9
Imaging and Structural Informatics

James F. Brinkley and Robert A. Greenes

9.1 Introduction

As is evident to anyone who has had an X-ray, a magnetic resonance imaging (MRI) exam, or a biopsy, images play a central role in the health care process. In addition, images play important roles in medical communication and education, as well as in research. In fact much of our recent progress, particularly in diagnosis, can be traced to the availability of increasingly sophisticated images that not only show the structure of the body in incredible detail but also show the function.

Although there are many **imaging modalities**, images of all types are increasingly being converted to or initially acquired in digital form. This form is more or less the same across all imaging modalities. It is therefore amenable to common image-processing methodologies for enhancement, analysis, display, and storage.

Because of the ubiquity of images in biomedicine, the increasing availability of images in digital form, the rise of high-powered computer hardware and networks, and the commonality of image-processing solutions, digital images have become a core data type that must be considered in many biomedical informatics applications. Therefore, this chapter is devoted to a basic understanding of this core data type and many of the image-processing operations that can be applied to it. Chapter 18, on the other hand, describes the integration of images and image processing in various applications, particularly those in radiology since radiology places the greatest demands on imaging methods.

The topics covered by this chapter and Chapter 18 are generally part of biomedical imaging informatics (Kulikowski, 1997), a subfield of biomedical informatics that has arisen in recognition of the common issues that pertain to all image modalities and applications once the images are converted to digital form. By trying to understand these common issues, we can develop general solutions that can be applied to all images, regardless of the source.

The common tasks addressed by imaging informatics can be roughly classified as **image generation, image manipulation, image management**, and **image integration**. Image generation is the process of generating the images and converting them to digital form if they are not intrinsically digital. Image manipulation uses preprocessing and post-processing methods to enhance, visualize, or analyze the images. Image management includes methods for storing, transmitting, displaying, retrieving, and organizing images. Image integration is the combination of images with other information needed

for interpretation, management, and other tasks. Because radiology places the greatest demand on image management, and because radiology represents the primary application of imaging methods, Chapter 18 is primarily concerned with the latter two tasks whereas this chapter concentrates on the former two.

A major purpose of image processing is to extract information about the structure of the body. As such, imaging informatics overlaps **structural informatics**, which is the study of methods for representing, organizing, and managing diverse sources of information about the physical organization of the body and other physical structures, both for its own sake, and as a means for organizing other information (Brinkley, 1991). Many of the topics in this chapter therefore have to do with how to represent, extract, and characterize the anatomic information that is present in images.

The examples for this chapter, particularly those for three-dimensional and functional imaging, are primarily taken from brain imaging, which is part of the growing field of **neuroinformatics** (Koslow and Huerta, 1997). We choose brain imaging because: (1) brain imaging is a strong area of interest of one of the authors (JB), (2) the national Human Brain Project (HBP) (Human Brain Project, 2003) is generating substantial results in the area of brain imaging, (3) a large portion of current medical imaging work is in brain imaging, and (4) some of the most advanced image-related work in informatics is currently being done in this area. Thus, in addition to introducing the concepts of digital images and image processing, this chapter represents an intersection of many of the concepts in imaging informatics, structural informatics, and neuroinformatics.

We first introduce basic concepts of digital images, and then describe methods for imaging the structure of the body in both two dimensions and three dimensions. We then describe two-dimensional and three-dimensional methods for processing structural images, primarily as a means for visualizing, extracting, and characterizing anatomy. The chapter ends with a discussion of methods for imaging the function of the body, virtually all of which involve mapping or registering the functional data onto the structural representations extracted using the techniques described in earlier sections.

9.2 Basic Concepts

9.2.1 Digital Images

A **digital image** typically is represented in a computer by a two-dimensional array of numbers (a **bit map**). Each element of the array represents the intensity of a small square area of the picture, called a **pixel**. If we consider the image of a volume, then a three-dimensional array of numbers is required; each element of the array in this case represents a volume element, called a **voxel**.

We can store any image in a computer in this manner, either by converting it from an analog to a digital representation or by generating it directly in digital form. Once an image is in digital form, it can be handled just like all other data. It can be transmitted over communications networks, stored compactly in databases on magnetic or optical media, and displayed on graphics monitors. In addition, the use of computers has created an entirely new realm of capabilities for image generation and analysis; images

can be computed rather than measured directly. Furthermore, digital images can be manipulated for display or analysis in ways not possible with film-based images.

9.2.2 Imaging Parameters

All images can be characterized by several parameters of image quality. The most useful of these parameters are spatial resolution, contrast resolution, and temporal resolution. These parameters have been widely used to characterize traditional X-ray images; they also provide an objective means for comparing images formed by digital imaging modalities.

- **Spatial resolution** is related to the sharpness of the image; it is a measure of how well the imaging modality can distinguish points on the object that are close together. For a digital image, spatial resolution is generally related to the number of pixels per image area.
- **Contrast resolution** is a measure of the ability to distinguish small differences in intensity, which in turn are related to differences in measurable parameters such as X-ray attenuation. For digital images, the number of bits per pixel is related to the contrast resolution of an image.
- **Temporal resolution** is a measure of the time needed to create an image. We consider an imaging procedure to be a real-time application, if it can generate images concurrent with the physical process it is imaging. At a rate of at least 30 images per second, it is possible to produce unblurred images of the beating heart.

Other parameters that are specifically relevant to medical imaging are the degree of invasiveness, the dosage of ionizing radiation, the degree of patient discomfort, the size (portability) of the instrument, the ability to depict physiologic function as well as anatomic structure, and the availability and cost of the procedure at a specific location.

A perfect imaging modality would produce images with high spatial, contrast, and temporal resolution; it would be low in cost, portable, free of risk, painless, and noninvasive; it would use nonionizing radiation; and it would depict physiologic functions as well as anatomic structure.

9.3 Structural Imaging

Imaging the structure of the body has been and continues to be the major application of medical imaging, although, as described in Section 9.6, functional imaging is a very active area of research. The development of the various structural imaging modalities can be seen partly as a search for the perfect imaging modality; a primary reason for the proliferation of modalities is that no single modality satisfies all the desiderata. Another reason for the proliferation of image-generation methods is that progress has occurred in parallel in four main areas, and researchers have developed new methods quickly by combining elements from each of these areas. The four areas of development are energy source, reconstruction method, higher dimensionality, and contrast agents.

9.3.1 Energy Source

Light

The earliest medical images used **light** to create photographs, either of gross anatomic structures or, if a microscope was used, of histologic specimens. Light is still an important source for creation of images, and in fact optical imaging has seen a resurgence of late for areas such as molecular imaging (Weissleder and Mahmood, 2001) and imaging of brain activity on the exposed surface of the cerebral cortex (Pouratian et al., 2003). Visible light, however, does not allow us to see more than a short distance beneath the surface of the body.

X-Rays

X-rays were first discovered in 1895 by Wilhelm Conrad Roentgen, who was awarded the 1901 Nobel Prize in Physics for this achievement. The discovery caused worldwide excitement, especially in the field of medicine; by 1900, there were already several medical radiological societies. Thus, the foundation was laid for a new branch of medicine devoted to imaging the structure and function of the body (Kevles, 1997).

Film-based **radiography** is the primary modality used in radiology departments today, although this emphasis is changing rapidly as digital or **computed radiography (CR)** services are installed. We produce a typical X-ray image by projecting an X-ray beam—one form of ionizing radiation—from an X-ray source through a patient's body (or other object) and onto an X-ray-sensitive film. Because an X-ray beam is differentially absorbed by the various body tissues, the X-rays produce shadows on the radiographic film. The resultant **shadowgraph** is a superposition of all the structures traversed by each beam. **Digital radiography (DR)** applies the same techniques, but nonfilm detectors are used. In a technique known as CR, a latent image is recorded on a specially coated cassette that is scanned by a computer to capture the image in digital form; in other techniques, detectors capture the data directly in digital form. Although the images obtained by these techniques may be printed subsequently on film, they do not need to be.

Both film and fluoroscopic screens were used initially for recording X-ray images, but the fluoroscopic images were too faint to be used clinically. By the 1940s, however, television and image-intensifier technology were used to produce clear real-time fluorescent images. Today, a standard procedure for many types of examinations is to combine real-time television monitoring of X-ray images with the creation of selected higher resolution film images. Until the early 1970s, film and **fluoroscopy** were the only X-ray modalities available.

Traditional X-ray images have high spatial resolution and medium cost. Furthermore, they can be generated in real time (fluoroscopy) and can be produced using portable instruments. Their limitations are their relatively poor contrast resolution, their use of ionizing radiation, and their inability to depict physiologic function. Alternate imaging principles have been applied to increase contrast resolution, to eliminate exposure to X-ray radiation, and so on. For example, in nuclear-medicine imaging, a **radioactive isotope** is chemically attached to a biologically active compound (such as iodine) and then is injected into the patient's peripheral circulation. The compound collects in the specific

body compartments or organs (such as the thyroid), where it is stored or processed by the body. The isotope emits radiation locally, and the radiation is measured using a special detector. The resultant nuclear-medicine image depicts the level of radioactivity that was measured at each point. Because the counts are inherently digital, computers have been used to record them. Multiple images also can be processed to obtain dynamic information, such as the rate of arrival or of disappearance of isotope at particular body sites.

Ultrasound

Another common energy source is **ultrasound** (echosonography), which developed out of research performed by the Navy during World War II. **Ultrasonography** uses pulses of high-frequency sound waves rather than ionizing radiation to image body structures. As each sound wave encounters tissues in a patient's body, a portion of the wave is reflected and a portion continues. The time required for the echo to return is proportional to the distance into the body at which it is reflected; the amplitude (intensity) of a returning echo depends on the acoustical properties of the tissues encountered and is represented in the image as brightness. The system constructs two-dimensional images by displaying the echoes from pulses of multiple adjacent one-dimensional paths. Such images can be stored in digital memories or recorded on videotape and then displayed as television (raster-display) images.

Nuclear Magnetic Resonance

Creation of images from **magnetism** grew out of **nuclear magnetic resonance (NMR) spectroscopy**, a technique that has long been used in chemistry to characterize chemical compounds. Many atomic nuclei within the body have a net magnetic moment, so they act like tiny magnets. When a small chemical sample is placed in an intense, uniform magnetic field, these nuclei line up in the direction of the field, spinning around the axis of the field with a frequency dependent on the type of nucleus, on the surrounding environment, and on the strength of the magnetic field.

If a radio pulse of a particular frequency is applied at right angles to the stationary magnetic field, those nuclei with rotation frequency equal to that of the radiofrequency pulse resonate with the pulse and absorb energy. The higher energy state causes the nuclei to change their orientation with respect to the fixed magnetic field. When the radiofrequency pulse is removed, the nuclei return to their original aligned state, emitting a detectable radiofrequency signal as they do so. Characteristic parameters of this signal—such as intensity, duration, and frequency shift away from the original pulse—are dependent on the density and environment of the nuclei.

In the case of traditional NMR spectroscopy, different molecular environments cause different frequency shifts (called chemical shifts), which we can use to identify the particular compounds in a sample. In the original NMR method, however, the signal is not localized to a specific region of the sample, so it is not possible to create an image. Creation of images from NMR signals known as MRI had to await the development of computer-based reconstruction techniques, which represent one of the most spectacular applications of computers in medicine.

