics Association.

Modayur, B., Ojemann, G., Lettich, E., Jakobovits, R., Sundsten, J., and Brinkley, J. (1997b). Information system for mapping and studying cortical language organization. In *Abstracts, Society for Neuroscience Annual Meeting*, page 2228, New Orleans.

Modayur, B., Prothero, J., Ojemann, G., Maravilla, K., and Brinkley, J. (1997c). Visualization-based mapping of language function in the brain. *Neuroimage*, 6:245–258.

Modayur, B., Prothero, J., Rosse, C., Jakobovits, R., and Brinkley, J. (1996). Visualization and mapping of neurosurgical functional data onto a 3-d mr-based model of the brain surface. In *Proceedings, 1996 AMIA Fall Symposium*, pages 304–308, Washington, D.C. American Medical Informatics Symposium.

Myers, L. and Brinkley, J. (1995). Visualization of brain surface features using registered partially segmented mri scans. In *SPIE Medical Imaging: Image Display*, pages 43–52, San Diego, CA. SPIE.

Ojemann, G., Ojemann, J., Lettich, E., and Berger, M. (1989). Cortical language localization in left, dominant hemisphere. *J. Neurosurgery*, 71:316–326.

Ono, M., Kubik, S., and Abernathy, C. (1990). *Atlas of the Cerebral Sulci.* Thieme Medical Publishers, New York.

Prothero, J. and Prothero, J. (1989). A software package in c for interactive 3d reconstruction and display of anatomical objects from serial section data. In *NCGA Proceedings*, pages 187–192.

Rosse, C., Ben Said, M., Eno, K., and Brinkley, J. (1995). Enhancements of anatomical information in umls knowledge sources. In *Proceedings, 19th Annual Symposium on Computer Applications in Medical Care*, pages 873–877, New Orleans. American Medical Informatics Association.

Rosse, C., Mejino, J., Modayur, B., Jakobovits, R., Hinshaw, K., and Brinkley, J. (1998). Motivation and organizational principles for anatomical knowledge representation: the digital anatomist symbolic knowledge base. *Journal of the American Medical Informatics Association*, 5(1):17–40.

Talairach, J. and Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain.* Thieme Medical Publishers, New York.

Tu, S., Eriksson, H., Gennari, J., Shahar, Y., and Musen, M. (1995). Ontology-based configuration of problem-solving methods and generation of knowledge-acquisition tools: application of protege-ii to protocol-based decision support. *Artificial Intelligence in Medicine*, 7:257–289.