9.3.2 Reconstruction Methods

Reconstruction techniques were first applied to X-ray images aimed at addressing the problem of superposition of structures in standard projection imaging. An X-ray image at a given point represents the total attenuation due to all the overlaid structures traversed by a beam as that beam passes through the body; shadows cast by surrounding structures may obscure the object that the clinician wishes to visualize. **Contrast radiography**—the use of radiopaque contrast material to highlight the areas of interest (e.g., stomach, colon, urinary tract)—was used as early as 1902 to address this problem. The first clinical experiments with **angiography**—imaging of blood vessels performed by the injection of opacifying agents into the bloodstream—were conducted in 1923.

The desire to separate superimposed structures also led to the development of a variety of analog tomographic techniques. In these methods, the X-ray source and detector were moved in opposite arcs, thereby causing a thin tomographic (planar) section to remain in focus while other planes were blurred. This method, however, exposes the patient to a relatively high X-ray dose because the blurred areas are exposed continuously.

Mathematical methods for reconstructing images from projections were first developed by Radon in 1917 and later were improved by other researchers. These methods were used in the 1950s and 1960s to solve scientific problems in many fields, including radio astronomy and electron microscopy. In the late 1960s, Cormack used the techniques to reconstruct phantoms (objects with known shape) using X-rays. In the early 1970s, Hounsfield led a team at the London-based EMI Corporation, which developed the first commercially viable computed tomography (CT) scanner.

Instead of depicting a directly measurable parameter (the absorption of X-ray beams as they pass through the body), CT mathematically reconstructs an image from X-ray-attenuation values that have been measured from multiple angles. As a result, it is possible to view cross-sectional slices through the body rather than two-dimensional projections of superimposed structures. Thus, CT images provide a precise mapping of the internal structures of the body in three-dimensional space—a function not provided by standard X-ray images. They also greatly improve contrast resolution.

In the basic CT imaging technique, the patient is placed between an X-ray-sensitive detector and an X-ray source that produces a collimated (pencil-like) beam. The measured difference between the source and detector X-ray intensities represents the amount of X-ray attenuation due to the tissues traversed by the beam; this measured attenuation is a superposition, or **projection**, of the attenuations of all the individual tissue elements traversed by the beam. In the simplest reconstruction method, called **back-projection**, the measured intensity is distributed uniformly over all the pixels traversed by the beam. For example, if the measured attenuation is 20, and 10 pixels were traversed, then the CT number of each of the 10 pixels is incremented by 2 units.

The attenuation measured from a single projection is not sufficient to reconstruct an image. The same back-projection computation, however, can be applied to the attenuations measured from multiple projections. The source and detector are rotated about the patient, and the X-ray attenuation is measured along each path. Because each pixel is traversed by multiple projection paths, its computed attenuation is the sum of the

contributions from each path. The total sum provides a reasonable first approximation of the X-ray attenuation of the individual pixel. The image is further refined using a mathematical edge-enhancement technique called **convolution**. In effect, convolution removes shadows that result from the back projection, thus sharpening the blurry image.

The development of the CT scanner dramatically improved our ability to visualize adjacent structures; for the first time, physicians were able to see inside a living human being clearly, but noninvasively. This ability led to a revolution in medicine almost as great as the one occasioned by the invention of X-ray imaging. As a result, Cormack and Hounsfield were awarded the 1979 Nobel Prize in Medicine.

After the invention of the CT scanner, this basic method of reconstruction from projections was applied to other energy sources, including magnetism (MRI), ultrasound (ultrasound-transmission tomography), and variants of nuclear-medicine imaging called **positron-emission tomography (PET)** and single-photon-emission computed tomography (SPECT).

The most dramatic example of reconstruction from projections other than CT is MRI, which is based on NMR (Oldendorf, 1991). As described in the previous section, NMR takes advantage of magnetic properties of nuclei to characterize the distribution and chemical environment of nuclei within a chemical sample. To create an image using these parameters, we need a way to restrict this sample to a small volume within a larger tissue. With this restriction, the parameters of the NMR signal from each small tissue volume can be mapped to voxel intensities depicting different tissue characteristics.

The restriction to a small sample volume is accomplished by taking advantage of the fact that the resonant frequency of atomic nuclei varies with the magnetic field. If the field can be made different for each small tissue volume, then a radiofrequency pulse with a given frequency will excite only those nuclei in the small volume that have the resonant frequency of that pulse. The basic method uses electromagnetic coils to superimpose a varying magnetic field on a large fixed magnetic field, thereby setting up a gradient in the magnetic field.

This gradient is changed electronically, setting the location of the sample volume. For example, we use one gradient to set the plane of section (the z direction, although the orientation of this section may be arbitrary with respect to the patient), and a second gradient sets a line within a single section (the x,y plane). As in CT, the signal detected along this line is a summation of the signals from all voxels along the line. Therefore, the x,y gradient is electronically rotated, rotating the plane of section and generating additional lines within a given plane. The same reconstruction techniques developed for CT then reconstruct the values for the individual voxels within the given plane. Because there are many different parameters that can be measured for each sampled voxel, many different types of images can be constructed, many of which are still being developed.

9.3.3 Higher Dimensionality

Most routine images in radiology are still two-dimensional. Because the body is a three-dimensional object that changes over time, however, there will always be a drive to create three-dimensional time-varying images. In recent years, advances in digital hardware

have provided the storage and throughput to manage large time-varying voxel-based data sets. Reconstruction modalities—such as CT, PET, and MRI—all are either inherently three-dimensional or can be made three-dimensional by acquisition of a series of closely spaced parallel slices (see Section 9.5). Thus, the only drawbacks of these techniques are the time and expense required to acquire a series of parallel slices, both of which are becoming smaller.

Ultrasound images, on the other hand, cannot be acquired as parallel slices because sound does not pass through bone or air. For this reason, we usually obtain three-dimensional ultrasound information by attaching a three-dimensional locating device to the transducer. The locator gives the position and orientation of the slice plane in space. Before the availability of hardware that could store large numbers of volume data, the ultrasound images were first processed in two dimensions to extract relevant anatomy as two-dimensional contours or regions; the two-dimensional contours were converted to three-dimensional contours based on the location information and then were displayed with vector graphics (Brinkley et al., 1978). Such an approach was useful for quantitation, but did not provide a realistic three-dimensional view of the object.

9.3.4 Contrast Agents

As noted above, one of the major motivators for development of new imaging modalities is the desire to increase contrast resolution. We have already discussed the use of radiologic contrast agents and reconstruction techniques as examples of highly successful attempts to increase contrast resolution among the different energy sources. In addition, histologic staining agents such as hematoxylin and eosin (H&E) have been used for years to enhance contrast in tissue sections, and magnetic contrast agents such as gadolinium have been introduced to enhance contrast in MR images.

Although these methods have been very successful, they generally are somewhat nonspecific. In recent years, advances in molecular biology have led to the ability to design contrast agents that are highly specific for individual molecules. In addition to radioactively tagged molecules used in nuclear medicine, molecules are tagged for imaging by magnetic resonance and optical energy sources. Tagged molecules are imaged in two dimensions or three dimensions, often by application of reconstruction techniques developed for clinical imaging. Tagged molecules have been used for several years *in vitro* by such techniques as immunocytochemistry (binding of tagged antibodies to antigen) (Van Noorden, 2002) and *in situ* hybridization (binding of tagged nucleotide sequences to DNA or RNA) (King et al., 2000). More recently, methods have been developed to image these molecules in the living organism, thereby opening up entirely new avenues for understanding the functioning of the body at the molecular level.

9.3.5 New and Emerging Structural Imaging Methods

Many new imaging techniques have been developed in recent years. Most of these techniques can be seen as a combination of an energy source, a computer-based processing or reconstruction technique, increased dimensionality due to advances in digital

hardware, and, increasingly, use of molecular contrast agents. The remainder of this section describes a few examples of these techniques.

At the gross anatomic level, **charge-coupled device (CCD) cameras** can be used to convert existing film-based equipment to units that can produce images in digital form. Storage phosphor, or CR, systems replace film by substituting a reusable phosphor plate in a standard film cassette. The exposed plate is processed by a reader system that scans the image into digital form, erases the plate, and packages the cassette for reuse. An important advantage of CR systems is that the cassettes are of standard size, so they can be used in any equipment that holds film-based cassettes (Horii, 1996). More recently, CR uses CCD arrays to capture the image directly.

Many new modalities are being developed based on magnetic resonance. For example, *magnetic resonance arteriography* (MRA) and *magnetic resonance venography* (MRV) image blood flow (Lee, 2003), and *diffusion tensor imaging* (DTI) is increasingly being used to image white matter fiber tracts in the brain (Figure 9.1) (Le Bihan et al., 2001).

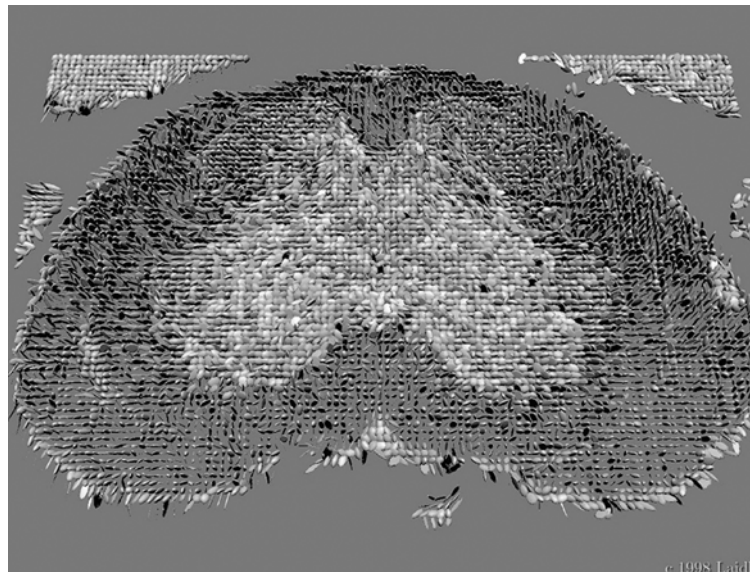

Figure 9.1. Diffusion tensor image (DTI) of the mouse spinal cord. At each pixel the DTI technique outputs a 3×3 diffusion tensor describing the measured diffusion of water in each of the six principle directions in three-dimensional space. In gray matter, the diffusion is generally uniform (isotropic) in all directions, but in white matter diffusion is reduced in the direction perpendicular to the fibers. Thus, DTI is used to visualize white matter fiber tracts. Since each pixel (or voxel in three-dimensions) is described by a 3×3 matrix, visualization techniques from computer graphics are needed in order to represent the information at each pixel. In this figure, the diffusion tensors are represented by ellipsoids with axes along the principle directions of the diffusion tensor. Photograph courtesy of David Laidlaw (Ahrens et al., 1998), http://www.gg.caltech.edu/~dhl/images.html.

Ultrasound machines have essentially become specialized computers with attached peripherals, with active development of three-dimensional imaging. The ultrasound transducer now often sweeps out a three-dimensional volume rather than a two-dimensional plane, and the data are written directly into a three-dimensional array memory, which is displayed using volume or **surface-based rendering** techniques (Figure 9.2) (Ritchie et al., 1996).

At the microscopic level, the *confocal microscope* uses electronic focusing to move a two-dimensional slice plane through a three-dimensional tissue slice placed in a microscope. The result is a three-dimensional voxel array of a microscopic, or even submicroscopic, specimen (Wilson, 1990; Paddock, 1994). At the electron microscopic level *electron tomography* generates three-dimensional images from thick electron-microscopic sections using techniques similar to those used in CT (Perkins et al., 1997).

At the molecular level tagged molecules are increasingly introduced into the living organism, and imaged with optical, radioactive or magnetic energy sources, often using reconstruction techniques and often in three dimensions. The combination of these various methods with highly specific tagged molecules has given rise to the field of **molecular imaging** (Weissleder and Mahmood, 2001; Massoud and Gambhir, 2003), which in addition to functional brain imaging (Section 9.6) represents some of the most exciting new developments in biomedical imaging. It is now becoming possible to combine gene sequence information, gene expression array data, and molecular imaging to determine not only which genes are expressed but also where they are expressed in the organism. These capabilities will become increasingly important in the **post-genomic era** for determining exactly how genes generate both the structure and function of the organism.

9.4 Two-Dimensional Image Processing

The rapidly increasing number and types of digital images has created many opportunities for image processing, since one of the great advantages of digital images is that

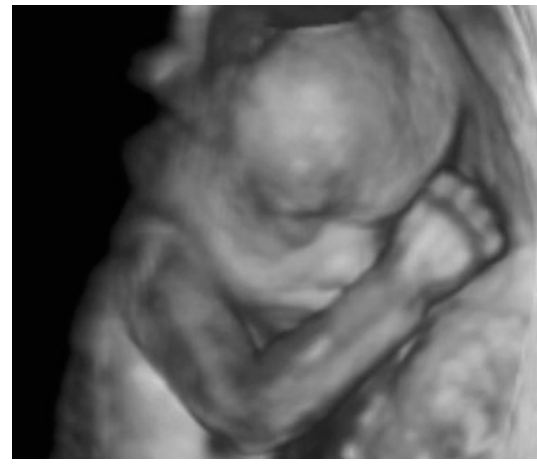

Figure 9.2. Three-dimensional ultrasound image of a fetus, *in utero*. The ultrasound probe sweeps out a three-dimensional volume rather than the conventional two-dimensional plane. The volume can be rendered directly using volume-rendering techniques, or as in this case, fetal surfaces can be extracted and rendered using surface-rendering techniques. Photograph courtesy of Per Perm, "GE Healthcare", http://www.gemedicalsystems.com/rad/us/education/msucme3d.html.

they can be manipulated just like any other kind of data. This advantage was evident from the early days of computers, and success in processing satellite and spacecraft images generated considerable interest in biomedical image processing, including automated image analysis for interpretation. Beginning in the 1960s, researchers devoted a large amount of work to this end, with the hope that eventually much of radiographic image analysis could be automated.

One of the first areas to receive attention was automated interpretation of chest X-ray images, because, previously, most patients admitted to a hospital were subjected to routine chest X-ray examinations. (This practice is no longer considered cost effective except for selected subgroups of patients.) Subsequent research, however, confirmed the difficulty of completely automating radiographic image interpretation, and much of the initial enthusiasm has long ago worn off. Currently, there is less emphasis on completely automatic interpretation and more on systems that aid the user, except in specialized areas such as brain imaging.

9.4.1 Basic Concepts in Two-Dimensional Image Processing

Digital image manipulation, or **image processing**, generally involves the transformation of one or more input images either into one or more output images or into some abstract representation of the contents of the input images. For example, the intensity values can be modified to improve contrast resolution, or a set of terms (*pleural effusion, lung nodule*) can be attached to specific regions of interest.

Images can be enhanced to permit human viewing, to show views not present in the original images, to flag suspicious areas for closer examination by the clinician, to quantify the size and shape of an organ, and to prepare the images for integration with other information. Most of these applications require one or more of the four basic image-processing steps: global processing, segmentation, feature detection, and classification. These steps are generally performed in order, although later steps may feed back to earlier ones, and not all steps are required for each application. Most steps generalize from two-dimensional to three-dimensional images, but three-dimensional images give rise to additional image-processing opportunities and challenges that are discussed in Section 9.5.

Global processing involves computations on the entire image, without regard to specific local content. The purpose is to enhance an image for human visualization or for further analysis by the computer. A simple but important example is *gray-scale windowing* of CT images. The CT scanner generates pixel values (Hounsfield numbers, or CT numbers) in the range of −1,000 to +3,000. Humans, however, cannot distinguish more than about 100 shades of gray. To appreciate the full precision available with a CT image, the operator can adjust the midpoint and the range of the displayed CT values. By changing the level and width (i.e., intercept and slope of the mapping between pixel value and displayed gray scale or, roughly, the brightness and contrast) of the display, radiologists enhance their ability to perceive small changes in contrast resolution within a subregion of interest.

Segmentation involves the extraction of regions of interest (ROIs) from the overall image. The ROIs usually correspond to anatomically meaningful structures, such as organs or parts of organs. The structures may be delineated by their borders, in which

case **edge-detection techniques** (such as edge-following algorithms) are used, or by their composition on the image, in which case **region-detection techniques** (such as texture analysis) are used (Haralick and Shapiro, 1992). Neither of these techniques has been completely successful; regions often have discontinuous borders or nondistinctive internal composition. Furthermore, contiguous regions often overlap. These and other complications make segmentation the most difficult subtask of the medical image-analysis problem. Because segmentation is difficult for a computer, it is often performed manually by a human operator or through a combination of automated and operator-interactive approaches. It therefore remains a major bottleneck that prevents more widespread application of image-processing techniques.

Feature detection is the process of extracting useful parameters from the segmented regions. These parameters may themselves be informative—for example, the volume of the heart or the size of the fetus. They also may be used as input into an automated **classification** procedure, which determines the type of object found. For example, small round regions on chest X-ray images might be classified as tumors, depending on such features as intensity, perimeter, and area.

Mathematical models often are used to aid in the performance of image-analysis subtasks. In classic pattern-recognition applications, the subtasks of global processing, segmentation, feature detection, and classification usually are performed sequentially. People, however, appear to perform pattern-recognition iteratively. For example, radiologists can perceive faint images and can trace discontinuous borders, in part because they know which features they are searching for. Many researchers have applied artificial intelligence techniques to imitate such interaction among subtasks. The computer is programmed with some of the higher-level anatomic knowledge that radiologists use when they interpret images. Thus, high-level organ models provide feedback to guide the lower-level process of segmentation.

The nature of the application determines which of these subtasks is performed, the choice of technique for each subtask, and the relative order of the subtasks. Because image understanding is an unsolved problem, and because many applications are possible, there is a wealth of image-processing techniques that can be applied to digital images.

9.4.2 Examples of Two-Dimensional Image Processing

Although completely automated image-analysis systems are still in the future, the widespread availability of digital images, combined with image management systems such as picture archiving and communication systems (PACS; Chapter 18) and powerful workstations, has led to many applications of image-processing techniques. In general, routine techniques are available on the manufacturer's workstations (e.g., an MR console or an ultrasound machine), whereas more advanced image-processing algorithms are available as software packages that run on independent workstations.

The primary uses of two-dimensional image processing in the clinical environment are for image enhancement, screening, and quantitation. Software for such image processing is primarily developed for use on independent workstations. Several journals are devoted to medical image processing (e.g., *IEEE Transactions on Medical Imaging,*

Journal of Digital Imaging, Neuroimage), and the number of journal articles is rapidly increasing as digital images become more widely available. We describe just a few examples of image-processing techniques in the remainder of this section.

Image enhancement uses global processing to improve the appearance of the image either for human use or for subsequent processing by computer. All manufacturers' consoles and independent image-processing workstations provide some form of image enhancement. We have already mentioned CT windowing. Another technique is **unsharp masking**, in which a blurred image is subtracted from the original image to increase local contrast and to enhance the visibility of fine-detail (high-frequency) structures. **Histogram equalization** spreads the image gray levels throughout the visible range to maximize the visibility of those gray levels that are used frequently. **Temporal subtraction** subtracts a reference image from later images that are registered to the first. A common use of temporal subtraction is digital-subtraction angiography (DSA) in which a background image is subtracted from an image taken after the injection of contrast material.

Screening uses global processing, segmentation, feature detection, and classification to determine whether an image should be flagged for careful review by a radiologist or pathologist. In such an approach, the computer is allowed to flag a reasonable number of normal images (false positives) as long as it misses very few abnormal images (false negatives). If the number of flagged images is small compared with the total number of images, then automated screening procedures can be economically viable. Screening techniques have been applied successfully to mammography images for identifying mass lesions and clusters of microcalcifications, to chest X-rays for small cancerous nodules, and to Papanicolaou (Pap) smears for cancerous or precancerous cells (Giger and MacMahon, 1996), as well as to many other images (Figure 9.3; see also Color Plate I).

Quantitation uses global processing and segmentation to characterize meaningful regions of interest. For example, heart size, shape, and motion are subtle indicators of heart function and of the response of the heart to therapy (Clarysse et al., 1997). Similarly, fetal head size and femur length, as measured on ultrasound images, are valuable indicators of fetal well-being (Brinkley, 1993b). Although the literature describes a wealth of automatic or semiautomatic techniques for segmenting images of the heart or of the fetus, the most common clinical scenario continues to be manual outlining by trained technicians. This situation should change, however, as semiautomatic techniques (those that let the user correct segmentation errors by the computer) become widely available on independent workstations that are custom-tailored for particular applications.

9.5 Three-Dimensional Image Processing

The growing availability of three-dimensional and higher dimensionality structural and functional images (Section 9.3.3) leads to exciting opportunities for realistically observing the structure and function of the body. Nowhere have these opportunities been more widely exploited than in brain imaging. Therefore, this section concentrates on three-dimensional brain imaging, with the recognition that many of the methods developed for the brain have been or will be applied to other areas as well.

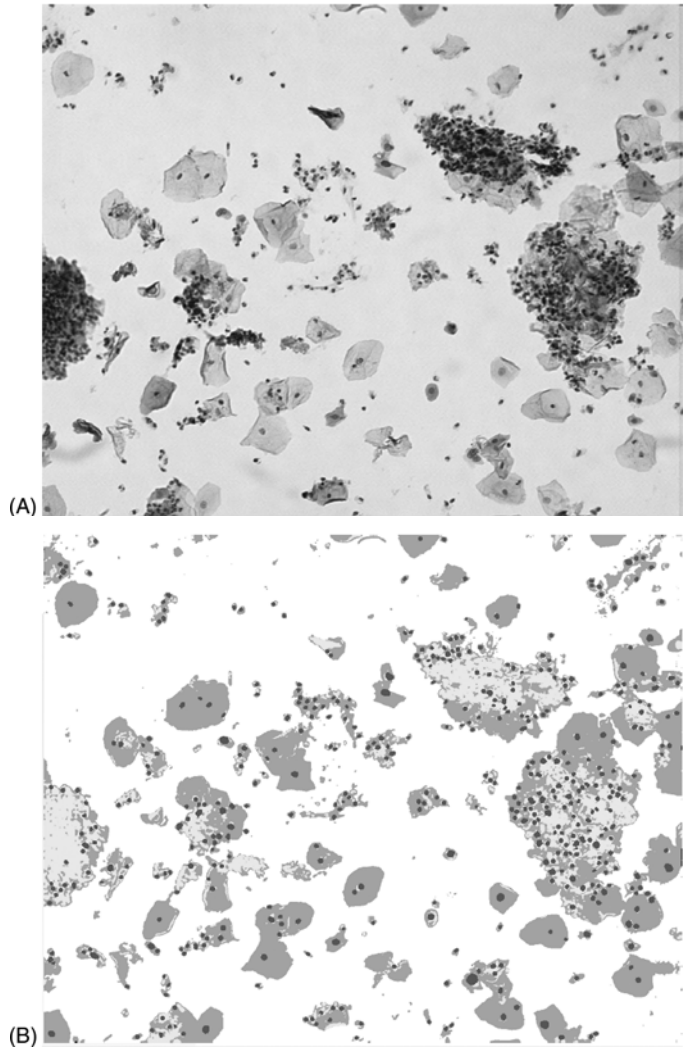

Figure 9.3. Automated screening of Papanicolaou (PAP) smears. Since large numbers of PAP smears are acquired routinely, there is a need to reduce the cost and potential errors associated with purely manual interpretation. (A) Raw microscopic image of cervical cells. (B) Segmented image. The program has segmented the cells and nuclei from the raw image, prior to feature detection and classification based on the features. Abnormally classified cells are flagged for review by the human operator. Photograph courtesy of Peter Locht, Visible Diagnostics, http://www.imm.dtu.dk/visiondag/VD03/medicinsk/pl.pdf.

The basic two-dimensional image-processing operations of global processing, segmentation, feature detection, and classification generalize to higher dimensions, and are usually part of any image-processing application. However, three-dimensional and higher dimensionality images give rise to additional informatics issues, which include

image **registration** (which also occurs to a lesser extent in two dimensions), *spatial* representation of anatomy, *symbolic* representation of anatomy, integration of spatial and symbolic anatomic representations in *atlases*, anatomic *variation*, and *characterization* of anatomy. All but the first of these issues deal primarily with anatomic structure, and therefore could be considered part of the field of structural informatics. They could also be thought of as being part of imaging informatics and neuroinformatics.

The following sections discuss these additional informatics issues.

9.5.1 Registration

As noted previously, three-dimensional image volume data are represented in the computer by a three-dimensional volume array, in which each *voxel* (volume element, analogous to the pixel in two dimensions) represents the image intensity in a small volume of space. In order to accurately depict anatomy, the voxels must be accurately registered (or located) in the three-dimensional volume (*voxel registration*), and separately acquired image volumes from the same subject must be registered with each other (*volume registration*).

Voxel Registration

Technologies such as CT, MRI, MRV, MRA, and confocal microscopy (Section 9.3) are inherently three-dimensional; the scanner generally outputs a series of image slices that can easily be reformatted as a three-dimensional volume array, often following alignment algorithms that compensate for any patient motion during the scanning procedure. For this reason, almost all CT and MR manufacturers' consoles contain some form of three-dimensional reconstruction and visualization capabilities.

As noted in Section 9.3.3, two-dimensional images can be converted to three-dimensional volumes by acquiring a set of closely spaced parallel sections through a tissue or whole specimen. In this case the problem is how to align the sections with each other. For whole sections (either frozen or fixed), the standard method is to embed a set of thin rods or strings in the tissue prior to sectioning, to manually indicate the location of these *fiducials* on each section, then to linearly transform each slice so that the corresponding fiducials line up in three dimensions (Prothero and Prothero, 1986). A popular current example of this technique is the Visible Human, in which a series of transverse slices were acquired, then reconstructed to give a full three-dimensional volume (Spitzer and Whitlock, 1998).

It is difficult to embed fiducial markers at the microscopic level, so intrinsic tissue landmarks are often used as fiducials, but the basic principle is similar. However, in this case tissue distortion may be a problem, so nonlinear transformations may be required. For example, Fiala and Harris (2001) have developed an interface that allows the user to indicate, on electron-microscopy sections, corresponding centers of small organelles such as mitochondria. A nonlinear transformation (warp) is then computed to bring the landmarks into registration.

An approach being pursued (among other approaches) by the National Center for Microscopy and Imaging Research (http://ncmir.ucsd.edu/) combines reconstruction

from thick serial sections with electron tomography (Soto et al., 1994). In this case the tomographic technique is applied to each thick section to generate a three-dimensional digital slab, after which the slabs are aligned with each other to generate a three-dimensional volume. The advantages of this approach over the standard serial section method are that the sections do not need to be as thin, and fewer of them need be acquired.

An alternative approach to three-dimensional voxel registration from two-dimensional images is stereo-matching, a technique developed in computer vision that acquires multiple two-dimensional images from known angles, which finds corresponding points on the images, and uses the correspondences and known camera angles to compute three-dimensional coordinates of pixels in the matched images. The technique is being applied to the reconstruction of synapses from electron micrographs by a HBP collaboration between computer scientists and biologists at the University of Maryland (Agrawal et al., 2000).

Volume Registration

A related problem to that of aligning individual sections is the problem of aligning separate image volumes from the same subject, i.e., *intrasubject* alignment. Because different image modalities provide complementary information, it is common to acquire more than one kind of image volume on the same individual. This approach has been particularly useful for brain imaging because each modality provides different information. For example, PET (Section 9.3.2) provides useful information about function, but does not provide good localization with respect to anatomy. Similarly, MRV and MRA (Section 9.3.5) show blood flow but do not provide the detailed anatomy visible with standard MRI. By combining images from these modalities with MRI, we can show functional images in terms of the underlying anatomy, thereby providing a common neuroanatomic framework.

In our own (JB) HBP work, we acquire an MRI volume dataset depicting cortical anatomy, an MRV volume depicting veins, and an MRA volume depicting arteries (Modayur et al., 1997; Hinshaw et al., 2002). By "fusing" these separate modalities into a single common frame of reference (anatomy, as given by the MRI dataset), it is possible to gain information that is not apparent from one of the modalities alone. In our case the fused datasets are used to generate a visualization of the brain surface as it appears at neurosurgery, in which the veins and arteries provide prominent landmarks (Figure 9.9).

The primary problem to solve in **multimodality image fusion** is volume registration—that is, the alignment of separately acquired image volumes. In the simplest case, separate image volumes are acquired during a single sitting. The patient's head may be immobilized, and the information in the image headers may be used to rotate and resample the image volumes until all the voxels correspond.

However, if the patient moves, or if examinations are acquired at different times, other registration methods are needed. When intensity values are similar across modalities, registration can be performed automatically by intensity-based optimization methods (Woods et al., 1992; Collins et al., 1994). When intensity values are not similar (as is the case with MRA, MRV, and MRI), images can be aligned to templates of the same modalities that are already aligned (Woods et al., 1993; Ashburner and Friston,

1997). Alternatively, landmark-based methods can be used. The landmark-based methods are similar to those used to align serial sections, but in this case the landmarks are three-dimensional points. The Montreal Register Program (MacDonald, 1993) (which can also do nonlinear registration, as discussed in Section 9.5.5) is an example of such a program.

9.5.2 Spatial Representation of Anatomy

The reconstructed and registered three-dimensional image volumes can be visualized directly using **volume rendering** techniques (Foley et al., 1990; Lichtenbelt et al., 1998) (Figure 9.2) which project a two-dimensional image directly from a three-dimensional voxel array by casting rays from the eye of the observer through the volume array to the image plane. Because each ray passes through many voxels, some form of segmentation (usually simple thresholding) often is used to remove obscuring structures. As workstation memory and processing power have advanced, volume rendering has become widely used to display all sorts of three-dimensional voxel data—ranging from cell images produced by confocal microscopy to three-dimensional ultrasound images, or to brain images created from MRI or PET.

Volume images can also be given as input to image-based techniques for warping the image volume of one structure to other, as described in Section 9.5.5. However, more commonly the image volume is processed in order to extract an explicit *spatial* (or quantitative) representation of anatomy. Such an explicit representation permits improved visualization, quantitative analysis of structure, comparison of anatomy across a population, and mapping of functional data. It is thus a component of most research involving three-dimensional image processing.

Extraction of spatial representations of anatomy, in the form of three-dimensional surfaces or volume regions, is accomplished by a three-dimensional generalization of the segmentation techniques discussed in Section 9.4.1. As in the two-dimensional case, fully automated segmentation is an unsolved problem, as attested to by the number of papers about this subject in *IEEE Transactions on Medical Imaging*. However, because of the high quality of MRI brain images, a great deal of progress has been made in recent years for brain imaging in particular; in fact, several software packages do a credible job of automatic segmentation, particularly for normal macroscopic brain anatomy in cortical and subcortical regions (Collins et al., 1995; Friston et al., 1995; Subramaniam et al., 1997; Dale et al., 1999; MacDonald et al., 2000; Brain Innovation B.V., 2001; FMRIDB Image Analysis Group, 2001; Van Essen et al., 2001; Hinshaw et al., 2002). The HBP-funded Internet Brain Segmentation Repository (Kennedy, 2001) is developing a repository of segmented brain images to use in comparing these different methods.

Popular segmentation and reconstruction techniques include reconstruction from serial sections, region-based methods, edge-based methods, model- or knowledge-based methods, and combined methods.

Reconstruction from Serial Sections

The classic approach to extracting anatomy is to manually or semiautomatically trace the contours of structures of interest on each of a series of aligned image slices, then to

tile a surface over the contours (Prothero and Prothero, 1982). The tiled surface usually consists of an array of three-dimensional points connected to each other by edges to form triangular facets. The resulting three-dimensional *surface mesh* is then in a form where it can be further analyzed or displayed using standard three-dimensional surface rendering techniques such as those applied in the computer-generated film industry (Foley, 2001).

Neither fully automatic contour tracing nor fully automatic tiling has been satisfactorily demonstrated in the general case. Thus, semiautomatic contour tracing followed by semiautomatic tiling remains the most common method for reconstruction from serial sections, and reconstruction from serial sections itself remains the method of choice for extracting microscopic three-dimensional brain anatomy (Fiala and Harris, 2001).

Region-Based and Edge-Based Segmentation

This and the following sections primarily concentrate on segmentation at the macroscopic level.

In region-based segmentation, voxels are grouped into contiguous regions based on characteristics such as intensity ranges and similarity to neighboring voxels (Shapiro and Stockman, 2001). A common initial approach to region-based segmentation is to first classify voxels into a small number of tissue classes such as gray matter, white matter, cerebrospinal fluid, and background, then to use these classifications as a basis for further segmentation (Choi et al., 1991; Zijdenbos et al., 1996). Another region-based approach is called region growing, in which regions are grown from seed voxels manually or automatically placed within candidate regions (Davatzikos and Bryan, 1996; Modayur et al., 1997). The regions found by any of these approaches are often further processed by mathematical morphology operators (Haralick, 1988) to remove unwanted connections and holes (Sandor and Leahy, 1997).

Edge-based segmentation is the complement to region-based segmentation; intensity gradients are used to search for and link organ boundaries. In the two-dimensional case, contour-following methods connects adjacent points on the boundary. In the three-dimensional case, isosurface-following or marching-cubes (Lorensen and Cline, 1987) methods connect border voxels in a region into a three-dimensional surface mesh.

Both region-based and edge-based segmentation are essentially low-level techniques that only look at local regions in the image data.

Model- and Knowledge-Based Segmentation

The most popular current method for medical image segmentation, for the brain as well as other biological structures, is the use of **deformable models**. Based on pioneering work called "Snakes" by Kass et al. (1987), deformable models have been developed for both two dimensions and three dimensions. In the two-dimensional case the deformable model is a contour, often represented as a simple set of linear segments or a *spline*, which is initialized to approximate the contour on the image. The contour is then deformed according to a cost function that includes both intrinsic terms limiting how much the contour can distort, and extrinsic terms that reward closeness to image

borders. In the three-dimensional case, a three-dimensional surface (often a triangular mesh) is deformed in a similar manner. There are several examples of HBP-funded work that use deformable models for brain segmentation (Davatzikos and Bryan, 1996; Dale et al., 1999; MacDonald et al., 2000; Van Essen et al., 2001).

An advantage of deformable models is that the cost function can include knowledge of the expected anatomy of the brain. For example, the cost function employed in the method developed by MacDonald (MacDonald et al., 2000) includes a term for the expected thickness of the brain cortex. Thus, these methods can become somewhat knowledge-based, where knowledge of anatomy is encoded in the cost function.

An alternative knowledge-based approach explicitly records shape information in a geometric constraint network (GCN) (Brinkley, 1992), which encodes local shape variation based on a training set. The shape constraints define search regions on the image in which to search for edges. Found edges are then combined with the shape constraints to deform the model and reduce the size of search regions for additional edges (Brinkley, 1985, 1993a). One potential advantage of this sort of model over a pure deformable model is that knowledge is explicitly represented in the model, rather than implicitly represented in the cost function.

Combined Methods

Most brain segmentation packages use a combination of methods in a sequential pipeline. For example, in our own recent work (JB) we first use a GCN model to represent the overall cortical "envelope", excluding the detailed gyri and sulci (Hinshaw et al., 2002). The model is semiautomatically deformed to fit the cortex, then used as a mask to remove noncortex such as the skull. Isosurface following is then applied to the masked region to generate the detailed cortical surface. The model is also used on aligned MRA and MRV images to mask out noncortical veins and arteries prior to isosurface following. The extracted cortical, vein, and artery surfaces are then rendered to produce a composite visualization of the brain as seen at neurosurgery (Figure 9.9).

MacDonald et al. (2000) describe an automatic multiresolution surface deformation technique called anatomic segmentation using proximities (ASP), in which an inner and outer surface are progressively deformed to fit the image, where the cost function includes image terms, model-based terms, and proximity terms. Dale et al. (1999) describe an automated approach that is implemented in the FreeSurfer program (Fischl et al., 1999). This method initially finds the gray-white boundary, then fits smooth gray-white (inner) and white-CSF (outer) surfaces using deformable models. Van Essen et al. (2001) describe the SureFit program, which finds the cortical surface midway between the gray-white boundary and the gray-CSF boundary. This mid-level surface is created from probabilistic representations of both inner and outer boundaries that are determined using image intensity, intensity gradients, and knowledge of cortical topography. Other software packages also combine various methods for segmentation (Davatzikos and Bryan, 1996; Brain Innovation B.V., 2001; FMRIDB Image Analysis Group, 2001; Sensor Systems Inc., 2001; Wellcome Department of Cognitive Neurology, 2001).

9.5.3 Symbolic Representation of Anatomy

Given segmented anatomic structures, whether at the macroscopic or microscopic level, and whether represented as three-dimensional surface meshes or extracted three-dimensional regions, it is often desirable to attach labels (names) to the structures. If the names are drawn from a controlled terminology they can be used as an index into a database of segmented structures, thereby providing a qualitative means for comparing structures from multiple subjects.

If the terms in the vocabulary are organized into symbolic qualitative models (**ontologies**) of anatomic concepts and relationships, they can support systems that manipulate and retrieve segmented structures in "intelligent" ways. If the anatomic ontologies are linked to other ontologies of physiology and pathology, they can provide increasingly sophisticated knowledge about the *meaning* of the various images and other data that are increasingly becoming available in online databases. It is our belief that this kind of knowledge (by the computer, as opposed to the scientist) will be required in order to achieve the seamless integration of all forms of imaging and nonimaging data.

At the most fundamental level, *Nomina Anatomica* (International Anatomical Nomenclature Committee, 1989) and its successor, *Terminologia Anatomica* (Federative Committee on Anatomical Terminology, 1998) provide a classification of officially sanctioned terms that are associated with macroscopic and microscopic anatomical structures. This canonical term list, however, has been substantially expanded by synonyms that are current in various fields, and has also been augmented by a large number of new terms that designate structures omitted from *Terminologia Anatomica*. Many of these additions are present in various controlled terminologies (e.g., MeSH (National Library of Medicine, 1999), SNOMED (Spackman et al., 1997), Read Codes (Schultz et al., 1997), GALEN (Rector et al., 1993)). Unlike *Terminologia* these vocabularies are entirely computer-based, and therefore lend themselves for incorporation in computer-based applications.

The most complete primate *neuroanatomical* terminology is *NeuroNames*, developed by Bowden and Martin at the University of Washington (Bowden and Martin, 1995). *NeuroNames*, which is included as a knowledge source in the National Library of Medicine's Unified Medical Language System (UMLS; see Chapter 7) (Lindberg et al., 1993), is primarily organized as a part-of hierarchy of nested structures, with links to a large set of ancillary terms that do not fit into the strict part-of hierarchy. Other neuroanatomical terminologies have also been developed (Paxinos and Watson, 1986; Swanson, 1992; Bloom and Young, 1993; Franklin and Paxinos, 1997). A challenge for biomedical informatics is either to come up with a single consensus terminology or to develop Internet tools that allow transparent integration of distributed but commonly agreed-on terminology, with local modifications.

Classification and ontology projects to date have focused primarily on arranging the terms of a particular domain in hierarchies. As we noted with respect to the evaluation of *Terminologia Anatomica* (Rosse, 2000), insufficient attention has been paid to the relationships between these terms. *Terminologia*, as well as anatomy sections of the controlled medical terminologies, mix *-is a-* and *-part of-*relationships in the anatomy segments of their hierarchies. Although such heterogeneity does not interfere with using

these term lists for keyword-based retrieval, these programs will fail to support higher-level knowledge (reasoning) required for knowledge-based applications.

In our own Structural Informatics Group at the University of Washington we (JB and co-workers) are addressing this deficiency by developing a *Foundational Model of Anatomy* (FMA), (Figure 9.4) which we define as a comprehensive symbolic description of the structural organization of the body, including anatomical concepts, their preferred names and synonyms, definitions, attributes, and relationships (Rosse et al., 1998; Rosse and Mejino, 2003).

The FMA is being implemented in Protégé-2000, a frame-based knowledge acquisition system developed at Stanford (Musen, 1998; Mejino et al., 2001). In Protégé anatomical concepts are arranged in class–subclass hierarchies, with inheritance of defining attributes along the *is-a* link, and other relationships (e.g., parts, branches, spatial adjacencies) represented as additional slots in the frame. The FMA currently consists of over 70,000 concepts, represented by about 100,000 terms, and arranged in over 1.2 million links using 110 types of relationships. These concepts represent structures at all levels: macroscopic (to 1 mm resolution) cellular, and macromolecular. Brain

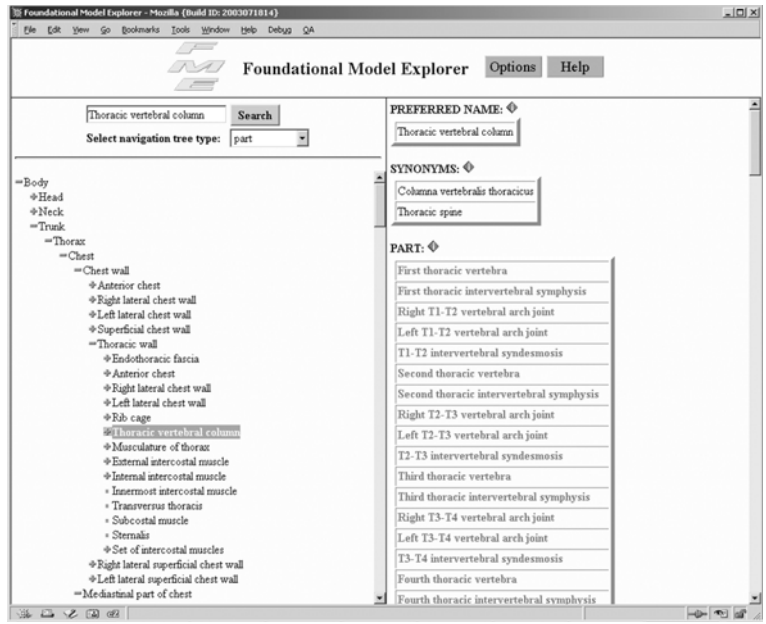

Figure 9.4. The Foundational Model Explorer, a Web viewer for the frame-based University of Washington Foundational Model of Anatomy (FMA). The left panel shows a hierarchical view along the part of link. Hierarchies along other links, such as is-a, branch-of, tributary-of, can also be viewed in this panel. The right hand panel shows the detailed local and inherited attributes (slots) associated with a selected structure, in this case the thoracic vertebral column. See also Figure 9.5. Photograph courtesy of the Structural Informatics Group, University of Washington.

structures have been added by integrating *NeuroNames* with the FMA as a *Foundational Model of Neuroanatomy* (FMNA) (Martin et al., 2001).

Our belief is that the FMA will prove useful for symbolically organizing and integrating biomedical information, particularly that obtained from images. But in order to answer nontrivial queries in neuroscience and other basic science areas, and to develop "smart tools" that rely on deep knowledge, additional ontologies must also be developed, among other things, for physiologic functions mediated by neurotransmitters, and pathologic processes and their clinical manifestations, as well as for the radiologic appearances with which they correlate. The relationships that exist between these concepts and anatomical parts of the body must also be explicitly modeled. Next-generation informatics efforts that link the FMA and other anatomical ontologies with separately developed functional ontologies will be needed in order to accomplish this type of integration.

9.5.4 Atlases

Spatial representations of anatomy, in the form of segmented regions on two-dimensional or three-dimensional images, or three-dimensional surfaces extracted from image volumes, are often combined with symbolic representations to form digital atlases. A digital atlas (which for this chapter refers to an atlas created from three-dimensional image data taken from real subjects, as opposed to artists' illustrations) is generally created from a single individual, which therefore serves as a "canonical" instance of the species. Traditionally, atlases have been primarily used for education, and most digital atlases are used the same way.

As an example in two dimensions, the Digital Anatomist Interactive Atlases (Sundsten et al., 2000) were created by outlining ROIs on two-dimensional images (many of which are snapshots of three-dimensional scenes generated by reconstruction from serial sections) and labeling the regions with terminology from the FMA. The atlases, which are available on the Web, permit interactive browsing where the names of structures are given in response to mouse clicks; dynamic creation of "pin diagrams", in which selected labels are attached to regions on the images; and dynamically generated quizzes, in which the user is asked to point to structures on the image (Brinkley et al., 1997).

As an example in three dimensions, the Digital Anatomist Dynamic Scene Generator (DSG, Figure 9.5; see also Color Plate II) creates interactive three-dimensional atlases "on-the-fly" for viewing and manipulation over the Web (Brinkley et al., 1999; Wong et al., 1999). In this case the three-dimensional scenes generated by reconstruction from serial sections are broken down into three-dimensional "primitive" meshes, each of which corresponds to an individual part in the FMA. In response to commands such as "display the branches of the coronary arteries" the DSG looks up the branches in the FMA, retrieves the three-dimensional model primitives associated with those branches, determines the color for each primitive based on its type in the FMA is-a hierarchy, renders the assembled scene as a two-dimensional snapshot, then sends it to a Web-browser, where the user may change the camera parameters, add new structures, or select and highlight structures. The complete scene may also be downloaded for viewing in a VRML browser.

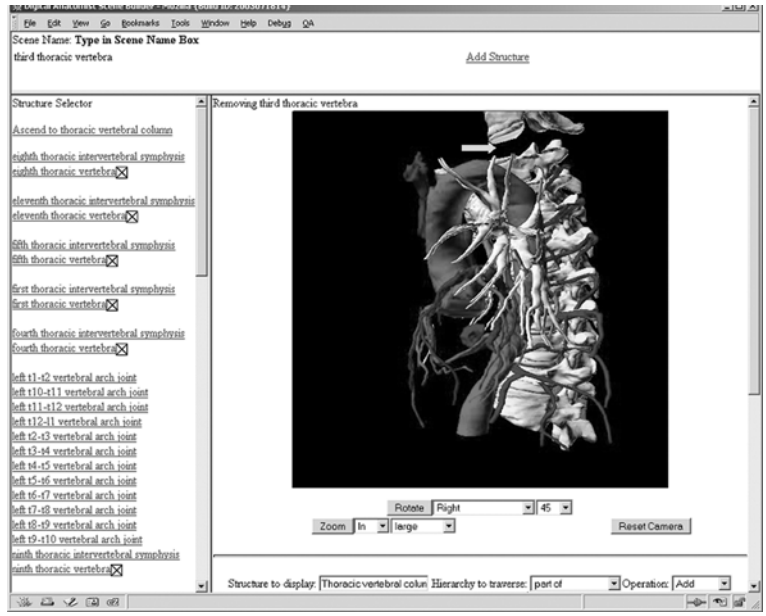

Figure 9.5. The Digital Anatomist Dynamic Scene Generator (see text). This-scene was created by requesting the following structures from the scene generator server: the parts of the aorta, the branches of the ascending aorta, the tributaries of the right atrium, the branches of the tracheobronchial tree, and the parts of the thoracic vertebral column. The server was then requested to rotate the camera 45 degrees, and to provide the name of a structure selected with the mouse, in this case the third thoracic vertebra. The selected structure was then hidden (note the gap indicated by the arrow). The left frame shows a partial view of the FMA part of hierarchy for the thoracic vertebral column. Checked structures are associated with three-dimensional "primitive" meshes that were loaded into the scene. Photograph courtesy of the Structural Informatics Group, University of Washington.

An example of a three-dimensional brain atlas created from the Visible Human is Voxelman (Hohne et al., 1995), in which each voxel in the Visible Human head is labeled with the name of an anatomic structure in a "generalized voxel model" (Hohne et al., 1990), and highly detailed three-dimensional scenes are dynamically generated. Several other brain atlases have also been developed primarily for educational use (Johnson and Becker, 2001; Stensaas and Millhouse, 2001).

In keeping with the theme of anatomy as an organizing framework, atlases have also been developed for integrating functional data from multiple studies (Bloom and Young, 1993; Toga et al., 1994, 1995; Swanson, 1999; Fougerousse et al., 2000; Rosen et al., 2000; Martin and Bowden, 2001). In their original published form these atlases permit manual drawing of functional data, such as neurotransmitter distributions, onto hard-copy printouts of brain sections. Many of these atlases have been or are in the process of being converted to digital form. The Laboratory of Neuroimaging (LONI)

at UCLA has been particularly active in the development and analysis of digital atlases (Toga, 2001b), and the Caltech HBP has released a Web-accessible three-dimensional mouse atlas acquired with micro-MR imaging (Dhenain et al., 2001).

The most widely used human brain atlas is the Talairach atlas, based on postmortem sections from a 60-year-old woman (Talairach and Tournoux, 1988). This atlas introduced a proportional coordinate system (often called "Talairach space") which consists of 12 rectangular regions of the target brain that are piecewise affine transformed to corresponding regions in the atlas. Using these transforms (or a simplified single affine transform based on the anterior and posterior commissures), a point in the target brain can be expressed in Talairach coordinates, and thereby related to similarly transformed points from other brains. Other human brain atlases have also been developed (Hohne et al., 1992; Caviness et al., 1996; Drury and Van Essen, 1997; Schaltenbrand and Warren, 1977; Van Essen and Drury, 1997).

9.5.5 Anatomic Variation

Brain information systems often use atlases as a basis for mapping functional data onto a common framework, much like geographic information systems (GISs) use the earth as the basis for combining data. However, unlike GISs, brain information systems must deal with the fact that no two brains are exactly alike, especially in the highly folded human cerebral cortex. Thus, not only do brain-imaging researchers have to develop methods for representing individual brain anatomy, they must also develop methods for relating the anatomy of multiple brains. Only by developing methods for relating multiple brains will it be possible to generate a common anatomic frame of reference for organizing neuroscience data. Solving this problem is currently a major focus of work in the HBP and in imaging informatics in general.

Two general approaches for quantitatively dealing with anatomic variation can be defined: (1) warping to a **template atlas**, and (2) **population-based atlases**. Variation can also be expressed in a qualitative manner, as described in the section on qualitative classification.

Warping to a Template Atlas

The most popular current quantitative method for dealing with anatomic variation is to deform or warp an individual target brain to a single brain chosen as a template. If the template brain has been segmented and labeled as an atlas (Section 9.5.4), and if the registration of the target brain to the template is exact, then the target brain will be automatically segmented, and any data from other studies that are associated with the template brain can be automatically registered with the target brain by inverting the warp (Christensen et al., 1996; Toga and Thompson, 2001). Such a procedure could be very useful for surgical planning, for example, since functional areas from patients whose demographics match that of the surgical patient could be superimposed on the patient's anatomy (Kikinis et al., 1996).

The problem, of course, comes with the word "exact". Since no two brains are even topologically alike (sulci and gyri are present in one brain that are not present in

another), it is impossible to completely register one brain to another. Thus, the research problem, which is very actively being pursued by many HBP researchers (Toga and Thompson, 2001), is how to register two brains as closely as possible. Methods for doing this can be divided into volume-based warping and surface-based warping.

Volume-based warping. Pure volume-based registration directly registers two image volumes, without the preprocessing segmentation step. Whereas intra (single)-patient registration (see Section 9.5.1) establishes a linear transformation between two datasets, inter (multiple)-patient registration establishes a nonlinear transformation (warp) that relates voxels in one volume to corresponding voxels in the other volume. Because of the great variability of the cerebral cortex pure volume-based registration is best suited for subcortical structures rather than the cortex. As in the linear case there are two basic approaches to nonlinear volume registration: *intensity-based* and *landmark-based*, both of which generally use either physically based approaches or minimization of a cost function to achieve the optimal warp.

The intensity-based approach uses characteristics of the voxels themselves, generally without the segmentation step, to nonlinearly align two image volumes (Gee et al., 1993; Collins et al., 1995; Christensen et al., 1996; Kjems et al., 1999). Most start by removing the skull, which often must be done manually.

The *landmark-based* approach is analogous to the two-dimensional case; the user manually indicates corresponding points in the two datasets (usually with the aid of three orthogonal views of the image volumes). The program then brings the corresponding points into registration while carrying along the intervening voxel data. The Montreal Register program (MacDonald, 1993) can do nonlinear three-dimensional warps, as can the 3-D Edgewarp program (Bookstein and Green, 2001), which is a generalization of the 2-D Edgewarp program developed by Bookstein (1989).

A variation of landmark-based warping matches curves or surfaces rather than points, then uses the surface warps as a basis for interpolating the warp for intervening voxels (Thompson and Toga, 1996; Davatzikos, 1997).

Surface-based warping. Surface-based registration is primarily used to register two cortical surfaces. The surface is first extracted using techniques described in Section 9.5.2, then image-based or other functional data are "painted" on the extracted surface where they are carried along with whatever deformation is applied to the surface. Since the cortical surface is the most variable part of the brain, yet the most interesting for many functional studies, considerable research is currently being done in the area of surface-based registration (Van Essen et al., 1998).

It is very difficult if not impossible to match two surfaces in their folded up state, or to visualize all their activity. (The cerebral cortex gray matter can be thought of as a two-dimensional sheet that is essentially crumpled up to fit inside the skull). Therefore, much effort has been devoted to "reconfiguring" (Van Essen et al., 2001) the cortex so that it is easier to visualize and register (Figure 9.6). A prerequisite for these techniques is that the segmented cortex must be topologically correct. The programs FreeSurfer (Dale et al., 1999), Surefit (Van Essen et al., 2001), ASP (MacDonald et al., 2000), and others all produce surfaces suitable for reconfiguration.

Common reconfiguration methods include *inflation*, *expansion to a sphere*, and *flattening*. *Inflation* uncrumples the detailed gyri and sulci of the folded surface by partially

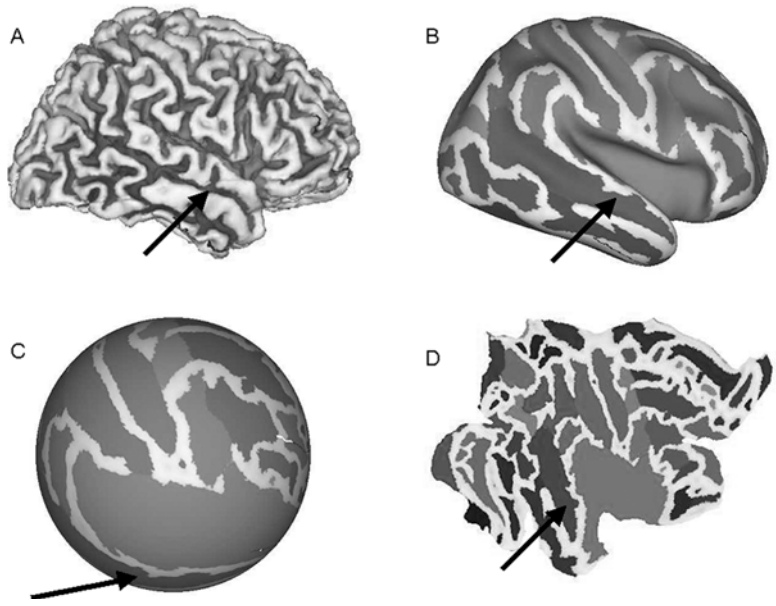

Figure 9.6. Brain surface reconfiguration using the Caret software suite developed by the David Van Essen laboratory at Washington University (Van Essen et al., 2001). The brain surface is first segmented (A), then inflated (B), expanded to a sphere (C), and flattened (D). At all stages any structural or functional data painted on the surface are carried along with the reconfiguration. In this case the brain sulci are painted onto the surface. The arrows point to the superior temporal sulcus in each configuration. Photograph courtesy of the SUMS database at Washington University (Van Essen, 2002), http://brainmap.wustl.edu:8081/sums/directory.do?dir_id=636032.

blowing the surface up like a balloon (Fischl et al., 1999; Brain Innovation B.V., 2001; Van Essen et al., 2001). The resulting surface looks like a lissencephalic (smooth) brain, in which only the major lobes are visible, and the original sulci are painted on the surface as darker intensity curves. These marks, along with any functional data, are carried along in the other reconfiguration methods as well.

Expansion to a sphere further expands the inflated brain to a sphere, again with painted lines representing the original gyri and sulci. At this point it is simple to define a surface-based coordinate system as a series of longitude–latitude lines referred to a common origin. This spherical coordinate system permits more precise quantitative comparison of different brains than three-dimensional Talairach coordinates because it respects the topology of the cortical surface. The surface is also in a form where essentially two-dimensional warping techniques can be applied to deform the gyri and sulci marked on the sphere to a template spherical brain.

The third approach is to *flatten* the surface by making artificial cuts on the inflated brain surface, then spreading out the cut surface on at two-dimensional plane while minimizing distortion (Fischl et al., 1999; Hurdal et al., 2000; Van Essen et al., 2001). Since it is impossible to eliminate distortion when projecting a sphere to a plane,

multiple methods of projection have been devised, just as there are multiple methods for projecting the earth's surface (Toga and Thompson, 2001). In all cases, the resulting flat map, like a two-dimensional atlas of the earth, is easier to visualize than a three-dimensional representation since the entire cortex is seen at once. Techniques for warping one cortex to another are applicable to flat maps as well as spherical maps, and the warps can be inverted to map pooled data on the individual extracted cortical surface.

The problem of warping any of these reconfigured surfaces to a template surface is still an active area of research because it is impossible to completely match two cortical surfaces. Thus, most approaches are hierarchical, in which larger sulci such as the lateral and central sulcus are matched first, followed by minor sulci.

Population-Based Atlases

The main problem with warping to a template atlas is deciding which atlas to use as a template. Which brain should be considered the "canonical" brain representing the population? As noted previously, the widely used Talairach atlas is based on a 60-year-old woman. The Visible Human male was a 38-year-old convict and the female was an older women. What about other populations such as different racial groups? These considerations have prompted several groups to work on methods for developing brain atlases that encode variation among a population, be it the entire population or selected subgroups. The International Consortium for Brain Mapping (ICBM), a collaboration among several brain mapping institutions spearheaded at UCLA (http://www.loni.ucla.edu/ICBM), is collecting large numbers of normal brain-image volumes from collaborators around the world (Mazziotta et al., 2001). To date several thousand brain-image volumes, many with DNA samples for later correlation of anatomy with genetics, are stored on a massive file server. As data collection continues methods are under development for combining these data into population-based atlases.

A good high-level description of these methods can be found in a review article by Toga and Thompson (2001). In that article three main methods are described for developing population-based atlases: *density-based*, *label-based*, and *deformation-based* approaches.

In the *density-based* method, a set of brains is first transformed to Talairach space by linear registration. Corresponding voxels are then averaged, yielding an "average" brain that preserves the major features of the brain but smoothes out the detailed sulci and gyri (Figure 9.7; see also Color Plate III). The Montreal average brain, which is an average of 305 normal brains (Evans et al., 1994), is constructed in this way. Although not detailed enough to permit precise comparisons of anatomic surfaces, it is nevertheless useful as a coarse means for relating multiple functional sites. For example, in our own work (JB) we have mapped cortical language sites from multiple patients onto the average brain, allowing a rough comparison of their distribution for different patient subclasses (Martin et al., 2000).

In the *label-based* approach, a series of brains are segmented, and then linearly transformed to Talairach space. A probability map is constructed for each segmented structure, such that at each voxel the probability can be found that a given structure is present at that voxel location. This method has been implemented in the Talairach Demon, an Internet server and Java client, developed by Fox et al. as part of the ICBM project

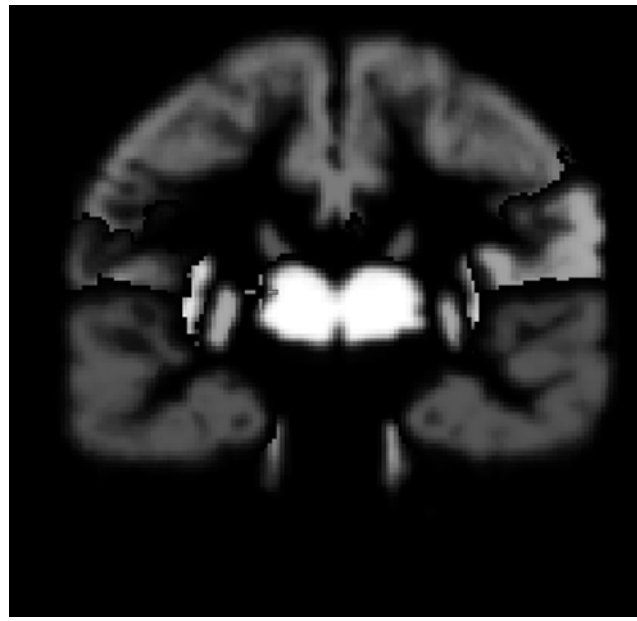

Figure 9.7. Probabilistic brain atlas, coronal section. Individual MRI image volumes from 53 subjects were linearly aligned, and each subject's lobes and deep nuclei were manually delineated. These delineations were averaged across the subjects and used to create probability maps for the likelihood of finding the specified lobe or nuclei at a given voxel position. Each structure is depicted in a different color in the color version of this image. The intensity of the color is proportional to the probability of finding that structure at the specified location. Photograph courtesy of Arthur Toga, Laboratory for Neuro Imaging, UCLA. http://www.loni.ucla.edu/NCRR/NCRR.Probabilistic.html.

(Lancaster et al., 2000). A Web user inputs one or more sets of Talairach coordinates, and the server returns a list of structure probabilities for those coordinates.

In the *warp-based* method, the statistical properties of deformation fields produced by nonlinear warping techniques (see Section 9.5.5) are analyzed to encode anatomic variation in population subgroups (Christensen et al., 1996; Thompson and Toga, 1997). These atlases can then be used to detect abnormal anatomy in various diseases.

9.5.6 Characterization of Anatomy

The main reason for finding ways to represent anatomy is to examine the relationship between structure and function in both health and disease. For example, how does the branching pattern of the dendritic tree influence the function of the dendrite? Does the pattern of cortical folds influence the distribution of language areas in the brain? Does the shape of the corpus callosum relate to a predisposition to schizophrenia?

Can subtle changes in brain structure be used as a predictor for the onset of Alzheimer's disease? These kinds of questions are becoming increasingly possible to answer with the availability of the methods described in the previous sections. However, in order to examine these questions, methods must be found for characterizing and classifying the extracted anatomy. Both qualitative and quantitative approaches are being developed.

Qualitative Classification

The classical approach to characterizing anatomy is for the human biologist to group individual structures into various classes based on perceived patterns. This approach is still widely used throughout science since the computer has yet to match the pattern-recognition abilities of the human brain.

An example classification at the cellular level is the 60 to 80 morphologic cell types that form the basis for understanding the neural circuitry of the retina (which is an outgrowth of the brain) (Dacey, 1999). At the macroscopic level, Ono has developed an atlas of cerebral sulci that can be used to characterize an individual brain based on sulcal patterns (Ono et al., 1990).

If these and other classifications are given systematic names and are added to the symbolic ontologies described in Section 9.5.3, they can be used for "intelligent" indexing and retrieval, after which quantitative methods can be used for more precise characterization of structure–function relationships.

Quantitative Classification

Quantitative characterization of anatomy is often called *morphometrics* (Bookstein, 1997) or *computational neuroanatomy* (Ascioli, 1999). Quantitative characterization permits more subtle classification schemes than are possible with qualitative methods, leading to new insights into the relation between structure and function, and between structure and disease (Toga, 2001a; Toga and Thompson, 2001).

For example, at the ultrastructural level, *stereology*, which is a statistical method for estimating from sampled data the distribution of structural components in a volume (Weibel, 1979), is used to estimate the density of objects such as synapses in image volumes reconstructed from serial electron micrographs (Fiala and Harris, 2001).

At the cellular level Ascoli et al. are engaged in the L-neuron project, which attempts to model dendritic morphology by a small set of parameterized generation rules, where the parameters are sampled from distributions determined from experimental data (Figure 9.8) (Ascioli, 1999). The resulting dendritic models capture a large set of dendritic morphologic classes from only a small set of variables. Eventually, the hope is to generate virtual neural circuits that can simulate brain function.

At the macroscopic level landmark-based methods have shown changes in the shape of the corpus callosum associated with schizophrenia that are not obvious from visual inspection (DeQuardo et al., 1999). Probabilistic atlas-based methods are being used to characterize growth patterns and disease-specific structural abnormalities in diseases such as Alzheimer's and schizophrenia (Thompson et al., 2001). As these techniques

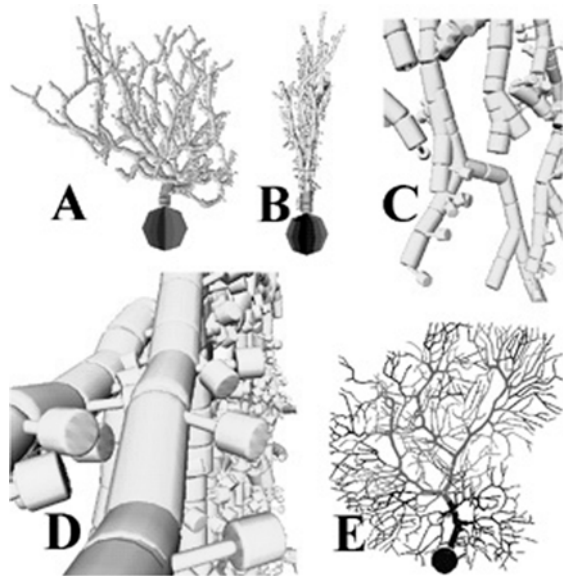

Figure 9.8. The L-neuron project. The branching pattern of a given population of neurons is modeled by a small set of parameters measured from experimental data. These parameters are used to generate synthetic neurons (A to D) that look very similar to experimentally reconstructed neurons (E). The virtual cell displayed in panels A to D was created with only 11 lines of stochastic rules to resemble the real 2107-compartment Purkinje cell shown in panel E. (A) Front view; (B) side view; (C) detail on dendritic branches; (D) detail on dendritic spines. Photograph courtesy of Georgio Ascioli, George Mason University, http://www.krasnow.gmu.edu/ascoli/CNG/index.htm. (Reprinted from Ascoli, GA. Progress and Perspectives in Computational Neuroanatomy. Anat Rec. 257(6): 195-207. Copyright © 1999, Wiley. Reprinted by Permission of Wiley-Liss Inc., A Subsidiary of John Wiley & Sons, Inc.)

become more widely available to the clinician, they should permit earlier diagnosis and hence potential treatment for these debilitating diseases.

9.6 Functional Imaging

Many imaging techniques not only show the structure of the body but also the function. For imaging purposes, function can be inferred by observing changes of structure over time. In recent years this ability to image function has greatly accelerated. For example, ultrasound and angiography are widely used to show the functioning of the heart by depicting wall motion, and ultrasound doppler can image both normal and disturbed blood flow (Mehta et al., 2000). **Molecular imaging** (Section 9.3.5) is increasingly able to depict the expression of particular genes superimposed on structural images, and thus can also be seen as a form of functional imaging.

A particularly profound application of functional imaging is the understanding of cognitive activity in the brain. It is now routinely possible to put a normal subject in a scanner, to give the person a cognitive task, such as counting or object recognition, and

to observe which parts of the brain light up. This unprecedented ability to observe the functioning of the living brain opens up entirely new avenues for exploring how the brain works.

Functional brain-imaging modalities can be classified as *image-based* or *nonimage-based*. In both cases it is taken as axiomatic that the functional data must be mapped to the individual subject's anatomy, where the anatomy is extracted from structural images using techniques described in the previous sections. Once mapped to anatomy, the functional data can be integrated with other functional data from the same subject, and with functional data from other subjects whose anatomy has been related to a template or probabilistic atlas. Techniques for generating, mapping, and integrating functional data are part of the field of Functional Brain Mapping, which has become very active in the last few years, with several conferences (Organization for Human Brain Mapping, 2001) and journals (Fox, 2001; Toga et al., 2001) devoted to the subject.

9.6.1 Image-Based Functional Brain Mapping

Image-based functional data generally come from scanners that generate relatively low-resolution volume arrays depicting spatially localized activation. For example, PET (Heiss and Phelps, 1983; Aine, 1995) and magnetic resonance spectroscopy (MRS) (Ross and Bluml, 2001) reveal the uptake of various metabolic products by the functioning brain; and **functional magnetic resonance imaging (fMRI)** reveals changes in blood oxygenation that occur following neural activity (Aine, 1995). The raw intensity values generated by these techniques must be processed by sophisticated statistical algorithms to sort out how much of the observed intensity is due to cognitive activity and how much is due to background noise.

As an example, one approach to fMRI imaging is the boxcar paradigm applied to language mapping (Corina et al., 2000). The subject is placed in the MRI scanner and told to silently name objects shown at 3-second intervals on a head-mounted display. The actual objects ("on" state) are alternated with nonsense objects ("off" state), and the fMRI signal is measured during both the on and the off states. Essentially the voxel values at the off (or control) state are subtracted from those at the on state. The difference values are tested for significant difference from nonactivated areas, then expressed as *t*-values. The voxel array of *t*-values can be displayed as an image.

A large number of alternative methods have been and are being developed for acquiring and analyzing functional data (Frackowiak et al., 1997). The output of most of these techniques is a low-resolution three-dimensional image volume in which each voxel value is a measure of the amount of activation for a given task. The low-resolution volume is then mapped to anatomy by linear registration to a high-resolution structural MR dataset, using one of the linear registration techniques described in Section 9.5.1.

Many of these and other techniques are implemented in the SPM program (Friston et al., 1995), the AFNI program (Cox, 1996), the Lyngby toolkit (Hansen et al., 1999), and several commercial programs such as Medex (Sensor Systems Inc., 2001), and BrainVoyager (Brain Innovation B.V., 2001). The FisWidgets project at the University of Pittsburgh is developing a set of Java wrappers for many of these programs that allow customized creation of graphical user interfaces in an integrated desktop environment

(Cohen, 2001). A similar effort (VoxBox) is underway at the University of Pennsylvania (Kimborg and Aguirre, 2002).

9.6.2 Nonimage-Based Functional Mapping

In addition to the image-based functional methods there are an increasing number of techniques that do not directly generate images. The data from these techniques are generally mapped to anatomy, and then displayed as functional overlays on anatomic images.

For example, cortical stimulation mapping (CSM) is a technique for localizing functional areas on the exposed cortex at the time of neurosurgery (Figure 9.9; see also Color Plate IV). In our own work (JB) the technique is used to localize cortical language areas so that they can be avoided during the resection of a tumor or epileptic focus (Ojemann et al., 1989). Following removal of a portion of the skull (craniotomy) the patient is awakened and asked to name common images shown on slides. During this

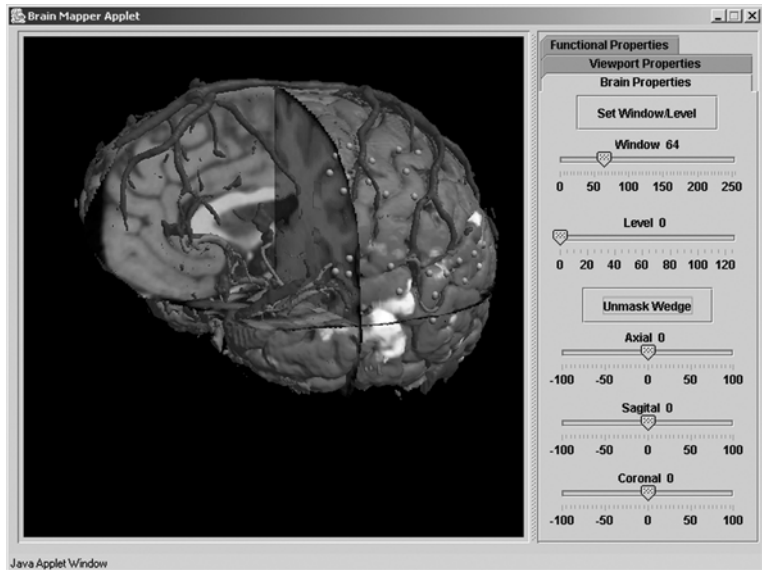

Figure 9.9. Remote visualization of integrated structural and functional brain data mapped onto a single patient's brain. MRI, MRV (veins), and MRA (arteries) brain-image volumes are acquired and registered, then segmented to generate the cortical surface, arteries, and brains. fMRI data representing areas of language processing are registered to the structural volumes, then projected to the surface as the light-colored regions. Cortical stimulation mapping (CSM) data (small spheres) acquired during neurosurgery are also registered to the patient's anatomy. The integrated data are rendered on a visualization server, which can be accessed from a web browser using a simple Java applet. Photograph courtesy of the Structural Informatics Group, University of Washington.

time the surgeon applies a small electrical current to each of a set of numbered tags placed on the cortical surface. If the patient is unable to name the object while the current is applied the site is interpreted as essential for language and is avoided at surgery. In this case, the functional-mapping problem is how to relate these stimulation sites to the patient's anatomy as seen on an MRI scan.

Our approach, which we call visualization-based mapping (Modayur et al., 1997; Hinshaw et al., 2002), is to acquire image volumes of brain anatomy (MRI), cerebral veins (MRV), and cerebral arteries (MRA) prior to surgery, to segment the anatomy, veins, and arteries from these images, and to generate a surface-rendered three-dimensional model of the brain and its vessels that matches as closely as possible the cortical surface as seen at neurosurgery. A visual-mapping program then permits the user to drag numbered tags onto the rendered surface such that they match those seen on the intraoperative photograph. The program projects the dragged tags onto the reconstructed surface, and records the x-y-z image-space coordinates of the projections, thereby completing the mapping.

The real goal of functional neuroimaging is to observe the actual electrical activity of the neurons as they perform various cognitive tasks. fMRI, MRS, and PET do not directly record electrical activity. Rather, they record the results of electrical activity, such as (in the case of fMRI) the oxygenation of blood supplying the active neurons. Thus, there is a delay from the time of activity to the measured response. In other words these techniques have relatively poor temporal resolution (Section 9.2.2). **Electroencephalography (EEG)** or **magnetoencephalography (MEG)**, on the other hand, are more direct measures of electrical activity since they measure the electromagnetic fields generated by the electrical activity of the neurons. Current EEG and MEG methods involve the use of large arrays of scalp sensors, the output of which are processed in a similar way to CT in order to localize the source of the electrical activity inside the brain. In general this "source-localization problem" is under-constrained, so information about brain anatomy obtained from MRI is used to provide further constraints (George et al., 1995).

9.7 Conclusions

This chapter focuses on methods for processing images in biomedicine, with an emphasis on brain imaging and the extraction and characterization of anatomic structure, both for its own sake, and as a substrate on which to map function. Other than the interest and expertise of the authors, an important reason to concentrate on brain imaging is that a large part of the most advanced image-processing work is currently in this area. As these techniques develop, and as new imaging modalities increasingly become available for imaging other and more detailed body regions, the techniques will increasingly be applied in all areas of biomedicine. For example, the development of molecular-imaging methods is analogous to functional-brain imaging, in that functional data, in this case from gene expression rather than cognitive activity, are mapped to an anatomic substrate. Since the same basic principles apply in both functional-brain mapping and molecular imaging, the same techniques apply: spatial and symbolic representation of

anatomy, dealing with anatomic variation, characterization of anatomy, visualization, and multimodality image fusion.

Thus, these general methods will increasingly be applied to diverse areas of biomedicine. As they are applied, and as imaging modalities continue to proliferate, an increasing demand will be placed on methods for managing the images and for storing and accessing them. At the same time, imaging will continue to play an increasing role in integrated biomedical information systems. These two subjects, management of images, and integration in biomedical applications, are the subjects of Chapter 18.

Questions for Discussion

1. What is the general principle that underlies computed axial tomography (CT)? What are the advantages of CT images over conventional X-ray images?
2. Explain the general principle underlying magnetic resonance imaging (MRI)? What are the advantages of this method compared to older methods of imaging?
3. Explain the differences among contrast, spatial, and temporal resolution.
4. Describe the four standard image-processing steps, and suggest how these might be applied by an image-analysis program looking for abnormal cells in a PAP smear.
5. What is the segmentation step in image analysis? Why is it so difficult to perform? Give two examples of ways by which current systems avoid the problem of automatic segmentation. Give an example of how knowledge about the problem to be solved (e.g., local anatomy) could be used in future systems to aid in automatic segmentation.
6. What additional informatics issues arise when going from two-dimensional to three-dimensional image processing? What are the three-dimensional versions of two-dimensional image-processing operations such as region growing and edge finding?
7. What is a three-dimensional brain atlas? What are the methods for registering a patient image volume to that atlas? What is the use of a brain atlas?
8. Give some example techniques for imaging the function of the brain.

Suggested Readings

Brinkley J.F. (1991). Structural informatics and its applications in medicine and biology. *Academic Medicine*, 66(10):589–591.
Short introduction to the field.

Brinkley J.F., Rosse C. (2002). Imaging and the Human Brain Project: A review. *Methods of Information in Medicine*, 41:245–260.
Review of image processing work related to the brain. Much of the brain-related material for this chapter was taken from this article.

Potchen E.J. (2000). Prospects for progress in diagnostic imaging. *Journal of Internal Medicine*, 247(4):411–424.
Nontechnical description of newer imaging methods such as cardiac MRI, diffusion tensor imaging, fMRI, and molecular imaging. Current and potential use of these methods for diagnosis.

Robb R.A. (2000). *Biomedical Imaging, Visualization, and Analysis*. New York: Wiley-Liss.
Overview of biomedical imaging modalities and processing techniques.

Rosse C., Mejino J.L.V. (2003). A reference ontology for bioinformatics: The Foundational Model of Anatomy. *Journal of Bioinformatics*, 36(6):478–500.
Description and principles behind a large symbolic ontology of anatomy.

Shapiro L.G., Stockman G.C. (2001). *Computer Vision*. Upper Saddle River, NJ: Prentice-Hall.
Detailed description of many of the representations and methods used in image processing. Not specific to medicine, but most of the methods are applicable to medical imaging.